



Technical Report

Guide to Transitioning Windows File Services from 7-Mode to Clustered Data ONTAP

Saurabh Singh, Brahmanna Chowdary, NetApp
June 2016 | TR-4522

Summary

This report is not a migration document, but it discusses the transition of Windows File Services from NetApp® Data ONTAP® operating in 7-Mode to clustered Data ONTAP. The report covers the migration of a CIFS information worker workload from Data ONTAP operating in 7-Mode to clustered Data ONTAP.

TABLE OF CONTENTS

1	Introduction	3
1.1	Scope	3
1.2	Out of Scope.....	3
1.3	Intended Audience.....	3
2	Overview of CIFS in Clustered Data ONTAP	3
2.1	Network Access to CIFS Server	4
2.2	CIFS with Global Name Space for each SVM	4
2.3	Network Access to CIFS Server	4
2.4	CIFS Server Management and Data Protection	5
3	Key Transition Considerations.....	5
4	Migration Methodology	6
5	Migration Tools	6
6	Windows File Services Solution Use Cases	7
6.1	Home Directory.....	7
6.2	Implementation of FPolicy Based Solutions.....	9
6.3	Implementation of Antivirus Solution.....	13
6.4	Implementation of CIFS Native Auditing	16
6.5	CIFS Events	17
	References	19

LIST OF TABLES

Table 1)	Feature comparison.	8
Table 2)	FPolicy configuration comparison.	11
Table 3)	Vscan configuration comparison.	15
Table 4)	Native auditing configuration comparison.	17

1 Introduction

Clustered Data ONTAP is a unified, scale-out storage system that solves the challenge of unpredictable data growth. Clustered Data ONTAP also provides efficient performance and supports multitenancy and data mobility.

This white paper provides best practices for Data ONTAP 7G and 7-Mode Common Internet File System (CIFS) customers who would like to transition to clustered Data ONTAP. This paper also provides a CIFS feature comparison of the two storage platforms and discusses the changes that have been made to clustered Data ONTAP. Additionally, the paper provides guidelines for configuring new environments that will either make the transition process easier or lead to a posttransition clustered Data ONTAP deployment that more closely matches the capabilities of your current environment.

NetApp assumes that readers of this paper already have a basic working knowledge of the capabilities of clustered Data ONTAP, including its architecture and value proposition. This paper provides only a brief overview of CIFS in cluster mode. For more information on clustered Data ONTAP, refer to the resources provided in the reference section at the end of this report.

1.1 Scope

This document provides information on the following topics:

- Overview of CIFS in clustered Data ONTAP
- Windows File Services features
- Differences and similarities in CIFS features between Data ONTAP operating in 7-Mode and clustered Data ONTAP

1.2 Out of Scope

This document does not include instructions for migrating your environment from 7-Mode to clustered Data ONTAP. Instead, it describes the changes in Windows File Services from 7-Mode to clustered Data ONTAP that must be considered during migration to clustered Data ONTAP.

1.3 Intended Audience

This report is intended for system engineers, professional service providers, storage administrators, system administrators, and data center managers. The report assumes basic familiarity with the following NetApp products and products from other vendors:

- NetApp FAS systems and the Data ONTAP operating system
- Windows File Services
 - CIFS protocol
 - Windows Active Directory (AD)
- Data ONTAP 7-Mode architecture

2 Overview of CIFS in Clustered Data ONTAP

Clustered Data ONTAP is designed to scale. It allows you to group pairs of physical nodes, share resources, and distribute work across the system while presenting a single entity to manage. When a cluster node becomes live, the cluster administrator decides whether to create a new cluster system, starting with the new node, or to join the node to an existing cluster. When a node joins an existing

cluster, the node shares its storage space, taking over the network traffic and serving the storage access requests.

2.1 Network Access to CIFS Server

Clustered Data ONTAP provides multitenancy with the storage virtual machine (SVM; formerly call Vserver). An SVM is not physically bound to one specific node, but spans the boundaries of physical nodes in the cluster, presenting a single view to client applications. An SVM allows a single authentication domain for each protocol, which enables data to be visible to—and accessible by—only the authorized owner. Volumes are a basic unit for storing datasets that are created on top of an aggregate, which resides on a single cluster node. A volume is always owned by a unique SVM, thus providing data security.

Note: The term “Vserver” has been replaced by “storage virtual machine” or “SVM.” However, “Vserver” continues to appear in GUI fields and wizard pages in various NetApp tools and applications and in the clustered Data ONTAP CLI.

2.2 CIFS with Global Name Space for each SVM

Clustered Data ONTAP provides a clusterwide name space by stitching volumes together through junctions. During SVM creation, a root volume is created. This root volume acts as the root of the global name space; it is the container that can join more volumes together. Volumes can be created on an as-needed basis and be linked together through junction paths under the root volume of an SVM. If you create a Common Internet File System (CIFS) share that is mapped to the root volume of an SVM, the end user will be able to access all the data saved on different volumes across the cluster through a single share.

Storage users often need to move their data from one location to another for various reasons, such as:

- Aggregate running out of space
- Load-balancing across hosts
- Moving data from higher- to lower-tier storage or from lower to higher

The volume-move activity in clustered Data ONTAP is achieved without disconnecting the existing CIFS clients. Starting with Data ONTAP 8.0, the cluster system provides the capability to seamlessly move volumes within an SVM with minimal impact on data access.

2.3 Network Access to CIFS Server

A CIFS server is created on an each-SVM basis. To access a CIFS server, one or more logical interfaces (LIFs) must be created and associated with the CIFS server and SVM. LIFs are a set of IP addresses assigned to network ports attached to cluster nodes. When needed, LIFs can be migrated to ports on different nodes. LIFs can be associated with failover rules. Doing so allows the cluster to reroute the network traffic to another available port when the original LIF encounters port failure, network interface failure, cable failure, and so on.

In the event of LIF migration, all CIFS connections are terminated; the older CIFS clients, such as Windows XP and Windows 2000, experience levels of service disruption. With improved SMB2.x and SMB 3.0 implementation, together with newer CIFS clients that support SMB 2.x or later, a CIFS server can provide a certain degree of nondisruptive operation. This means that, upon LIF failover or migration, CIFS clients with SMB 2.0 and later capability are able to reclaim the disconnected CIFS session, reopen the CIFS files, and continue with “leftover” CIFS processing. As a result, the application running on those CIFS clients is *not* aware of network changes that are introduced by LIF migration.

2.4 CIFS Server Management and Data Protection

A CIFS server can be managed and accessed from any node in a cluster. The management of a given CIFS server can be restricted to an administrator without giving access to cluster context. An SVM administrator can manage a CIFS server without having detailed knowledge of the cluster or the location of the SVM. Clustered Data ONTAP allows administrators to provide storage through an SVM and CIFS server for multiple customers as in the Storage Service Provider (SSP) model. In that way, each customer can independently manage the company's data without worrying about a security breach because of the secure multitenancy feature of clustered Data ONTAP data security and access management.

Starting with Data ONTAP 8.2, clustered Data ONTAP supports off-box antivirus (AV). Access to file and directory information can be managed with export policy, volume access permissions using Storage Level Access Guard (SLAG), and share-level access control lists (ACLs). The access permissions are checked in the same order: export policy, volume-level access permissions, and share-level ACLs. Before granting access to an AD user, Windows users must be mapped to UNIX users if the volume security style is either UNIX or mixed.

3 Key Transition Considerations

The key transition considerations for CIFS are:

- File folding in Data ONTAP operating in 7-Mode is superseded by the use of deduplication in clustered Data ONTAP.
 - Kerberos authentication using DES and RC4-HMAC encryption is available in clustered Data ONTAP.
 - AD-based LDAP authentication is available in clustered Data ONTAP. Non-AD LDAP and workgroup authentication is not available in clustered Data ONTAP.
 - Features added in clustered Data ONTAP 8.2.1 include:
 - LDAP over SSL (using Start-TLS)
 - Off-box AV scanning
 - Multiple domain search for user mapping
 - Features added in clustered Data ONTAP 8.2.2 include:
 - `CIFS.restrict_anonymous` option. The default for clustered Data ONTAP is `no_restriction`.
 - Features added in clustered Data ONTAP 8.3 include:
 - NetBIOS aliases
 - Storage Level Access Guard (SLAG)
 - Character mapping
 - Auditing log-on and log-off events
 - Home directory visibility for administrators
 - Viewing and management of shares, open files, and sessions in the Microsoft Management Console (MMC)
- Note:** Windows PowerShell scripts allow clusterwide enumeration of CIFS objects and provide additional functionality for clustered Data ONTAP systems.
- Microsoft GPOs are supported for Kerberos settings, refresh time interval, and refresh time interval offset in clustered Data ONTAP.
 - Features that are available in 7-Mode but not available in clustered Data ONTAP include:
 - Live View Audit
 - Perfmon.exe support

- LDAP (secure) using signing and sealing
- Display of CIFS client activities

4 Migration Methodology

Migration is not just a script to run and finish. Migration is a multistep procedure and requires proper planning to be achieved perfectly. Assess different parameters to understand the environment and requirements before you proceed with the migration.

Figure 1 shows the migration methodology from Data ONTAP operating in 7-Mode to clustered Data ONTAP.

Figure 1) Migration methodology from Data ONTAP operating in 7-Mode to clustered Data ONTAP.



1. **Evaluate.** Evaluate clustered Data ONTAP. Identify the required feature set. Check 7vC feature mapping.
2. **Plan.** Create migration plans, provision destination storage, and configure migration tools.
3. **Execute.** There are three components to the execution phase.
 - a. Configuration migration
 - b. Data migration
 - c. Feature transition
4. **Verify.** Verify new system configurations and provide documentation.

The comprehensive four-step process is designed to handle the most complex migration scenarios and execute the right migration strategy for your specific environment.

- **Planning.** The NetApp team will work with the client's IT infrastructure team to understand the technical, business, and operational requirements to minimize the impact on operations during the migration.
- **Project management.** NetApp project managers will work with the customer's staff throughout the entire migration project, creating the initial statement of work (SOW) and project plan, providing progress reports, communicating issues, and offering alternative solutions and optimizations.
- **Experience.** NetApp PS will fine-tune procedures, recognize and avoid common pitfalls, and develop workarounds. NetApp Professional Services has developed proficiency, field-proven methodologies, and tools that help to ensure a successful migration project.
- **Methodology.** The NetApp data migration methodology is a comprehensive four-step process designed to handle the most complex migration scenarios and execute the right migration strategy for the client's specific environment.

5 Migration Tools

When the plan is ready, use tools to migrate the data to the destination storage platform. The following tools can be used to migrate the data, depending on the customer requirements:

- 7-Mode Transition Tool (7MTT)

- NDMPCopy

The following non NetApp host-based tools are commonly used for data migration:

- Logical volume managers (LVMs) from various vendors
- ScriptLogic Secure Copy
- Rsync
- Robocopy/Richcopy
- PEER Software PeerSync
- Data Dynamics StorageX
- XCP–NFSv3 migration (Linux CLI only)

6 Windows File Services Solution Use Cases

6.1 Home Directory

The home directory is a feature that is deployed extensively in a CIFS environment. A CIFS home directory is a dynamic share in memory that maps to a specific path on a virtual server that serves as the home directory for a specific user. Data ONTAP finds this path automatically using the patterned share name and search paths that you provide.

The home directory is a special type of share that is different from static shares in the following ways:

- You cannot change the share-level ACL and the comment for a home directory.
- The `CIFS shares show` command does not display the home directories.
- A home directory share name can be a patterned name.
- The share path for a home directory is a relative path; the share path for a regular static share is an absolute path.

Comparison of Home Directory Between 7-Mode and Clustered Data ONTAP

The way that home directory works in Data ONTAP 7-Mode is quite different from the way it works in clustered Data ONTAP.

- Data ONTAP 7-Mode uses `options CIFS.home_dir_namestyle` to determine the home directory pattern and search path. In clustered Data ONTAP you have more flexibility in terms of defining home share names and share paths.
- Data ONTAP 7-Mode uses the `CIFS homedir showuser` command to find a user's home directory path and shows the path to the admin. In clustered Data ONTAP you can use `CIFS home-directory show-user`.
- Data ONTAP 7-Mode uses `options CIFS.home_dirs_public_for_admin` and `options CIFS.home_dirs_public` to control the publicity of a user's home directory and allow someone to mount a share that is not his or her own. Starting in Data ONTAP 8.3, the `BUILTIN\Administrators` group can connect to CIFS HDs if the `-is-home-dirs-access-for-admin-enabled` option is enabled.
- Data ONTAP 7-Mode is able to use mapped UNIX user names in the share pattern through `options CIFS.home_dir_namestyle mapped`. The same can be achieved in clustered Data ONTAP by mapping the Windows/UNIX user.

In clustered Data ONTAP, there are four variables that determine how a user is mapped to a directory:

Share Name

This is the name of the share that you create that the user connects to. You must set the home directory property for this share.

The share name can use the following dynamic names:

- %w (the user's Windows user name)
- %d (the user's Windows domain name)
- %u (the user's mapped UNIX user name)

To ensure that the share name is unique across all home directories, the share name must contain either the %w or the %u variable. The share name can contain both the %d and the %w variable (for example, %d/%w) or can contain both a static portion and a variable portion to the share name (for example, home_%w).

Share Path

This is the relative path, defined by the share. It is associated with one of the share names that is appended to each search path. The path generates the user's entire home directory path from the root of the SVM. It can be static (for example, home), dynamic (for example, %w), or a combination of the two (for example, eng/%w).

Search Paths

This is the set of absolute paths from the root of the specified SVM that directs the clustered Data ONTAP search for home directories. You specify one or more search paths by using the SVM CIFS home-directory `search-path add` command. If you specify multiple search paths, clustered Data ONTAP tries them in the order specified until it finds a valid path.

Directory

This is the home directory that you create for the user. It is usually the user's name. You must create this directory in one of the directories defined by the search paths. Table 1 lists the share name format and compares 7G/7-Mode and clustered setups. Assume that the search path is set to /home, the user's login domain is NTDom, the user's Windows login name is NTUser, and the user's mapped UNIX name is UnixUser.

Table 1) Feature comparison.

Share Name	Login User	Data ONTAP 7-Mode	Clustered Data ONTAP
NTUser	me	Original or ntnamestyle; path is /home/NTUser	Share name is %w.
~NTDom~NTUser	me	Domain style; path is /home/NTDom/NTUser	Share name is %d%w.
~UnixUser	me	Mapped style; path is /home/UnixUser	Not available.
CIFS.HOMEDIR	me	Any name style; path is decided based on name style	Share name is CIFS.HOMEDIR.

Considerations During Migration Planning

Before you perform a migration, be aware of the workload you plan to migrate and make note of the following details that will help to size the environment for migration:

- Determine the number of home directories or users.
- Check `CIFS_homedir.cfg` to pull information regarding the search paths (volumes used for home directories).
- Determine if the users are classified based on I/O profiles such as Power User or Normal User.

Based on these three factors, determine how to segregate the capacity and load on the cluster.

Optimal Leveraging of Clustered Architecture

Storage Performance

To achieve better storage performance, distribute the workload across different nodes in the cluster. This distribution is possible because of the clusterwide namespace in clustered Data ONTAP.

Example scenario: 4,000 users have home directories on a single volume in 7-Mode

You need to migrate this environment to a four-node clustered Data ONTAP cluster. You can create one volume per node (four volumes in total) and assign them to the SVM hosting the CIFS server for the `home` directory setup. Now you can create 1,000 home directories per volume. Make sure you add the four volumes in the `home` directory search path.

You now have 1,000 users per node and have provided scalability and a performance boost to the overall solution.

Network Performance

You can load balance the network by creating data LIFs and nodes for the SVM and create DNS load balancing to use different LIFs every time a new request comes in.

From this example, create four LIFs with one per node for the SVM. Use DNS load balancing to send access requests to different nodes and load balance the network load from one node in 7-Mode to multiple nodes in clustered Data ONTAP.

Migration Methods

The 7-Mode Migration Tool (7MTT) supports migration of the HOMEDIR environment. 7MTT transfers both the path to the HOMEDIR and the pattern used to create the user's dynamic shares. Many-to-one consolidation can be done very easily because of the clusterwide namespace. There are third-party tools that can also be used for migration. One such solution is PeerSync by Peer Software.

6.2 Implementation of FPolicy Based Solutions

NetApp FPolicy® software is a file-access notification system that enables an administrator to monitor file access in Network File System (NFS) and CIFS configured storage. Introduced for the scaled-out architecture in NetApp clustered Data ONTAP 8.2, FPolicy enables a rich set of use cases with our selected partners. FPolicy requires that all nodes in a cluster run Data ONTAP 8.2 or later. FPolicy supports all SMB versions and major NFS versions, including NFSv3 and NFSv4.

FPolicy natively supports a simple file-blocking use case that allows administrators to prevent end users from storing unwanted files. For example, an administrator can block the storage of audio and video files in data centers and thus save precious storage resources. This feature blocks files based only on extension; for more advanced features, consider partner solutions.

This system enables partners to develop applications that cater to a diverse set of use cases, including:

- File screening
- File-access reporting
- User and directory quotas

- Hardware security module and archiving solutions
- File replication
- Data governance

Enhancements in Clustered Data ONTAP Compared to 7-Mode

Clustered Data ONTAP replicates most of the 7-Mode features pertaining to FPolicy with enhanced versions of those features. Important feature enhancements include:

- **FPolicy request/response schema.** The FPolicy request response schema has been enhanced in clustered Data ONTAP. It now uses the xml-over-tcp protocol for communication rather than the RPC-based request/response schema used in 7-Mode.
- **Filters.** Filters are used along with a given file operation to reduce the noise between the appliance and the FPolicy server to achieve better performance.
- **Security.** To enhance the security between the FPolicy server and clustered Data ONTAP, we introduced the concept of a privileged user. This user must be added to the security login so that the FPolicy server can communicate with the appliance using this user.
- **SSL connection.** SSL connection is introduced in clustered Data ONTAP for secure communication between the FPolicy server and clustered Data ONTAP.
- **Counters and parameters.** The FPolicy performance counters and parameters have changed. Refer to vendor-specific TRs for FPolicy, available in the [NetApp Library](#).
- **Connection initiation.** Connection is initiated from the FPolicy server in 7-Mode, but in clustered Data ONTAP it is initiated from the clustered Data ONTAP side.

Considerations Before Planning a Migration

Following are the aspects you need to consider when planning a migration for an environment in which an FPolicy based solution—such as auditing, quota, archiving, and so on—is used.

1. Collect the following information about the existing setup:
 - a. Use case (for example, auditing, quota, file screening, and archiving).
 - b. Since the FPolicy framework in clustered Data ONTAP is different from that in 7-Mode, the behavior of the vendor software and the operations it monitors also changes. Accordingly, there are filters for different use cases that help fine-tune performance.
 - c. Protocols being monitored (NFS, CIFS). Which protocols are being monitored dictates which file operations need to be monitored and which filters need to be applied. Not all operations and filters apply to both CIFS and NFS.
 - d. Volumes/shares being monitored. Make a note of volumes and shares being monitored and also the amount of data. This information will help size the environment to achieve optimal performance.
 - e. Number of users in the environment.
 - f. File operations being monitored. This information will provide insight into which filters can be applied for the given set of file operations to achieve optimal performance.
2. When segregating the volume load, make sure that the FPolicy server can be reached through at least one LIF on the node to process FPolicy requests for the file operations performed on that node.
3. Make sure that the security login is configured based on vendor recommendations for sending NetApp ONTAPI® calls.
4. Refer to FPolicy vendor software and clustered Data ONTAP solution guides for detailed deployment steps.

Comparison Between Data ONTAP 7-Mode and Clustered Data ONTAP for FPolicy Configuration

Table 2 shows the FPolicy configuration comparison between 7-Mode and clustered Data ONTAP to make it easier to map the different configuration parameters.

Table 2) FPolicy configuration comparison.

Feature	Data ONTAP 7-Mode	Clustered Data ONTAP
Events	Specified using FPolicy monitor: Add <i>PolicyName</i> -p <protocol> -f <operation>	Specified by creating an event: fpolicy policy event create - SVM <SVM name>
External FPolicy server	Don't need to specify it explicitly because the FPolicy server registers with the server.	Need to create an external engine and add the FPolicy server using fpolicy policy external-engine create
Policy	Created by fpolicy create <policy name> screen	Need to create FPolicy policy explicitly, which acts as a placeholder to combine all components of FPolicy setup.
Scope	The command syntax to work with file volumes is as follows: fpolicy vol[ume] {inc[lude] exc[lude]} {add remove set eval} <i>PolicyName</i> vol-spec The command syntax to work with file extension is as follows: fpolicy extensions { include exclude } { set add remove } <i>PolicyName</i> ext-list	Scope defines the boundaries under which the FPolicy would act. -shares-to-include -shares-to-exclude -volumes-to-include -volumes-to-exclude -export-policies-to-include -export-policies-to-exclude -file-extensions-to-include -file-extensions-to-exclude
Filters	No concept of filters in 7-Mode	Filters help reduce the chatter between the FPolicy server and the appliance by filtering and sending only required fields to the FPolicy server. fpolicy policy -SVM -event- name <event name> -filters monitor-ads close-with-modification close-without-modification first-read first-write offline-bit open-with-delete-intent open-with-write-intent write-with-size-change
Enabling	fpolicy enable <i>PolicyName</i> [- f]	fpolicy policy enable <policy_name> -sequence <number>

Feature	Data ONTAP 7-Mode	Clustered Data ONTAP
Scan Mandatory feature	<p>This option mandates every I/O to go through the FPolicy framework.</p> <p>Note: If the FPolicy server is not available, the I/Os would be blocked.</p> <p>To enable: <code>fpolicy options PolicyName required on</code></p> <p>To disable: <code>fpolicy options PolicyName required off</code></p>	Scan mandatory is configured when creating FPolicy policy.
Notification type	Notifications are sent using named pipes.	Notifications are sent using xml-over-tcp protocol.

Clustered Data ONTAP Best Practices for FPolicy

NetApp recommends following FPolicy best practices for server hardware, operating systems, patches, and so on.

Configuration of FPolicy External Engine for SVM

Providing additional security comes with a performance cost. Enabling SSL communication will have a performance effect on CIFS.

Configuration of FPolicy Events for SVM

Monitoring file operations has an effect on the overall user experience. In fact, filtering unwanted file operations on the storage side improves the overall user experience. NetApp recommends monitoring the minimum number of file operations and enabling the maximum number of filters without breaking the use case. The CIFS home directory environment has a high percentage of `getattr`, `read`, `write`, `open`, and `close` operations. NetApp recommends using filters for these operations.

Configuration of FPolicy Scope for SVM

Restrain the scope of the policies to relevant storage objects, such as shares, volumes, and exports, rather than enabling them throughout the SVM. NetApp recommends checking directory extensions. If `is-file-extension-check-on-directories-enabled` is set to `True`, directory objects are subjected to the same extension checks as regular files.

Network Configuration

Network connectivity between the FPolicy server and the controller should be of low latency. NetApp recommends separating FPolicy traffic from client traffic by using a private network.

Note: In a scenario in which the LIF for FPolicy traffic is configured on a different port than the LIF for client traffic, the FPolicy LIF might fail over to another node because of a port failure. This action will make the FPolicy server unreachable from the node and the FPolicy notifications for the file operations on the node will fail.

Note: Make sure that the FPolicy server can be reached through at least one LIF on the node to process FPolicy requests for the file operations done on that node.

Hardware Configuration

The FPolicy server can be on either a physical server or a virtual server. If the FPolicy server is in a virtual environment, make sure to allocate dedicated resources (CPU, network, and memory) to the virtual server.

Multiple Policy Configuration

The FPolicy policy for native blocking has the highest priority, irrespective of the sequence number. Decision-altering policies have a higher priority than others. Policy priority depends on use cases. To determine the appropriate priority, NetApp recommends working with partners.

Managing FPolicy Workflow and Dependency on Other Technologies

NetApp recommends disabling an FPolicy policy before making any configuration changes. For example, if you want to add or modify an IP address in the external engine configured for the enabled policy, then first disable the policy.

If you configure FPolicy to monitor NetApp FlexCache® volumes, NetApp recommends that you do not configure FPolicy to monitor `read` and `getattr` file operations. Monitoring these operations in Data ONTAP requires the retrieval of inode-to-path (I2P) data. Because I2P data cannot be retrieved from FlexCache volumes, it must be retrieved from the origin volume. Therefore, monitoring these operations eliminates the performance benefits that FlexCache can provide.

When both FPolicy and an off-box AV solution are deployed, the AV solution receives notifications first. FPolicy processing starts only after AV scanning is complete. A slow AV scanner could affect overall performance, so AV solutions must be sized properly.

Add all shares that you want to monitor or audit into the share-include list during scope definition. Turn off monitoring on the file server if you do not want to monitor it. Disabling FPolicy on the SVM is not helpful, because the Varonis probe service probes the file server and automatically disables or enables FPolicy if it notices a disconnection.

Sizing Considerations

FPolicy performs inline monitoring of CIFS operations, sends notifications to the external server, and waits for a response, depending on the mode of external engine communication (synchronous or asynchronous). This process affects the performance of CIFS access and CPU resources. To mitigate any issues, NetApp recommends assessing and sizing the environment before enabling FPolicy. Performance is affected by the number of users; workload characteristics, such as operations per user and the data size; and network latency.

Migration Methods

Currently 7MTT doesn't support FPolicy configuration migration. Please gather the required data from the 7-Mode appliance and refer to Table 3 to get the configuration comparison. Accordingly, configure the settings in clustered Data ONTAP.

6.3 Implementation of Antivirus Solution

The off-box AV feature provides virus-scanning support to the NetApp clustered Data ONTAP operating system. In this architecture, virus scanning is performed by external servers that host AV software from third-party vendors. This feature offers AV functionality that is similar to the functionality currently available in Data ONTAP operating in 7-Mode.

The off-box AV feature provides virus-scanning support by triggering in-band notifications to the external virus-scanning servers during various file operations, such as `open`, `close`, `rename`, and `write` operations. Because of the in-band nature of these notifications, the client's file operation is suspended

until the file scan status is reported back by the virus-scanning server, a Windows Server instance that is referred to as a Vscan server.

The Vscan server, upon receiving a notification for a scan, retrieves the file through a privileged CIFS share and scans the file contents. If the AV software encounters an infected file, the software attempts to perform remedial operations on the file. The remedial operations are determined by the settings that are configured in the AV software.

After completing all necessary operations, the Vscan server reports the scan status to clustered Data ONTAP. Depending on the scan status, clustered Data ONTAP allows or denies the file operation requested by the client. In clustered Data ONTAP 8.2.1, virus scanning is available only for CIFS-related traffic.

Enhancements in Clustered Data ONTAP as Compared to 7-Mode for Virus Scanning

Clustered Data ONTAP replicates most of the 7-Mode features pertaining to Vscan with an enhanced version of those features.

- **Granular scan exclusion.** Clustered Data ONTAP gives you the ability to exclude files from virus scanning based on file size and location (path) or to scan only the files that are opened with execute permissions.
- **Support for updates to the AV software.** Clustered Data ONTAP supports rolling updates of the AV software and maintains information about the software running version along with the scan status of files. If the AV software running in a single server in a scanner pool is updated to a later version, the scan status of all files that were already scanned is not discarded.
- **Security enhancements.** Clustered Data ONTAP validates incoming connection requests sent by the Vscan server. Before the server is allowed to connect, the connection request is compared to the privileged users and IP addresses defined in the scanner pools to verify that the request originates from a valid Vscan server.
- **Use of AV connector software.** In clustered Data ONTAP, the appliance doesn't directly connect to the Vscan server. An additional piece of software, called an AV connector, is required for clustered Data ONTAP. The connector acts as an interface between SVMs and Vscan servers and pulls scan requests from the cluster.

Considerations Before Planning a Migration

Below are aspects to consider when planning a migration for an environment in which Vscan is being used for antivirus protection of CIFS shares.

1. The existing Vscan server software should support clustered Data ONTAP.
2. Make sure that the Vscan server can be reached through at least one LIF on the node to process Vscan requests for the file operations performed on that node.
3. Use the SPM Tool to size the environment considering the AV workload.
4. Refer to the Interoperability Matrix Tool (IMT) to check the overall supportability of the components in question.

Comparison Between Data ONTAP 7-Mode and Clustered Data ONTAP for FPolicy Configuration

Table 3 shows a Vscan configuration comparison between Data ONTAP operating in 7-Mode and clustered Data ONTAP to make it easier to map the different configuration parameters.

Table 3) Vscan configuration comparison.

Feature	Data ONTAP 7-Mode	Clustered Data ONTAP
Vscan server	Controller IP added to the Vscan server configuration and the Vscan server initiates the connection. Detected automatically. Must specify secondary scanners explicitly.	Must create scanner pool so that the controller initiates the connection to the Vscan server.
File extension exclusion	Specified using <code>vscan extension exclude</code> .	Specified while creating <code>on-access-policy</code> .
File extension inclusion	Specified using <code>vscan extension include</code> .	Starting with 8.3.1, you have the option to add an include list while creating <code>on-access-policy</code> .
Shares to scan	Specified as a share property. Options: <code>novscan</code> – Disable scans for a share <code>vscan</code> – Enable scans for a share <code>vscanread</code> – Turn on scanning for read-only access <code>novscanread</code> – Turn off scanning for read-only access	Specified as a share property using <code>-vscan-fileop-profile</code> . Options: <code>no_scan</code> – None <code>standard</code> – Open, close, and rename <code>strict</code> – Open, read, close, and rename <code>writes_only</code> – Close (only for newly created or modified files)
Time-out values	Specified using <code>vscan options timeout set value</code> .	Specified using <code>SVM vscan scanner-pool modify</code> .
host_scanners	Can be configured to use the same scanners for all NetApp vFiler® units.	No concept of host scanners because of secure multitenancy setup. Same scanners can be used for all SVMs only if all SVMs are joined to the same domain/trusted domain.

Clustered Data ONTAP Best Practices for Off-Box Antivirus

Consider the following recommendations for configuring the off-box AV functionality in clustered Data ONTAP:

- Restrict privileged users to virus-scanning operations. Normal users should be discouraged from using privileged user credentials. This restriction can be implemented by turning off login rights for privileged users on AD.
- Privileged users are not required to be part of any user group that has a large number of rights in the domain, such as the administrators group or the backup operators group. Privileged users must be validated only by the storage system so that they are allowed to create Vscan server connections and access files for virus scanning.
- Use the computers running Vscan servers only for virus-scanning purposes. To discourage general use, disable the Windows terminal services and other remote access provisions on these machines and grant the right to install new software on these machines only to administrators.
- Dedicate Vscan servers for virus scanning and do not use them for other operations, such as backups. You might decide to run the Vscan server as a virtual machine (VM). If this is the case, make sure that the resources allocated to the VM are not shared and are enough to perform virus scanning. Consult the vendor for guidance on AV engine requirements.

- Provide adequate CPU, memory, and disk capacity to the Vscan server to avoid resource bottlenecks. Most Vscan servers are designed to use multiple CPU core servers and to distribute the load across the CPUs. Consult the vendor for guidance on AV engine requirements.
- NetApp recommends using a dedicated network with a private VLAN for the connection from the SVM to the Vscan server so that the scan traffic is not affected by other client network traffic. Create a separate NIC that is dedicated to the AV VLAN on the Vscan server and to the data LIF on the SVM. This step simplifies administration and troubleshooting if network issues arise.

Note: In a scenario in which the LIF for Vscan traffic is configured on a different port than the LIF for client traffic, the Vscan LIF might fail over to another node because of a port failure. This action will make the Vscan server unreachable from the node, causing the scan notifications for the file operations on the node to fail. Ensure that the Vscan server can be reached through at least one LIF on the node to process Vscan requests for the file operations performed on that node.

- Connect the NetApp storage system and the Vscan server by using at least a 1GbE network.
- For an environment with multiple Vscan servers, connect all servers that have similar high-performing network connections. Connecting the Vscan servers improves performance by allowing load sharing.
- For remote sites and branch offices, NetApp recommends using a local Vscan server rather than a remote Vscan server because the ROBO scenario is a perfect candidate for high latency. If cost is a factor, use a laptop or PC for moderate virus protection. You can schedule periodic complete file system scans by sharing the volumes or qtrees and scanning them from any system in the remote site.
- Use multiple Vscan servers to scan the data on the SVM for load-balancing and redundancy purposes. The amount of CIFS workload and resulting AV traffic varies per SVM. Monitor CIFS and virus-scanning latencies on the storage controller. Trend the results over time. If CIFS latencies and virus-scanning latencies increase because of CPU or application bottlenecks on the Vscan servers beyond trend thresholds, CIFS clients might experience long wait times. Add additional Vscan servers to distribute the load.
- Install the latest version of Antivirus Connector. For detailed information about supportability, refer to the NetApp [Interoperability Matrix Tool](#) (IMT).
- Always keep AV engines and definitions up to date. Consult the vendor for recommendations on update frequency.
- In a multitenancy environment, a scanner pool (pool of Vscan servers) can be shared with multiple SVMs provided the Vscan servers and the SVMs are part of the same domain or in a trusted domain.

Migration Methods

Currently, 7MTT doesn't support Vscan configuration migration. Gather the required data from the 7-Mode appliance and refer to Table 3 to get the configuration comparison. Accordingly, configure the settings on clustered Data ONTAP.

6.4 Implementation of CIFS Native Auditing

The native auditing framework in clustered Data ONTAP supports both CIFS and NFS protocols. Auditing in CIFS is based on New Technology File System (NTFS) ACLs (SACLs) or NFS 4.x ACLs.

The native auditing infrastructure provides features to securely generate and manage audit logs in a timely fashion along with file access monitoring support. Auditing is mainly used in organizations to meet compliance requirements.

Auditing for NAS events is a security measure that enables you to track and log certain CIFS and NFS events on storage virtual machines (SVMs) with NetApp FlexVol® volumes. This measure helps you to track potential security problems and provides evidence of any security breaches.

For more information, refer to [TR-4189: Clustered Data ONTAP CIFS Auditing Quick Start Guide](#).

6.5 CIFS Events

You can audit the following events:

- **SMB file and folder access events.** You can audit SMB file and folder access events on objects stored on FlexVol volumes belonging to the auditing-enabled SVMs.
- **CIFS log-on and log-off events.** You can audit CIFS log-on and log-off events for CIFS servers on SVMs with FlexVol volumes.
- **Central access policy staging events.** You can audit the effective access of objects on CIFS servers using permissions applied through proposed central access policies. Auditing through the staging of central access policies enables you to see what the effects are of central access policies before they are deployed. Auditing of central access policy staging is set up using AD GPOs; however, the SVM auditing configuration must be configured to audit central access policy staging events. Although you can enable central access policy staging in the auditing configuration without enabling the Windows Server Dynamic Access Control (DAC) feature on the CIFS server, central access policy staging events are generated only if DAC is enabled. DAC is enabled through a CIFS server option. It is not enabled by default.

Native Auditing Enhancements in Clustered Data ONTAP as Compared to 7-Mode

The following native-auditing enhancements are available in clustered Data ONTAP:

- **Event auditing.** Most of the events that could be audited in 7-Mode can also be audited in clustered Data ONTAP. There are more events added. For more details on event auditing, refer to the [Clustered Data ONTAP CIFS Auditing Quick Start Guide](#).
- **Licensing.** In 7-Mode, CIFS must be licensed for auditing to work. In clustered Data ONTAP, CIFS and NFS licenses are not required for auditing to work.
- **Unified view for logs.** Logs are generated and kept on a per-node basis and are then consolidated in one location.
 - **Staging files.** These files are the intermediate binary files on individual nodes where audit records are stored before consolidation and conversion. Staging files are contained in staging volumes.
 - **Consolidation task.** A consolidation task gets created when auditing is enabled. This long-running task on each SVM takes the audit records from staging files across the member nodes of the SVM. This task merges the audit records 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.

Native Auditing Configuration Comparison Between Data ONTAP 7-Mode and Clustered Data ONTAP

Table 4 shows the native auditing configuration comparison between 7-Mode and clustered Data ONTAP to make it easier to map the features.

Table 4) Native auditing configuration comparison.

Feature	Data ONTAP 7-Mode	Clustered Data ONTAP
Events	<ul style="list-style-type: none">• CIFS file and directory access• Log-on and log-off events• Local user and group account management	<ul style="list-style-type: none">• CIFS file and directory access• Log-on and log-off events• Central access policy staging events

Feature	Data ONTAP 7-Mode	Clustered Data ONTAP
Log file size and format	By default in .EVT format. Log file can be between 512K and 64GB. The default size is 524,288 bytes.	Can be .EVTX or .XML format. By default a single file is 100MB and multiple files can be created based on rotation policy.
Accessing audit logs	By default, the external event log is the <code>/etc/log/adtlog.evt</code> file. You can specify another file as the event log. If the specified file does not already exist, Data ONTAP creates the file when it saves information to the file. The directory containing the file, however, must exist; otherwise, an error message appears when you specify the file.	Access to audit logs is through a pull mechanism and the logs are retrieved over NFS, CIFS, or another file access protocol method. Audit logs are not integrated with the syslog framework so logs cannot be accessed through the push mechanism.
Log file rotation	Audit events can be saved automatically to the event log based on a time interval or the size of the internal log file; that is, how full the <code>CIFSaudit.alf</code> file is. If you specify both a size threshold and a time interval, audit events are saved to the event log whenever the size threshold or the time interval is reached. The default for the size threshold is 75%. The default for the time interval is one day. Each time the internal log file is automatically saved to the external event file, an extension is added to the base name of the event file. You can select one of the following types of extensions to be added: <ul style="list-style-type: none"> • counter • timestamp 	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 then creates a new active converted event log file.
Guaranteed auditing	Not present in 7-Mode	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 be completed 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.
Live view	Can be configured in 7-Mode	NA

Considerations Before Planning a Migration

Before you configure and enable auditing on an SVM, be aware of the following requirements and considerations:

- Make a note of the SACLS configured in the existing 7-Mode environment configured for auditing on the files on directories.
- The maximum number of auditing-enabled SVMs supported in a cluster is 50.

- Auditing is not tied to CIFS or NFS licensing, unlike 7-Mode. You can configure and enable auditing even if CIFS and NFS licenses are not installed on the cluster.
- 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 depends 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 depends on having available space in the volume containing the directory where converted 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 event logs to retain in the auditing directory by using the `-rotate-limit` parameter when creating an auditing configuration. Doing so can help to ensure that there is enough available space for the event logs in the volume.
- Staging volume is 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 SVMs to store audit records of data access for data volumes in that particular aggregate. Each SVM audit record is stored in a separate directory within the staging volume. Cluster administrators can view information about staging volumes, but most other volume operations are not permitted. Only clustered Data ONTAP can create staging volumes. Clustered Data ONTAP automatically assigns a name to staging volumes. All staging volume names begin with `MDV_aud_` followed by the UUID of the aggregate containing that staging volume (for example, `MDV_aud_1d0131843d4811e296fc123478`).
- DAC must be enabled to generate central access policy staging events. DAC is not enabled by default.

Migration Methods

Currently the 7MTT doesn't support auditing configuration migration. Please gather the required data from the 7-Mode appliance and refer to Table 4 to get the feature comparison. Configure the settings accordingly on clustered Data ONTAP.

References

- [Best Practices Guide for Clustered Data ONTAP 8.2.x and 8.3.x Windows File Services](#)
- [TR-4189: Clustered Data ONTAP CIFS Auditing Quick Start Guide](#)
- [File Access Management Guide for CIFS – 8.3](#)
- [File Access and Protocol Management Guide](#)

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>. TR-4522-0616