Technical Report

Best Practices Guide for Clustered Data ONTAP 8.2.x and 8.3.x Windows File Services

Brahmanna Chowdary Kodavali, Marc Waldrop, NetApp
August 2015 | TR-4191

Abstract

Windows File Services on NetApp® clustered Data ONTAP® 8.3.x or 8.2./8.2.x offer new use cases and features. This technical report covers those 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 the Data ONTAP storage operating system. 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 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.x and 8.3.

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

Note: For information about feature details and services on NetApp storage systems, see the clustered Data ONTAP documentation “File Access and Protocols Management Guide for CIFS” on the NetApp Support website.

2 Overview of Windows File Services in Clustered Data ONTAP 8.3.1

With clustered Data ONTAP 8.3.1, a few features were added. The following details are just the key features. For a complete list of the features, see Table 9 in the appendix. For details about features that are not outlined here, consult the File Access and Protocols Management Guide.

2.1 SMB Encryption

SMB encryption is a feature that is part of the SMB3 protocol that clustered Data ONTAP now supports. The feature allows for the securing of data during transit (that is, in flight) from client to server. This feature is not securing data at rest where the data is secured as it sits on disk. If you are looking for details about storage encryption or at-rest disk encryption, consult the clustered Data ONTAP documentation titled “Physical Storage Management Guide.”

SMB encryption has the following requirements:

- Client support for SMB3
- Clustered Data ONTAP 8.3.1 (or later)

SMB encryption in clustered Data ONTAP can be configured globally at the SVM level or on a per-share basis. Enabling the feature at the SVM level causes SMB encryption to be required for all clients that are attempting to establish connections to shares. This requirement can create a problem for those clients that do not support SMB encryption. If you have a mixed environment of clients that support SMB3 and those that don’t, then you should consider using the per-share setting when enabling SMB encryption. This type of implementation flexibility fits inline with what Microsoft did as well when it introduced SMB encryption.

Unlike other encryption methods with which you may have to maintain and/or manage a set of keys, with SMB encryption the key management is done within the protocol itself. It's a self-contained security mechanism utilizing the AES-CCM algorithm to secure and sign the data. No separate key management is required or necessary. If you are interested in how the keys are generated, consult the SMB specification "Server Message Block Protocol Versions 2 and 3" (sections 3.1.4.2 and 3.2.5.3).

Enabling the feature is again done at the SVM or per-share level. The following are how to enable the feature for both:

- Enable at the SVM level:
  `vserver cifs security modify –vserver <SVM> -is-smb-encryption-required <true|false> (default: false)`
- Enable at the share level:
  `vserver cifs share properties add –vserver <svm> -share-name <share> -share-property encrypt-data`

Table 1 outlines the various combinations when enabling SMB encryption.

Table 1) SMB Encryption combinations

SVM Setting: -is-smb-encryption-required	**Share Property Setting:** encrypt-data	**Encryption Behavior**
False	False	No encryption enabled.
False	True	Encryption used only on the shares where the property 'encrypt-data' has been added.
True	False	Encryption is required for all share connections. The SVM setting supersedes the share setting.

Recommendations

This feature is a security feature and is recommended only for use in environments where, from a protocol perspective, in-flight data security is a requirement. Recommendation is to use the per-share level setting vs. the per-SVM setting. This recommendation is made for the following reasons:

- Allows clients that don't support SMB encryption to still connect to shares that don't contain information that needs to be secured while in flight.
- If not all data needs to be secured in flight, provides flexibility in enabling the feature on an as-needed basis for only the data access points (that is, shares) that are necessary.
- Not recommended for use on SMB3 connections over SMB shares for SQL Server or Hyper-V. There is a notable impact to performance, and those types of enterprise applications have latency dependencies that may not tolerate an increase in latency when using SMB encryption.

2.2 SVMDR (SMB Details Only)

Simply put, SVM is a disaster recovery solution at the SVM granular level. It encompasses details and information far beyond just Windows File Services. The following information is just a portion of the overall SVMDR solution. If you are interested in learning more about SVMDR, see TR-4015, "SnapMirror Configuration and Best Practices Guide for Clustered Data ONTAP." SVMDR is an asynchronous version of MetroCluster® using SnapMirror® to accomplish the replication instead of RAID SyncMirror® used by MetroCluster.

For those familiar with vFiler® DR in 7-Mode, the SVMDR solution is quite similar. It enables the recovery of an SVM and its configuration onto a separate cluster. It utilizes the NetApp SnapMirror technology in order to replicate the information between the source and destination clusters. It can be configured in two different modes: ID preserve and ID discard. This is important when it comes to Windows File Services because it dictates any potential impact in a DR scenario.

With ID preserve, the SVM configuration is replicated in its entirety to the DR SVM that resides on a separate cluster. This means things such as protocol configurations, services, and networking are replicated to the DR SVM. For Windows File Services, this means there is near minimal impact when a DR scenario is activated. The cutover to a DR site means that all of the settings are the same, so clients can quickly and easily reconnect to their data using the same CIFS server names, IP addresses, and so on. The clients need to reestablish connections to any open files (there are no locks transitioned between source SVM and the DR SVM) at the time of a DR cutover, which is what is meant by “near minimal” impact.

With ID discard, most but not all of the SVM details are replicated over to the DR SVM. The key items that are not replicated over when using ID discard are name services settings (DNS, LDAP, NIS, and so on), CIFS server domain setup, and network settings (LIFs, routes, firewall). This type of DR setup could have a more increased impact to clients upon activation of a DR scenario.

Table 2 shows the features that are replicated depending on the DR configuration.

Table 2) CIFS SVMDR ID preserve vs. discard

CIFS-Specific Feature	ID Preserve	ID Discard
CIFS local users/groups	Y	Y
Audit	Y	Y
BranchCache	Y	Y
CIFS server options	Y	Y
Home directory	Y	Y
Shares	Y	Y
CIFS server security	Y	Y
CIFS symlink settings	Y	Y
FPolicy® policies	Y	Y
Fsecurity policies	Y	Y
Name/group mapping	Y	Y
Name services (DNS, LDAP, NIS, local UNIX users/groups)	Y	**N**
CIFS server/domain details	Y	**N**
Network settings (LIFs, routes, firewall)	Y	**N**

Recommendations

The configuration of SVMDR goes well beyond just Windows File Services, and there is no clear-cut recommendation on whether to use ID discard vs. ID preserve to be made from this guide. You should

take into consideration Table 2 and understand the impact that the features *not* replicated have. If SVMDR is set up using ID discard, then understand the following:

- Network settings are different, and this may have an impact on configurations external to clustered Data ONTAP. For example, if you have applications that refer to shares by IP address (for example, \\192.168.1.100\ProjectXshare), with ID discard the IP address on the DR SVM is different.
- CIFS server and domain details are not replicated. Prior to configuring and establishing an SVMDR relationship, CIFS needs to be configured on the DR SVM. This means the source CIFS server and destination CIFS server have different names and register themselves with Active Directory with different names. This creates a situation to that of the preceding IP address scenario. Care needs to be taken to make sure applications and clients can connect to an SVM during a DR scenario. This may involve the following:
  - Changing client and/or application connection details due to the different server names.
  - DNS modifications to make sure of a more seamless failover to the DR site (that is, create DNS records upon failing over that alias for the source CIFS server name to route to the DR CIFS server name). This could be done to eliminate the need to change client and/or application connection details.
  - Clients may need to flush their DNS resolver caches (for example, Windows clients issuing "ipconfig /flushdns").
  - Kerberos authentication and modifications: If DNS aliasing is used to make the DR situation more seamless due to different CIFS server names, then further adjustments need to be considered to make sure of Kerberos authentication. Kerberos is an authentication mechanism. It functions using service principal names and tickets that are initially generated using the name of the CIFS server when it was created. The tickets are generated and encrypted/decrypted using an algorithm that involves the principal's password.

    If name resolution allows for a client to locate the CIFS server, it may then attempt to obtain a Kerberos ticket for the old name and aligned to the original source SVM CIFS server. Depending on the environment setup, it is possible that upon presentation of the ticket, the DR CIFS server might be unable to decrypt the ticket. This can occur if the source and DR CIFS servers are members of the same Active Directory domain. Depending on the environment, additional steps might need to be taken:

    - DR CIFS server member of same Active Directory domain as source CIFS server:
      - May need to consider deleting the original CIFS server computer object from Active Directory during a DR scenario. This would require that CIFS setup be run again on the original source SVM, prior to returning back to service the original CIFS server, in order for the computer object to be created again.
      - Make sure DNS aliasing reroutes clients to DR CIFS server.
      - Set up SPN alias using the "setspn" tool from a Windows host that is a member of the domain.
    - DR CIFS server not member of same Active Directory domain as source CIFS server:
      - Make sure DNS aliasing routes clients to DR CIFS server.
      - Set up SPN alias using the "setspn" tool from a Windows host that is a member of the domain.

Additional considerations that do not depend on the configuration chosen:

- No user data should be stored on the SVM root volume. User data stored on the root volume is not replicated.
- This is again an asynchronous replication and DR implantation. Therefore, there is an impact to CIFS clients and open files. No open file locks are replicated between source and DR.

2.3 Dynamic DNS (DDNS)

Dynamic DNS means having the ability to automatically register networking details with a DNS server. This saves time when managing network information such as IP addresses. When configured, this updates the IPv4 A and PTR records in DNS for network interfaces configured in an SVM. Where CIFS comes into the DDNS conversation is how or what information is updated for a CIFS server.

There are two settings that control DDNS: one that enables DDNS globally and another that controls whether a particular LIF has its information updated.

- Enable DDNS globally:
  ```
  dns dynamic-update modify -vserver <SVM> -is-enabled true -use-secure
  true -domain-name <SVM_FQDN>
  ```

- Enable DDNS per LIF:
  ```
  net int modify –vserver <SVM> –lif <LIF> -is-dns-update-enabled true
  ```

After you have enabled DDNS, it provides the ability of CIFS to register with DNS. The CIFS server name registration is the CIFS server name appended with the Active Directory domain name. It registers those interfaces that are configured as LIFs with CIFS protocol defined. Those interfaces with the “data-protocol” set to “none” are not registered. The TTL is 24 hours by default.

Recommendations

- Enable DDNS on the interfaces you intend to use for CIFS.
- DDNS is not supported if you use the on-box DNS feature in clustered Data ONTAP.
- If your DNS server only allows secure updates, you need to specify “use-secure” when enabling DDNS on the SVM.

3 Overview of Windows File Services in Clustered Data ONTAP 8.3

Clustered Data ONTAP 8.3 brought forward many features to help bring parity with Data ONTAP operating in 7-Mode. Several other new features were introduced as well. Following is a list of all the features, and, where necessary, best practices are outlined.

3.1 Windows File Services Features in Clustered Data ONTAP 8.3

Following is a list of the major new features that clustered Data ONTAP 8.3 provides. For a complete list of features introduced in version 8.3, see Table 10 in the appendix.

- Microsoft Management Console (MMC) support for viewing and managing open files, open sessions, and shares
- NetBIOS aliases
- Storage-Level Access Guard (SLAG)
- Native file access auditing for user logon and logoff
- Group policy object (GPO) security policy support
- NetApp FPolicy pass-through read support
- Offloaded data transfer (ODX) enhancements
- Support for Microsoft dynamic access control (DAC)

3.2 MMC Support for Viewing and Managing Open Files, Open Sessions, and Shares

Prior to version 8.3, clustered Data ONTAP support for the MMC was mainly for viewing only. You could view shares, for example, but could not manage them. Starting in clustered Data ONTAP 8.3, you can now do the following with the MMC:

- Create a share.
- Stop sharing a share.
- View details about current open sessions.**
- Close out an outstanding session.**
- View details about currently open files.**
- Close out an open file.**

Note: The view displayed by the preceding capabilities marked with a double asterisk (**) are node specific instead of cluster specific. So, when you use the MMC to connect to the CIFS server host name (that is, `cifs01.domain.local`), you are routed, based on how you have set up DNS, to a single LIF within your cluster.

For example, a 4-node cluster with a single data LIF per node (4 total LIFs) has each LIF defined in DNS to service the name `cifs01.domain.local`. Based on a DNS round robin, a connection to the MMC when using `cifs01.domain.local` can be routed to any 1 of the 4 data LIFs. After the MMC is connected, the view for files and sessions reflects what the node on which the LIF lives is aware of. Viewing shares within the MMC is a global cluster representation regardless of the LIF to which the connection is routed.

Recommendations

When you attempt to manage or view open files and you want to have a clusterwide view of the open files or sessions, you have options:

- **Option 1.** Within the MMC snap-in in Windows, you can open multiple Computer Management snap-ins within the same window. When you provide the details in the snap-in about which computer to connect to, specify a separate data LIF IP.

- **Option 2.** The alternative to the MMC to manage open files and sessions is to use the NetApp PowerShell Toolkit. If you have downloaded and installed at least the 3.2 version of the NetApp PowerShell Toolkit, then you have the available cmdlets to view and manage open files and sessions. The cmdlets with the NetApp toolkit provide a clusterwide view of open files and sessions. The cmdlets to view and manage sessions are:
  - Current sessions: `get-nccifssession`
  - Current open files: `get-nccifssessionfile`
  - Close outstanding CIFS session: `close-nccifssession`
  - Close an open file: `close-nccifssessionfile`

  For more details about these cmdlets, download and install the latest PowerShell Toolkit (the preceding cmdlets are available starting with the 3.2 version of the toolkit).

3.3 NetBIOS Aliases

A NetBIOS alias is a unique name that consists of 15 characters, with a 16th character that is a unique identifier. The unique identifier is used to define a service (for example, workstation, messenger, and so on). This allows a CIFS server to be referenced by more than one NetBIOS name. These additional names are advertised by the server. A NetBIOS alias should not be confused with a DNS alias, which is intended to allow a host to be discovered by more than a single name through DNS.

The general use for this feature is when you are consolidating equipment. This consolidation can be down to a single clustered Data ONTAP storage virtual machine (SVM, formerly known as Vserver). If you have client applications that use NetBIOS names for access to their data, then you may have to create additional NetBIOS aliases within clustered Data ONTAP.

Recommendations

If the architecture needs more than a single NetBIOS name, you can create up to 200 per SVM. From the command-line interface (CLI) within clustered Data ONTAP, use the following commands to create the NetBIOS aliases:

- To create NetBIOS aliases during CIFS server creation: `vserver cifs server create –vserver <SVM> –cifs-server <netbiosName> –netbios-aliases <comma_separated_list>…`
- To add NetBIOS aliases: `vserver cifs server add-netbios-aliases –vserver <SVM> -netbios-aliases <comma_separated_list>`
- To remove NetBIOS aliases: `vserver cifs server remove-netbios-aliases –vserver <SVM> -netbios-aliases <comma_separated_list>`

3.4 Storage-Level Access Guard (SLAG)

SLAG is a third layer of security or auditing that you can set at the storage level. It cannot be seen or managed by a client (that is, it is not seen by way of the Permissions tab from a Windows client). SLAG security is applied before any file- or folder-level permissions are evaluated. In Windows permission terms, SLAG is a separate DACL and is generally referred to as a *SLAG DACL*. It can be used only to remove permissions, not add them.

SLAG applies only to environments that have New Technology File System (NTFS) security styles. In a pure UNIX security-style environment, SLAG is ignored. If an environment has both Windows and UNIX users accessing NTFS security-style data, to deploy SLAG, UNIX users must map to a valid Windows user.

For more details about SLAG and how to configure it, see the “File Access Management Guide for CIFS” on the NetApp Support website.

3.5 Native File Access Auditing for User Logon and Logoff

Before version 8.3, clustered Data ONTAP included support for auditing file access. With the introduction of 8.3 comes the ability to audit the success or failure of logging on to a share, as well as a successful logoff from a share. Table 3 shows the event IDs that are supported.

Table 3) Native auditing support

Event ID	Message Summary
4624	An account was successfully logged on.
4625	An account failed to log on.
4634	An account was logged off.

Recommendations

Enabling logon and logoff generates a considerable amount of additional events. You must exercise caution when you enable this feature. These events provide additional security capability by being helpful in:

- Monitoring unauthorized access to shares
- Monitoring users' CIFS sessions and their active session time

For more details about auditing and clustered Data ONTAP, see TR-4189: "Clustered Data ONTAP CIFS Auditing Quick Start Guide" or the product documentation File Access Management Guide for CIFS.

3.6 Group Policy Object (GPO) Security Policy Support Additions

This feature is just an enhancement to what is already supported in clustered Data ONTAP. Table 4 illustrates what has been added.

Table 4) GPO policy support additions.

Group Policy
Event Audit: Audit account logon events
Event Audit: Audit logon events
Event Audit: Audit directory service request
Event Audit: Audit object access
Event Log: Maximum-security log size
Event Log: Retain security log days
Event Log: Retention method for security log
Privilege Rights: TakeOwnership
Privilege Rights: SecurityPrivilege
Privilege Rights: Bypass Traverse Checking

3.7 NetApp FPolicy Pass-Through Read Support

Clustered Data ONTAP 8.3 introduces the *pass-through* read feature for FPolicy. Pass-through read is the ability to have a file read from archived storage in an HSM implementation without the file's being recalled back to primary storage.

In an HSM implementation, after a file is archived to a lower tier of storage, a stub file is left in its place. When a user encounters these stub files and double-clicks them, it kicks off a process that requires FPolicy to recall the file back to primary storage. Without pass-through read support, when the recall occurs, the file is read back into primary storage, and the stub file is reconstituted to its original state.

With pass-through read support, the file is recalled and given to the requesting client. However, instead of reconstituting the file to primary storage, the file is read directly off of the secondary storage location.

Recommendations

Best practices and recommendations for using this feature depend on the implementation of the FPolicy architecture and the FPolicy software vendors. For more details regarding FPolicy, see section 5.8, FPolicy.

3.8 Offloaded Data Transfer (ODX) Enhancements

ODX support was originally introduced in clustered Data ONTAP 8.2. In clustered Data ONTAP 8.3, NetApp implemented changes to enhance performance. The key enhancements are as follows:

- **Direct-copy.** ODX uses a point-in-time (PIT) file to assist in copying data so that the contents of the file don't change in flight. With direct-copy, you can use the semantics by which the file was opened to avoid using the PIT copy, and you can copy straight from the source to the destination.

  Direct-copy is used only if the source of the copy starts on an SMB share. For direct-copy to work, the source file has to be opened in such a way that no changes occur to the file. To prevent changes to the file, it has to be opened in a manner that allows only reading of the file by anyone else who opens it. The concept of "opening only for read" does not mean that permissions on the file are set to read-only, however.

  The CIFS protocol includes the concept of "share access" mode. This is something controlled at a level lower than the setting of permissions. A client application can ask that a file be opened, regardless of what permissions are set, in a manner that permits access different than what is allowed by the defined NTFS file permissions when subsequent openers of the file request access (that is, users have read/write but share access mode can state "only read access is allowed while I have the file open"). If the first opener of the file can open it in a way that allows no openers of the file to modify the file and then initiates a copy offload operation, then direct-copy can be used.
- **Increased token size.** When copy offload is used, the data itself is copied between the source and the destination, with only control (that is, progress) data sent to the client. In a traditional copy involving a client, the data must be copied up to the client from the source and then written back down to the destination. With copy-offload, however, the client is removed from the up-and-down process of moving data. So that the correct data gets moved, there is an exchange of tokens between the client and the source or destination.

  The tokens represent a set amount of data for the destination to read from the source and to write into the destination location. In clustered Data ONTAP earlier than 8.3, the token sizes were limited to 8MB. So, when using copy-offload on versions earlier than 8.3, the maximum amount of data that could be pulled by the destination was 8MB at a time. The introduction of 32MB token sizes increases the amount of data that can be pulled and decreases the amount of back-and-forth between the client and the source.

Recommendations

The new enhancements to ODX do not change any recommendations. Follow the same guidelines that are outlined in the overview section that covers clustered Data ONTAP 8.2, section 5.3, Copy Offload (ODX).

3.9 Support for Microsoft Dynamic Access Control (DAC)

DAC is a feature that was introduced in Windows 2012. Starting with clustered Data ONTAP 8.3, NetApp has introduced support for this feature. DAC uses a combination of data classification and user or device claims to reduce the need for a large number of groups to maintain security of file resources. It is not a replacement for the traditional file and share permissions, however; rather, it is more of an enhancement to them. DAC includes the following:

- **Resource property lists.** These are used for classifying resources or data (for example, by country, department, or job title).
- **Central access rules.** This is a series of access control lists (ACLs) that define the criteria that must be met for access.
- **Central access policies (CAPs).** These are essentially a collection of central access rules and are what is enabled by the Security tab in the properties of a folder.

The bulk of the configuration for DAC occurs on a Windows 2012 server. After you have configured DAC, it is deployed by group policy from Microsoft Active Directory. You must include the clustered Data ONTAP SMB server in the GPO in order for it to be available. After it has been deployed by group policy,

DAC is used by enabling a CAP on user data with the Windows GUI and the Security tab on the properties of a folder.

This feature provides a number of benefits, including:

- **Central management of security settings.** DAC is configured from the Windows 2012 Active Directory Administrative Center. After it has been set within the Active Directory Administrative Center, the remaining configuration is performed through group policy deployment and the Security tab in the properties of a folder.
- **Permissions can now use "and" instead of "or."** In current permissions settings, a user can be a member of groupX or groupY to obtain permissions. Now with DAC, you can set an ACL that states that a user must be a member of groupX and have a title of "Engineer" to obtain access to data.
- **Reduced number of groups of which a user needs to be a member.** The reduction in the total number of groups of which a user can be a member helps in many ways. For example, it helps reduce the size of Kerberos tokens. Large Kerberos tokens can create issues with logging into resources. An additional benefit of a reduction in group membership is from an overall administrative standpoint. The fewer groups you have, most likely the easier it is to manage an environment.

The following example uses one of the preceding benefits to illustrate the impact that DAC can have. The example shows how DAC can assist in lowering the number of groups needed:

- Company X has offices in 20 countries = 20 groups to define user locations by country.
- Company X has 10 branch offices per country = 200 groups to define users by either branch office or geographic location.
- Company X has 100 customers across all its branch offices = 20,000 groups to define customers by geographic and branch office location.

With DAC, you could implement conditional expressions and claims, helping to potentially reduce the number from 20,000 to as few as 130:

- Groups to represent countries: 20.
- Groups to represent branch offices: 10.
- Groups to represent customers: 100.
- **Total: 130.** The conditional expression would provide access based on something such as `MemberOf(Country) AND MemberOf(Branch) AND MemberOf(Customer).`

This is just a small example of how this feature can help reduce group membership and administrative efforts in a Windows File Services environment. One note about the feature for NetApp's implementation is in regard to the ability to perform automatic file classification. DAC has a feature that allows datasets to be automatically scanned and classified with certain permissions or restrictions based on the files themselves. This particular feature is not one that the NetApp implementation supports. Clustered Data ONTAP supports manual classification of files, but not the automated process in the Microsoft version of this feature.

Recommendations

NetApp recommends using DAC in a NetApp environment as if it were implemented on Microsoft. There are no specific best practices or recommendations that are exclusive to NetApp. For details about configuring DAC in a NetApp environment, consult the "File Access Management Guide for CIFS" on the NetApp Support website.

4 Overview of Windows File Services in Clustered Data ONTAP 8.2.1

4.1 LDAP over SSL (start-TLS)

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 cleartext. The communications in cleartext 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 of by the application or user with SSL. SSL can provide encryption of the data in transit, as well as mutual authentication. In clustered Data ONTAP 8.2.1, NetApp supports only 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. 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 uses 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 allows 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 <SVM_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 <SVM_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 <SVM_name> … -use-start-tls true
```

LDAP over SSL for Active Directory LDAP:

```
vserver cifs security modify –vserver <SVM_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 need to succeed by 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 about 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 are beyond the scope of this technical report.

If you decide to use LDAP over SSL, you must keep the LDAP port in your SVM client configuration at 389. The clustered Data ONTAP implementation uses start-TLS to provide the LDAP over SSL security. Start-TLS relies on the LDAP conversation occurring over port 389. This is a slight difference over using LDAPS, which uses port 636. You can confirm the port setting for the SVM with the following command:

```
vserver services ldap client show –vserver <SVM_name> –client-config <configName> –fields port
```

Changing the port to 636 results in an inability to communicate with your LDAP server. This is not a NetApp limitation; it is how start-TLS works. The expectation with start-TLS is that the conversation occurs over the unsecure port of 389.

4.2 Multidomain Name Mapping

In a multiprotocol environment, it is possible that users from both Windows and UNIX might 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 with a Windows user accessing UNIX security-style data. The same applies to a UNIX user accessing UNIX data compared with 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 relevant only 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 NTFS 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 is 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 attempts 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 maps 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 invokes the multidomain name-mapping process. The search looks 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 is the account used to complete name mapping.

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 limits the searches in the trusted domains for the mapped Windows account. See the “Recommendations” section 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 starts a process in clustered Data ONTAP to reach out and discover trust relationships.

Performance

The impact to the end-user experience occurs 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 and 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 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 is 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 <SVM_name>
```

Recommendations

A few best practices can be explored:

- **Establish a preferred domain search list.** A preferred list limits the search to those domains defined. The user must exist in one of the domains defined in the list, or an error results while attempting to map the Windows user. The search is conducted in the order in which the domains are defined in the list. If a user exists 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 fully qualified domain name (FQDN) for the domain.
- **Order of the preferred list.** The preferred list as mentioned is searched based on the order in which 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 and modifying the preferred trust list:

- **Setting a preferred list.** The following example adds `domainB.local` to the preferred search list. This adds the entry to the end if an existing list is 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 changes 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
```

4.3 Separate Active Directory Authentication

This feature allows you create an Active Directory account without a CIFS license. It also provides the ability to join, modify, or 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 uses domain accounts to run applications to manage their LUNs on clustered Data ONTAP. Without a full CIFS license, you are unable to create shares to use the CIFS server as a file server.

The CIFS server can communicate with Active Directory to authenticate users, security identifier (SID) lookups, and so on. This allows environments that already use Active Directory to keep their existing security management policies in place and allow 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 before 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.

4.4 Off-Box Antivirus

The off-box 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 operating in 7-Mode.

The off-box 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. Because of

the in-band nature of these notifications, the client's file operation is suspended until the scan status is reported back by the external virus-scanning server. The Vscan servers, upon receiving a notification for a scan, retrieve the file over a privileged CIFS share and scan the file contents. If the scanner encounters 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 depends on the configuration defined on the virus scan servers.

After completing all necessary operations, the virus scan server responds with the scan status to clustered Data ONTAP. Depending on the status sent from the scan, clustered Data ONTAP allows or denies the requested file operation by the client. In clustered Data ONTAP 8.2.1, virus scanning is only for CIFS-related traffic. For more details about off-box antivirus, see section 6, "Antivirus Architecture," or the following NetApp technical reports:

Antivirus Solution Guide for Clustered Data ONTAP 8.2.1: Symantec

Antivirus Solution Guide for Clustered Data ONTAP 8.2.1: McAfee

Antivirus Solution Guide for Clustered Data ONTAP 8.2.1: Trend Micro

Antivirus Solution Guide for Clustered Data ONTAP 8.2.1: Sophos

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.
  - The option to scan only files that are opened with execute permissions.
- **AV engine version:**
  - Rolling updates of the AV scan engine support. Clustered Data ONTAP maintains the current running version of a virus scan server along with the scan status of a file. If a single server in a pool updates 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 confirm that they are valid scanners. A comparison is performed against defined scanner pools to confirm that the privileged user and IP address are allowed to connect.

5 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.

5.1 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)
- Microsoft BranchCache
- Local users and groups
- File-directory (FSecurity)
- FPolicy native file blocking and partner use case support

- Roaming profiles and folder redirection
- Access-based enumeration (ABE)
- Support for Microsoft previous versions
- Offline folders (client-side caching)
- SMB signing
- Remote VSS
- Native file access auditing

5.2 Server Message Block (SMB) 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 Microsoft 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 invalidates 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 that the client connection can be proactively moved to the partner node before the actual failover event.
- **Cluster client failover.** This enables cluster-capable applications to close the stale “opens” and “locks” during failure recovery and enables 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 is 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 whether 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 the witness protocol include:
  - At least one data LIF must be present on each node per SVM.
  - CA shares should not be mapped by using an IP address. They should be mapped by 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 falls back to NTLM. Hyper-V deployment relies heavily on Kerberos authentication.

5.3 Copy Offload (ODX)

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

Traditional host operating systems are designed to run or handle applications efficiently, but are 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 that 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 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 a block (FC or iSCSI) by using the copy offload functionality.

lists the scenarios supported by copy offload.

Table 5) 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.

Following 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

To troubleshoot copy offload:

1. Confirm that SMB 3.0 is enabled on the SVM.
2. Confirm that both CIFS copy offload and SVM scoped subfile-sisclone (enabled by default) options are enabled.
3. Check that the client supports copy offload (Windows 8 or Windows Server 2012).
4. Check the source data volume. It must be at least 2GB and cannot be a read-only, compressed, or 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 of less than 6GB. For platforms with greater than 6GB of memory, the PIT file size can go up to 65GB:

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

5.4 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 increases 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.

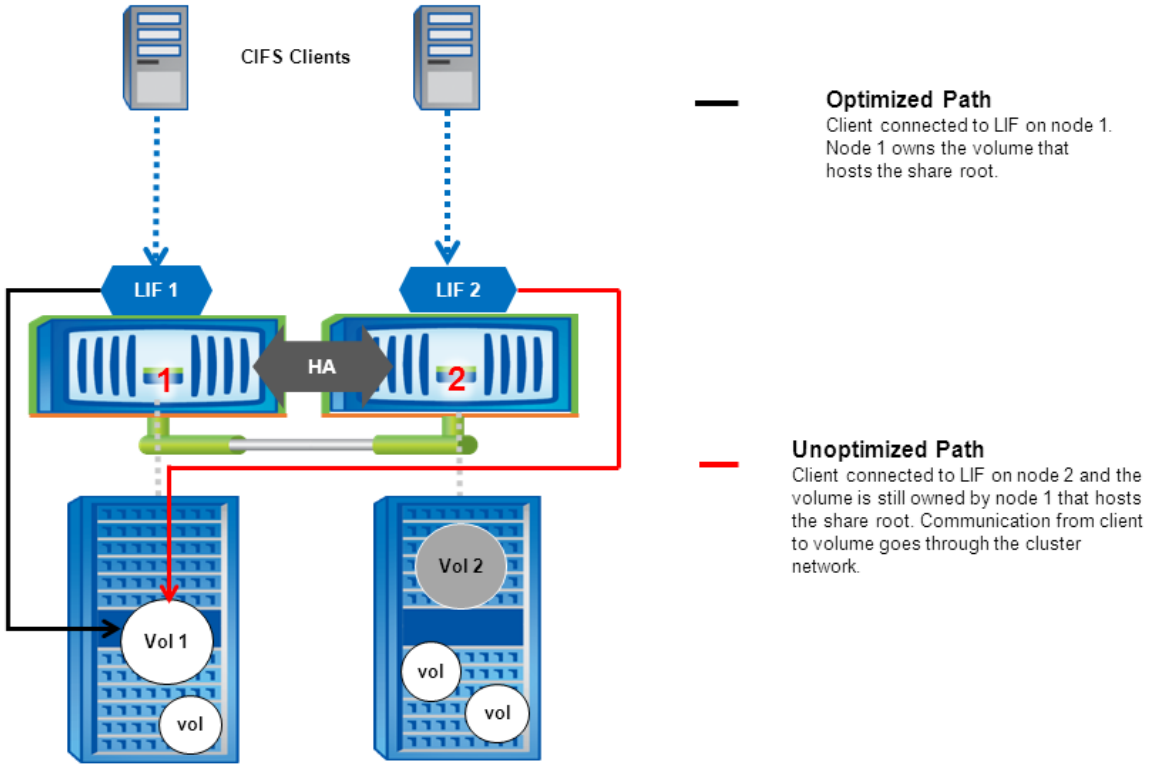

Node referrals address this challenge by redirecting the client connection to the appropriate node that owns the volume that hosts the root of the CIFS share. In that way, the data is accessed locally instead of being accessed over the cluster network. Node referral happens during the connection initiation 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

To troubleshoot node referral issues:

1. Check whether the LIF has migrated to a different node than the node that owns the data volume.
2. Check whether the data volume has moved to the different node but no data LIF has been configured.
3. Use the following performance counters to identify the node referral statistics:
   a. `node_referral_local` displays the number of clients connected by using a LIF hosted by the same node that hosts the share root.
   b. `node_referral_remote` displays the number of clients connected by 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 causes all clients on that share to connect to one node, thus lowering performance through the cluster interconnect.
- Shares are mapped by using specific IP addresses for applications. Clients might get referred to an IP address different from the IP address mapped in the application.
- The share is a CA share. Hyper-V over SMB does not support node referrals.
- The environment is one in which Kerberos authentication is enforced. Node referral provides an IP address to the client to connect to the controller. This forces Windows clients to fall back onto NTLM authentication, and the authentication fails.

5.5 Microsoft 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 can be slow because of 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 6) 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
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 performed on Windows clients. Windows 7 Ultimate or Professional clients and Windows 8 clients support BranchCache.
- A hosted cache mode has an advantage over a 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:

- The hash is flushed frequently.
- The client is accessing the data from the source instead of accessing it from the cache.

To view the current setting applied, run the command `vserver cifs branchcache show`:

```
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. The minimum size of the hash store is 1GB, and each terabyte of content requires a 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 reduces the effort in configuring each share with a BranchCache property. The default value is `per-share`.

5.6 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 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 associates the group privileges to the member domain users.

Use cases for this feature include:

- Backup applications running with the 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 7) 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 the security log.
SeCreateSymbolicLinkPrivilege	Create symlinks (available with Microsoft 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 of authentication and authorization, two levels of cache are introduced to store user-specific information:

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

  This caching option stores a 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 are processed with the cache information available locally on the controller.

Verification

Following are steps for troubleshooting authentication and privilege issues.

Authentication Failures with Local Users or Group Privileges

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

Authentication Failures for Domain Account

1. Make sure the domain account is not disabled. If the account is disabled, an attempt to look up the user locally is performed. This results in a local authentication error because the user is not available in the local database. This is typical Windows behavior.
2. Check whether the user account is from a trusted domain. Validate the trust between the domain to which the controller is joined and the user's domain.

Recommendations

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

5.7 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 for Storage-Level Access Guard (SLAG) (until clustered Data ONTAP 8.3). 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 it on the controller by using FSecurity.

Performance

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

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

Verification

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

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

Recommendations

When 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 (for example, volume mount and unmount operations)
- During high workloads
- During volume move operations

Working in a Multiprotocol Scenario

Exercise caution when applying permissions on mixed-mode volumes and folders because that changes the existing access permission and converts it into NTFS security types. That might cause access disruptions.

Applying Advanced Inheritance Options

The ACE inheritance field specifies how ACEs are 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 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 modify these values.

  For the values of control flags, see 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.

5.8 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 have 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

- The policy scope should be properly configured. Leaving the `Extensions Included` and `Extensions Excluded` options blank results 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 is not blocked. The policy is not enforced until one of the following scope options is configured:
  - `shares-to-include`
  - `shares-to-exclude`
  - `volumes-to-include`
  - `volumes-to-exclude`
- The `-is-file-extension-check-on-directories-enabled` (advanced privilege) option enables 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 is monitored. NetApp recommends changing the setting to `True`.
- File operations that should be monitored for file blocking are `open`, `create`, and `rename`.

Partner Solutions

FPolicy supports advanced features by working with the partner solution (for example, storage management, access management and data governance, quotas, 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, the overall client experience is affected, and monitoring performance counters can help identify resource bottlenecks:

- Increasing FPolicy latency has 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 or 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 the performance of the FPolicy subsystem:
  - `Statistics show –object fpolicy_server –instance SVM:servername`
  - `request_sent_rate`
  - `request_latency`
  - `response_received_rate`

Verification

Verify whether policy events are configured properly:

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

Verify whether the policy scope is configured properly:

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

Verify whether the safeguards are configured properly and executed in advanced mode:

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

Recommendations

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

Policy Configuration

Regarding configuration of the FPolicy external engine for the SVM (Vserver):

- Providing additional security comes with a performance cost. Enabling SSL communication has a performance impact on CIFS.
- If the FPolicy server resources are underused, NetApp recommends increasing the value of `Request Queue Length`. This improves CIFS throughput when the external engine is of sync type. Exercise caution when increasing this value. A very high value leads to CIFS timeouts and high CIFS latencies. You must find the optimal value by working with partners. The value can be changed in the advanced mode:

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

Regarding configuration of the FPolicy event for the SVM:

- Monitoring file operations has 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 optimal 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.

Regarding configuration of the FPolicy scope for the SVM:

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

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. Find the optimal value by working with partners.

Multiple Policy Configuration

- The FPolicy policy for native blocking has the highest priority regardless of the sequence number.
- Decision-altering policies should have higher priority than other policies do.

- Policy priority depends on use cases. To determine the appropriate priority, NetApp recommends working with partners.

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 tunable parameters are part of the FPolicy external engine object, and they can be configured in advanced mode.

Table 8) Fpolicy tunable parameters

Tunable Parameter	Default Value (Sec)	Description	When Applicable
`request-cancel-timeout`	20	Measures the time that the SVM waits for a response from the FPolicy server. Beyond this timeout, the client operation is forwarded to an alternate server, if it exists, or access is allowed or denied based on the mandatory attribute.	Manages CIFS latency when you have a slow FPolicy server.
`request-abort-timeout`	40	Measures the time that the CIFS client request spends on the SVM. Beyond this timeout, the client operation is allowed or denied based on the mandatory attribute.	Manages CIFS latency when you have a slow SVM.
`status-request-interval`	10	Sets intervals at which the SVM sends queries on pending requests to the FPolicy server.	Works with partners to set the optimal 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.
`max-server-requests`	50	Sets the maximum number of outstanding screen requests that are queued for an FPolicy server.	An optimal value should be found by working with partners. Increase the value when the 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 the SVM and FPolicy server is disconnected.	This is required to minimize client disruption due to a slow FPolicy server.
`keep-alive-interval`	120	The SVM sends keep-alive messages to the FPolicy server to detect half-open connections.	This is useful when the client traffic is nonexistent and the FPolicy server is disconnected.

5.9 Managing FPolicy Workflow and Dependency on Other Technologies

- NetApp recommends that you disable the FPolicy policy before you make any configuration changes. 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 NetApp FlexCache® volumes, NetApp recommends that you do not configure FPolicy to monitor read and getattr file operations on the FlexCache volumes. Monitoring these operations in Data ONTAP requires retrieving inode-to-path (I2P) data. Because I2P data cannot be retrieved from the FlexCache volume, it has to be retrieved from the origin volume. As a result, some performance benefits of FlexCache are not realized.
- When both FPolicy and an off-box AV solution are deployed, the AV solution gets the notification first. FPolicy processing starts only after AV scanning has been completed. A slow AV scanner could affect overall performance, hence the AV has to be sized properly.

5.10 Roaming Profiles and Folder Redirection

User profiles on Windows machines separate each user's settings from other users' 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. 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, then multiple clients are downloading the user settings from the network share. This is called a *logon storm*. This can be addressed by using NetApp Flash Cache™ intelligent caching. 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

NetApp has no specific recommendations because this is a Windows configuration. For more information, see Microsoft Best Practices for User Profiles.

5.11 Access-Based Enumeration (ABE)

ABE displays only the files and folders that a user has permission 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 has a performance impact on the content enumeration because an additional permissions check is performed to hide the content from users who do not have access to the data.

Verification

To troubleshoot the ABE feature:

1. Verify that the ABE feature is enabled on the share through share properties.
2. Verify that 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 result in authentication failures.

Recommendations

NetApp recommends that you avoid having too many objects under any folder. File enumeration takes longer if too many objects are under the folder.

5.12 Support for Microsoft Previous Versions

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

To troubleshoot issues with “previous versions” issues:

1. Verify that NetApp Snapshot® copies are scheduled and are available for the volume.
2. Verify that Snapshot directory access is enabled on the volume.
3. Verify that `volume modify –volume datavol1 –snapdir-access` is set to `True`.
4. If you are restoring from a path that crosses junctions, check whether the directories have junctions underneath them. Restore operations can fail with an error if directories have junctions underneath them.

Recommendations

NetApp has no specific recommendations.

5.13 Offline Folders (Client-Side Caching)

Offline folders is a Windows client-side feature that uses the `client-side caching` feature. 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 are 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 are 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 you connect to the share. If offline files are not configured on the share by using `-offline-files`, then by default manual file caching is set on the share.

This is a Windows client feature, and NetApp has no specific recommendations.

For more information, see http://technet.microsoft.com/en-us/library/cc784484(v=ws.10).aspx.

5.14 SMB Signing

SMB signing prevents the transmission and reception of data across a network from being 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 is a performance impact because each message has to be signed and verified by either the server or the client to confirm that the message has originated from a rightful source.

Recommendations

NetApp recommends that you enable SMB signing if security is a key requirement in your organization.

5.15 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, see TR-4172: Microsoft Hyper-V over SMB 3.0 with Clustered Data ONTAP: Best Practices.

5.16 Native File Access Auditing

Native file access auditing provides a file-auditing framework that supports both CIFS and NFS protocols. Auditing in CIFS is based on NTFS system ACLs (SACLs) or NFS 4.x ACLs.

The native auditing infrastructure provides Data ONTAP capability to securely generate and manage audit logs. Auditing is used in organizations mainly to meet compliance requirements.

Performance

Enabling native auditing affects CIFS latency and CPU utilization. Follow these recommendations to mitigate overall performance impact:

- Audit only those files and folders that are of interest to you.
- Enable SACLs for only those file operations that are of interest to you. NetApp recommends setting SACLs only for create and delete file operations, if that meets your requirements.
- Enable success and failure auditing based on your requirements.

- Choose either the XML or EVTX file format because there is no performance overhead. EVTX is supported on Data ONTAP 8.2.1 and later versions.
- Monitor audit event latencies to examine the effect of auditing on overall SMB latency.
- Use QoS on volumes and SVMs on which auditing is enabled. QoS helps define and monitor workloads serving to identify potential problems.

Recommendations

To improve the client experience when you enable native auditing, address the following recommendations:

- Every aggregate in the cluster, including aggr0, requires reserving 2GB of free space to enable auditing. Make sure that enough free space is available to create a staging volume.
- NetApp recommends using the *vserver security file-directory (FSecurity)* utility to configure SACLS on a large home directory environment. It is faster than applying SACLs over Windows Explorer.
- NetApp recommends the following to configure a destination volume:
  - The destination volume holds the consolidated audit log files. The destination volume must be set while configuring the audit policy.
  - The destination volume should never be filled up to more than 90% at any point in time. This rule should be followed in each of the SVMs in the cluster. To prevent the volumes from filling up, NetApp recommends active monitoring of destination volumes.
  - Configure the `rotate-limit` to an optimal value to prevent the destination volumes from filling up.
  - The optimal size of the destination volume depends on the generated log size, which in turn depends on:
    - **Load.** The greater the CIFS traffic, the greater the number of generated audit records.
    - **Duration.** How long the audit logs are kept before rotating.
    - **SACL type.** The number of records generated depends on the SACLs enabled.

    Destination volume size = generated log size x [rotate limit + 1].

    NetApp recommends keeping an additional buffer of 10% to 15% in the destination volume.
- NetApp recommends the following to configure guaranteed auditing:
  - Guaranteed auditing is a new feature enabled by default in clustered Data ONTAP. With this feature, an audit event is recorded for every client operation, thus providing a highly reliable audit framework for regulatory and compliance requirements.
  - When guaranteed auditing is enabled and the destination volume or staging volume is full, CIFS client operations are blocked. NetApp recommends turning off this feature if regulatory requirements do not mandate guaranteed auditing. You can turn off the feature with the following diagnostic mode command:

```
vserver audit modify -vserver <vserver_name> -destination <unix path> -rotate-size 100MB -
rotate-limit 0 -audit-guarantee true|false
```

- The following limits apply when you enable auditing in a multi-SVM cluster:
  - Auditing is supported in up to 10 SVMs in Data ONTAP 8.2.
  - Auditing is supported in up to 50 SVMs in Data ONTAP 8.2.1.

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

6 Antivirus Architecture

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

6.1 Components of the Vscan or 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 opens 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.
- The same service account must be used to run the AV scan engine service.
- You must 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 antivirus software installation guidance and best practices, see the respective vendor documentation.

6.2 Components of a 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 is 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 more SVMs and antivirus servers helps to scale out the solution:

- Maximum for Data ONTAP 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 and their status, and to view information about all the scanner pools belonging to all the SVMs (Vservers) or information about 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 that serve an SVM added to the scanner pool. NetApp recommends having at least two scanners per scanner pool, because having more than one scanner helps deal with fault tolerance and regular maintenance issues.
- The number of scanners to be connected per SVM depends 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 the AV connector service, AV scan engine, and privileged user. In the case of secure multi-tenancy, the privileged user must be different for different SVMs to meet 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 is 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 that 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

An on-access policy defines the scope of scanning of files when a client accesses them. 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 the performance impact of antivirus scanning, file type, size, and path can be excluded. Make sure that all file types that are required to be scanned are configured for scanning:

- A maximum of 10 on-access policies is allowed per SVM.
- A maximum of 100 paths and file extensions is allowed on the exclusion list per on-access policy.

Verification

To verify the on-access policy setting and to view information about all on-access policies belonging to all SVMs or information about one on-access policy belonging to an SVM, use the SVM `vscan on-access-policy show` command. This information helps you manage the on-access policies.

Recommendations

- You might want to exclude large files (file size can be specified) because they might result in slow response or a scan request timeout for CIFS users. The 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 a quarantine directory or some paths where only virtual hard drives or databases are stored.
- Make sure that all the exclusions are specified in one policy, because only one policy can be enabled at a time. NetApp also highly recommends having the same set of exclusions specified on the AV scanners. 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 when you create or modify 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 standard default profile.
- 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 affects performance.
- If you are looking for maximum performance with liberal scanning, you may select writes only. Using this profile scans only files that have been modified and closed.

Other General Infrastructure Recommendations

- Use an AV scanner server that is dedicated to antivirus scanning and is not used for other jobs such as backup. The reason is that any application running on the machine shares the CPU cycle and memory on the server. This increases the CPU latency (cycle) for the AV process and reduces 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 must 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 that you adhere to the hardware specifications provided by the antivirus software vendors.
- Use a dedicated network with a private VLAN from the SVM to the AV scanner to prevent the scan traffic from being 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 if network issues arise.
- Connect the NetApp storage system and AV scanner by using at least a 1GbE network, and for virtualized (shared) AV scanners, use a 10GbE network. 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 improves the performance by load sharing.
- For remote sites and branch offices, use local AV scanners rather than remote AV scanners because of the high latency. If cost is a factor, then customers can rely on laptop or PC virus protection for moderate antivirus security. They can also schedule periodic complete file system scans by sharing the volumes or qtrees and scanning them from any system in the remote site.
- For load-balancing and redundancy purposes, use multiple AV scanners to scan the data on an SVM. The amount of CIFS workload and resulting antivirus traffic varies 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 because of CPU or application bottlenecks on the antivirus servers beyond trend thresholds, CIFS clients might experience long wait times. Add more AV servers to distribute the load.
- Install the latest version of the AV connector.
- Constantly update antivirus engines and definitions, and the update frequency should be in accordance with recommendations from AV vendors.
- If you choose to use a pod architecture, consider the following:
  - For secure multi-tenancy, sharing a scanner between two or more SVMs is not possible because the SVM and the scanner must 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 similarly to that of a pod.

For more details about off-box antivirus, see the specific vendor's antivirus software deployment guide. The details necessary to cover this feature are beyond the scope of this technical report.

Appendix: Charts of Features Released

Windows File Services Features in Clustered Data ONTAP 8.3.1

Table 9) Clustered Data ONTAP 8.3.1 new features.

CIFS Feature
SMB 3 encryption
SVMDR (global feature not just for CIFS but has CIFS component)
Dynamic DNS (DDNS)
Guest Account Security
FSCTLs for DBCC for SQL over SMB

Windows File Services Features in Clustered Data ONTAP 8.3

Table 10) Clustered Data ONTAP 8.3 new features

CIFS Feature
Support for Microsoft DAC
AES 128/256 for CIFS Kerberos authentication
ODX direct-copy
MMC support for viewing and managing open files and sessions
NetBIOS aliases
SLAG
Native auditing for logon and logoff to shares
UNIX character mapping
GPO security policy support
FPolicy pass-through read support
CIFS restrict anonymous capability
Control bypass traverse checking
CIFS home directory show user command
Control of CIFS home directory access for admins

Windows File Services Features in Clustered Data ONTAP 8.2.1

Table 11) Clustered Data ONTAP 8.2.1 new features

CIFS Feature
Multidomain user mapping
LDAP over SSL (start-TLS)
Off-box antivirus
Separate AD licensing

Windows File Services Features in Clustered Data ONTAP 8.2

Table 12) Clustered Data ONTAP 8.2 new features

CIFS Feature
SMB 3.0
Copy offload
SMB autolocation
BranchCache
Local users and groups
FSecurity
FPolicy
Roaming profiles and folder redirection
Access-based enumeration (ABE)
Microsoft previous version support
Offline folders
SMB signing
Remove VSS
File access auditing or file access monitoring

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
Version 1.0.2	November 2014	Added 8.3 feature best practices
Version 1.0.3	August 2015	Added 8.3.1 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.

Copyright Information

Trademark Information