Technical Report

Best Practices Guide for Clustered Data ONTAP 8.2 Windows File Services

Brahmanna Chowdary Kodavali, Marc Waldrop, Sharyathi Nagesh, Dhaval Bhadeshiya, NetApp

December 2013 | TR-4191

Abstract

Windows® File Services on clustered NetApp® Data ONTAP® 8.2 or 8.2.x brings new use cases and features after the release of version 8.1. This technical report covers these new features and best practices.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1 Introduction

Clustered Data ONTAP was introduced to provide more reliability and scalability to the applications and services hosted on Data ONTAP. Windows File Services are one of the key value propositions of clustered Data ONTAP because they provide services through the Server Message Block (CIFS/SMB) protocol.

Clustered Data ONTAP 8.2.x brings added functionality and features to Windows File Services. This technical report presents an overview of the new features for Windows File Services in the latest versions of clustered Data ONTAP.

1.1 Intended Audience

This technical report is for IT administrators, solution architects, technical architects, professional service engineers, and system engineers.

1.2 Purpose and Scope

This technical report provides a brief overview of SMB implementation and other Windows File Services features with recommendations and basic troubleshooting information for clustered Data ONTAP 8.2 and 8.2.x.

Note: For configuration and best practices for features introduced prior to Data ONTAP 8.2, refer to TR-3967: Deployment and Best Practices Guide for Clustered Data ONTAP 8.1 Windows File Services.

Note: For information about feature details and services procedures on NetApp storage systems, refer to the Data ONTAP 8.2 File Access and Protocol Management Guide on the NetApp Support website.

2 Overview of Windows File Services in Clustered Data ONTAP 8.2

Clustered Data ONTAP 8.2 complements its earlier version with additional features and added capacity for new use cases.

2.1 Windows File Services Features in Clustered Data ONTAP 8.2

Following are the new features introduced in Windows File Services in 8.2:

- Server Message Block (SMB) 3.0
- Copy offload (ODX)
- Node referrals (SMB autolocation)
- Remote VSS
- BranchCache®
- Local users and groups
- FPolicy® native file blocking and partner use case support
- File access auditing
- File-directory (FSecurity)
- Access-based enumeration
- Microsoft® previous versions support
- Roaming profiles and folder redirection
- Offline folders (client-side caching)

- SMB signing

2.2 Server Message Block Version 3.0

SMB 3.0 is the revised version of the SMB 2.x protocol, introduced by Microsoft in Windows 8 and Windows Server® 2012. The SMB 3.0 protocol offers significant enhancements to the SMB protocol in terms of availability, scalability, reliability, and protection. Enhancements to the SMB protocol with version 3.0 open up new use cases in the enterprise application segments and support applications such as Hyper-V® and SQL Server®.

Clustered Data ONTAP 8.2 implements SMB protocol version 3.0 and the following optional protocol features to support the Hyper-V over SMB use case. Without these optional features, SMB 3.0 functions as a minor protocol revision of SMB 2.1.

The optional SMB 3.0 features for supporting the Hyper-V over SMB use case are:

- **Continuously available shares (CA shares).** Enable high availability for file shares that can be accessible during failures and controller failover scenarios. For this feature, apply the following property to the file share:

```
-share-properties continuously-available
```

- **Persistent handles.** Persistent handles are an enhancement to the durable handle that was introduced in SMB 2.0. In the case of durable handles, the server preserves the file handle and allows the client to reconnect to the file after a brief network outage. The challenge with the durable handle is that if any other client attempts to access the same file, the file will invalidate the previous durable handle opened by the first client.
  Persistent handles solve this challenge by allowing the server to preserve the file handle opened during the file open for a predetermined time period after a network failure. During the predetermined time period, any client other than the client that has the persistent handle cannot get a handle on the file. After the client reestablishes the connection with the controller, the client can reclaim the file handle.
- **Witness protocol.** This allows notification to the client about storage-side failovers so the client connection can be proactively moved to the partner node prior to the actual failover event.
- **Cluster client failover.** Enables cluster-capable applications to close the stale “opens” and “locks” during failure recovery and applications to reclaim their file handles after moving over to a new node.
- **Request replay.** Request replays are protocol extensions that handle replay of nonidempotent requests in the event of a network failure.

Performance

- To achieve better resiliency, the lock states of the persistent handles are mirrored to the storage failover partner. When the lock state of a persistent handle changes, that state will be mirrored to the partner node. Because each node maintains the lock state with persistent handles, maintaining a copy of the partner node’s state reduces the number of locks to half to accommodate the partner node’s lock state information.
- SMB 3.0 is an enhancement or extension to the SMB 2.x protocol. SMB 2 counters can be used for troubleshooting protocol issues.

Verification

The following command helps to verify if a file is opened with a persistent handle:

```
vserver locks show -smb-attrs
```

This command lists all the locks in the storage virtual machine (SVM, formerly known as a Vserver). In the command output, if the value for “`Open Type`” is “`persistent,`” it means the file is opened with a persistent handle.

Recommendations

- The CA property should be enabled only on shares hosting Hyper-V virtual machines.
- Troubleshooting and best practices for witness protocol:
  - At least one data LIF must be present on each node per SVM.
  - CA shares should not be mapped using an IP address. They should be mapped using NetBIOS or the fully qualified domain name (FQDN).
  - The node referrals (SMB autolocation) feature should be turned off. It is not supported because this option refers the client to an IP address, and authentication will fall back to NTLM. Hyper-V deployment heavily relies on Kerberos authentication.

2.3 Copy Offload

One of the challenging tasks for most data center administrators is moving or copying data across servers. The effect of this task is high if the dataset is large.

Traditional host operating systems are designed to run or handle applications efficiently, but not designed for data movement. Storage devices or systems are designed for managing data efficiently. Microsoft introduced the offloaded data transfer (ODX, or “copy offload”) feature with SMB 3.0 to leverage the data management capabilities of the storage systems. This feature offloads the copy operation to the storage so the storage system can efficiently move or copy more quickly than the host-side copy operation.

As part of SMB 3.0, the copy offload feature is implemented in clustered Data ONTAP 8.2. The unique advantage with NetApp’s implementation of copy offload is that it supports cross-protocol copy operations. That means that only clustered Data ONTAP provides the capability of copying data between CIFS and block (FC or iSCSI) using the copy offload functionality.

Table 1 lists the scenarios supported by copy offload.

Table 1) Copy offload types and technology used.

Type	Technology
Intravolume	SIS clone
Intervolume within same node	Internal replication (block copy) engine
Intervolume on different node	Internal replication (block copy) engine over cluster network

Performance

Performance of copy offload varies depending on the scenario. However, its performance is better than that of traditional host-side copy in any case.

Here is an order based on the copy throughput. Copy within the same data volume is the fastest.

- Performance of file copy on the same data volume
- Performance of file copy across data volumes on the same node
- Performance of file copy across data volumes on different nodes
- File copy without copy offload

Verification

Steps for troubleshooting copy offload:

1. SMB 3.0 is enabled on the SVM.
2. Both CIFS copy offload and SVM scoped subfile-sisclone (enabled by default) options should be enabled.

3. Client must support copy offload (Windows 8 or Windows Server 2012).
4. Source data volume must be at least 2GB and cannot be a read-only/compressed/sparse volume.

Recommendations

The first copy offload operation creates a scratch space called a point-in-time (PIT) file on the source data volume. The PIT file uses 6.25% of the total volume and a maximum of up to 16GB on platforms with memory less than 6GB. For platforms with greater than 6GB memory, the PIT file size can go up to 65GB.

- Copy offload performs best with data volumes of size greater than 260GB due to the maximum number of tokens (2,048) available to perform the copy operation.
- Max PIT file size (16GB)/token size (8MB) = 2,048.
- The minimum data volume size for copy offload function to work is 2GB.
- To copy the same file multiple times, first copy the file from source to destination data volume and then create the subsequent copies within the data volume. This will increase the copy throughput.

2.4 Node Referrals (SMB Autolocation)

Because SVMs span multiple nodes, clustered Data ONTAP allows clients to access the data from any node in the cluster. This is to increase resource availability to the CIFS clients. Applying this feature opens the possibility that the CIFS clients may access the resources from an unoptimized path, as shown in Figure 1.

Figure 1) Data access path.

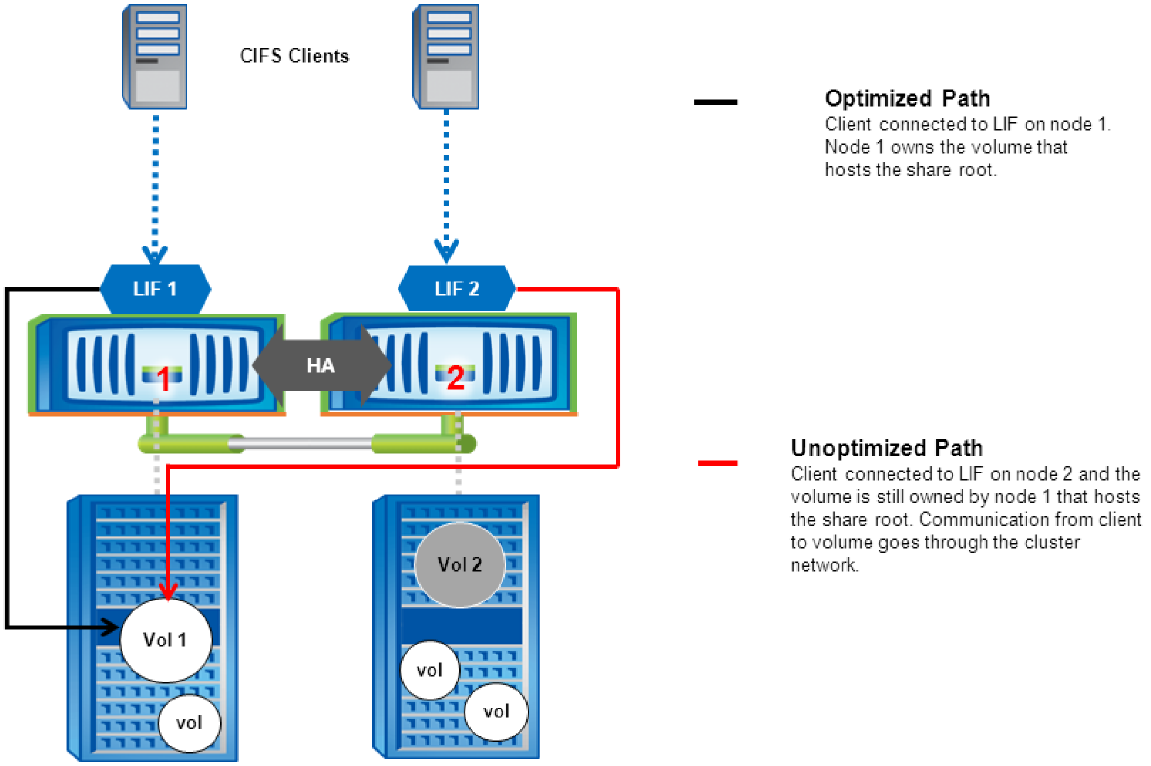
CIFS Clients
LIF 1
LIF 2
1
HA
2
Vol 1
vol
Vol 2
vol
vol
Optimized Path
Client connected to LIF on node 1. Node 1 owns the volume that hosts the share root.
Unoptimized Path
Client connected to LIF on node 2 and the volume is still owned by node 1 that hosts the share root. Communication from client to volume goes through the cluster network.

Node referrals address this challenge by redirecting the client connection to the appropriate node owning the volume that hosts the root of the CIFS share. This is to make sure the data is accessed locally instead of being accessed over the cluster network. Node referral happens during the connection initiation phase.

Performance

Clients are given referrals based on the data volume that hosts the root of the share. Node referrals perform best if the complete data under the CIFS share is hosted on a single node.

Verification

Steps for troubleshooting node referral issues:

1. Check if LIF is migrated to different node than the node that owns the data volume.
2. Check if data volume is moved to the different node but no data LIF is configured.
3. Use the following performance counters to identify the node referral statistics:
    a. node_referral_local displays the number of clients connected using a LIF hosted by the same node that hosts the share root.
    b. node_referral_remote displays the number of clients connected using a LIF hosted by a node different from the node that hosts the share root.
    c. node_referral_issued displays the number of clients that were provided a referral to a LIF local to the node that hosts the share root when the client is connected from a remote node.
    d. node_referral_not_possible displays the number of clients that were not issued a referral in spite of using a nonoptimal LIF because no active data LIF was found on the node that hosts the share root.

Recommendations

Although node referrals allow clients to connect to an optimized path, NetApp recommends avoiding the use of node referrals when one or more of the following are true:

- A share that contains volume junctions, symlinks, and widelinks points to the data volumes or shares hosted on different nodes. This will cause all clients on that share to connect to one node, thus lowering performance through the cluster interconnect.
- Shares are mapped using specific IP addresses for applications. Clients might get referred to an IP address different than the IP address mapped in the application.
- Share is a CA share. Hyper-V over SMB does not support node referrals.
- Environment is one in which Kerberos authentication is enforced. Node referral provides an IP address to the client to connect to the controller. This will force Windows clients to fall back onto NTLM authentication, and the authentication will fail.

2.5 BranchCache

BranchCache was originally introduced with Windows Server 2008 R2. BranchCache addresses the challenges that enterprise users face in accessing data over the WAN.

Data access over the WAN is slow due to higher network latencies or slow links. Branch office or remote office environments typically face problems such as slow WAN links to the main office.

BranchCache was first introduced in Data ONTAP 8.1.1 for 7-Mode. With this release, BranchCache is introduced in clustered Data ONTAP with a few enhancements over the 7-Mode implementation.

Data ONTAP acts as a content server in the overall BranchCache solution. Data ONTAP does not implement any components of hosted cache or distributed cache in the remote office itself.

Table 2) BranchCache implementation differences between 7-Mode and clustered Data ONTAP.

Configuration Option	BranchCache in 7G/7-Mode	BranchCache in Clustered Data ONTAP
Hash location	In memory	On permanent storage

Configuration Option	BranchCache in 7G/7-Mode	BranchCache in Clustered Data ONTAP
Hash store location	In memory	Configurable on user-specified path
Hash store size	Depends on storage memory	Configurable from 1GB to 1TB
Manual hash generation	No	Yes

Performance

- Client-side configuration of BranchCache must be configured on Windows clients. Windows 7 Ultimate/Professional and Windows 8 clients support BranchCache.
- Hosted cache mode will have an advantage over distributed cache mode because a dedicated server is available to serve the cache information.

Verification

Check the hash store path and hash store size configuration for the following issues:

- Hash is flushed frequently.
- Client is accessing the data from the source instead of accessing from the cache.

Run the command `vserver cifs branchcache show` to view the current setting applied:

```
vserver cifs branchcache show

               Operating  Allowed     Max
Vserver        Mode       Versions    Size   Path
-------------- ---------- ----------- ------ ------------------------
cifsvs1        per_share  enable_all     1GB /datavol1
```

Recommendations

- Configure the hash store size based on the data change rate. Minimum size of the hash store is 1GB, and each terabyte (TB) of content will require 512MB cache size.
- Configuration on the client is appropriately set based on the caching mode (hosted or distributed) used in the organization.
- Configure the BranchCache operating mode to “All-Shares” if all the shares under the CIFS server are intended for providing BranchCache content. This will reduce the effort in configuring each share with BranchCache property. Default value is “per-share.”

2.6 Local Users and Groups

Active Directory® is a centralized database for providing authentication and authorization services for enterprise users. Active Directory is a standard LDAP server for managing the users and systems in Microsoft Windows environments.

Certain applications demand specific privileges for data access locally on the data source. To address this demand, Windows provides users and groups with management locally on the system to restrict the permissions for the domain user at the file server level.

Data ONTAP acting as a Windows file server also provides this capability through the local users and groups (LUG) feature. With clustered Data ONTAP 8.2, an administrator can create or customize a user or a group with specific privileges. After a group is configured with specific privileges, domain users can be added as members of the group. That will associate the group privileges to the member domain users.

Use cases for this feature are:

- Backup applications running with domain account need access to the resources (shares, files, and folders) on the controller.
- Large numbers of domain users need access to the shares, and each user group needs different permissions on the shares.
- Enterprise applications such as Microsoft SQL Server need specific privileges such as SeSecurityPrivilege on the local system.

Table 3) Privileges supported by clustered Data ONTAP 8.2.

Privilege	Description
SeTCBPrivilege	Act as part of the operating system
SeBackupPrivilege	Back up files and directories, overriding any access control lists (ACLs)
SeRestorePrivilege	Restore files and directories, overriding any ACLs
SeTakeOwnershipPrivilege	Take ownership of files or other objects
SeSecurityPrivilege	Manage auditing and security log
SeCreateSymbolicLinkPrivilege	Create symlinks (available with Windows Vista® and later Windows clients)

Performance

Local authentication is faster than domain authentication because all the necessary information is local and does not need to communicate with an external server for authentication.

To improve the performance authentication and authorization, two levels of cache are introduced to store user-specific information:

- **Level 1**: `ad-sid-to-local-membership`

  This caching option will store domain SID to its local group membership mappings and its cumulative privileges.
- **Level 2**: `username-to-creds`

  This caching option stores the domain user to its cumulative local group membership and cumulative privileges. It is populated only if the domain user is a member of 50 or more groups.

When a domain user authenticates for the first time, there might be a very slight performance degradation (depending on the local group memberships) to build credentials. Subsequent authentication requests will be processed with the cache information available locally on the controller.

Verification

Steps for troubleshooting of authentication and privilege issues:

Authentication Failures with Local Users or Group Privileges

- Enable options `-is-local-auth-enabled` and `-is-local-users-and-groups-enabled`.
- Make sure all the nodes are upgraded to clustered Data ONTAP 8.2 that supports local users and groups.

Authentication Failures for Domain Account

- Make sure the domain account is not disabled. If the account is disabled, an attempt to look up the user locally will be performed. This will result in a local authentication error because the user is not available in the local database. This is typical Windows behavior.

- Check if the user account is from a trusted domain. Validate the trust between the domain to which the controller joined and the user's domain.

Recommendations

- Local users and groups are not a direct alternative to Active Directory domain authentication. For larger numbers of users, it is recommended to use Active Directory instead of local users and groups for better manageability.
- If customization is required for a large number of users with a specific set of privileges, then:
  - Configure a group with the specific set of privileges.
  - Add the users to the group. It will save the effort of having to customize individual users and privileges that instead can be granted to "groups."

2.7 File-Directory (FSecurity)

File-directory is equivalent to the 7-Mode feature FSecurity. File-directory implements all the features of FSecurity in clustered Data ONTAP except storage-level access guard (SLAG). File-directory in clustered Data ONTAP allows administrators to create permission sets and apply them from the controller. It has removed the dependency of the external application to create the job description file and apply on the controller using FSecurity.

Performance

While applying bulk security settings repetitively on directory and files, using file-directory instead of any client-side tool will significantly enhance performance. File-directory makes it easier to manage and standardize the user permissions on the storage through the security descriptor template.

Current implementation does not support configuring NFSv4 access control entries (ACEs).

Verification

Run the following command to view current status of the file-directory jobs:

```
vserver security file-directory job show
```

Recommendations

Working with other features and commands:

FSecurity or file-directory should not be used in these potential race conditions:

- During changes to SVM (Vserver) global namespaces
  Example: Volume mount and unmount operations
- During high workload
- During volume move operations

Working in Multiprotocol Scenario

Exercise caution when applying permissions on mixed-mode volumes and folders because this will change the existing access permission and convert it into NTFS security types. This might cause access disruptions.

Applying Advanced Inheritance Options

The ACE inheritance field specifies how ACEs will be propagated to subfolders, folders, and files. This can be specified with *–apply-to* in the security descriptor.

NetApp recommends the following inheritance modes:

- This folder, subfolders, and files
- This folder, subfolders

NetApp does not recommend using the following combination of inheritance modes:

- Subfolders and files only
- Subfolders only

Using Advanced Features

- **Control flags.** Control flags are useful when more control is required on the ACE inheritance. Control flags will control both:
  - Inherited from the parent folder
  - Propagated to child objects (files, folders)

  The values for these control flags can be found in Microsoft documentation. Because this is an advanced option, NetApp recommends that only an advanced user should modify these values.

  For the values of control flags, refer to security_descriptor_control.
- **Rights raw.** This allows hexadecimal values specified by Microsoft to define access rights to files and folders. The values (for example, managing the security descriptor) can be obtained from Microsoft documentation. This feature is available under advanced mode.

2.8 FPolicy

Native File Blocking

FPolicy native file blocking is equivalent to the 7-Mode native file blocking feature. This feature allows administrators to screen the files stored by end users based on file extensions and block them (based on the corporate data policy).

For example, if corporate policy defines that MP3 or MP* files cannot be stored on the storage, it can be implemented using native file blocking.

Performance

Native blocking processes all requests on the storage itself instead of sending and processing on the external server. This improves the user experience by reducing response time.

User data can be monitored or blocked based on the policy scope configuration. The following parameters describe the policy and define the data:

- shares-to-include
- shares-to-exclude
- volumes-to-include
- volumes-to-exclude
- export-policies-to-include (NFS only)
- export-policies-to-exclude (NFS only)
- file-extensions-to-include
- file-extensions-to-exclude

Verification

To check the policy information, execute the following command:

```
vserver fpolicy policy show

Vserver      Policy       Events     Engine        Is Mandatory  Privileged
               Name                                                Access
-------------- -----------  ---------- ------------- ------------  -----------
cifsvs1        fps1         evnt1      native        true          yes
```

To verify the extensions included in the scope, use the following command:

```
vserver fpolicy policy scope show

Vserver           Policy              Extensions           Extensions
Name              Name                Included             Excluded
----------------- ------------------- -------------------- -------------------
cifsvs1           fps1                mp?                  doc
```

To verify the file operations that are being monitored or filters applied, use the following command:

```
vserver fpolicy policy event show

          Event                        File                       Is Volume
Vserver   Name               Protocols Operations    Filters      Operation
--------- ------------------ --------- ------------- ------------ ------------
cifsvs1   evnt1              cifs      create,       -            false
                                       create_dir
```

Recommendations

- Policy scope should be properly configured. Leaving the options “Extensions included” and “Extensions excluded” blank will result in monitoring every file and blocking the user access to the file.
- “Extensions excluded” has priority over “Extensions included.” If a file extension is mentioned in both, then the file will not be blocked. Policy will not be enforced until one of the following scope options is configured:
  - shares-to-include
  - shares-to-exclude
  - volumes-to-include
  - volumes-to-exclude
- The option “-is-file-extension-check-on-directories-enabled” (advanced privilege) will enable the policy to monitor the files under the directories. The default setting for the options is “False,” and that means that no file under the directories will be monitored. Changing the setting to “True” is recommended.
- File operations that should be monitored for file blocking are “open, create, and rename.”

Partner Solutions

FPolicy supports advanced features working with the partner solution (for example, storage management, access management and data governance, quota, archiving, file replication, and others). These solutions require deploying external FPolicy servers along with clustered Data ONTAP. These are generic best practices independent of FPolicy servers, but many of these recommendations require end users to work with FPolicy partners.

Performance

Partner solutions require deploying an FPolicy server to process FPolicy notifications generated on the storage. Because of this round trip, overall client experience will be affected, and monitoring performance counters can help to identify resource bottlenecks:

- Increasing FPolicy latency will have a cascading effect on CIFS latency. Monitor both workload (CIFS) and FPolicy latency.

- Use Data ONTAP quality of service (QoS) to set up a workload for each volume/SVM that has FPolicy enabled.
  - Display these workload statistics: statistics show –object workload
  - Monitor these counters: average, read, and write latencies; total number of operations, reads, and writes
- Use Data ONTAP FPolicy counters to monitor performance of FPolicy subsystem:
  - Statistics show –object fpolicy_server –instance SVM:servername

  a. request_sent_rate

  b. request_latency

  c. response_received_rate

Verification

User needs to verify if policy events are configured properly:

```
fpolicy policy event show -vserver <vserver name> -event-name <event name> -instance
```

User needs to verify if the policy scope is configured properly:

```
fpolicy policy scope show -vserver <vserver name> -policy-name <policy name> -instance
```

User needs to verify if the safeguards are configured properly and executed in advanced mode:

```
fpolicy policy external-engine show -vserver <vserver name> -engine-name <engine name> -instance
```

Recommendations

Follow FPolicy application best practices for server hardware, operating system, patches, and so on.

Policy Configuration

Configuration of FPolicy external engine for the SVM (Vserver):

- Providing additional security comes with a performance cost. Enabling SSL communication will have a performance impact on CIFS.
- If the FPolicy server resources are underutilized, NetApp recommends increasing the value of `‘Request Queue Length’`. This will improve CIFS throughput when the external engine is of sync type. Caution has to be exercised while increasing this value. A very high value leads to CIFS timeouts and high CIFS latencies. Users have to find the optimum value working with partners. The value can be changed in the advanced mode:

```
fpolicy policy external-engine modify –vserver <vserver> -engine-name <engine>  -max-server-
requests  <new value>
```

Configuration of FPolicy event for the SVM:

- Monitoring file operations have an effect on the overall client experience.
- Filtering unwanted file operations on the storage side improves the overall client experience.
- NetApp recommends monitoring minimum file operations and enabling the maximum number of filters without breaking the use case. Work with a partner for optimum value.
- 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 the SVM:

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

Hardware Configuration

- Networking:
  - Network connectivity between the FPolicy server and the controller should be of low latency.
- FPolicy server:
  - The FPolicy server can be on either a physical server or a virtual server.
  - If the FPolicy server is on a virtual server, make sure that enough CPU, network, memory, and disk resources are allocated. Users have to find the optimum value when working with partners.

Multiple Policy Configuration

- FPolicy policy for native blocking has the highest priority irrespective of the sequence number.
- Decision-altering policies should have higher priority than other policies.
- Policy priority depends on use cases. To determine the appropriate priority, working with partners is recommended.

FPolicy Safeguards

Safeguards are tunable and are provided within the FPolicy framework to handle performance and connection disruption issues between the SVM (Vserver) and the FPolicy application. These act as levers to change FPolicy behavior. The tunables are part of the FPolicy external engine object, and they can be configured in advanced mode.

Table 4) FPolicy tunable parameters.

Tunable	Default Value (Sec)	Description	When Applicable
`request-cancel-timeout`	20	Measures the time that SVM waits for a response from FPolicy server. Beyond this timeout, the client operation is forwarded to an alternate server, if it exists, or access will be allowed/denied based on the mandatory attribute.	Manages CIFS latency when you have a slow FPolicy server.
`request-abort-timeout`	40	Measures the time CIFS client request spends on SVM. Beyond this timeout, the client operation is allowed/denied based on the mandatory attribute.	Manages CIFS latency when you have a slow SVM.
`status-request-interval`	10	Intervals at which SVM sends queries on pending requests to the FPolicy server.	Works with partner to set the optimum value to manage a slow FPolicy server.
`max-connection-retries`	5	Storage appliance (SVM) attempts sending keep-alive requests before deciding that the FPolicy server has gone bad.	If disruption is due to network issues, increase the value.

Tunable	Default Value (Sec)	Description	When Applicable
`max-server-requests`	50	Maximum number of outstanding screen requests that will be queued for an FPolicy server.	An optimum value to be found working with partners. Increase the value when FPolicy server can handle more notifications, but increasing it beyond a particular value increases CIFS latency and CIFS timeouts.
`server-progress-timeout`	60	When internal FPolicy specific queues are full and no response is received from the FPolicy server for this time, the connection between SVM and FPolicy server will be disconnected.	Required to minimize client disruption due to slow FPolicy server.
`keep-alive-interval`	120	SVM sends keep-alive messages to FPolicy server to detect half-open connections.	This is useful when the client traffic is nonexistent and FPolicy server is disconnected.

2.9 Managing FPolicy Workflow and Dependency on Other Technologies

- Disabling the FPolicy policy before making any configuration changes is recommended. For example, if you want to add or modify an IP address in the external engine configured for the enabled policy, you should first disable the policy.
- If you configure FPolicy to monitor FlexCache® volumes, it is recommended that you do not configure FPolicy to monitor read and getattr file operations on the FlexCache volumes. Monitoring these operations in Data ONTAP needs to retrieve inode-to-path (I2P) data. Because I2P data cannot be retrieved from the FlexCache volume, it has to be retrieved from the origin volume. This results in performance benefits from FlexCache not being realized.
- When both FPolicy and an offbox AV solution are deployed, the AV solution will get the notification first. FPolicy processing will start only after AV scanning is completed. A slow AV scanner could affect overall performance, and hence the AV has to be sized properly.

2.10 Roaming Profiles and Folder Redirection

User profiles on Windows machines separate each user’s settings from another user’s settings and the local computer. Each user profile is stored locally on the system and keeps each user’s settings in a separate user profile folder.

Roaming profiles are a type of user profile that allow enterprises to store the user profile in a remote and centralized location so that a user can get access anywhere to profile configuration information such as desktop settings, documents, and so on.

“Folder redirection” is a client-side feature along with roaming profiles to keep all the user data (pictures, documents, videos, and similar formats) away from the local desktop so that the data can be accessed from anywhere.

Performance

There might be a delay while the user logs into the system because the user settings have to be opened from the remote network share.

If multiple users are configured with roaming profiles and they are logging in at the same time, this results in multiple clients downloading the user settings from the network share. This is called a logon storm. This can be addressed by using Flash Cache™. Flash Cache keeps the cache of frequently accessed files so that the data can be served quickly without adding much load onto the physical disk.

Recommendations

There are no specific recommendations from NetApp because it is a Windows configuration, and it is recommended to refer to Microsoft Best Practices for User Profiles.

2.11 Access-Based Enumeration

Access-based enumeration (ABE) displays only the files and folders that a user has permissions to access. If a user does not have "read" (or equivalent) permissions for a folder, Windows hides the folder from the user's view. This is beneficial for large directories with many people accessing them.

Performance

Enabling ABE will have a performance impact on the content enumeration because an additional permissions check is done to hide the content from users who do not have access to the data.

Verification

Steps for troubleshooting ABE feature:

1. Verify ABE feature is enabled on the share through share properties.
2. Verify the specific user has read permission on the folder or file.
3. Verify the trust between the domains if the user is trying to access the resources on an SVM (Vserver) that is part of a different domain. Any issues with the trust will result in authentication failures.

Recommendations

Avoid having too many objects under any folder. File enumeration takes longer if too many objects are under the folder.

2.12 Microsoft Previous Versions Support

Microsoft previous versions support is a feature that makes previous versions of files or folders on a network drive available to the CIFS user. The user can choose to browse through the previous versions or to restore from them.

This feature allows end users to restore their data from the folder properties tab without the storage administrator's intervention.

Performance

There are no performance implications.

Verification

Steps to troubleshoot "previous versions" issues:

1. Snapshot™ copies are scheduled and are available for the volume.
2. Snapshot directory access is enabled on the volume.
3. Volume modify –volume datavol1 –snapdir-access true.

4. Restoring from the path that crosses junctions: restore operation can fail with an error if directories have junctions underneath them.

Recommendations

There are no specific recommendations.

2.13 Offline Folders

“Offline folders” is a Windows client-side feature that uses the `client-side caching` feature. The offline file caching allows a user to work with network files and programs even when the user is not connected to the network.

When the user makes network files available offline, the corresponding files will be cached locally on the user’s computer so that the files can be accessible during a network outage as well. When the network connection is restored, any changes to the data in the share that is marked offline will be synchronized with the server copy.

Performance

Offline folders improve the user experience by caching the network files locally so that the user can access the data even when the network is not available. After the data is cached, the data is served from a local machine.

Recommendations

Configure the offline files options on the share before connecting to the share. If offline files are not configured on the share using -offline-files, then by default manual file caching will be set on the share.

This is a Windows client feature, and there is no specific recommendation from NetApp.

For more information, refer to http://windows.microsoft.com/en-IN/windows7/Understanding-offline-files.

2.14 SMB Signing

SMB signing makes sure that the transmission and reception of data across a network are not altered by any method. Traditional SMB authentication without SMB signing is vulnerable to man-in-the-middle attacks. To avoid these kinds of issues, secure transmission of SMB traffic might be required.

Implementing mutual authentication, SMB signing protects data over the network from these attacks by adding a digital signature to each SMB packet.

Performance

There will be a performance impact because each message has to be signed and verified by either the server or the client to make sure the message is originated from the rightful source.

Recommendations

Enable SMB signing if security is a key requirement in the organization.

2.15 Remote VSS

Remote VSS is a new feature introduced to protect the data on a remote share. Remote VSS allows VSS-aware backup applications to create a volume shadow copy of VSS-aware applications that stores data on remote SMB 3.0 file shares.

In a Hyper-V over SMB scenario, the virtual machines are stored on an SMB 3.0 CA share hosted on clustered Data ONTAP. Without remote VSS, it is not possible to take a backup of virtual machines on the remote SMB share.

For more information, refer to TR-4172: Microsoft Hyper-V over SMB 3.0 with Clustered Data ONTAP: Best Practices.

2.16 File Access Auditing or File Access Monitoring

The native auditing framework provides a file auditing framework that supports both CIFS and NFS protocols. Auditing in CIFS is based on New Technology File System (NTFS) system 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.

For more information, refer to TR-4189: Clustered Data ONTAP CIFS Auditing Quick Start Guide.

3 Overview of Windows File Services in Clustered Data ONTAP 8.2.1

3.1 LDAP over SSL

Clustered Data ONTAP needs to communicate with external systems for the purpose of completing user authentication. This communication may be with Active Directory domain controllers or LDAP servers. In many cases this communication occurs in clear text. The communications in clear text may include user credentials and other critical information about an environment. This means it is possible to use a network monitoring device to view the communication between LDAP clients and servers. This is particularly of issue when LDAP simple bind is used, because the credentials (user name and password) are passed over the network unencrypted. This type of exchange can quickly lead to the compromise of credentials and is a security vulnerability.

Secure Sockets Layer (SSL) is a secure protocol developed for sending information securely over the Internet. The data integrity can be taken care by the application or user using SSL. SSL can provide encryption of the data in transit, as well as mutual authentication. In clustered Data ONTAP 8.2.1, NetApp only supports securing authentication (not mutual authentication) and encrypting the exchange of data.

In clustered Data ONTAP 8.2.1, two options are available for enabling LDAP over SSL. The two options are mutually exclusive of each other. One can be enabled without the other, or both can be enabled, depending on the needs of the environment:

- **LDAP over SSL for user mapping.** All Windows accounts must map to a UNIX user. If an environment utilizes an LDAP server to house UNIX user accounts and clustered Data ONTAP is configured to use LDAP servers for user mapping, then you can enable SSL for communication to those LDAP servers.
- **LDAP over SSL for Active Directory LDAP.** After you set up a CIFS server and make it a member of an Active Directory domain, clustered Data ONTAP uses LDAP for CIFS server metadata work. Enabling an additional option in clustered Data ONTAP will allow this metadata work to be securely exchanged.

Performance

There is no expected performance degradation by the use of this feature.

Verification

Verify that a certificate is properly installed:

```
security certificate show –vserver <vserver_name> -type server-ca
```

LDAP over SSL requires certificate services to exist in the environment. Without a valid certificate from a certificate authority, installing LDAP over SSL is not possible.

Confirm whether the option for LDAP over SSL is enabled for user mapping:

```
vserver services ldap client show -client-config <config_name> -fields use-start-tls
```

Confirm whether the option for LDAP over SSL is enabled for Active Directory communication:

```
cifs security show -vserver <vserver_name> -fields use-start-tls-for-ad-ldap
```

Recommendations

If your environment requires a more secure exchange of LDAP data, modify the following options:

LDAP over SSL for user mapping:

```
vserver services ldap client create  vserver <vserver_name> … -use-start-tls true
```

LDAP over SSL for Active Directory LDAP:

```
vserver cifs security modify –vserver <vserver_name>  -use-start-tls-for-ad-ldap true
```

By default, both of the preceding referenced options are set to disabled. Note that if you configure LDAP over SSL, your connections to the LDAP server will need to succeed using SSL connections. There is no fallback to a non-SSL connection if you install a certificate and then enable the options.

For complete details on setting up LDAP and installing the SSL certificate, review the “File Access and Protocols Management Guide” for clustered Data ONTAP 8.2.1. Setting up and configuring LDAP over SSL is beyond the scope of this technical report. You may also find additional details in TR-4073: Secure Unified Authentication with NetApp Storage Systems.

3.2 Multidomain Name Mapping

In a multiprotocol environment, it is possible that users from both Windows and UNIX will need to access data that is secured with a security style unlike the type that matches the type of the client from which they are accessing the data, for example, a Windows user accessing NTFS-style data compared to a Windows user accessing UNIX® security-style data. The same applies to a UNIX user accessing UNIX data compared to a UNIX client accessing NTFS-secured data. To accomplish this, a process called user mapping occurs during the initial connection by the client. Multidomain name mapping is only relevant when a UNIX user attempts to access datasets that are secured with an NTFS-style ACL.

When a UNIX user accesses a file or folder with an NT ACL, the UNIX user name must be mapped to a corresponding Windows account. If local name mapping has been specified in the name-mapping switch, it will be possible to define a mapping from a UNIX user or a regular expression to a Windows account with a wildcard (*) for the domain name and the lookup to be successful. If this type of entry is encountered, the name-mapping engine will attempt to locate the first instance of the mapped Windows account in the home domain (the domain of which the CIFS server is a member) and then locate the two-way trusted domains of the home domain.

For example, you can have a name-mapping rule that is similar to the following:

```
cluster1::> name-mapping show -vserver cifs01 -direction unix-win -instance
                      Vserver: cifs01
               Name Mapping Direction: unix-win
                  Position: 1
               Pattern: bobbyj
                  Replacement: *\\bobbyjwin
```

The name mapping in this example will map an incoming UNIX user of *bobbyj* to a Windows user *bobbyjwin*. The asterisk (*) in the rule is what makes it a wildcard mapping. When this rule is encountered, it will invoke the “multidomain name-mapping” process. The search will look in the home domain of which the CIFS server is a member and by default all trusted domains. If the Windows name exists in multiple domains, the first domain where the Windows user name is encountered will be the account used to complete name mapping.

In order to reduce the number of domains in which to search, this feature has an option to define a preferred list of trusted domains. Setting up a preferred list will limit the searches in the trusted domains for the mapped Windows account. See the “Recommendations” for further explanation on the preferred list.

You can view the list of trusted domains by using a command from the CLI. The command is `vserver cifs domain trusts show`. The following is an example:

```
Cluster1::> vserver cifs domain trusts show -node cluster1-01 -vserver cifs01 -home-domain
domainA.local -instance

               Node: cluster1-01
               Vserver: cifs01
   Home Domain Name: DOMAINA.LOCAL
   Trusted Domain Name: DOMAINB.LOCAL, DOMAINA.LOCAL
```

As the preceding example shows, SVM (Vserver) “cifs01” is a member of DOMAINA.LOCAL. A trust exists between DOMAINA.LOCAL and DOMAINB.LOCAL. In this example, an individual node was specified in the command syntax; however, the command can be run without specifying a node. The results returned would reflect what each node is aware of regarding trusted domains. The trusted domain list is built per cluster node. Each node becomes aware of trusts when one of the following occurs:

- **On demand.** A client connects to an SVM and discovers a wildcard-mapping rule for a similar user.
- **Periodic rediscovery.** After trust has been established, the cached information about trusts is rediscovered every four hours.
- **Manually.** From the CLI, a storage administrator can issue the command “`vserver cifs domain trusts rediscover`”. This will start a process in clustered Data ONTAP to reach out and discover trust relationships.

Performance

The impact to the end-user experience will occur at the very beginning of an attempt by a UNIX user to access data that has an NTFS ACL. The performance impact is during the initial authentication/user lookup, which is necessary to complete the user-mapping process. After user mapping completes, there is no impact to the actual exchange of data.

Keep in mind the following when multidomain user mapping needs to be completed:

- **Home domain DC connection.** The CIFS server needs a connection to a DC in the home domain. The default settings will cause the user lookup in the home domain and all trusted domains. If a connection has timed out to the home domain, it might be necessary to issue a DC rediscovery for the home domain.
- **Trusted domains have not already been discovered.** An on-demand discovery of trusted domains occurs if there are no previously cached trusted domains.
- **Worst case: all domains are consulted.** If the environment has a large number of Windows trust relationships established, searching all of them to discover a user can take time. Depending on the location and responsiveness of the DCs, the user experience will be affected.

Verification

Manually display the trust information:

```
vserver cifs domain trusts show [-node <node_name>]
```

Manually rediscover trust relationships:

```
vserver cifs domain trusts rediscover -vserver <vserver_name>
```

Recommendations

There are a few best practices that can be explored:

- **Establish a preferred domain search list.** A preferred list will limit the search to those domains defined. The user will need to exist in one of the domains defined in the list, or an error will result while attempting to map the Windows user. The search is conducted in the order in which the domains are defined in the list. Should a user exist in more than one of the domains defined, the search stops upon first discovery of that user. When entering domain names, make sure to use the FQDN for the domain.
- **Order of the preferred list.** The preferred list as mentioned is searched based on the order the domains are listed. If you have domains that contain a larger number of your mapped Windows accounts, those domains should be placed near the beginning of the list.

The following commands are examples of setting, modifying, and viewing the preferred trust list:

- **Setting a preferred list.** The following example will add domainB.local to the preferred search list. This will add the entry to the end should an existing list be established:

```
cluster1::> vserver cifs domain name-mapping-search add -trusted-domains domainB.local -vserver
cifs01
```

- **Modify (or reorder) an established list.** The `name-mapping modify` option can be used not only to reorder an existing list, but also to add to the current list. The command accepts a comma-separated list of trusted domains. The following command will change the order to make domainC.local first, make domainB.local second, and add domainD.local to the end of the list:

```
cluster1::> vserver cifs domain name-mapping-search modify -trusted-domains
domainC.local,domainB.local, domainD.local -vserver cifs01
```

3.3 Separate AD Authentication

This feature allows you create an Active Directory account without a CIFS license. It also provides the ability to join/modify/unjoin a CIFS server to a Windows Active Directory domain without the need for a CIFS license.

This feature is useful for allowing domain users to manage the cluster when those users have no CIFS file-sharing needs. Think of a SAN-only environment that utilizes domain accounts to run applications to manage their LUNs on clustered Data ONTAP. Without a full CIFS license, you will be unable to create shares in order to utilize the CIFS server as a file server.

The CIFS server will be able to communicate with Active Directory to authenticate users, security identifier (SID) lookups, and so on. This allows environments that already utilize Active Directory to keep their existing security management policies in place and allow for simpler management of their user accounts.

Performance

The overall performance impact is no different than would be present for normal CIFS file-sharing operations because the authentication of a user occurs prior to the exchange of data. The time to complete user authentication depends on the network latency, the ability of the Active Directory DCs to reply, and the number of incoming requests for authentication.

3.4 Offbox Antivirus

The offbox Vscan feature provides antivirus scanning support to clustered Data ONTAP, where the virus scanning is performed by third-party machines hosting virus scanners from various vendors. This feature provides a functionality similar to that currently used in Data ONTAP 7-Mode.

The offbox Vscan feature provides virus-scanning support by triggering in-band notifications to external virus-scanning servers during various file operations such as open, close, rename, and write. Due to the in-band nature of these notifications, the client's file operation will be suspended until the scan status is reported back by the external virus-scanning server. The vscan servers, upon receiving a notification for a scan, will retrieve the file over a privileged CIFS share and scan the file contents. Should the scanner encounter a situation in which it becomes necessary to take action on a file, the scanner may attempt to perform remedial operations on an infected file. The remedial action will depend on the configuration defined on the virus scan servers.

After completing all necessary operations, the virus scan server will respond with the scan status to clustered Data ONTAP. Depending on the status sent from the scan, clustered Data ONTAP will allow or deny the requested file operation by the client. In clustered Data ONTAP 8.2.1, virus scanning is only for CIFS-related traffic.

Although this feature is similar to the 7-Mode implementation, there are some key enhancements. A few of those enhancements are:

- **Granular scan exclusion.**
  - The ability to exclude files from being scanned based on the size and location (path) of the file
  - Option to only scan files that are opened with execute permissions
- **AV engine version.**
  - Rolling updates of the AV scan engine support. Clustered Data ONTAP will maintain the current running version of a virus scan server along with the scan status of a file. Should a single server in a pool update its version, it does not require discarding the scan status of all files already scanned.
- **Security enhancements.**
  - Clustered Data ONTAP validates incoming connection requests by an antivirus server to make sure they are valid scanners. A comparison is done against defined scanner pools to make sure the privileged user and IP address are allowed to connect.

4 Antivirus Architecture

For a successful antivirus solution configuration, we must understand the various components of the configuration such as external antivirus scanner (vendor software), Data ONTAP AV connector, and Data ONTAP Vscan settings.

4.1 Components of the Vscan/AV Scanner Server

Clustered Data ONTAP Antivirus Connector

The Data ONTAP antivirus connector must be installed on the antivirus scan server. The Data ONTAP antivirus connector communicates with the antivirus scan software and SVM to process the scan requests.

Performance

The Data ONTAP antivirus connector and antivirus scan software communicate with each other on the loopback address (127.0.0.1), so there are no performance implications.

Verification

To verify the connectivity:

1. Right-click the Configure Data ONTAP Management LIFs for Polling application shortcut created on the desktop during the installation and select run as administrator. This will open up the Configure Data ONTAP Management LIF configuration dialog box.
2. Click `test` to verify the connectivity and authenticate the connection.

Recommendations

- Credentials used as service accounts to run the AV connector service must be added as the privileged user in the scanner pool.
- Same service account must be used to run the AV scan engine service.
- Configure Data ONTAP management LIFs.
- Credentials used for polling must have at least read access to the network interface.

You might want to use a separate user to poll the Data ONTAP management LIFs for security purposes. Preferred accounts should be "cluster admin" or "vsadmin."

Antivirus Software

The antivirus software is installed and configured on the external Windows Server instance (referred to as Vscan server) to scan the files for viruses or any other malicious data. The antivirus software must be compliant with clustered Data ONTAP. You must also specify the remedial actions to be taken on the infected files in this software.

For the antivirus software installation guidance and best practices, refer to the respective vendor documentation.

4.2 Components of System Running Clustered Data ONTAP

Scanner Pool

A scanner pool is used to validate and manage the connection between the Vscan servers and the SVM (Vserver). You can create a scanner pool for an SVM and define the list of Vscan servers and privileged users that can access and connect to that SVM. You can also specify the scan request timeout period. If the scan response to a scan request is not received within this timeout period, access will be denied in mandatory scan cases.

Performance

Scanner pool performance depends on the performance of AV scanners in the pool and the network connecting the SVM and the AV scanner.

Adding additional SVMs and antivirus servers will help to scale out the solution:

- Maximum for 8.2.1: 20 scanner pools per SVM
- Maximum: 100 Vscan servers and privileged users per scanner pool

Verification

To verify the scanner pool settings, such as the list of scanners connected, status, and view information about all scanner pools belonging to all SVMs (Vservers) or one scanner pool belonging to an SVM, use the `vserver vscan scanner-pool show` command.

Recommendations

- Make sure that you have all the AV scanners for serving an SVM added to the scanner pool. NetApp recommends having at least two scanners per scanner pool because having more than one scanner will help deal with fault tolerance and regular maintenance issues.
- The number of scanners to be connected per SVM will depend on the size of the environment.
- It is mandatory to have an AV scanner and an SVM in the same security domain. The same user account must be used for running AV connector service, AV scan engine, and privileged user. In the case/event of secure multi-tenancy, the privileged user must be different for different SVMs to make sure of multi-tenancy compliance.
- The scan request timeout period should be less than the CIFS timeout. The default timeout value is in case of mandatory timeout, which might lead to access denial.

Scanner Pool Policy

A scanner pool policy defines when the scanner pool will be active. A Vscan server is allowed to connect to an SVM only if its IP and privileged user are part of the active scanner pool list for that SVM.

Note: The scanner policies are all system defined, and you cannot create a customized scanner policy.

A scanner policy can have one of the following values:

- **Primary.** Makes the scanner pool always active.
- **Secondary.** Makes the scanner pool active only when none of the primary Vscan servers are connected.
- **Idle.** Makes the scanner pool always inactive.

Verification

To verify the scanner pool policy, use the `vscan scanner-pool show` command.

Recommendations

Make sure you have applied a primary policy to a primary scanner pool and have applied a secondary policy to the backup scanner pool.

On-Access Policy

On-access policy defines the scope of scanning of files when accessed by a client. You can specify the maximum size of the file, which must be considered for virus scanning. File extensions and paths are to be excluded from scanning. You can also choose from the available set of filters to define the scope of scanning.

Performance

To reduce performance impact of antivirus scanning, file types, size, and paths can be excluded. Make sure that all file types that are required to be scanned are configured for scanning:

- Maximum: 10 on-access policies per SVM
- Maximum: 100 paths and file extensions in exclusion list per on-access policy

Verification

Verify the on-access policy setting and view information about all on-access policies belonging to all SVMs or one on-access policy belonging to an SVM to manage the on-access policies by using the SVM `vscan on-access-policy show` command.

Recommendations

- You might want to exclude large files (file size can be specified) because they might result in slow response or scan request timeout for CIFS users. Default exclusion size is 2GB.
- You might want to exclude file types and extensions such as .vhd or .tmp because they might not be appropriate for scanning.
- You might also want to consider excluding certain paths such as quarantine directory or some paths where only virtual hard drives or databases are stored.
- Make sure all the exclusions are specified in one policy because only one policy can be enabled at a time. Having the same set of exclusions specified on the AV scanners as well is highly recommended. For details about supported exclusions, contact your respective antivirus vendors.

Vscan File Operations Profile

The Vscan file operations profile (`–vscan-fileop-profile`) parameter defines which action on the CIFS share can trigger virus scanning. You must configure this parameter while creating or modifying a CIFS share.

This parameter can have one of the following values:

- **No scan.** Virus scans are never triggered for this share.
- **Standard.** Virus scans can be triggered by open, close, and rename operations. This is the default profile.
- **Strict.** Virus scans can be triggered by open, read, close, and rename operations.
- **Writes only.** Virus scans can be triggered only when a file that has been modified is closed.

Recommendations

- Use default profile that is standard.
- If you are looking at very strict scanning options, you can use the “strict” profile. However, using the strict profile generates more scan requests and will have an effect on performance.
- If you are looking for maximum performance with liberal scanning, you may select writes only. Using this profile will scan only files that have been modified and closed.

Others: General Infrastructure Recommendations

- Use an AV scanner server dedicated to antivirus scanning and not used for other jobs such as backup. The reason is that any application running on the machine will share the CPU cycle and memory on the server. This will increase the CPU latency (cycle) for the AV process and will reduce the number of AV requests being processed in any particular time interval.
- You may decide to run the AV scanner as a virtual machine as well. However, you need to make sure that the resources allocated to the virtual machines are not shared and are enough to perform scanning.
- Provide adequate CPU, memory, and disk to the antivirus server to avoid resource bottlenecks. Most antivirus servers are designed to use multiple CPU core servers and to distribute the load across the CPUs.
- Make sure you adhere to the hardware specifications provided by the antivirus software vendors.
- NetApp recommends using a dedicated network using a private VLAN from SVM to AV scanner to make sure that the scan traffic is not affected by other client network traffic. Create a separate NIC on the antivirus server and data LIF on the SVM dedicated to the antivirus VLAN. This simplifies administration and troubleshooting should network issues arise.
- Connect the NetApp storage system and AV scanner using at least network 1GbE for virtualized (shared) AV scanners 10GbE. This should help to avoid network bottlenecks.

- For an environment with multiple NetApp storage devices and multiple scanners, connect all AV scanners with similar high-performing network connections as primary to all the NetApp storage devices. This will improve the performance by load sharing.
- For remote sites/branch offices, using local AV scanners rather than remote AV scanners is recommended due to high latency. If cost is a factor, then customers can rely on laptop/PC virus protection for moderate virus protection. They can also schedule periodic complete file system scans by sharing the volumes/qtrees and scanning them from any system in the remote site.
- Use multiple AV scanners to scan the data on SVM for load-balancing and redundancy purposes. The amount of CIFS workload and resulting antivirus traffic will vary per SVM. Monitor CIFS and virus scan latencies on the storage controller. Trend the results over time. If CIFS latencies and virus scan latencies increase due to CPU or application bottlenecks on the antivirus servers beyond trend thresholds, CIFS clients might experience long wait times. Add additional AV servers to distribute the load.
- The latest version of the AV connector should be installed.
- Antivirus engines and definitions should constantly be up to date, and the update frequency should be in accordance with recommendations from AV vendors.
- Pod architecture can also be utilized; however, there are a few considerations.
  - For secure multi-tenancy, sharing a scanner between two or more SVMs is not possible, because the SVM and the scanner need to be part of the same security domain.
  - Because clustered Data ONTAP can have multiple nodes and SVMs spread across many nodes, the scanner pool more or less does work similar to that of a pod.

For more details about offbox antivirus, refer to the antivirus deployment guide. The details necessary to cover this feature are beyond the scope of this technical report.

Version History

Version	Date	Document Version History
Version 1.0	June 2013	8.2 feature best practices
Version 1.0.1	December 2013	Added 8.2.1 feature best practices

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.

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer’s responsibility and depends on the customer’s ability to evaluate and integrate them into the customer’s operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

Go further, faster®

NetApp®

www.netapp.com