



Technical FAQ

CIFS Technical FAQ

Clustered Data ONTAP

Marc Waldrop, NetApp
June 2015

TABLE OF CONTENTS

1	Purpose	5
2	General CIFS Terms	5
	What is an opportunistic lock, or oplock?	5
	What is a lease?	5
	What is a share?	5
	What does the acronym SMB stand for?	6
	What is NTLM?	6
	What is Kerberos?	6
	What are share permissions and file/folder (or NTFS) permissions?	6
	Where can I find more information about the CIFS protocol in general?	6
3	Support Considerations	6
	Which Kerberos encryption types are supported?	6
	Which versions of the SMB protocol are supported in Windows?	6
	Which version of clustered Data ONTAP supports which version of the SMB protocol?	7
4	Name Resolution and Time	7
	Are my storage virtual machine name and my CIFS server name the same thing?	7
	I added a DNS cname entry for my CIFS server. I can ping the cname, but I can't get access. What's wrong?	8
	Does the time set on the clustered Data ONTAP storage controller need to match the Active Directory domain time (within five minutes) in order for the CIFS servers created on SVMs to join the domain?	8
5	Authentication and Authorization	8
	What is the difference between authentication and authorization?	8
	Do I need a CIFS license if I never plan to deploy CIFS shares?	8
	Do I need to have user mapping functioning if I am only going to use CIFS or NTFS security style volumes or qtrees?	8
	What is the equivalent to "wcc" in clustered Data ONTAP?	8
	What happens if I don't specify a security style for a FlexVol volume or a qtree during creation?	9
	What happens to file/folder permissions when I change the security style of a volume?	9
	Why does a folder show up as owned by BuiltIn\Administrator, yet the user who created it was Domain\maym?	11
	I am trying to set up CIFS and get an error that references "Error: Strong(er) authentication required," and CIFS setup fails. Why does this happen?	11
	What is the maximum Kerberos token size support for CIFS authentication?	11
6	Features (Including Virus Scanning, FPolicy, Auditing, and So On)	12
	Which version of clustered Data ONTAP supports off-box antivirus?	12
	Is LDAP SSL the same thing as LDAP signing and sealing?	12
	Is SMB3 witness the same as file share witness in a Microsoft cluster?	12
	Does enabling SMB3 enable SMB3-capable client connections to survive storage events including LIF migrate, ARL, and storage failover?	12

What happens if I set up a home directory that has the same name as an existing static share: for example, home directory bobbyj and an actual share named bobbyj?	12
Is there a limit to the number of search paths that can be specified for home directories?	13
Can I set a group as the CIFS superuser?	13
Can a share on a CIFS SVM participate in a distributed file system (DFS)?	13
Is the Microsoft feature BranchCache supported in clustered Data ONTAP?	13
Does clustered Data ONTAP support both modes of BranchCache (hosted cache and distributed cache modes)?	13
Can I use the autolocation feature when using SMB3-CA shares and Hyper-V?	13
Is the SMB3 share property “continuously-available” supported for General Information worker type workloads (ie: home directories or project shares)?	14
Is SMB3 Encryption supported in clustered Data ONTAP?	14
Do I have to maintain separate encryption keys for SMB3 encryption?.....	14
Is there a performance impact to using SMB3 encryption?	14
Does SMB3 encryption secure my data at-rest on disk?	14
Is there NDO for CIFS when activating a DR scenario when using SVMDR?	14
7 Differences Between 7-Mode and Clustered Data ONTAP	15
What is “node” scope versus “cluster” scope?	15
What is the CIFS superuser?.....	15
Is there an equivalent to the 7-Mode option cifs.tcp_window_size in clustered Data ONTAP?	15
Why does clustered Data ONTAP not have the equivalent of the Data ONTAP 7-Mode feature for symlinks called “nosymlink_strict_security”?	15
In 7-Mode there was the option cifs.bypass_traverse_checking, but I don’t see it in clustered Data ONTAP. Is bypass traverse checking not supported?	15
Is there an equivalent to the 7-Mode CIFS option cifs.grant_implicit_exe_perms in clustered Data ONTAP?	16
In clustered Data ONTAP SMB signing only has the single option “-is-signing-required” vs. 7-Mode, which has “cifs.signing.enable” or “cifs.smb2.signing.required.” How do I turn off SMB signing in clustered Data ONTAP?	16
In 7-Mode a common concern was pblks. Do pblks exist in clustered Data ONTAP?	16
How are symlinks enabled in clustered Data ONTAP vs. 7-Mode?	16
I upgraded my 7-Mode system to clustered Data ONTAP and suddenly I am having troubles with clients accessing CIFS shares via clients who are using SMB2 and have either a WAN acceleration device from Riverbed or have removed the traverse/execute permission from folders?	17
8 How It Works	17
How is access to CIFS determined?.....	17
How can I check to see if CIFS is licensed?.....	18
How can I verify that a CIFS server exists?	18
How can I verify that the right CIFS version is enabled?	18
How can I verify that CIFS is allowed for the SVM?	18
How can I verify that CIFS is allowed on the network interface?	18
What happens when a client attempts to map a drive to a share located on a clustered Data ONTAP CIFS server (at a high level)?	18
What is a namespace, and how does it relate to CIFS?	18

What is a volume junction?	19
What external resources does my CIFS server need to be able to reach using the LIFs assigned to the SVM?	19
What are export policies, and do I need one?	19
What is the end-user experience during a storage failover event?	19
I changed the path to which my CIFS share points (referring to a new volume or qtree); however, the client still sees data from the old location.	20
What is the maximum character length supported for a path in CIFS?	20
Is it possible to control to which LIF the CIFS feature “autolocation” refers a client?	20
What does the “status” column information mean in the output of “cifs domain discovered-servers show”?	20
8.1 Support Considerations	21
Which network ports might need to be reachable by a CIFS server?	21
9 Miscellaneous	22
I have folders in a share that one client can see but not another. It occurs on Windows 7 clients and later when using SMB2 to connect.	22
Can I use the MMC to manage a clustered Data ONTAP CIFS server?	22
Does NetApp Support the Microsoft feature called “Work Folders”?	22
What does the symlink-properties option do when creating or modifying the particular setting on a cifs share?	22
I need to enable a version of the SMB protocol in my SVM that was previously disabled (whether by default or manually disabled), will it impact data access for existing clients?	23
Version History	24
 LIST OF TABLES	
Table 1) CIFS and Windows versions.	7
Table 2) CIFS and clustered Data ONTAP versions.	7
Table 3) Max Kerberos Token Size	11

1 Purpose

The purpose of this technical FAQ is to answer questions surrounding the use and implementation of CIFS in the NetApp® clustered Data ONTAP® operating system. The FAQ is broken into several sections that are defined here. This is an internal NetApp document meant to assist NetApp SEs, PSEs, Support and other folks with common questions they have while troubleshooting or selling a solution. This document is meant to be kept internal as a reference document only.

- **General CIFS Terms.** This section contains questions and answers detailing common terms used when discussing CIFS. It covers general CIFS terms true for both Windows® servers and Data ONTAP.
- **Support Considerations.** Here you will find a few charts that cover which versions of the SMB protocol are supported in both Windows and Data ONTAP.
- **Name Resolution and Time.** This section deals with questions and answers that pertain to name resolution. In order to access data, you first need to be able to find the CIFS server.
- **Authentication and Authorization.** After you find the CIFS server, you then need to prove who you are. Then a determination needs to be made as to whether you have access.
- **Features.** This section deals with features including the FPolicy® file-screening policy, antivirus, and auditing.
- **Differences Between 7-Mode and Clustered Data ONTAP –** The questions here will cover details that reflect differences between Data ONTAP operating in 7mode and clustered Data ONTAP.
- **How It Works. –** Explanations on how CIFS and Windows File Services work in clustered Data ONTAP
- **Miscellaneous. –** This is a catch all section that details questions that don't fit into any of the above sections.

2 General CIFS Terms

What is an opportunistic lock, or oplock?

Answer: An oplock, despite its name, which contains the word *lock*, is actually a means through which a client caches data locally to reduce client-to-server chatter and allow the potential for better performance.

What is a lease?

Answer: A lease is very similar to an oplock, but it allows additional levels of caching. Leases were introduced starting in SMB version 2.1. In addition, an SMB lease can upgrade its lease level from a lower cache level to a higher cache level on subsequent opens on a file by the same client (oplocks can downgrade only). Leasing also addresses a limitation in oplocks in which a client that opens the same file a second time without closing the file causes a downgrade or break of an oplock. Leasing accomplishes this using what is known as a lease key.

What is a share?

Answer: A share is similar to an NFS export. It exposes a volume or qtree to SMB-capable clients in order to allow them to read and write data over the network.

What does the acronym SMB stand for?

Answer: It stands for Server Message Block. As the protocol has grown, a number has been added to the end to note the version: SMB is SMB1 or CIFS, SMB2 is SMB versions 2.0/2.1, and SMB3 is SMB version 3.0. Throughout this document, CIFS and SMB are used interchangeably, but when using exact definitions, SMB is the protocol, and CIFS is version 1.0 of the protocol.

What is NTLM?

Answer: NTLM is a security protocol used to authenticate users in SMB environments. It uses a challenge/response mechanism and a one-way hash of the user's password in order to validate the credentials provided by the client.

What is Kerberos?

Answer: Kerberos is a security protocol used to authenticate users in SMB environments. It utilizes a ticket-based system whereby the tickets are encrypted with the passwords of the systems to which the clients are connected. Kerberos is the preferred authentication method for modern Windows CIFS clients and is generally considered more secure than NTLM. Consult the following Microsoft TechNet article for more details: <http://technet.microsoft.com/en-us/library/cc961963.aspx>.

What are share permissions and file/folder (or NTFS) permissions?

Answer: Share-level permissions control what a client can do when accessing a share over the network. The permissions at this level are less granular than NTFS permissions and include the following: Full Control, Read, and Change. File/folder (NTFS) permissions are more extensive than share permissions and include, for example, Read, Write, Modify, Execute, and Full Control. These permissions are used to control local access to files directly from the location housing the data. Share permissions are often combined with file/folder permission in order to control access to and the rights provided to clients.

Where can I find more information about the CIFS protocol in general?

Answer: Consult the following third-party external links for more information about the CIFS/SMB protocols:

- SMB1: <http://msdn.microsoft.com/en-us/library/cc246231.aspx>
- SMB2/3: <http://msdn.microsoft.com/en-us/library/cc246482.aspx>

3 Support Considerations

Which Kerberos encryption types are supported?

Answer: Clustered Data ONTAP up to version 8.2.1 supports DES and RC4-HMAC Kerberos at this time for CIFS. Starting in clustered Data ONTAP 8.3, the introduction of AES encryption was added for Kerberos.

Which versions of the SMB protocol are supported in Windows?

Answer: The following table shows which SMB versions are supported in the Windows operating system.

Table 1) CIFS and Windows versions.

Protocol Version	Windows XP	Windows 2000	Windows 2003	Windows 2008	Windows 2008 R2	Windows 7	Windows 8	Windows 2012
SMB 1.0	✓	✓	✓	✓	✓	✓	✓	✓
SMB 2.0				✓	✓	✓	✓	✓
SMB 2.1					✓	✓	✓	✓
SMB 3.0							✓	✓

*****Note:** SMB 3.02 is another release of the SMB protocol from Microsoft. However, currently NetApp has no support included in clustered Data ONTAP as of 8.3.1. Many of the features in SMB 3.02 are for features specific to Microsoft only or are aligned to SMB features not currently included in our implementation of Clustered Data ONTAP.

Which version of clustered Data ONTAP supports which version of the SMB protocol?

Answer: The following table shows which version of clustered Data ONTAP supports which version of the SMB protocol.

Table 2) CIFS and clustered Data ONTAP versions.

Protocol Version	Clustered Data ONTAP 8.0	Clustered Data ONTAP 8.1	Clustered Data ONTAP 8.2 and Later
SMB 1.0	✓	✓	✓
SMB 2.0		✓	✓
SMB 2.1		✓	✓
SMB 3.0			✓

4 Name Resolution and Time

Are my storage virtual machine name and my CIFS server name the same thing?

Answer: Generally, no. The storage virtual machine (SVM, formerly known as a Vserver) is the “virtual storage server” within the overall cluster that you created. SVMs can serve different purposes, including serving CIFS, NFS, iSCSI, and/or FCP clients and hosts. In order to serve CIFS within the SVM, you must create a CIFS server, and this construct is what you reference when you need to access a CIFS share. The CIFS server name can match the SVM name, but it does not have to.

I added a DNS cname entry for my CIFS server. I can ping the cname, but I can't get access. What's wrong?

Answer: The issue is very likely that the environment disabled NTLM authentication and allows only Kerberos. When you created the DNS cname, it allowed the client to get routed to the CIFS server using name resolution discovery. However, the client was unable to obtain a Kerberos ticket to authenticate. You need to look into creating an SPN alias. See the following NetApp Support KB article for more details: <https://kb.netapp.com/support/index?page=content&id=2018231>.

Does the time set on the clustered Data ONTAP storage controller need to match the Active Directory domain time (within five minutes) in order for the CIFS servers created on SVMs to join the domain?

Answer: Yes. This requirement is for all servers joining an Active Directory® domain. Set the time and the time zone and set up NTP on the storage controller to point to the same time source server (or domain controller) that other member servers use in the domain to enable authentication. This requirement is a Microsoft requirement related to Kerberos.

5 Authentication and Authorization

What is the difference between authentication and authorization?

Answer: Authentication is the process by which you prove who you are. When clients connect to a CIFS server, they provide a set of credentials to the server to prove their identity. The authentication method used, NTLM or Kerberos, dictates the amount of information presented by the client in advance or gathered and sent off to a third party by the CIFS SVM.

Authorization is the level of access a user is allowed that has already been verified. Authorization includes share- and file-level permissions. After a user has been validated, that user's permission determines what the user can or cannot do during that session.

Do I need a CIFS license if I never plan to deploy CIFS shares?

Answer: Yes and no. Starting in 8.2.1 you can join a CIFS server to a Windows Active Directory domain, but you cannot create shares without a CIFS license. A sample use case for this is when customers who utilize domain accounts to manage the cluster need a way to authenticate those users. See the section "Configuring and Managing Active Directory Computer Accounts for SVMs..." in the NetApp "File Access Management Guide for CIFS."

Do I need to have user mapping functioning if I am only going to use CIFS or NTFS security style volumes or qtrees?

Answer: Yes. All Windows users need to map to a UNIX® user when connecting to shares regardless of the security style.

What is the equivalent to "wcc" in clustered Data ONTAP?

Answer: The command to show the equivalent to "wcc" is available from the diag-level CLI and is *diag secd authentication show-creds* –node <nodeName> –vserver <SVMName> –win-name <domain\winUserName>. The following is an example of looking up a Windows user:

```
o  diag secd authentication show-creds -node RTPClus-01 -vserver cifs01
   -win-name RTPNC\bobbyj
```

UNIX UID: bobbyj <> Windows User: RTPNC\bobbyj (Domain User)

GID: nistestgroup

Supplementary GIDs: <None>

Windows Membership:

RTPNC\Domain Users (Domain group)

BUILTIN\Users (Alias)

User is also a member of Everyone, Authenticated Users, and Network Users

Privileges (0x80)

What happens if I don't specify a security style for a FlexVol volume or a qtree during creation?

Answer: The volume/qtree will inherit a security style upon creation, depending on several factors. Those factors are:

- FlexVol® volumes inherit the security style of the root volume of the SVM where the FlexVol volumes are being created.
- Qtree inherits the security style of the parent FlexVol volume in which it is being created.

The exception is Infinite Volumes; they will always use the unified security style.

What happens to file/folder permissions when I change the security style of a volume?

Answer: It depends on the security type you are going from and to, along with the permissions on the actual data itself. Here are a few simple examples that show the impact of changing the volume security style.

Environment:

Active Directory domain = Domain

All Windows users map to a username that is equal to their UNIX account (example: Windows user Domain\waldrop has an equivalent UNIX username of Waldrop).

<u>ACL Details</u>	<u>Change From → To</u>	<u>Change to ACL After Style Change</u>	<u>Result</u>
<i>Everyone – Read and Execute</i> <i>Domain\riggoj – Full Control</i> <i>Domain\clarkg – Full Control</i> <i>Owner – Domain\riggoj</i>	NTFS → UNIX	No	The resulting UNIX mode bit permissions would be 755. The NTFS ACL will be converted to UNIX mode bits such that it doesn't provide any more access than would have been granted when the security style was NTFS. In this case, because Domain\riggoj had Full Control, the "Owner" bit is 7. The Everyone group is what will dictate the remainder of the mode bits. In this case the "Group" and "Other" bits will be set to allow Read and Execute. In this case the user Domain\clarkg will lose the ability to write to the file.
<i>Domain\riggoj – Full Control</i> <i>Domain\clarkg – Full Control</i> <i>Owner – Domain\riggoj</i>	NTFS → UNIX	No	The resulting UNIX mode bits will be 700. Again, the NTFS ACL is analyzed, and UNIX mode bit permissions are set to not allow any more permissions than would be allowed if it were still NTFS. In this situation, because there is no Everyone group, the UNIX ACL cannot assign permissions for group and other. As a result, the only person who has access to the file is the UNIX user riggoj. To restore access to other users, riggoj needs to modify the UNIX ACL accordingly. ***See following copy for the impact that takes place when you change the permissions on files after a volume security style change.

When you change the security style of a volume, the previous ACL is retained in the background. As long as there are no changes to the ACL for the file or directory, then the original ACL is retained in the background. For example:

File name: *foo.txt*

NTFS ACL: *Everyone – Read and Execute; domain\riggoj Full Control; domain\clarkg Full Control*

Security style change: *NTFS → Unix*

After security style change, the UNIX ACL would show: *755 (Everyone entry would limit access to group and other)*

Make no changes and change the security style back: *Unix → NTFS*

Postchange, NTFS ACL would show: *Everyone – Read and Execute; domain\riggoj – Full Control; domain\clarkg – Full Control*

The key is that the ACL was not modified while set to UNIX. If in the preceding example a user from a UNIX workstation mounted an NFS export that points to the location where *foo.txt* resides and typed *chmod 775 foo.txt*, the original NTFS ACL would be gone. After switching the security style back to NTFS, permissions to the file would need to be modified to restore Write access for *domain\clarkg*. To show the impact, here is the same scenario with the UNIX "chmod" inserted.

File name: *foo.txt*

NTFS ACL: *Everyone – Read and Execute; domain\riggoj Full Control; domain\clarkg Full Control*

Security style change: *NTFS → Unix*

After security style change, the UNIX ACL would show: *755 (Everyone entry would limit access to group and other)*

Issue the command “chmod” from UNIX user riggoj: *chmod 757 foo.txt*

Change the security style back: *Unix → NTFS*

Resultant NTFS ACL: *domain\riggoj – Full Control, domain\clarkg – Read and Execute, All others – Full Control*

In order to restore the Windows user *domain\clarkg* access to “write” to the file, you need to re-ACL it using manual modification or set inheritance on a parent folder. The ACL, when viewed, will contain entries that say something similar to “waldrop (UNIXPermUid\waldrop)” and/or “unixgroup (UNIXPerGid\unixgroup).” In order to either manually set the permissions or set inheritance, entries that contain the wording UNIX... will need to be removed prior to attempting to set a new NTFS ACL; otherwise it will fail.

Why does a folder show up as owned by Builtin\Administrator, yet the user who created it was Domain\maym?

Answer: The answer to this resides in the fact that the user *domain\maym* was very likely a member of the BUILTIN\Administrators group locally on the CIFS server. This membership is either direct or indirect using an associated group. This approach is used to match Microsoft behavior and is not NetApp specific.

I am trying to set up CIFS and get an error that references “Error: Strong(er) authentication required,” and CIFS setup fails. Why does this happen?

Answer: This is likely the result of using a version of clustered Data ONTAP that does not support LDAP signing and sealing. This error message indicates that the DC to which you are talking has a higher requirement for LDAP connection than is being provided.

What is the maximum Kerberos token size support for CIFS authentication?

Answer: The answer is depends on the release of clustered Data ONTAP running in the cluster. Please see the following table for more details:

Table 3) Max Kerberos Token Size

Clustered Data ONTAP version	Max Kerberos Token Size
Pre-8.2.2	16k
8.2.2 to pre-8.3	32k
8.3 and later	64k

6 Features (Including Virus Scanning, FPolicy, Auditing, and So On)

Which version of clustered Data ONTAP supports off-box antivirus?

Answer: Version 8.2.1 and later support this.

Is LDAP SSL the same thing as LDAP signing and sealing?

Answer: Yes and no. The following few bullets show where they are similar and where they are not.

- LDAP over SSL (support started in clustered Data ONTAP 8.2.1)
 - Generally used to secure the transmission of LDAP authentication information during the bind phase, particularly for those who use simple binds.
 - Utilizes a certificate to accomplish the securing of data in transit.
 - Requires that the communication to succeed over SSL or LDAP traffic is not exchanged; there is no fallback to non-SSL.
 - The connection to the LDAP server MUST use port 389. LDAP over SSL utilizes start-TLS, and that feature requires that the connection utilize port 389. If you set the LDAP port in the clustered Data ONTAP configuration to 636, the connection to the LDAP server will fail.
- LDAP signing
 - Adds a signature to each exchange, similar to what SMB signing does for CIFS exchanges, to provide a level of validation that the packet was not tampered with.
 - Provides no level of encryption.
- LDAP sealing
 - Similar to LDAP over SSL but not certificate based.
 - The connection in this case does use port 636, but as of the last published date of this article clustered Data ONTAP does not support this method of securing LDAP traffic. It is on the roadmap for consideration in a future release. See your local NDA contact for further details and messaging.

Is SMB3 witness the same as file share witness in a Microsoft cluster?

Answer: No. SMB3 witness is a feature within CIFS that allows faster failover of a CIFS connection between an SMB3 client and a server. A file share witness in a Microsoft® cluster is used to assist in establishing a quorum for Microsoft Windows clusters.

Does enabling SMB3 enable SMB3-capable client connections to survive storage events including LIF migrate, ARL, and storage failover?

Answer: No. Just enabling the SMB3 protocol on an SVM does not make the connection capable of nondisruptive operations. In order to provide complete NDO for the SMB3 protocol in addition to enabling SMB3, you need to specify the share property “continuously available.” Continuously available shares are supported as of Data ONTAP 8.2.1 for Hyper-V® and SQL Server® data workloads only.

What happens if I set up a home directory that has the same name as an existing static share: for example, home directory bobbyj and an actual share named bobbyj?

Answer: A user attempting to connect to the user’s home directory share will be directed to the static share. Home directory shares are “pseudo” shares that exist but are not created through a manual process. A static share is one created by an admin at some point. When faced with a conflict like this, the

client will be connected to the static share and not the home directory. Verify when you use home directories that you don't have conflicts between the pseudo shares that will be created for a home directory and actual storage admin-defined shares.

Is there a limit to the number of search paths that can be specified for home directories?

Answer: The short answer is no. However, keep in mind that if you configure home directories in clustered Data ONTAP, a request by a client to connect to a home directory path is tried on all paths specified by the search path (in the order specified) until a matching folder is found. As the search path grows to include more and more volumes, those users whose home directories reside near the end of the specified locations can see an impact on their workstation logins. The users' home directory will be mapped as part of their standard Windows client login.

Can I set a group as the CIFS superuser?

Answer: No. This only supports defining a user.

Can a share on a CIFS SVM participate in a distributed file system (DFS)?

Answer: The primary question usually asked here is "Can the SVM house the Microsoft DFS root?" and the answer to that is no. A share that is available on the SVM can, however, be the target of a DFS link configured on the DFS server.

Is the Microsoft feature BranchCache supported in clustered Data ONTAP?

Answer: Yes, starting in 8.2.

Does clustered Data ONTAP support both modes of BranchCache (hosted cache and distributed cache modes)?

Answer: The answer to that is yes, but with a clear distinction in that we are only a content server in the BranchCache architecture. An SVM is not participating as a caching server in a BranchCache setup. BranchCache works off of identifiers that are based on hashes. The first client to request access to the data is the client who initially caches the data (in a distributed model) or who sends the data to a central server for caching (in a hosted model). The second, third, and later clients who access the same data will be provided identifiers to locate the data on either its peers or a central caching server. For more details on BranchCache, see official Microsoft documentation on how the feature works and is architected.

Can I use the autolocation feature when using SMB3-CA shares and Hyper-V?

Answer: If the environment is running a release with a fix for bug 584472, then the answer is "yes". In clustered Data ONTAP releases prior to those containing the fix, autolocation created an issue involving machine account authorization. Hyper-V makes use of machine accounts to access SMB3-CA shares. When accessing a share created on a clustered Data ONTAP SVM, its possible a client might be routed, depending on name resolution configuration (ie: DNS records), to any node that has a LIF defined in DNS with the CIFS server name.

For example a CIFS server named "CIFS01" is created on an SVM residing on a cluster that has (4) nodes. Each of those nodes has a data LIF created to serve cifs traffic for "CIFS01". Within DNS you have created (4) separate A records that point to CIFS01 and IP address1, 2, 3, etc. A share is created that points to a volume on node1. The client attempts to access [\\cifs01.dom.local\share](#) and through name resolution gets routed to the LIF on node4. The SVM will see that the share is to the volume owned by node1 and will now send the client back a referral with the IP address for the LIF on node1. It's here that the authentication mechanism used by the client will change. When using the FQDN or hostname to access a CIFS server, Kerberos is generally the authentication mechanism used by clients. However, once the referral is issued the client falls back to using NTLM because they are sent the IP address as the referral. This causes the client to fallback to using NTLM due to how Kerberos works with

what are called SPNs or Service Principal Names. There is no method by which the client will attempt access to what in Kerberos is called a “service” using the IP address. Without the fix for bug 584472 clustered Data ONTAP CIFS servers will deny authentication by NTLM when the connection attempt is being made by a machine account. Please note that NTLM by general user accounts (ie: domainA\bobbyj) is supported even in absence of 584472. This is only for machine accounts accessing CIFS shares.

Is the SMB3 share property “continuously-available” supported for General Information worker type workloads (ie: home directories or project shares)?

Answer: As of the last update to this article the continuously-available share property is only supported for SMB3 connections for Hyper-V servers (starting in 8.2) and SQL server 2012 (starting in 8.2.1). The general information worker, those who use CIFS shares for project data or home directory storage, is not supported. Please also note that there is similar guidance from Microsoft on general information work and SMB3-CA type shares.

Is SMB3 Encryption supported in clustered Data ONTAP?

Answer: It is supported starting in 8.3.1

Do I have to maintain separate encryption keys for SMB3 encryption?

Answer: No, this is a self-contained encryption model. The keys are based on an AES-CCM algorithm and is all managed through the protocol itself. If you are interested in how the keys are generated, please consult the SMB specification located here [“Server Message Block Protocol Versions 2 and 3”](#) (sections 3.1.4.2 and 3.2.5.3).

You can enable SMB encryption at the SVM or share level through separate options. The recommendation is to enable it on a per share basis due to the nature of mixed clients versions in most environments. The feature is only supported by clients who support the SMB3 protocol. If the feature is enabled at the SVM level, then encryption will need to be negotiated for any share connection regardless of the per share setting (ie: the SVM setting supersedes the share setting).

Is there a performance impact to using SMB3 encryption?

Answer: Yes and it can be a notable impact. The impact will be similar to using SMB signing and could be higher. The impact is different because SMB signing is merely reading the data and computing a signature that will be added to the frame. SMB encryption is reading the data and then writing it back out in a secure encrypted format, this results in additional CPU usage. The performance changes will be seen in increased CPU despite the network traffic staying the same. The recommendation is to test using a workload that one expects to have when accessing data you intend to encrypt.

Does SMB3 encryption secure my data at-rest on disk?

Answer: No, this is securing the data in flight between source and destination.

Is there NDO for CIFS when activating a DR scenario when using SVMMDR?

Answer: Starting in 8.3.1 clustered Data ONTAP introduced support for SVMMDR. This is a very similar feature to vfiler DR for those familiar with 7-mode. The answer to the question is NO, there will be a client interruption to data access when you activate the DR side in a SVMMDR relationship. SVMMDR is an asynchronous replication DR technology that uses SnapMirror to replicate configuration details and user data. The SVMMDR relationship ensures the configurations information between source SVM and destination SVM is replicated.

****NOTE:** when using ID-Preserve configuration between source and destination match. When using ID-Discard CIFS domain, Name Resolution and Network settings are not replicated

Additionally, SVMDR ensures that user data is replicated from source SVM to the SVM in the DR cluster. However, no information about the current lock state for any open files are transferred over to the DR cluster. When activating the DR side in an actual DR scenario, the connected SMB clients prior to the DR scenario will have to re-establish connection to their previously opened files. Again this is an asynchronous DR technology between (2) separate clusters. For more information about SVMDR, please consult the product documentation for the release running in the cluster.

7 Differences Between 7-Mode and Clustered Data ONTAP

What is “node” scope versus “cluster” scope?

Answer: Node scope refers to data that represents the current node from which you are running commands or specified within the CLI itself. Cluster scope shows data across the cluster, no matter the node to which you are currently connected within the cluster.

What is the CIFS superuser?

Answer: Most file systems include a superuser account. Such an account enables system administrators to modify files and their attributes without reference to the usual file permission system. When a Windows user connects to the SVM and is listed as a superuser, that user's account is similar to a user defined as root.

Is there an equivalent to the 7-Mode option `cifs.tcp_window_size` in clustered Data ONTAP?

Answer: No. Prior to 8.2 there was no option to change the CIFS TCP window size. Starting in 8.2 the option was not necessary due to the implementation of TCP autotuning.

Why does clustered Data ONTAP not have the equivalent of the Data ONTAP 7-Mode feature for symlinks called “`nosymlink_strict_security`”?

Answer: This particular option from Data ONTAP 7-Mode was necessary because it allowed a mechanism to control access to symlink paths that didn't reside in the same volume path. In 7-Mode the path to a symlink could be constructed using an absolute path. This absolute path when constructed could link multiple volumes from the same controller into what appeared as a single namespace. To control whether access should be allowed across volumes, the CIFS share option “`nosymlink_strict_security`” was created. This option controls whether access is allowed if the target of a symlink travels outside the original entry point on which the client started (for example, a client connects to a CIFS share on volume HRData [/vol/HRdata] and through navigating encounters an absolute symlink that takes the client to /vol/ITdata).

The use of junction points in clustered Data ONTAP makes the use case for the 7-Mode option somewhat unnecessary. The use of junctions allows an environment to be designed to move in and out of volumes within a single share. However, for symlinks starting in clustered Data ONTAP 8.3, a new CIFS share option called “`is-symlink-strict-security-enabled`” was introduced. This option will control symlinks that redirect a client outside the share on which their connection originated. It is a per share setting and applies to both relative and absolute symlinks.

In 7-Mode there was the option `cifs.bypass_traverse_checking`, but I don't see it in clustered Data ONTAP. Is bypass traverse checking not supported?

Answer: Yes, it's supported. However, there is no option to configure this as of clustered Data ONTAP 8.2. Bypass traverse checking is available and enabled by default. For example, if a user is trying to reach [\\cifsservername\fooshare\subfoo1\subfoo2](#) but does not have actual permissions to subfoo1 (but does to fooshare and the subdirectory subfoo2), the user will be able to access the path if that user specifies the entire path in the request.

The exception to this is if the path to which the user connects includes mixed or UNIX security styles. If it does, the user needs to verify that that user has the X attribute on all folders along the path to the destination.

Starting in clustered Data ONTAP 8.3, a new privilege was created that will control whether bypass traverse checking security checks are done. This new privilege is called `SeChangeNotifyPrivilege`. This privilege is assigned by default to the following local groups: Administrators, Backup Operators, Power Users, Users, and Everyone. To control whether an account can still bypass security checks, you would need to remove that privilege from the local groups using a command like the following:

```
Cluster::*> cifs user-and-groups privilege remove-privilege -user-or-group-name BUILTIN\Users -privileges SeChangeNotifyPrivilege
```

This example command would remove the privilege for the BUILTIN\Users group.

Is there an equivalent to the 7-Mode CIFS option `cifs.grant_implicit_exe_perms` in clustered Data ONTAP?

Answer: Yes. First, here's a quick recap of what that option does in 7-Mode. In 7-Mode, this option gave a Windows client the capability of running an exe file located on a share backed by a UNIX security style volume or qtree where the UNIX permissions did not have the executable bit set. In clustered Data ONTAP that option is `-read-grants-exec` and is changed using the command `cifs options modify -read-grants-exec enabled -vserver <SVMName>`.

In clustered Data ONTAP SMB signing only has the single option “-is-signing-required” vs. 7-Mode, which has “`cifs.signing.enable`” or “`cifs.smb2.signing.required`.” How do I turn off SMB signing in clustered Data ONTAP?

Answer: In short, you don't. This behavior matches closely that of Windows starting in SMB2. The “is-signing-required” setting is for all SMB protocols. For more details around what Microsoft did, see the following article: <http://blogs.technet.com/b/josebda/archive/2010/12/01/the-basics-of-smb-signing-covering-both-smb1-and-smb2.aspx>.

In 7-Mode a common concern was pblks. Do pblks exist in clustered Data ONTAP?

Answer: The answer to this is no. If you are not familiar with pblks, locate the following KB article on the Support site that explains what they are and provides troubleshooting details: <https://kb.netapp.com/support/index?page=content&id=1013397>.

How are symlinks enabled in clustered Data ONTAP vs. 7-Mode?

Answer: In 7mode the use of symlinks was controlled through an option named “`cifs.symlinks.enable`”, which defaults to “enabled”. From there the type of symlink was controlled through the symlink translations file and allowed you to setup what are known as “widelinks” as well. For more details on symlinks in 7mode, see the File Access and Protocols Management Guide for the release of Data ONTAP you are running.

In clustered Data ONTAP there is no option like 7-mode, rather enabling symlinks is controlled via a cifs share option called “`symlink-properties`”. In releases up to 8.3 it had three possible settings: enable, hide and read_only. When set to “enable” it will allow read-write access to a symlink. If set to read_only, it will permit just read access to a symlink. The value of “hide” causes symlinks to not show up in a directory listing when viewed from a Windows or SMB capable client.

Starting in 8.3.1 due to changes for burt 881956 there are three additional possible settings: symlinks, symlinks_and_widelinks and disable. These were added to allow for additional granular control over how both symlinks (these are not NTFS symlinks) are accessible for SMB clients and also how an SMB feature called DFS is advertised.

For more information on setting up symlinks in clustered Data ONTAP, please see the File Access and Protocols Management guide for the release running on the cluster. Refer to the following KB article for further details on the impact of symlink settings in clustered Data ONTAP - <https://kb.netapp.com/support/index?page=content&id=3014510> .

I upgraded my 7-Mode system to clustered Data ONTAP and suddenly I am having troubles with clients accessing CIFS shares via clients who are using SMB2 and have either a WAN acceleration device from Riverbed or have removed the traverse/execute permission from folders?

Answer: This may be the result of changes in clustered Data ONTAP vs. Data ONTAP operating in 7-mode. There is a particular SMB feature called “DFS”, which stands for “Distributed File System”. In short DFS allows for multiple shares or data points, that may or may not live in the same file system (or on the same nodes), to be accessed from a central namespace. This feature is advertised to clients upon connection to a share and occurs during what is called the “Tree Connect”. A tree connect is the SMB call that is issued by the protocol when a client attempts to connect to a share on a server.

In 7-mode this particular DFS feature was only advertised when a particular share was configured with the “widelink” share setting. In cluster Data ONTAP, starting in 8.1 this particular feature was unconditionally set to “enabled”. This can create issues depending on the particular operation being attempted by a client or non-NetApp device.

- **WAN Accelerator from Riverbed:** With this particular DFS feature advertised, it will create a situation where it can no longer accelerate SMB2 traffic. As a Riverbed cannot accelerate SMB2 traffic that is housed on clustered Data ONTAP
- **Removed traverse/execute permissions:** When the DFS capability is advertised it appears to cause clients to ask for “execute” access to folders in a share. When this permission is removed, the client will be denied access to the resource. This is something outlined in a KB article from Microsoft - <http://support.microsoft.com/kb/2385108> .

For more details on this issue please see bug 746314 and the following Support Bulletin KB7010148 (which can be found on the NetApp Customer Support website). There is also a NetApp Support site KB article that explains the impact of symlinks and DFS. Please see “[What is the impact of a shares symlink settings and DFS advertisement in clustered Data ONTAP....](#)”

8 How It Works

The following section covers how CIFS works in clustered Data ONTAP.

For more information on CIFS, see the following RFCs:

- [SMB v1](#)
- [SMB v2 and 3](#)

How is access to CIFS determined?

Answer:

- Valid CIFS license
- CIFS server exists
- CIFS versions enabled
- CIFS allowed on the storage virtual machine (SVM)
- CIFS allowed on the data LIF
- CIFS share exists
- Required permissions on the share exist

- Required permissions on the files and folders exist

How can I check to see if CIFS is licensed?

Answer:

```
license show -package CIFS
```

How can I verify that a CIFS server exists?

Answer:

```
cifs server show -vserver [SVM_name]
```

How can I verify that the right CIFS version is enabled?

Answer:

```
cifs server options show -vserver <SVM_name> -fields smb2-enabled,smb3-enabled
```

How can I verify that CIFS is allowed for the SVM?

Answer:

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

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

How can I verify that CIFS is allowed on the network interface?

Answer:

```
network interface show -vserver [SVM_name] -fields data-protocol
```

If CIFS is not allowed on the data LIF, the LIF must be destroyed and recreated, or a new LIF needs to be created. The ports allowed on a data LIF depend on the allowed data protocol. For instance, if CIFS is not enabled on a data LIF, then ports 139 and 445 will not be available and as such not in a LISTENING state.

What happens when a client attempts to map a drive to a share located on a clustered Data ONTAP CIFS server (at a high level)?

Answer:

1. The client begins with a three-way TCP handshake to ports 139 and 445.
2. The CIFS server generally responds to both three-way TCP handshake requests.
3. The client generally breaks down the connection to port 139.
4. The client and server exchange SMB “Negotiate Protocol” calls to determine the SMB version to use for the connection.
5. The client and server exchange SMB “Session Setup” calls in which, among other things, the user is authenticated (that is, the user name or Kerberos ticket is exchanged here).
6. The client and server exchange SMB “Tree Connect” calls in which the client passes to the server the name of the share to which it intends to connect.

What is a namespace, and how does it relate to CIFS?

Answer: A namespace is a logical grouping of separate volumes joined together by junctions to create a single logical file system. It is logical in the sense that all the volumes might or might not be located on the

same node. When a volume is created, it can be assigned to a “junction path,” which helps to organize the volumes into a single namespace.

Shares, when created, are tied to a path in the namespace. For example, a volume called foovol is mounted in clustered Data ONTAP to /foovol. In order to access foovol, when a CIFS share is created it needs to be pointed to /foovol.

Later, if you create goovol and mount it under /foovol/goovol in clustered Data ONTAP, it will be available to those users who have access to the share pointing to foovol. When you then navigate using the share to foovol, users will see a folder call goovol. To a CIFS client, the junctions appear as an ordinary directory.

What is a volume junction?

Answer: When you create a volume in clustered Data ONTAP, a junction allows the volume to be mounted for later presentation using a CIFS share or an NFS export. A junction isn't required when you create a volume; however, a junction is necessary in order to make the volume available to NFS or CIFS clients.

What external resources does my CIFS server need to be able to reach using the LIFs assigned to the SVM?

Answer: The CIFS SVM needs to verify that it has a data LIF that can access any of the following:

- Active Directory servers
- NIS
- LDAP
- DNS

What are export policies, and do I need one?

Answer: Export policies for CIFS are effectively a third level of “security” beyond your typical share- and file-level permissions. In clustered Data ONTAP versions prior to 8.2, an export policy was mandatory. Starting in clustered Data ONTAP 8.2, they are optional and disabled by default for CIFS. For further information on export policies, consult the “File Access Management Guide for CIFS,” the section titled “Securing SMB Access Using Export Policies.”

What is the end-user experience during a storage failover event?

Answer: The answer depends on the protocol version that the client uses to connect to the storage controller. Depending on the protocol version running and the configuration of CIFS shares, during a storage failover event the node that does the takeover might have no record of open files on the failing node.

When a file is open over a CIFS share, the lock state for that file is held in memory. During an SFO event, the lock state is lost when the surviving node takes over. Depending on the version of the SMB protocol that the connection uses, it dictates the ability (or lack thereof) for the clients to survive an SFO event.

- **SMB1.** This version of the SMB protocol has no way to recover any previously open file handles; the protocol just doesn't have this capability. The open file handles are held in memory, and thus when the controller fails over to its partner, the lock state for open files is not present. A client sees an interruption to open files/folders at the time of a failover event.
- **SMB2/2.x.** This version of the SMB protocol introduced durable handles and an increased capability to survive some storage events. Durable handles allow brief network interruptions to occur between client and storage. During an SFO event, the client will be disconnected from the

controller while the LIF is migrated. After the LIF is migrated and brief network interruption resolved, a client will attempt to reconnect to a previously opened file from the failed node. However, a downfall of durable handles is that they are not guaranteed to survive a “reboot” of either client or server. Despite the ability to survive the brief network disruption, the loss of the state of the file between storage cluster nodes is why SMB2.x cannot survive an SFO event. Although they are a giant step forward for CIFS and resiliency, SMB2/2.x have limitations.

They do not survive server reboots, and multiuser access to the same file can cause the loss of durability because durable handles rely on certain oplock or lease levels to retain durability. In addition, durable handles are passive. A client needs to issue I/O against an open file so that the connection to the durable handle is maintained after a survivable event. Again, however, because lock state is kept in memory on the node where the file resides and because durable handles are designed for intermittent network drops versus full-scale reboots of hosts, the clients see a disruption in access to open files.

- **SMB3-non-CA (noncontinuously available).** The SMB3-non-CA (big distinction) acts pretty much like the SMB2 client.
- **SMB3-CA (continuously available).** Clients connecting to the storage controller using SMB3 and shares that have the continuously available share property set can survive a storage failover event. Generally the clients see a slight pause in any active I/O that occurred at the time. The survivability is made possible by features such as persistent handles, witness, and lock mirroring that were introduced in SMB3 and CA-capable shares/connections.

For more details on maintaining operational ability during storage events, consult TR-4100. That report details nondisruptive operations and the SMB protocol.

I changed the path to which my CIFS share points (referring to a new volume or qtree); however, the client still sees data from the old location.

Answer: If you change the path of the CIFS share to point to a new location and don't have your client disconnect and reconnect to the share, the client will still reference the old path. The reconnect is necessary due to the client tracking its connection using what's called a “tree ID,” or TID. The TID is provided during the initial connection to a CIFS share during the “tree connect” exchange. This ID is used to associate a client's connection to a share path.

What is the maximum character length supported for a path in CIFS?

Answer: This is traditionally 255 characters and is not controlled by clustered Data ONTAP. See the following Microsoft link for more details about this limitation: <http://msdn.microsoft.com/en-us/library/aa365247.aspx>.

Is it possible to control to which LIF the CIFS feature “autolocation” refers a client?

Answer: Currently there is no way to control to which LIF the autolocation feature refers a client. The autolocation referral works by looking at the SVM's data LIFs relevant to the location of the volume itself. If the data LIFs where the volume resides are assigned the NFS and CIFS protocols, then those LIFs are candidates for referral. If an interface is down but would normally meet the criteria for referral, it will not be used. The interface needs to be in the up state; it must also have both CIFS and NFS as allowed protocols.

What does the “status” column information mean in the output of “cifs domain discovered-servers show”?

Answer: The following are the various status outputs you might see.

Status	What It Means
OK	The connection to this server is usable.
Unavailable	The server is currently unavailable for use.
Slow	The connection to the server is usable but has been determined to be slow.
Expired	The connection to this server expired.
Undetermined	This server is one discovered but not yet used.
Unreachable	This server was discovered but is no longer reachable.

8.1 Support Considerations

Which network ports might need to be reachable by a CIFS server?

Answer: The following list shows the ports used for CIFS by clustered Data ONTAP.

Table 3) Network ports for CIFS.

Protocol	Destination Port	Source Port	Port Details
TCP/UDP	53	1024-65535	DNS
TCP/UDP	88	1024-65535	Kerberos
TCP	135	1024-65535	RPC
TCP/UDP	137	1024-65535	NetBIOS name server
UDP	138	1024-65535	NetBIOS datagram
TCP	139	1024-65535	NetBIOS session services
TCP/UDP	389	1024-65535	LDAP
TCP	445	1024-65535	SMB over TCP

Protocol	Destination Port	Source Port	Port Details
TCP/UDP	464	1024-65535	Kpasswd

9 Miscellaneous

I have folders in a share that one client can see but not another. It occurs on Windows 7 clients and later when using SMB2 to connect.

Answer: This issue could very well be due to a client side introduced with the SMB2 protocol. See the following article for Microsoft that explains the client-side caches and how to disable them:

<http://technet.microsoft.com/en-us/library/ff686200%28v=ws.10%29.aspx>.

Can I use the MMC to manage a clustered Data ONTAP CIFS server?

Answer: Yes and no. Prior to clustered Data ONTAP 8.3, much of what you see in the MMC is view only. Starting with 8.3 we introduced the ability to manage the following: create/delete CIFS shares, manage open files, and manage client sessions. When using the MMC to manage open files and client sessions, the view of those will be node specific. This means that when you use the MMC and add in the “Computer Management” snap-in specifying the CIFS server as the computer to which you want to connect, if you use an FQDN or short name, then when the connection is made, DNS is consulted to resolve the name. On whichever cluster node the connection is made, the display of sessions and open files will be what that node knows.

For example: An SVM has four nodes, and each has a data LIF. Each data LIF is registered in DNS to serve traffic for a CIFS server named CIFS01. When you specify in the Computer Management snap-in the hostname “CIFS01,” DNS will be consulted, and the management workstation will connect to one of four nodes. Let’s say all the open files and sessions are on node 4 of the cluster, but the snap-in connects to node 1. When viewing the open files and session, the MMC will show nothing or just the session established by the admin workstation. In order to see all the open files and session, you will need to add another Computer Management snap-in and specify the IP address associated with the data LIF that resides on node 4.

The other alternative would be to use the Data ONTAP PowerShell Toolkit, version 3.2. That version of the toolkit contains cmdlets that will provide a clusterwide enumeration of open files and sessions. The cmdlets are get-nccifssession, get-nccifssessionfile, close-nccifssession, and close-nccifssessionfile.

Does NetApp Support the Microsoft feature called “Work Folders”?

Answer: The answer to this is No. The reason for this has to do with the fact that the feature Work Folders requires the use of locally attached storage. So the lack of support is not a NetApp limitation but rather how Microsoft implemented the feature.

What does the symlink-properties option do when creating or modifying the particular setting on a cifs share?

Answer: This particular option controls whether symlink resolution for CIFS is enabled on a share. The setting is a per-share setting and depending on the release of clustered Data ONTAP you are running, it

has different possible settings:

Clustered Data ONTAP 8.3 and earlier	Clustered Data ONTAP 8.3.1 and later
Enable	Enable
Read_only,enable	Read_only,enable
Hide	Hide
"" (null) or "-"	"" (null) or "-"
	Symlinks
	Symlinks_and_widelinks
	disable

The settings for symlink properties is important for multiple reasons. The first is it controls how symlinks are advertised and accessible to SMB clients (note these are not NTFS symlinks, which are not supported). For all versions of clustered Data ONTAP as the chart above shows enable, read_only,enable and hide / "" (null) / - are all the same. Starting in 8.3.1 several additional settings were introduced to provide additional control over symlinks and a feature called "DFS advertisement". This additional granularity is possible by specifying symlinks, symlinks_and_widelinks or disabled. The explanation of each of these is covered in the "File Access and Protocols Management guide" for the relative clustered Data ONTAP releases. Another resource to consult on the importance of these settings is NetApp Customer Support KB article titled ["What is the impact of a shares symlink settings and DFS advertisement in clustered Data ONTAP on SMB2.x/3 client traffic"](#)

I need to enable a version of the SMB protocol in my SVM that was previously disabled (whether by default or manually disabled), will it impact data access for existing clients?

Answer: No, currently connected clients will continue to use the protocol version negotiated with the cDOT CIFS server during the initial connection. New connections that are established will attempt to make use of the newly enabled SMB protocol versions depending on the clients capability. When clients attempt initiates connections to a CIFS server, they will present a list of SMB dialect versions they support. This is done through what is called the "Negotiate Protocol" exchange. The server will then look at the list, selecting from the provided client list the highest version it supports AND has enabled. So this is why a client with an existing connection will continue to work just fine and why new clients can potentially take advantage of the different protocol versions enabled.

Version History

Version	Date	Document Version History
Version 1.0	April 2014	First publishing
Version 1.1	November 2014	Update to include new questions since first publishing
Version 1.2	January 2015	Update to include questions since November
Version 1.3	March 2015	Update to include questions since January 2015
Version 1.4	June 2015	Updates for clustered Data ONTAP 8.3.1 release

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

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

© 2015 NetApp, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. Specifications is subject to change without notice. NetApp, the NetApp logo, Data ONTAP, FlexVol, and FPolicy are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. UNIX is a registered trademark of The Open Group. Microsoft, Active Directory, Hyper-V, SQL Server, and Windows are registered trademarks of Microsoft Corporation. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.