



Technical FAQ

NFS Technical FAQ

Updated for ONTAP 9.1

Justin Parisi, NetApp
October 2016

Abstract

This document is the Technical FAQ for Network File Systems (NFS). The document covers frequently asked questions about the NFS protocol.

Document Classification

Internal Only

Version History

Version	Date	Document Version History
Version 1.0	August 2013	First publishing
Version 2.0	September 2014	Updated to include ONTAP 8.3 information
Version 3.0	July 2015	Added 8.3.1 information
Version 3.1	July 2016	Added ONTAP 9.0 information
Version 3.2	October 2016	Added ONTAP 9.1 information

TABLE OF CONTENTS

Version History	2
1 Purpose	5
1.1 Frequently Asked Questions	5
What is NFS?	5
What is an NFS export?	5
What is an NFS mount?	5
What is pNFS?	5
What is an ACL?	5
What is an ACE?	5
What is authentication?	5
What is authorization?	6
What are <code>chown</code> and <code>chgrp</code> ?	6
What is <code>chmod</code> ?	6
What is a delegation?	6
What is an export policy?	6
What is an export policy rule?	6
What is a file handle?	6
What is a file lock?	6
What is a junction path?	6
What is a lease?	7
What are mode bits?	7
What is name mapping?	7
Are group mappings supported in ONTAP?	7
What is the NAS module (<code>nblade</code>)?	7
What is the network lock manager?	7
What is the network status monitor?	8
What is a referral?	8
What is the replicated database (RDB)?	8

What is the root user?	8
What is the security daemon (SecD)?	8
What is security style?	8
What are a UID and a GID?	8
What is the user space?	8
1.2 NFS Options Explained	9
1.3 Export Policy Rule Options Explained	17
1.4 How It Works	18
NFS in ONTAP	19
1.5 Support Considerations	24
Which versions of NFS are supported by ONTAP?	24
Does NetApp support NFSv4.2?	24
Are TCP and UDP NFS supported?	24
Which NFS ports are used by ONTAP?	25
Which Kerberos encryption types are supported?	25
Which Kerberos security modes are supported for NFS mounts in ONTAP?	25
Is krb5p process offloading supported for AES-NI capable processors?	25
How many ACEs are supported?	25
Are qtree exports currently supported?	25
Can qtrees or folders inside volumes be used as mount points?	26
What is the maximum number of nested netgroups allowed?	26
How many export policies and export policy rules are supported?	26
Why is there a maximum number of export policies and rules?	26
How many AUTH_SYS and AUTH_GSS groups are supported in ONTAP?	27
Is <code>exportfs</code> currently supported?	28
Is <code>showmount</code> supported?	29
Is Dynamic DNS supported?	29
Which bit levels of file handles are supported?	29
What is the difference between a session and a connection?	29
How many simultaneous TCP connections are supported?	30
Is there an <code>nfsstat -d</code> command to detect misalignment on VMware datastores?	30
Is the <code>.snapshot</code> directory visible in NFSv4?	30
What are the best practices for VMware vSphere?	30
Is Windows NFS supported?	30
Is Oracle dNFS supported?	30
Which LDAP schemas are supported?	30
Is the Data ONTAP operating in 7-Mode option <code>nfs.hide_snapshot</code> supported?	30
Are the following tunable variables supported?	30
Can data volumes be accessed through NAS without mounting them into a namespace?	30
Can volumes be mounted to paths that do not exist?	31
Why don't Data ONTAP operating in 7-Mode NFS options (such as in node shell) work in ONTAP?	31

Are LDAPS and LDAP signing and sealing supported in ONTAP?	31
Is cluster administration authentication through LDAP supported?	31
Is LDAP the same as Active Directory?	31
What is the Kerberos token size limit for NFS in ONTAP?	31
2 Decision Points	32
2.1 Why Use NFS?	32
Why should I use NFS?	32
Why should I use NFSv3?	32
Why should I use NFSv4.x?	32
Why should I use pNFS?	33
2.2 Considerations for Making the Transition from NFS in Data ONTAP Operating in 7-Mode	34
2.3 Common Uses of NFS	34
3 Common Issues	35
3.1 Cannot Mount	35
3.2 Permission Denied or Access Issues	37
3.3 Files Written as “Nobody”	38
3.4 Stale File Handle or NFS Mount	39
4 References	40

LIST OF TABLES

Table 1) NFS options in ONTAP.	9
Table 2) Export policy rule options.	17
Table 3) Supported NFS versions in ONTAP.	24
Table 4) Common mount failures.	35
Table 5) Common access issues.	37
Table 6) Files written as “nobody” in NFSv4.....	38

LIST OF FIGURES

Figure 1) Namespaces.	20
----------------------------	----

1 Purpose

The purpose of this technical FAQ is to answer questions surrounding the use and implementation of NFS in the NetApp® ONTAP® operating system.

1.1 Frequently Asked Questions

What is NFS?

Answer: The Network File System (NFS) allows a client computer to access files on shared storage over the network. Clients are generally Linux or UNIX, but Microsoft Windows also supports NFS by using service for UNIX packages. NFS builds on the [Open Network Computing Remote Procedure Call](#) system. Its technical and implementation details are defined in [RFC 1094](#). (RFC, or request for comments, is a memorandum published by the [Internet Engineering Task Force](#).)

What is an NFS export?

Answer: This export is a file system that is made available through NFS by an NFS server.

What is an NFS mount?

Answer: This mount is a connection through TCP or UDP to an NFS export on an NFS server. NFS mounts can have a variety of options and flags when attaching to an NFS server. For details, see the [mount man pages](#).

What is pNFS?

Answer: pNFS is parallel NFS. The pNFS protocol gives clients direct access to files across two or more data servers. By accessing multiple data servers in parallel, clients achieve significant I/O acceleration. The pNFS protocol has been designed to deliver graceful performance scaling on both a per-client and a per-file basis without sacrificing backward compatibility with the standard NFS protocol. Clients without the pNFS extension are still able to access data. For more information about pNFS, see [TR-4063: Parallel Network File System Configuration and Best Practices for ONTAP 8.2 and Later](#). File-based pNFS is currently the only supported mechanism in ONTAP.

What is an ACL?

Answer: An ACL is an access control list. An ACL is used to provide a list of users and their access level to a file system. ACLs can be Windows style or NFSv4 style. Data ONTAP does not support POSIX ACLs.

What is an ACE?

Answer: An ACE is an access control entry, or an entry in an access control list.

What is authentication?

Answer: Authentication is the process of identifying an individual, usually based on a user name and a password. In security systems, authentication is distinct from authorization. Authentication states that the individual is who he or she claims to be, but says nothing about the access rights of the individual.

What is authorization?

Answer: Authorization takes place after authentication. After it is determined that a user is valid, the user's credentials, permissions, group memberships, and so on are gathered to determine the level of access that is allowed on a volume or file.

What are `chown` and `chgrp`?

Answer: These commands are UNIX shell commands used to change the owning user and group of a file or directory.

What is `chmod`?

Answer: This command is a UNIX shell command to change the permissions and special flags on a file or directory.

What is a delegation?

Answer: A delegation is an NFSv4.x feature that can be used to improve read or write performance by caching metadata and data operations on the NFS client. For more information about NFSv4.x delegations, see [TR-4067: ONTAP NFS Implementation Guide](#).

What is an export policy?

Answer: An export policy is a container for export policy rules. Policies get applied to individual volumes. Policies are assigned at volume creation. When no policy is specified, the policy assigned is the default policy.

What is an export policy rule?

Answer: An export policy rule is a rule that specifies the access that is allowed to an NFS export.

Multiple rules can exist in a single export policy. The number of export policy rules is limited to the [number allowed per cluster](#). For more information about export policy rules, see [TR-4067: NFS Implementation and Best Practice Guide](#). The default export policy does not contain a rule by default. Verify that the default policy has a rule defined.

What is a file handle?

Answer: A file handle is a unique numerical identifier used by NFS to identify exported files. The file handle is constructed by the NFS server. In ONTAP, the file handle is created by using the volume identifiers and inode numbers within the cluster subsystem. If a file handle changes while an NFS mount is active, the file handle becomes stale and the file system requires a remount.

What is a file lock?

Answer: A file lock is a mechanism used to verify that a file that is open can be modified only by the owner of the file lock.

What is a junction path?

Answer: A junction path is the logical path to a volume in ONTAP. A junction path can have the same name as the volume that is using it, or it can use a different name altogether. A junction path can be mounted like a regular exported file system. Junction paths can be defined upon volume creation or by using the `volume mount` command; they can be destroyed by using the `volume unmount` command or through NetApp OnCommand® System Manager. The junction path, along with the export policy rule, creates an analog to an export of Data ONTAP operating in 7-Mode.

For more information and examples, see [TR-4067: NFS Implementation and Best Practice Guide](#).

What is a lease?

Answer: A lease is a feature in NFSv4.x that is the time negotiated to honor a lock. Leases guarantee how long a lock will stay on a server. During the lease, a client must renew the lease (similar to keepalive packets) to maintain its hold on a lease. If the client does not renew the lease, the server has the option to revoke the lock. For more information about leases, see [TR-4067: NFS Implementation and Best Practice Guide](#).

What are mode bits?

Answer: Mode bits enable UNIX systems to define permissions for files and directories. These apply to volumes and qtrees as well. Permission mode bits are set in four different octets:

- The first octet sets the special permission bits (such as the sticky bit).
- The second octet sets permissions for the file owner.
- The third octet sets permissions for the group owner.
- The fourth octet sets permissions for everyone else.

Permission bits are used by Data ONTAP for volumes and qtrees. They can be modified through the command line by using the `volume modify` or `qtree modify` commands with the `-unix-permissions` flag or through OnCommand System Manager.

Note: Mode bits cannot be modified on NTFS security-style volumes.

For more information about Linux permission bits, see the [File Permissions article on Linux.org](#).

What is name mapping?

Answer: This is the process by which one style of user account is mapped to another. Valid name mappings include:

- UNIX user to Windows user
- Windows user to UNIX user
- Kerberos Service Principal Name (SPN) to UNIX user

Name mapping rules are set in ONTAP through the `vserver name-mapping` command set.

For more information about name mapping, see the [Multiprotocol Technical FAQ](#) and [TR-4067: NFS Implementation and Best Practice Guide](#).

Are group mappings supported in ONTAP?

Answer: Group mappings are not supported in ONTAP.

What is the NAS module (nblade)?

Answer: The NAS module (nblade) is a per-node entity in ONTAP that translates CIFS and NFS requests into packets the cluster can understand, as well as back into CIFS and NFS packets when the request has been processed by the cluster. All NAS traffic passes through the NAS module local to the data LIF on which the request arrived. The NAS module also provides caching mechanisms similar to SecD to enable efficiency in NAS requests.

What is the network lock manager?

Answer: To help manage file access conflicts and protect NFS connections during failures, NFS offers a file- and record-locking service called the *network lock manager* (NLM). NLM is a separate service (running on port 4045) that NFS makes available to user applications. To use the locking service, applications must make calls to standard lock routines. NLM is used only with NFSv3.

What is the network status monitor?

Answer: This monitor is a service specific to NFSv3 used to notify NFS clients and servers of reboots for lock management purposes.

What is a referral?

Answer: A referral is an NFSv4.x feature that can be used to provide direct traffic for NFS clients by referring them to nodes in a cluster that own the file systems being accessed. For more information about NFSv4.x referrals, see [TR-4067: NFS Implementation and Best Practice Guide](#).

What is the replicated database (RDB)?

Answer: In ONTAP, this is what makes a cluster a cluster. The replicated database is a series of tables that are replicated across all nodes in a cluster through the cluster network. These tables remain in sync by way of a set of four user-space RDB applications (`mgwd`, `vifmgr`, `vldb`, and `bcomd`). In that way, each node in the cluster is aware of all entries in the cluster while providing a single entry point for management.

What is the root user?

Answer: Also known as the “superuser,” the root user provides administrator access to a UNIX operating system. The root user is typically assigned a UID of 0.

What is the security daemon (SecD)?

Answer: This daemon is the user-space application responsible for authentication and authorization requests in ONTAP. SecD makes connections to domain controllers and name service servers, sends and receives authentication and authorization requests, and translates name mappings. SecD also works over remote procedure calls (RPCs) to commit configuration changes to the replicated database (RDB) and caches credentials so that there is efficient access to NAS protocols. SecD is a per-node process in which requests use the SecD application on the node that owns the data LIF on which the request arrived.

What is security style?

Answer: Security style refers to the type of permissions used on a volume in Data ONTAP. If a volume has UNIX security style, then mode bits or NFSv4-style ACLs are used. If a volume has NTFS security style, then only Windows-style ACLs are used. If a volume has mixed security style, then the ACL used depends on which type of client last modified the ACL. For further details on security styles, see the [Multiprotocol Technical FAQ](#) and [TR-4067: NFS Implementation and Best Practice Guide](#).

Clustered Data ONTAP 8.2 also provides the Unified security style for NetApp Infinite Volume. For more information about Unified security, see [TR-4178: Infinite Volume Deployment and Implementation Guide](#).

What are a UID and a GID?

Answer: These acronyms stand for user ID and group ID, respectively. These acronyms are numeric identifiers for a user in a UNIX operating system. The UID and the GID are used to determine access to files based on mode-bit settings or ACLs.

What is the user space?

Answer: The user space is a concept in ONTAP through which dedicated cluster applications run to provide end-user services, such as access, volume location, and command-line availability.

1.2 NFS Options Explained

The following table covers the various options used for NFS servers, the version of ONTAP in which they are available, the privilege level, and their use. You can view all NFS server options by using the `nfs server show` command or through OnCommand System Manager.

For further details on many of these options, see [TR-4067: NFS Implementation and Best Practice Guide](#).

Also see the CLI output from `man nfs modify` to get manual page definitions.

Table 1) NFS options in ONTAP.

NFS Option	Version	Privilege Level	Definition
NFS Access (-access)	All	Admin	Access allowed or not.
NFSv3 (-v3)	All	Admin	Enable/disable NFSv3.
NFSv4 (-v4.0)	8.1 and later	Admin	Enable/disable NFSv4.
UDP (-udp)	All	Admin	Enable/disable UDP.
TCP (-tcp)	All	Admin	Enable/disable TCP.
Spin Authentication (-spinauth)	All	Admin	Enable/disable spinauth; deprecated.
Default Windows User (-default-win-user)	All	Admin	Specify the default Windows user for multiprotocol access.
NFSv4.0 ACL Support (-v4.0-acl)	8.1 and later	Admin	Enable/disable NFSv4 ACL support.
NFSv4.0 Read Delegation (-v4.0-read-delegation) NFSv4.0 Write Delegation (-v4.0-write-delegation)	8.1 and later	Admin	Enable/disable NFSv4 read and write delegations.
NFSv4 ID Mapping Domain (-v4-id-domain)	8.1 and later	Admin	Specify the NFSv4 ID domain.
NFSv4.1 Minor Version Support (-v4.1)	8.1 and later	Admin	Enable/disable NFSv4.1 minor version support.
Rquota Enable (-rquota)	8.1 and later	Admin	Enable/disable rquota support.

pNFS Support (-v4.1-pnfs)	8.1 and later	Admin	Enable/disable pNFS support.
NFSv4.1 ACL Support (-v4.1-acl)	8.1 and later	Admin	Enable/disable NFSv4.1 ACL support.
NFS vStorage Support (-vstorage)	8.1 and later	Admin	Enable/disable NFS vStorage support.
Default Windows Group (-default-win-group)	8.2 and later	Admin	Set the default Windows group, Infinite Volume only. See TR-4037: "Introduction to NetApp Infinite Volume" for details about Infinite Volume.
NFSv4.1 Read Delegation (-v4.1-read-delegation) NFSv4.1 Write Delegation (-v4.1-write-delegation)	8.2 and later	Admin	Enable/disable NFSv4.1 read and write delegations.
NFS Mount Root Only (-mount-rootonly)	8.2 and later	Admin	Allow mounts only from the root user.
NFS Root Only (-nfs-rootonly)	8.2 and later	Admin	Allow NFS only from the root user.
RPC GSS Context Cache (-rpcsec-ctx-high) RPC GSS Context Idle (-rpcsec-ctx-idle)	All	Advanced	Specifies the max number of RPCSEC_GSS contexts and the idle timeout of the cache. See RFC 2203 for information about RPCSEC_GSS contexts.
NFSv2 (-v2)	8.0 and 8.1.x	Advanced	Enable/disable NFSv2 (removed from ONTAP 8.2).
NFSv3 EJUKEBOX Error (-enable-ejukebox)	All	Advanced	Enables/disables NFSv3 EJUKEBOX errors. Jukebox errors are "NFS not responding" errors. This option is enabled by default.
Include EJUKEBOX Replies in the NFS Replay Cache (-cache-ejukebox)	8.0 only	Advanced	Specifies whether EJUKEBOX errors are cached for NFSv3.
Require All NFSv3 Reads to Return Read Attributes (-v3-require-read-attributes)	All	Advanced	Specifies whether NFSv3 read operations are required to return read attributes.

Show Change in FSID as NFSv3 Clients Traverse File Systems (-v3-fsid-change)	All	Advanced	Specifies whether Data ONTAP shows changes in file system identifiers (FSIDs) as NFSv3 clients traverse file systems. If you change the value of this parameter, clients must remount any paths over which they use NFSv3.
Enable the Dropping of a Connection When an NFSv3 Request Is Dropped (-v3-connection-drop)	All	Advanced	Option to enable/disable connection drops if an NFSv3 connection is dropped; useful for older clients that might not handle connection drops properly; enabled by default.
NFS Response Trace Enabled (-trace-enabled) NFS Response Trigger (-trigger)	All	Advanced	Specifies whether Data ONTAP logs NFS requests when they exceed the NFS response trigger time, which is defined by the <code>-trigger</code> parameter.
UDP Maximum Transfer Size (-udp-max-xfer-size) TCP Maximum Transfer Size (-tcp-max-xfer-size)	All	Advanced	Specifies the maximum transfer size (in bytes) that the storage system negotiates with the client for TCP and UDP connections. Range for UDP is 8,192 to 57,344. Range for TCP is 8,192 to 65,536.
NFSv3 TCP Maximum Read Size (-v3-tcp-max-read-size) NFSv3 TCP Maximum Write Size (-v3-tcp-max-write-size)	8.1 and later	Advanced	Specifies the maximum transfer size (in bytes) that the storage system negotiates with the client for TCP transport of data for NFSv3 read and write requests. The range is 8,192 to 1,048,576 for reads and 8,192 to 65,536 for writes. The default setting is 65,536 when created. Number specified will depend on the application vendor's recommendation.
Show Change in FSID as NFSv4 Clients Traverse File Systems (-v4-fsid-change)	8.1 and later	Advanced	Specifies whether Data ONTAP shows changes in file system identifiers (FSIDs) as NFSv4 clients traverse file systems. If you change the value of this parameter, clients must remount any paths over which they use NFSv4.
NFSv4.0 Referral Support (-v4.0-referrals)	8.1 and later	Advanced	Specifies whether Data ONTAP supports NFSv4.0 referrals. The default setting is disabled when created. You can set this parameter to Enabled only if the <code>-v4-fsid-change</code> option is also set to Enabled. If clients accessing the node do not support NFSv4.0 referrals, set this option to Disabled; otherwise, those clients

			will not be able to access the file system. This parameter is not supported for SVMs with Infinite Volume.
NFSv4 Validate UTF-8 Encoding of Symbolic Link Data (-v4-validate-symlinkdata)	8.1 and later	Advanced	Specifies whether Data ONTAP validates the UTF-8 encoding of symbolic link data. The default setting is Disabled when created. This is useful for international implementations of NFSv4.
NFSv4 Lease Timeout Value (-v4-lease-seconds) NFSv4 Grace Timeout Value (-v4-grace-seconds)	8.1 and later	Advanced	Specifies the locking lease and grace reclaim timeouts.
Preserves and Modifies NFSv4 ACL (-v4-acl-preserve)	8.1 and later	Advanced	Enables/disables preservation of the NFSv4 ACL in the event a <code>chmod</code> is performed.
NFSv4.1 Implementation ID Domain (-v4.1-implementation-domain) NFSv4.1 Implementation ID Name (-v4.1-implementation-name) NFSv4.1 Implementation ID Date (-v4.1-implementation-date)	8.1 and later	Advanced	Specifies the NFSv4.1 implementation information.
NFSv4.1 Referral Support (-v4.1-referrals)	8.1 and later	Advanced	Specifies whether Data ONTAP supports NFSv4.1 referrals. The default setting is Disabled when created. You can set this parameter to Enabled only if the <code>-v4-fsid-change</code> option is also set to Enabled. If clients accessing the node do not support NFSv4.1 referrals, set this option to Disabled; otherwise, those clients will not be able to access the file system. This parameter is not supported for SVMs with Infinite Volume.
Number of Slots in the NFSv4.x Session Slot Tables (-v4.x-session-num-slots) Size of the Reply That Will Be Cached in Each NFSv4.x Session Slot (-v4.x-session-slot-reply-cache-	8.2 and later	Advanced	Specifies the number of session slot tables and the size of those slots. Adjusting these values depends on the application and OS using them and can affect performance (positively or negatively). For details on slot tables, see RFC 5661 .

size)			
Maximum Number of ACEs per ACL (-v4-acl-max-aces)	8.2 and later	Advanced	Specifies the maximum number of allowed ACEs per ACL. Default is 400. The range is 192 to 1,024. Setting the value higher than the default can affect performance, so set only when necessary.
Validation of Qtree IDs for Qtree File Operations (-validate-qtree-export)	8.2.1 and later	Advanced	This optional parameter specifies whether ONTAP performs an additional validation on qtree IDs. The default setting is Enabled. This parameter is only effective when a qtree is assigned an export policy.
Showmount Enabled (-showmount)	8.3 and later	Advanced	This optional parameter specifies whether to allow or disallow clients to see all of the Vserver's NFS exports list using the <code>showmount -e</code> command. The default setting is Disabled.
AUTH_SYS and RPCSEC_GSS Auxiliary Groups Limit (-extended-groups-limit)	8.3 and later	Advanced	This optional parameter specifies the maximum number of auxiliary groups supported over RPC security options AUTH_SYS and RPCSEC_GSS in Data ONTAP. The range is 32 to 1,024. The default value is 32.
AUTH_SYS Extended Groups Enabled (-auth-sys-extended-groups)	8.3 and later	Advanced	This optional parameter specifies whether Data ONTAP supports fetching auxiliary groups from a name service rather than from the RPC header. The default setting is Disabled.
Set the Protocol Used for Name Services Lookups for Exports (-name-service-lookup-protocol)	8.3 and later	Advanced	This optional parameter specifies the protocol to use for doing name service lookups. The allowed values are TCP and UDP. The default setting is UDP.
Permitted Kerberos Encryption Types (-permitted-enc-types)	8.3 and later	Advanced	This optional parameter specifies the permitted encryption types for Kerberos over NFS. The default setting is des,des3,aes-128,aes-256.
NFS Mount Daemon Port (-mountd-port)	8.3 and later	Advanced	This optional parameter specifies which port the NFS mount daemon (mountd) uses. The default setting is 635.
Network Lock Manager Port	8.3 and later	Advanced	This optional parameter specifies which port the network lock manager (NLM) uses.

(-nlm-port)			The default setting is 4045.
Network Status Monitor Port (-nsm-port)	8.3 and later	Advanced	This optional parameter specifies which port the network status monitor (NSM) uses. The default setting is 4046.
NFS Quota Daemon Port (-rquotad-port)	8.3 and later	Advanced	This optional parameter specifies which port the NFS quota daemon (rquotad) uses. The default setting is 4049.
Set the Protocol Used for Name Services Lookups for Exports (-name-service-lookup-protocol)	8.3.1 and later	Advanced	This optional parameter specifies the protocol to use for doing name service lookups. The allowed values are TCP and UDP. The default setting is UDP.
NFSv3 MS-DOS Client Support (-v3-ms-dos-client)	8.2.3 and later; 8.3.1 and later	Admin	This optional parameter specifies whether to enable access for NFSv3 MS-DOS clients. The default setting is disabled at the time of creation. This parameter is not supported for Vservers with Infinite Volume.
Time to Live Value (in msec) of a Positive Cached Credential (-cached-cred-positive-ttl)	8.3.1 and later	Advanced	This optional parameter specifies the age of the positive cached credentials after which they will be cleared from the cache. The value specified must be from 60,000 through 604,800,000. The default setting is 86,400,000.
Time to Live Value (in msec) of a Negative Cached Credential (-cached-cred-negative-ttl)	8.3.1 and later	Advanced	This optional parameter specifies the age of the negative cached credentials after which they will be cleared from the cache. The value specified must be from 60,000 through 604,800,000. The default setting is 7,200,000.
Skip Permission Check for NFS Write Calls from Root/Owner (-skip-root-owner-write-perm-check)	8.3.1 and later	Advanced	This optional parameter specifies if permission checks are to be skipped for NFS WRITE calls from root/owner. For copying read-only files to a destination folder that has inheritable ACLs, this option must be enabled. Warning: When enabled, if an NFS client does not make use of an NFS ACCESS call to check for user-level permissions and then tries to write onto read-only files, the operation will succeed. The default setting is Disabled.
Display Maximum NT ACL	8.3.1 and	Advanced	This optional parameter controls the

Permissions to NFS Client (-ntacl-display-permissive-perms)	later		permissions that are displayed to NFSv3 and NFSv4 clients on a file or directory that has an NT ACL set. When True, the displayed permissions are based on the maximum access granted by the NT ACL to any user. When False, the displayed permissions are based on the minimum access granted by the NT ACL to any user. The default setting is False.
Trust No-Match Result from Any Name Service Switch Source During Netgroup Lookup (-netgroup-trust-any-ns-switch-no-match)	8.3.1 and later	Advanced	This optional parameter specifies if you can consider a no-match result from any of the netgroup ns-switch sources to be authoritative. If this option is enabled, then a no-match response from any of the netgroup ns-switch sources is deemed conclusive even if other sources could not be searched. The default setting is Disabled, which causes all netgroup ns-switch sources to be consulted before a no-match result is deemed conclusive.
DNS Domain Search Enabled During Netgroup Lookup (-netgroup-dns-domain-search)	8.3.1 and later	Advanced	If you enable this optional parameter, during client access check evaluation in a netgroup, Data ONTAP performs an additional verification so that the domain returned from DNS for that client is listed in the DNS configuration of the Vserver. Doing so enables you to validate the domain when clients have the same short name in multiple domains. The default setting is Enabled.
Map Unknown UID to Default Windows User (-map-unknown-uid-to-default-windows-user)	8.3.1 and later	Advanced	If you enable this optional parameter, unknown UNIX users that do not have a name mapping to a Windows user are mapped to the configured default Windows user. Doing so allows all unknown UNIX users to have access with the credentials of the default Windows user. If you disable this parameter, all unknown UNIX users without name mapping are always denied access. By default, this parameter is enabled.
Ignore the NT ACL Check for NFS User 'root' (-ignore-nt-acl-for-root)	8.3.1 and later	Advanced	This optional parameter specifies whether Windows ACLs affect root access from NFS. If this option is enabled, root access from NFS ignores the NT ACL set on the file or directory. If auditing is enabled for the Vserver and there is no name-mapping

			present, then a default SMB credential (Builtin\administrator) is used for auditing and an EMS warning is generated. The default setting is Disabled, which causes NFS “root” to be mapped to a Windows account, like any other NFS user.
Use 64 Bits for NFSv3 FSIDs and File IDs (-v3-64bit-identifiers)	ONTAP 9.0 and later	Advanced	This optional parameter specifies whether Data ONTAP uses 64 bits (instead of 32 bits) for file system identifiers (FSIDs) and file identifiers (file IDs) that are returned to NFSv3 clients. If you change the value of this parameter, clients must remount any paths over which they use NFSv3. When -v3-fsid-change is disabled, enable this parameter to avoid file ID collisions.
Ignore Client Specified Mode Bits and Preserve Inherited NFSv4 ACL When Creating New Files or Directories (-v4-inherited-acl-preserve)	ONTAP 9.0 and later	Advanced	This optional parameter specifies whether the client-specified mode bits should be ignored and the inherited NFSv4 ACL should be preserved when creating new files or directories. The default setting is Disabled.
Lookup for the filename in unconverted language if converted language lookup fails (-v3-search-unconverted-filename)	ONTAP 9.1 and later	Advanced	This optional parameter specifies whether to continue the search with unconverted name while doing lookup in a directory.

1.3 Export Policy Rule Options Explained

The following table lists the various options for export policy rules and what they are used for. All export policy rule options can be viewed by using the `export-policy rule show` command or through OnCommand System Manager.

For further details on many of these options, see [TR-4067: NFS Implementation and Best Practice Guide](#).

Also see the CLI output from `man export-policy rule modify` to get manual page definitions.

Table 2) Export policy rule options.

Export Policy Rule Option	Privilege Level	What It Does
Policy Name (-policyname)	Admin	Shows/sets policy name.
Rule Index (-ruleindex)	Admin	Sets the rule index for the export policy rule. This is the order in which the rules are applied.
Access Protocol (-protocol)	Admin	Shows/sets the protocols allowed by an export policy rule.
Client Match (-clientmatch)	Admin	Shows/sets the host name, IP, netgroups, subnet, or domain allowed to access through the export policy rule.
RO Access/RW Access Rule (-rorule) (-rwrule)	Admin	Shows/sets the values for which AUTH flavor (such as krb, sys, and so on) is allowed RW and RO access. None will cause users to come in as anon. Never denies access.
User ID to Which Anonymous Users are Mapped (-anon)	Admin	Shows/sets the user ID to which anonymous users are mapped.
Superuser (-superuser)	Admin	Determines the AUTH flavor that allows root/superuser access. None squashes root to anon.
Honor SetUID Bits in SETATTR (-allow-suid)	Admin	Specifies whether set user ID (suid) and set group ID (sgid) access is enabled by the export rule.

Allow Creation of Devices (-allow-dev)	Admin	Specifies whether the creation of devices is enabled by the export rule.
NTFS Unix Security Options (-ntfs-unix-security-ops)	Advanced	Specifies whether UNIX-type permission changes on NTFS (Windows) volumes are prohibited (fail) or allowed (ignored) when the request originates from an NFS client.
Vserver NTFS Unix Security Options (-ntfs-unix-security-ops-vs)	Advanced	Shows the SVM-wide setting for this option; can be modified only at the NFS server level.
Change Ownership Mode (-chown-mode)	Advanced	Controls whether or not users other than the superuser can <code>chown</code> .
Vserver Change Ownership Mode (-chown-mode-vs)	Advanced	Shows the SVM-wide setting for this option; can be modified only at the NFS server level. Must be set to Unrestricted to allow the export policy rule to apply in NFSv4.

1.4 How It Works

The following section covers how NFS works in ONTAP, from mount to unmount, in both NFSv3 and NFSv4.x. For details on how specific NFS features work, see [TR-4067: NFS Implementation and Best Practice Guide](#).

For more information about NFS, see the following RFCs:

- [RFC 1094: NFS RFC](#)
- [RFC 1813: NFS Version 3](#)
- [RFC 3530: NFS Version 4](#)
- [RFC 7530: NFS Version 4 \(updated\)](#)
- [RFC 5661: NFS Version 4.1](#)
- [RFC 5663: pNFS](#)

NFS in ONTAP

NFS Access

Access to NFS is determined by the following:

- Valid NFS license
- NFS server existence
- NFS versions enabled
- NFS allowed on the storage virtual machine (SVM)
- NFS allowed on the data LIF
- Export policies and rules
- Valid junction path
- Access to the parent volumes in the path

Checking the NFS license

```
license show -feature NFS
```

Verifying that the NFS server exists

```
nfs server show -vserver [vservname]
```

Verifying that the NFS version is enabled

```
nfs server show -vserver [vservname] -fields v3,v4.0,v4.1,v4.1-pnfs*
```

*v2 is only for ONTAP 8.0.x and 8.1.x.

Verifying that NFS is allowed on the SVM

```
vserver show -vserver [vservername] -fields allowed-protocols
```

If NFS is not allowed, modify the SVM to allow it.

Verifying that NFS is allowed on the data LIF

```
network interface show -vserver nfs -fields data-protocol
```

If NFS is not allowed on the data LIF, the LIF must be destroyed and recreated or a new LIF must be created. The ports allowed on an NFS data LIF depend on the allowed data protocol. For instance, if NFS is not enabled on a data LIF, then ports 111, 2049, and so on will not be open and in LISTENING status.

Junction paths and export policy rules

In ONTAP, the mount point is based on the volume's junction path. Permission to mount and access the file system is controlled through export policy rules as well as the permission settings on the volume. To check these, use the following commands in the CLI:


```
volume show -vserver [vservername] -fields policy,junction-path,unix-  
permissions,user,group,security-style  
export-policy rule show -policy [policyname] -instance
```

You can also use System Manager to verify NFS settings.

Parent volumes

NFS access can be affected by restrictive permissions on parent volumes in a junction path. In ONTAP, the concept of a single namespace allows volumes to be mounted to other volumes in the namespace.

Figure 1) Namespaces.



In Figure 1, multiple levels of junctions are being used. The top level is the root junction, which the SVM root volume is mounted to. This is always /.

Example:

```
cluster::> vol show -vserver nfs -volume rootvol -fields junction-path
(volume show)
vserver volume junction-path
-----
nfs      rootvol /
```

The other levels can be a combination of namespace and/or data volumes. These junctions contain each parent volume's junction path. All junctions contain / as their top-level junction, meaning that every data volume in a namespace is subjected to the SVM root volume's export policy and UNIX permissions.

As a result, you should verify that the SVM root volume allows access to the client attempting to mount.

NFSv3 in ONTAP

When a client mounts an NFS file system hosted on a ONTAP storage system, the mount version used depends on the client OS and the version of NFS that the server supports. In most cases, this default can be adjusted through a configuration file, such as `/etc/nfsmount.conf` in RHEL, CentOS, and Fedora.

The mount version can also be controlled by using mount options. By specifying `-o vers=3` or `-o nfsvers=3`, the client can be forced to mount NFS version 3.

Examples of mounting a ONTAP file system through NFSv3 are:

```
[nfsclient] # mount -o nfsvers=3 cluster:/junction /mountpoint_on_client
[nfsclient] # mount -o vers=3 cluster:/junction /mountpoint_on_client
```

When a mount command is issued to a ONTAP NFS server from an NFS client, the following occurs:

1. RPC is made to port 111 (portmapper) of the NFS server to attempt a TCP connection through the portmapper.
2. When the RPC call has been acknowledged, portmapper issues a GETPORT call to port 111 of the NFS server data LIF to obtain which port NFS is allowed on.
3. The NFS server returns the port 2049 (nfs) to the client.
4. The client then closes the connection to port 111.
5. A new RPC call is made to port 2049 of the NFS server data LIF.
6. The NFS server returns the call successfully and the client sends an NFS NULL call to port 2049 of the NFS server's data LIF. This action checks whether the parent volume allows access to the mount. In this case, the parent volume is mounted to /, or the SVM root.
7. The NFS NULL call is returned successfully and the client proceeds with the mount attempt.
8. Portmapper sends a GETPORT call to the NFS server's data LIF asking for the `mountd` port and provides the credentials, NFS version number, and whether the mount will use TCP or UDP.
9. The cluster checks the NFS settings and verifies whether the credentials supplied are allowed to mount based on the export policy rules. This operation is carried out through an RPC call from the NAS blade to the SecD process. If SecD is not functioning properly, this check fails and the mount gets `access denied`. If the NFS version or TCP/UDP is not allowed, the client reports the error.
10. The NFS server replies successfully if the version provided is supported and if the mount can use the specified TCP or UDP connection. It also replies if the AUTH security provider is supported (AUTH_SYS or AUTH_GSS, for example).
11. When the GETPORT call passes, the client issues a V3 MNT call to the junction path specified in the `mount` command through port 635 (`mountd`) of the NFS server data LIF.
12. The cluster uses the junction path provided by the client and searches for the path in the Volume Location Database (vldb). If the entry exists, the cluster gathers the file handle information from the vldb.

13. The NFS server returns the file handle to the client, as well as replies which AUTH types are supported by the export policy rule. If the AUTH type provided by the server matches what the client sent, the mount succeeds.
14. Portmapper from the client then sends another GETPORT call for NFS, this time providing the client's host name.
15. The NFS server replies with port 2049 and the call succeeds.
16. Another NFS NULL call is made from the client over port 2049 of the NFS data LIF and is acknowledged by the NFS server.
17. A series of NFS packets with FSINFO and PATHCONF information is traded between the client and the NFS server.

When a read or write is attempted to a ONTAP NFS server from an NFS client through NFSv3, the following occurs:

1. The client attempts to traverse the volume and the cluster resolves the UID/GID so that the user has access. If the volume is NTFS security style, the cluster attempts to map the UNIX user to a valid Windows user. See the [Multiprotocol FAQ](#) for details on this process.
2. All traffic travels over the NFS port specified by the NFS server. In the case of ONTAP, the port is 2049.
3. Reads and writes use a series of V3 calls, including:
 - ACCESS
 - LOOKUP
 - CREATE
 - SETATTR
 - GETATTR
 - REaddir/REaddirplus
- Reads and writes of files in NFS depend on file handles to identify which files are being accessed.
- Reads and writes of files depend on lock assignment to prevent file conflicts.

NFSv4 in ONTAP

When a client mounts an NFS hosted on a ONTAP storage system, the mount version used depends on the client OS and the version of NFS that the server supports. In most cases, this default can be adjusted through a configuration file, such as `/etc/nfsmount.conf` in RHEL, CentOS, and Fedora.

The mount version can also be controlled through mount options. By specifying the `-t nfs, -o vers=4` or `-o nfsvers=4` option, the client can be forced to mount by using NFSv4.

The following are examples of mounting a ONTAP file system by using NFSv4:

```
[nfsclient] # mount -o nfsvers=4 cluster:/junction /mountpoint_on_client
[nfsclient] # mount -t nfs cluster:/junction /mountpoint_on_client
[nfsclient] # mount -o vers=4 cluster:/junction /mountpoint_on_client
```

When a mount command is issued to a ONTAP NFS server from an NFS client, the following occurs:

1. The NFS client makes a TCP connection attempt to the NFS server data LIF through port 2049 (nfs). NFSv4 uses only TCP and only port 2049 for all calls.
2. The NFS server either allows or rejects the TCP connection, based on the NFS server settings.
3. When the TCP connection has been established, a V4 NULL call is made by the NFS client to the NFS server data LIF. This call is a check to see whether v4 is allowed on the NFS server through the AUTH type provided by the client (AUTH_SYS or AUTH_GSS, for example).

4. The cluster checks the NFS settings and verifies whether the credentials supplied are allowed to mount based on the export policy rules. This is done through an RPC call from the NAS blade to the SecD process. If SecD is not functioning properly, this check fails and the mount gets `access denied`. If the NFS version or TCP is not allowed, the client reports the error.
5. The NFS server replies successfully if the version provided is supported and if the mount can use the specified TCP connection, and also if the AUTH security provider is supported.
6. The NFS client then makes a compound v4 call for PUTROOTFH, GETFH, and GETATTR for the specified mount point's parent volume. This call supplies the UID and GID of the user attempting access.
7. The cluster checks the UID and GID through SecD to verify that they exist in the provided NFSv4 domain and does a translation so that the user has access to mount the parent volume. In this case, the parent volume is `/`, or the SVM root.
8. The NFS server replies with the file handle of the SVM root volume and the attribute information of the parent volume.
9. The NFS server and client exchange a series of GETATTR calls to check access to the SVM root volume.
10. If the access is allowed to the user attempting access, then the client issues a LOOKUP DH call to the NFS server's data LIF for the SVM root volume's file handle/junction path; for example, `0x871aba87/vol11`.
11. The cluster uses the junction path provided by the client and searches for the path in the Volume Location Database (vldb). If the entry exists, the cluster gathers the file handle information from the vldb.
12. The NFS server returns the file handle to the client, as well as the attributes of that volume.
13. The NFS client and server exchange a series of GETATTR calls on the new file handle so that access is allowed for the user attempting mount access.
14. If access is allowed, the mount is successful.

When a read or write is attempted to a ONTAP NFS server from an NFS client through NFSv4, the following occurs:

1. All traffic travels over the NFS port specified by the NFS server. In the case of ONTAP, the port is 2049.
2. The client issues an ACCESS call to the NFS server by using the file handle retrieved during the mount process to determine whether the user attempting access has permission to do what he or she is trying to do (read or write).
3. The cluster checks the volume permissions and export policy by using the SecD process so that access is allowed. The user credentials are also checked for access to the volume. If the volume is NTFS security style, the cluster attempts to map the UNIX user to a valid Windows user. See the [Multiprotocol FAQ](#) for details on this process.
4. The NFS server returns the results of the ACCESS call to the client.
5. The NFS client sends a SETCLIENTID call to the NFS server to help manage the state associated with this client on the server.
6. The NFS server assigns a unique ID to the client and returns the ID to the NFS client. This client ID is used in stateful NFSv4 operations, such as locking and leasing mechanisms.
7. The client then sends a SETCLIENTID_CONFIRM to the NFS server so that the client receives the correct client ID.
8. The server responds with the client ID confirmation.
9. When a write occurs, the client sends an OPEN call to the NFS server with the file handle of the volume/file name; for example, `0x821795d6/filename`.

10. The NFS server replies to the client with a stateID. This is the lock/lease given to the file while it is being modified.
11. The cluster reserves a connection for this stateID in the NAS blade so that the cluster can keep track of the file while it is being modified.
12. After saving the file, the client sends WRITE and SETATTR calls to the NFS server.
13. The NFS server replies with an acknowledgment of the WRITE and SETATTR calls.
14. When the save is complete, the client sends a CLOSE call with the stateID information.
15. The cluster releases the stateID and reserved connection in the NAS blade and replies to the client with a CLOSE.

1.5 Support Considerations

Which versions of NFS are supported by ONTAP?

Answer: The following table shows which versions of NFS are supported by which versions of ONTAP.

Table 3) Supported NFS versions in ONTAP.

NFS Version	ONTAP 8.0.x	ONTAP 8.1.x	ONTAP 8.2 and later
v2	Yes	Yes	No
v3	Yes	Yes	Yes
v4	No	Yes	Yes
v4.1	No	Yes*	Yes

*** 8.1.1 and later**

Does NetApp support NFSv4.2?

Answer: NFSv4.2 is still not an official NFS protocol standard. NetApp will implement this version in a future release.

Are TCP and UDP NFS supported?

Answer: Both TCP and UDP are supported for NFS access. By default, UDP is not enabled. To enable UDP, use the `nfs modify` command or OnCommand System Manager.

Which NFS ports are used by ONTAP?

Answer: The following list shows the ports used for NFS by ONTAP. Some of these ports can be modified in ONTAP 8.3 and later:

program	vers	proto	port	service
100000	2	udp	111	portmapper
100000	2	tcp	111	portmapper
100000	3	udp	111	portmapper
100000	3	tcp	111	portmapper
100000	4	udp	111	portmapper
100000	4	tcp	111	portmapper
100003	3	udp	2049	nfs
100003	3	tcp	2049	nfs
100003	4	tcp	2049	nfs
400010	1	tcp	2049	
100005	1	udp	635	mountd
100005	2	udp	635	mountd
100005	3	udp	635	mountd
100005	1	tcp	635	mountd
100005	2	tcp	635	mountd
100005	3	tcp	635	mountd
100021	4	udp	4045	nlockmgr
100021	4	tcp	4045	nlockmgr
100024	1	udp	4046	status
100024	1	tcp	4046	status
100011	1	udp	4049	rquotad

Which Kerberos encryption types are supported?

Answer: ONTAP versions earlier than 8.3 support only DES and 3DES Kerberos. ONTAP 8.3 and later support AES-128 and AES-256.

Which Kerberos security modes are supported for NFS mounts in ONTAP?

Answer: The following chart shows the supported Kerberos security modes in ONTAP.

ONTAP Version(s)	krb5	krb5i	krb5p
8.0.x–8.2.x	Yes	No	No
8.3–8.3.x	Yes	Yes	No
9.0	Yes	Yes	Yes

Is krb5p process offloading supported for AES-NI capable processors?

Answer: This is currently not supported for krb5p, but is supported for SMB signing and sealing.

How many ACEs are supported?

Answer: Four hundred are supported in 8.0.x and 8.1.x; 1,024 are supported in 8.2. The option is configurable in 8.2 and later only through the `-v4-acl-max-aces` option.

Are qtree exports currently supported?

Answer: Qtree exports were not supported in 8.2 and earlier versions. These exports are supported in 8.2.1 for NFSv3 only. In 8.3, support for NFSv4.x has been added to qtree exports. Qtree exports in 8.2.1 and later operate exactly like volume exports: A policy must contain a rule and be applied to the qtree.

Qtrees include an option to inherit/not inherit the parent export policy. This option must be adjusted if different export policies are used for qtrees than are used by the parent volume.

Can qtrees or folders inside volumes be used as mount points?

Answer: Yes. As long as the folder or qtree exists, it can be used as a mount point. But only volumes and qtrees can have export policy rules applied to them.

What is the maximum number of nested netgroups allowed?

Answer: There is no limit defined for nested netgroups, but NetApp recommends attempting to keep nested netgroup numbers as low as possible and confirming that all host names that are in netgroups can be resolved through DNS. For more information on name service limits and best practices, see [TR-4379: Name Services Best Practice Guide](#).

How many export policies and export policy rules are supported?

Answer: Policy and rule limits depend on platform and cluster size and are documented in the ONTAP Maximum Configurations guides, which can be [found on the Field Portal by searching for “maximums.”](#)

Why is there a maximum number of export policies and rules?

Answer: Using an inordinate number of policies can add memory pressure to the replicated database, so the number of policies should be limited to only those in use. It is not considered a best practice to create policies without rules.

Additionally, export policies and rules should be designed in a way that maximizes the overall improvements in ONTAP. In Data ONTAP operating in 7-Mode, exports had limitations, such as:

- Lived in a flat file
- Not persistent across reboots unless the exports file was modified
- Potential outages on failovers because of nonpersistent exports
- One line per volume, potentially making large and unwieldy files

In ONTAP, multiple rules can be applied per policy and multiple volumes can use the same policy. That means potentially that only one policy needs to be managed for an entire cluster.

In Data ONTAP operating in 7-Mode, if a storage system had 300 volumes and 600 qtrees, the exports file likely had more than 900 lines. If you are thinking 1:1 in ONTAP, you could assume that means 900 policies, each with one rule. However, that's unnecessary.

ONTAP provides a vast improvement in export management over Data ONTAP operating in 7-Mode because Data ONTAP is no longer limited to a single exports file. There is now an entire database at the system's disposal.

In many cases, exports in Data ONTAP operating in 7-Mode had the same settings for security. Let's assume that of the aforementioned 900 exports, 500 of those all used the default export entries of:

```
-sec=sys,rw,nosuid
```

In ONTAP, rather than repeating that same line over and over, as seen in the Data ONTAP operating in 7-Mode exports file, a **single** export policy with a **single** rule can be created and applied to 500 volumes. So, 500 lines become this:

```
cluster::> export-policy rule show -vserver nfs -policyname 7mode -instance
(vserver export-policy rule show)

Vserver: nfs
Policy Name: 7mode
Rule Index: 1
Access Protocol: nfs
Client Match Hostname, IP Address, Netgroup, or Domain: 0.0.0.0/0
RO Access Rule: sys
RW Access Rule: sys
User ID To Which Anonymous Users Are Mapped: 65534
Superuser Security Types: none
Honor SetUID Bits in SETATTR: false
Allow Creation of Devices: true
```

That leaves 400 exports. Those also are likely not unique, but, even if they are, a **single** policy with 400 rules can be created. Then the `clientmatch` field can be used to tell the rule to apply to the desired client or subnet. The limit for export policy rules is limited per cluster and is based on the amount of memory in a system.

How many AUTH_SYS and AUTH_GSS groups are supported in ONTAP?

Answer: There are 16 AUTH_SYS groups and 32 AUTH_GSS groups in ONTAP 8.2.x and earlier. ONTAP 8.3 and later offer support for 1,024 groups (SYS and GSS) by way of the equivalent to the Data ONTAP operating in 7-Mode option `nfs.authsys.extended_groups_ns.enable` called `auth-sys-extended-groups`. This option is covered in detail in [TR-4067: NFS Best Practice and Implementation Guide](#).

Is exportfs currently supported?

Answer: exportfs is not supported in ONTAP. However, in ONTAP 8.3 and later, flushing the export rule cache is possible with the following command, found at the advanced level (exportfs -f equivalent):

NAME	<code>vserver export-policy cache flush --</code> Flush the export caches
AVAILABILITY	This command is available to cluster administrators at the admin privilege level.
DESCRIPTION	The vserver export-policy cache flush command clears out the contents of the export policy caches for a Vserver. You might need to flush the caches to allow the changes to immediately take effect for your NFS clients because of: - o A change to your export policy rules. - o Modifying a host name record in a name server (i.e., local hosts or DNS). - o Modifying the entries in a netgroup in a name server (i.e., local netgroup, LDAP, or NIS). - o Recovering from a network outage that resulted in a netgroup being partially expanded. To flush the caches, you must specify the following items: - o Vserver: either a specific Vserver or use "*" to flush all of them. You can optionally specify the following items: - o Node: if flushing the access cache, you can also specify which node to flush it on. - o Cache to flush: by default all but showmount will be flushed. Note that the showmount cache is not used to determine NFS client access and as such is only flushable explicitly.
PARAMETERS	<code>-vserver <vserver name></code> - Vserver This parameter specifies the name of the Vserver on which you want to flush the caches. <code>[-node <nodename>]</code> - Node This parameter specifies the node on which you want to flush the access cache. <code>[-cache {all access host id name netgroup showmount}]</code> - Cache Name This parameter specifies the name of the cache which you want to flush. Possible values include the following: - o all - All caches but showmount. This is the default. - o access - The export-policy rules access cache. - o host - The host name to IP cache. - o id - The ID to credential cache. - o name - The name to ID cache. - o netgroup - The netgroup cache. - o showmount - The showmount caches.

Additionally, checking export policy access is supported in ONTAP 8.2.3 and 8.3 with the `export-policy check-access` command (previously known as `exportfs -c` in 7-Mode).

```
vserver export-policy check-... Data ONTAP 8.3 vserver export-policy check-...

NAME
vserver export-policy check-access -- Given a Volume And/or a Qtree, Check to See If the Client
Is Allowed Access

AVAILABILITY
This command is available to cluster and Vserver administrators at the admin privilege level.

DESCRIPTION
The vserver export-policy check-access command checks whether a specific client is allowed access
to a specific export path. This enables you to test export policies to ensure they work as
intended and to troubleshoot client access issues.

The command takes the volume name (and optionally the qtree name) as input and computes the
export path for the volume/qtree. It evaluates the export policy rules that apply for each path
component and displays the policy name, policy owner, policy rule index and access rights for
that path component. If no export policy rule matches the specified client IP address access is
denied and the policy rule index will be set to 0. The output gives a clear view on how the
export policy rules are evaluated and helps narrow down the policy and (where applicable) the
specific rule in the policy that grants or denies access. This command is not supported on
Infinite Volumes.
```

Is `showmount` supported?

Answer: ONTAP versions earlier than 8.3 do not currently support `showmount`. In those versions, the result of the command is `/` (everything). NetApp added `showmount` support in ONTAP 8.3 and later, but it is limited in functionality as compared to 7-Mode. For details, see [TR-4067](#). For `showmount` functionality in versions earlier than 8.3, [use the `showmount` plug-in](#).

Is Dynamic DNS supported?

Answer: Yes. Support for Dynamic DNS (DDNS) was added to ONTAP 8.3.1.

Which bit levels of file handles are supported?

Answer: ONTAP supports both 32-bit and 64-bit file handles. Some applications require that NFS servers offer support for legacy 32-bit file ID lengths. RFC-1813 requires that NFSv3 return 8 bytes for the file ID (aka inode number) because it's defined in the specification as `uint64` (XDR unsigned hyper). All RFC-1813-compliant NFSv3 vendors will return 8 bytes.

In ONTAP 8.3.x and earlier, the operating system never returns anything in the upper 4 bytes of the NFSv3 file ID, so legacy 32-bit applications can operate normally. NetApp WAFL[®] (Write Anywhere File Layout) still uses 32-bit inode numbers. As for 64-bit applications, the 32-bit architecture still fits nicely. Thus, neither 32-bit nor 64-bit applications will have issues with the current ONTAP architecture. In ONTAP 9, a new NFS server option, `-v3-64bit-identifiers`, has been added to offer the ability to use only 64-bit FSID and file IDs for NFSv3 operations. The option is disabled by default, so ONTAP 9 will operate as previous releases did. If [disabling FSID changes in NFSv3](#), be sure to enable this option to avoid file ID collisions. This option is recommended when using the [NetApp FlexGroup volume](#) feature.

What is the difference between a session and a connection?

Answer: A connection operates at the transport layer of the OSI model (layer 4). A session operates at the session layer (layer 5). NFS operates at layers 5 through 7 (session, presentation, and application). Connections are the initial handshake between client and server. Sessions are the NFS RPC calls initiating the client-server relationship. In NFSv3, a session and a connection have a 1:1 relationship. NFSv4.x has an *n*:1 relationship in which a session can have multiple TCP connections for multipathing. However, ONTAP does not presently support multipathed NFS sessions.

How many simultaneous TCP connections are supported?

Answer: The maximum number of TCP connections is limited per node/HA pair and based on the amount of memory in a system. These connections are shared with all NAS connections. In a cluster, the limit is the sum of the HA pair limits in the cluster. For more information on maximums in ONTAP, refer to the [Configuration Maximums guides](#) found on the Field Portal.

Is there an `nfstat -d` command to detect misalignment on VMware datastores?

Answer: No. Use the NetApp Virtual Storage Console to detect and correct misaligned I/O on NFS datastores.

Is the `.snapshot` directory visible in NFSv4?

Answer: No. In NFSv3, the `.snapshot` directory is shown only at the mount point. Because there is no MOUNT protocol in NFSv4, the `.snapshot` directory is never populated. However, the `.snapshot` directory is still accessible through directory traversal. If looking for a way to hide snapshot directories at the NFS level (ie, the `nfs.hide_snapshot` 7-mode option), NFSv4.x would be a viable solution.

What are the best practices for VMware vSphere?

Answer: See [TR-4068: VMware vSphere on ONTAP](#) and [TR-4333: VMware vSphere 6 on NetApp ONTAP](#).

Is Windows NFS supported?

Answer: Yes. As of ONTAP 8.3.1 and 8.2.3, Windows NFS support was added and is controlled through a new NFS server option called `-v3-ms-dos-client`. For details, see [TR-4067: NFS Implementation and Best Practice Guide](#).

Is Oracle dNFS supported?

Answer: Yes. See [TR-3961: Oracle Database 11g Release 2 Performance Using ONTAP](#) for details.

Which LDAP schemas are supported?

Answer: All LDAP schemas are supported except for RFC-2307bis in ONTAP 8.2.x and earlier versions. RFC-2307bis support has been added to ONTAP 8.3. [TR-4073: Secure Unified Authentication with NetApp Storage Systems](#) covers this in detail. This support includes any and all variations of these schemas, as ONTAP allows you to modify LDAP schemas.

Is the Data ONTAP operating in 7-Mode option `nfs.hide_snapshot` supported?

Answer: The option `nfs.hide_snapshot` is not supported in ONTAP. See [bug 622226](#) for details. As a workaround, either turn snapshot directories off at the volume level or use NFSv4.x.

Are the following tunable variables supported?

- `nfs.tcp.recvwindowsize`
- `nfs.ifc.rcv.high`
- `nfs.ifc.xmt.high`

Answer: ONTAP does not have tunable variables for NFS because manual tuning is unnecessary. The network stack is self-tuning and contains flow-control mechanisms for NAS protocols. These mechanisms are covered in detail in [TR-4067: NFS Implementation and Best Practice Guide](#).

Can data volumes be accessed through NAS without mounting them into a namespace?

Answer: No. All data volumes must be mounted into a namespace to allow NAS access.

Can volumes be mounted to paths that do not exist?

Answer: No. Volumes must be mounted to paths that exist. For example, if mounting to `/vol/volname` in the namespace, then `/vol` should exist either as a volume called `vol` or as a folder called `vol` inside the `vsroot` volume, which is mounted to `/`.

Why don't Data ONTAP operating in 7-Mode NFS options (such as in node shell) work in ONTAP?

Answer: Data ONTAP operating in 7-Mode options for NFS don't work for ONTAP because those options exist at the d-blade/WAFL system level. NFS operations in ONTAP take place at the n-blade level. ONTAP operations that still take place at the d-blade level still honor the options in Data ONTAP operating in 7-Mode. Many of these options have been hidden, however.

Are LDAPS and LDAP signing and sealing supported in ONTAP?

Answer: LDAPS (LDAP over SSL) support was added in 8.2.1, but only for start-tls. LDAP signing and sealing are supported in ONTAP as of version 9.0. The option to control this is the `-session-security` option on the LDAP client configuration:

```
[-session-security {none|sign|seal}] - Client Session Security
This parameter specifies the level of security to be used for LDAP communications. If you do not
specify this parameter, the default is none.

LDAP Client Session Security can be one of the following:

o none - No Signing or Sealing.
o sign - Sign LDAP traffic.
o seal - Seal and Sign LDAP traffic.
```

Is cluster administration authentication through LDAP supported?

Answer: LDAP authentication for use in administering the cluster is supported in ONTAP 8.2 and later. This covers only non-Windows LDAP servers. To implement Windows LDAP servers for cluster administration authentication, use domain tunneling instead, which is covered in the product documentation.

Can I assign LDAP groups to manage the cluster?

Answer: As of ONTAP 9.0, LDAP groups can be used to manage clusters.

Is LDAP the same as Active Directory?

Answer: No. LDAP is used by Microsoft Active Directory, but it is not the same. Thus, cluster LDAP setup is needed only if Active Directory was configured as an LDAP server for use in UNIX attribute retrieval. See [TR-4073](#) for more information regarding Active Directory as an LDAP server.

What is the Kerberos token size limit for NFS in ONTAP?

Answer: The Kerberos token size is 4096kb in ONTAP 8.2. ONTAP 8.3 and later do not have a token size limit.

2 Decision Points

2.1 Why Use NFS?

Why should I use NFS?

Answer: NFS is useful in the following scenarios:

- File-based applications that are sensitive to queue depths and delays
- CPU-, disk-, or memory-intensive workloads for which the NetApp storage system can handle a majority of the workload instead of relying on hosts to do it—useful in virtualized workstation environments
- High-file-count environments
- Applications that are sensitive to storage failovers and disruptions
- Deep directory environments with ASCII-based data

Why should I use NFSv3?

Answer: NFSv3 has a slight performance advantage over current implementations of NFSv4.x, but NFSv4.x is catching up. NFSv3 also has the benefit of being simpler to implement than NFSv4.x. However, the version of NFS to choose is application dependent, because some applications fare better with the leasing and security that NFSv4 provides.

Why should I use NFSv4.x?

Answer: The following are benefits noted in [TR-4067: NFS Implementation and Best Practice Guide](#):

- Firewall-friendly, because NFSv4 uses only a single port (2049) for its operations
- Advanced and aggressive cache management, such as delegation in NFSv4.0
- Mandatory strong RPC security types that employ cryptography
- Internationalization
- Compound operations
- Works only with TCP
- Stateful protocol (not stateless like NFSv3)
- Kerberos configuration for efficient authentication mechanisms (uses 3DES for encryption)
- No replication support
- Migration (for dNFS) by using referrals
- Support of access control that is compatible with UNIX and Windows
- String-based user and group identifiers
- Parallel access to data (does not apply to NFSv4.0)

Why should I use pNFS?

Answer: The following are benefits noted in [TR-4063: Parallel Network File System Configuration and Best Practices for ONTAP 8.2](#):

- **Efficient load balancing for business-critical applications.** Because the computing requirements of high-performance commercial environments keep growing, storage capacity, CPU utilization, and network bandwidth need to scale seamlessly. Although adding incremental storage capacity in a ONTAP namespace is seamless, moving volumes to new storage for capacity growth, CPU, and network utilization rebalancing also has to be nondisruptive to business-critical applications running on the compute nodes. This requires near-zero downtime, and can apply to an oil and gas scenario in which hundreds of compute nodes are tasked to complete multiple jobs.

With pNFS, any volume that is transitioned live within the cluster namespace does not affect the compute nodes accessing the volumes from a new storage node. The NFS file handle does not change when the file system is served from a different cluster node after the move within the namespace; therefore, the file system does not need to be unmounted or remounted. This efficient way of load balancing helps to address highly utilized CPU, memory, and network in clustered storage. Although additional storage nodes in the scale-out architecture provide additional CPU and memory resources, the intelligent pNFS client is updated with the new location of the file system, helping to move the data access to the new location.

- **Effective resource utilization for applications with different workload characteristics.** With traditional NFS, metadata, data, and control share the same I/O path. As the data changes in its structure and grows in storage capacity, a single file system can have system resources saturated beyond a certain attainable performance. Applications implemented in electronic device automation (EDA) environments are very CPU-intensive. The CPU on a single node in a cluster could most likely prove to be a bottleneck. Although EDA demonstrates mostly sequential workloads that drive high CPU utilization, there can also be workloads such as user home directories and scratch spaces for software build environments on the same storage that are mostly random in nature. Redistributing these workloads within a cluster namespace and providing an optimal access path to the file systems are always challenges.

Clients supporting pNFS, in addition to following the necessary best practice recommendations for prevailing disk and aggregate sizing, can address heterogeneous workloads. A pNFS client can access file systems with different workload characteristics, from spinning disks or solid-state drives (SSDs), memory, and deduplication-aware extended cache. Regardless of where the file system resides in the namespace, the pNFS clients access the volume on a direct data path that is isolated from the metadata and the control path. This process provides parallelism to the application seeking access to file systems existing on different volumes through single or multiple NFS mounts, improving the throughput and optimizing latency by distributing heterogeneous workload across multiple nodes.

- **Scale and provision storage to achieve quick service-level agreements (SLAs) for IT as a service (ITaaS).** With today's changing and diverse business requirements, achieving SLAs in a timely manner and shortened time to market is critical. For cloud vendors who host ITaaS or platform as a service (PaaS), scaling, downsizing, and provisioning storage capacity require a quick turnaround time. Although scale-out architecture provides the flexibility to provision storage on demand, volumes in the namespace can reside in different tiers of storage (SSDs, Fibre Channel, or SATA disks) to support different workload requirements (frequently used to less frequently used). Traditional NFS clients that mount volumes from storage will have to choose the correct cluster node at mount time to provide an optimized data path to meet the performance SLAs. If multiple data LIFs exist in an SVM, then choosing the correct path can be challenging.

NFSv4.1 and pNFS can provide flexibility for mounting the file system or volumes from anywhere in the cluster namespace. The volumes in the namespace that are remote to the mount point specified by the client can still have a direct data path for optimal performance. Clustered applications can be mounted over pNFS, while monolithic ones can still be mounted over NFSv3. File systems that are exported from storage can have clients mount over different varieties of NFS, so they can coexist without making any significant changes to applications accessing the data or client-side infrastructure.

This feature reduces the overhead of frequent change management, which translates into time taken from SLAs. Multiple clients can simultaneously access different exported file systems. pNFS, an open source, does not require additional software or drivers on the client that are proprietary to enable it. As long as the client and server support pNFS, a simple mount suffices.

- **Nonconventional NFS clients in hypervisors and databases to improve accessibility, reliability, and performance.** Traditionally, some of the leading and most popular hypervisors such as those from VMware and databases such as those from Oracle have a built-in NFSv3 client that performs remote procedure call (RPC) communication similar to kernel NFS. This NFS client is lightweight and is integrated into the independent vendor stack. Today, the majority of NFSv3 traffic to a hypervisor datastore or a database volume uses a single TCP session. This implies that even if there is link aggregation at the network layer, clients are still limited by a single TCP session. The performance scales linearly as the number of connections is increased. There is no multipath support in NFSv3 either, so a client is unable to have multiple connections to a single NFSv3 server. If this single TCP session breaks (such as with a link failure), there must be another standby path for the RPC communication to resume. This need limits the accessibility, reliability, and performance between the client and the server.

NFSv4.1 and pNFS features will have a positive impact when these nonconventional NFS clients start adopting NFSv4.1 in their stack. NFSv4.1 provides a session model that allows multiple connections in one session. Presently, an NFSv4.1 client in a hypervisor and database can have multiple connections in one session.

2.2 Considerations for Making the Transition from NFS in Data ONTAP Operating in 7-Mode

For details on making the transition from NFS in Data ONTAP operating in 7-Mode, see [TR-4052: ONTAP Transition Guide](#) and [ONTAP for 7-Mode Administrators](#).

2.3 Common Uses of NFS

- Oracle databases
- VMware datastores
- Source-code repositories
- Compute farms
- Log repositories
- Video and imaging render farms
- TIBCO
- SAS grid software
- Home directories

3 Common Issues

3.1 Cannot Mount

The following section covers common errors and scenarios in which an NFS mount fails to a ONTAP system. The section also covers how to resolve the issue.

Table 4) Common mount failures.

Error	What to Check	How to Resolve
Access denied by server while mounting.	NFS client: • If using Kerberos, is the configuration correct? See TR-4073 for details. • If using AUTH_UNIX or AUTH_SYS, does the user resolve in the name service? NFS server: • NFS server options. • Can the user be resolved by the server into a UID? • Is SecD running? • Export policy rules of the volume. • Export policy rules of the parent volumes. Common mistakes in export policies include: • No rule defined in export policy. • Clientmatch is incorrect (such as 0.0.0.0 instead of 0.0.0.0/0 for all clients). • Client match does not allow client attempting access. • Access protocol does not allow NFS. • RO policy is set to incorrect value. • User is squashed to the anon user, which does not have permissions to the volume. NFS server options that could cause access-denied errors during mount include: • NFS mount root only.	Review the NFS client configuration. Review the NFS server configuration, export policies, and rules and make corrections.
Requested NFS version or transport protocol is	NFS server: • NFS server options for the NFS version.	Review the NFS server configuration to verify that the protocol and NFS

not supported.	- TCP/UDP settings. - NFS server is running. Network/firewall: - Verify that the firewall is not blocking NFS or related ports. - Verify that the data LIF allows NFS. SVM: - Verify that the SVM allows NFS as a protocol.	version are enabled and the server is running. Review the network settings and data LIFs to verify that NFS is allowed. Review the SVM to verify that NFS is allowed.
Mount hangs indefinitely.	Network: - Is the LIF up? Can it be pinged? - Is the DNS entry correct? - Is the LIF at home? Are failover groups configured properly? - Is the firewall blocking any of the NFS ports?	Review the network settings. Review the data LIFs on the cluster.
Mounting failed, reason given by server: No such file or directory.	Mount syntax: - Is the right mount path specified? NFS server: - Is the junction the same as the mount path? - Is the volume mounted? - If using LS mirrors, have they been updated? NFS client: - If volume permission changes have been made, has the volume been remounted? - Is the volume mounted with no access cache?	Review the mount syntax. Review the NFS volume.
Mount point is busy or already mounted.	NFS client: Is something already mounted to that mount point?	Review the output of the <code>mount</code> command.
Mount point/test does not exist.	NFS client: Does the mount point exist?	Use a valid mount point.
Only root can do that.	NFS client:	Check client and server

	- Does the user have permission to mount? NFS server: - Is root-only mount set?	configuration.
--	--	----------------

3.2 Permission Denied or Access Issues

The following section covers issues in which an NFS mount succeeds but accessing the mount fails; it also covers how to resolve the issue. Not all scenarios are covered.

Table 5) Common access issues.

Error	What to Check	How to Resolve
Permission denied (while accessing mounts and reading/writing).	NFS server: - Do the data volume's security settings permit the user access? - Do the parent volume's security settings permit the user access? - What is the volume's security style? Does the user attempting access map to a valid Windows user if the security style is NTFS? - If able to change directories but not able to read or write, but UNIX permissions seem to allow access to all users, does the cluster know the user attempting access?	Verify and modify the volume's security. Verify and modify the export policy rule to allow access. Verify that the user can map properly into name service. Verify that the user exists in name service.
Permission denied (while attempting <code>chmod/chown/chgrp</code>). Operation not permitted (while attempting <code>chown/chmod/chgrp</code>).	NFS server: - Is <code>chown</code> allowed by anyone other than root? - Is the user the owner of the file? NFS client: - Is the user root?	Change the NFS server and export policy rule options for <code>chown</code> to <code>unrestricted</code>.
Not a directory (when traversing <code>.snapshot</code> directory).	NFS client: Check the kernel version.	See Bugzilla 798809 .

3.3 Files Written as “Nobody”

The following section covers issues in which NFSv4 clients show file ownership as the “nobody” user; it also covers how to resolve the issue. Not all scenarios are covered.

A stale file handle error occurs when the server file system has changed and the file handle is no longer valid. For example: Client A opens file `xxx.yyy` for edit; Client B deletes this file; Client A tries to save the edit, and Client A gets a stale file handle error.

This can occur not just for operations on individual files, but also because of changes in the directory structure.

Table 6) Files written as “nobody” in NFSv4.

Error	What to Check	How to Resolve
No error; files written as the “nobody” user (or some other unexpected user).	NFS client: • <code>/var/log/messages</code> file. • Is the NFSv4 domain specified in <code>/etc/idmapd.conf</code>? • What is the user name attempting access? Can the client resolve the name in the name service? NFS server: • Can the cluster translate the user name to a UID? • Is the NFSv4 domain set? • If the user writing the file is root, does the export policy squash root to the anon user (<code>superuser = none</code>)? • Is the name service working properly?	Fix the NFSv4 domain ID. Verify that the user name matches the NFSv4 domain user name exactly (case-sensitive). Verify that the client and cluster can resolve the user name. Adjust the export policy rule to allow superuser access if desired.

3.4 Stale File Handle or NFS Mount

The following section covers scenarios in which a stale file handle error might occur and how to resolve them. Not all scenarios are covered.

Table 7) Troubleshooting stale file handle errors.

Error	What to Check	How to Resolve
<code>mount.nfs: Stale file handle.</code>	NFS server: • Did the junction path for the volume change? • Is the volume mounted? • Does the volume still exist? NFS client: • Is the mount already mounted somewhere else on the client?	Verify and remount the volume from the client. Mount the volume from the cluster if it is not mounted.
Cannot open directory: Stale file handle.	NFS server: • Was the <code>fsid-change</code> option modified?	Remount the volume from the client.
Was not found in <code>/proc/mounts</code> .	NFS client: • Does the mount show up in <code>/proc/mounts</code>? • Does the mount show <code>(deleted)</code> in the output?	Reboot the client.

4 References

Read the following references for additional information:

- [ONTAP Documentation](#)
- TR-3580: NFSv4 Enhancements and Best Practices Guide
www.netapp.com/us/system/pdf-reader.aspx?m=tr-3580.pdf
- TR-3961: Oracle Database 11g Release 2 Performance Using ONTAP
www.netapp.com/us/system/pdf-reader.aspx?m=tr-3961.pdf
- TR-4037: Introduction to NetApp Infinite Volume
www.netapp.com/us/system/pdf-reader.aspx?m=tr-4037.pdf
- TR-4052: ONTAP Transition Guide
www.netapp.com/us/system/pdf-reader.aspx?m=tr-4052.pdf
- TR-4063: Parallel Network File System Configuration and Best Practices for ONTAP
www.netapp.com/us/system/pdf-reader.aspx?m=tr-4063.pdf
- TR-4067: NFS Implementation and Best Practice Guide
www.netapp.com/us/system/pdf-reader.aspx?m=tr-4067.pdf
- TR-4068: VMware vSphere on ONTAP
www.netapp.com/us/system/pdf-reader.aspx?m=tr-4068.pdf
- TR-4073: Secure Unified Authentication with NetApp Storage Systems
www.netapp.com/us/system/pdf-reader.aspx?m=tr-4073.pdf
- TR-4129: Namespaces in ONTAP
www.netapp.com/us/media/tr-4129.pdf
- TR-4379: Name Services Best Practices
www.netapp.com/us/system/pdf-reader.aspx?m=tr-4379.pdf
- TR-4557: NetApp FlexGroup Technical Overview
www.netapp.com/us/system/pdf-reader.aspx?m=tr-4557.pdf
- Linux.org
linux.org
- Linux.die.net
linux.die.net
- IETF
<http://ietf.org>

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 1994–2016 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark Information

NetApp, the NetApp logo, Go Further, Faster, AltaVault, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANtricity, SecureShare, Simplicity, Simulate ONTAP, SnapCenter, Snap Creator, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, WAFL, and other names are trademarks or registered trademarks of NetApp Inc., in the United States and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netapptmlist.aspx>.