

Deploy SMB server-based services ONTAP 9

NetApp September 18, 2024

This PDF was generated from https://docs.netapp.com/us-en/ontap/smb-admin/dynamic-home-directories-concept.html on September 18, 2024. Always check docs.netapp.com for the latest.

Table of Contents

D	eploy SMB server-based services	1
	Manage home directories	1
	Configure SMB client access to UNIX symbolic links	. 14
	Use BranchCache to cache SMB share content at a branch office	. 22
	Improve Microsoft remote copy performance.	. 53
	Improve client response time by providing SMB automatic node referrals with Auto Location	. 60
	Provide folder security on shares with access-based enumeration	. 66

Deploy SMB server-based services

Manage home directories

How ONTAP enables dynamic home directories

ONTAP home directories enable you to configure an SMB share that maps to different directories based on the user that connects to it and a set of variables. Instead of creating separate shares for each user, you can configure one share with a few home directory parameters to define a user's relationship between an entry point (the share) and the home directory (a directory on the SVM).

A user who is logged in as a guest user does not have a home directory and cannot access other users' home directories. There are four variables that determine how a user is mapped to a directory:

Share name

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

The share name can use the following dynamic names:

- ° %w (the user's Windows user name)
- ° %d (the user's Windows domain name)
- %u (the user's mapped UNIX user name) To make the share name unique across all home directories, the share name must contain either the %w or the %u variable. The share name can contain both the %d and the %w variable (for example, %d/%w), or the share name can contain a static portion and a variable portion (for example, home_%w).

Share path

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

Search paths

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

• Directory

This is the user's home directory that you create for the user. The directory name is usually the user's name. You must create the home directory in one of the directories that are defined by the search paths.

As an example, consider the following setup:

- User: John Smith
- User domain: acme
- · User name: jsmith
- SVM name: vs1
- Home directory share name #1: home_%w share path: %w
- Home directory share name #2: %w share path: %d/%w
- Search path #1: /vol0home/home
- Search path #2: /vol1home/home
- Search path #3: /vol2home/home
- Home directory: /vol1home/home/jsmith

Scenario 1: The user connects to \\vs1\home_jsmith. This matches the first home directory share name and generates the relative path jsmith. ONTAP now searches for a directory named jsmith by checking each search path in order:

- /vol0home/home/jsmith does not exist; moving on to search path #2.
- /vollhome/home/jsmith does exist; therefore, search path #3 is not checked; the user is now connected to his home directory.

Scenario 2: The user connects to \\vs1\jsmith. This matches the second home directory share name and generates the relative path acme/jsmith. ONTAP now searches for a directory named acme/jsmith by checking each search path in order:

- /vol0home/home/acme/jsmith does not exist; moving on to search path #2.
- /vollhome/home/acme/jsmith does not exist; moving on to search path #3.
- /vol2home/home/acme/jsmith does not exist; the home directory does not exist; therefore, the connection fails.

Home directory shares

Add a home directory share

If you want to use the SMB home directory feature, you must add at least one share with the home directory property included in the share properties.

About this task

You can create a home directory share at the time you create the share by using the vserver cifs share create command, or you can change an existing share into a home directory share at any time by using the vserver cifs share modify command.

To create a home directory share, you must include the homedirectory value in the -share-properties option when you create or modify a share. You can specify the share name and share path using variables that are dynamically expanded when users connect to their home directories. Available variables that you can use in the path are w, d, and u, corresponding to the Windows user name, domain, and mapped UNIX user name, respectively.

Steps

1. Add a home directory share:

```
vserver cifs share create -vserver vserver_name -share-name share_name -path
path -share-properties homedirectory[,...]
```

-vserver vserver specifies the CIFS-enabled storage virtual machine (SVM) on which to add the search path.

-share-name share-name specifies the home directory share name.

In addition to containing one of the required variables, if the share name contains one of the literal strings w, u, or d, you must precede the literal string with a % (percent) character to prevent ONTAP from treating the literal string as a variable (for example, w).

- $^\circ\,$ The share name must contain either the \$w or the \$u variable.
- The share name can additionally contain the %d variable (for example, %d/%w) or a static portion in the share name (for example, home1_%w).
- If the share is used by administrators to connect to other users' home directories or to permit users to connect to other users' home directories, the dynamic share name pattern must be preceded by a tilde (~).

```
The vserver cifs home-directory modify is used to enable this access by setting the -is -home-dirs-access-for-admin-enabled option to true) or by setting the advanced option -is -home-dirs-access-for-public-enabled to true.
```

-path path specifies the relative path to the home directory.

-share-properties homedirectory[,...] specifies the share properties for that share. You must specify the homedirectory value. You can specify additional share properties using a comma delimited list.

1. Verify that you successfully added the home directory share by using the vserver cifs share show command.

Example

The following command creates a home directory share named *%w*. The oplocks, browsable, and changenotify share properties are set in addition to setting the homedirectory share property.



This example does not display output for all of the shares on the SVM. Output is truncated.

cluster1::> vserver cifs share create -vserver vs1 -share-name %w -path %w
-share-properties oplocks,browsable,changenotify,homedirectory

Related information

Adding a home directory search path

Requirements and guidelines for using automatic node referrals

Managing accessibility to users' home directories

Home directory shares require unique user names

Be careful to assign unique user names when creating home directory shares using the \Re_W (Windows user name) or \Re_U (UNIX user name) variables to generate shares dynamically. The share name is mapped to your user name.

Two problems can occur when a static share's name and a user's name are the same:

- When the user lists the shares on a cluster using the net view command, two shares with the same user name are displayed.
- When the user connects to that share name, the user is always connected to the static share and cannot access the home directory share with the same name.

For example, there is a share named "administrator" and you have an "administrator" Windows user name. If you create a home directory share and connect to that share, you get connected to the "administrator" static share, not to your "administrator" home directory share.

You can resolve the issue with duplicate share names by following any of these steps:

- Renaming the static share so that it no longer conflicts with the user's home directory share.
- Giving the user a new user name so that it no longer conflicts with the static share name.
- Creating a CIFS home directory share with a static name such as "home" instead of using the *w* parameter to avoid conflicts with the share names.

What happens to static home directory share names after upgrading

Home directory share names must contain either the w or the u dynamic variable. You should be aware of what happens to existing static home directory share names after upgrading to a version of ONTAP with the new requirement.

If your home directory configuration contains static share names and you upgrade to ONTAP, the static home directory share names are not changed and are still valid. However, you cannot create any new home directory shares that do not contain either the w or u variable.

Requiring that one of these variables is included in the user's home directory share name ensures that every share name is unique across the home directory configuration. If desired, you can change the static home directory share names to names that contain either the w or u variable.

Add a home directory search path

If you want to use ONTAP SMB home directories, you must add at least one home directory search path.

About this task

You can add a home directory search path by using the vserver cifs home-directory search-path add command.

The vserver cifs home-directory search-path add command checks the path specified in the -path option during command execution. If the specified path does not exist, the command generates a message prompting for whether you want to continue. You choose y or n. If you choose y to continue, ONTAP creates the search path. However, you must create the directory structure before you can use the search path in the home directory configuration. If you choose not to continue, the command fails; the search path is not created. You can then create the path directory structure and rerun the vserver cifs home-directory search-path add command.

Steps

- 1. Add a home directory search path: vserver cifs home-directory search-path add -vserver vserver -path path
- 2. Verify that you successfully added the search path using the vserver cifs home-directory search-path show command.

Example

The following example adds the path /home1 to the home directory configuration on SVM vs1.

The following example attempts to add the path /home2 to the home directory configuration on SVM vs1. The path does not exist. The choice is made to not continue.

Related information

Adding a home directory share

Create a home directory configuration using the %w and %d variables

You can create a home directory configuration using the Sw and Sd variables. Users can then connect to their home share using dynamically created shares.

Steps

- 1. Create a qtree to contain user's home directories: volume qtree create -vserver vserver_name -qtree-path qtree_path
- 2. Verify that the qtree is using the correct security style: volume qtree show
- 3. If the qtree is not using the desired security style, change the security style using the volume qtree security command.
- 4. Add a home directory share: vserver cifs share create -vserver vserver -share-name %w -path %d/%w -share-properties homedirectory\[,...\]

-vserver vserver specifies the CIFS-enabled storage virtual machine (SVM) on which to add the search path.

-share-name %w specifies the home directory share name. ONTAP dynamically creates the share name as each user connects to their home directory. The share name will be of the form *windows_user_name*.

-path %d/%w specifies the relative path to the home directory. The relative path is dynamically created as each user connects to their home directory and will be of the form *domain/windows_user_name*.

-share-properties homedirectory[,...] + specifies the share properties for that share. You must specify the homedirectory value. You can specify additional share properties using a comma delimited list.

- 5. Verify that the share has the desired configuration using the vserver cifs share show command.
- 6. Add a home directory search path: vserver cifs home-directory search-path add -vserver vserver -path path

-vserver vserver-name specifies the CIFS-enabled SVM on which to add the search path.

-path path specifies the absolute directory path to the search path.

- 7. Verify that you successfully added the search path using the vserver cifs home-directory search-path show command.
- 8. For users with a home directory, create a corresponding directory in the qtree or volume designated to

contain home directories.

For example, if you created a qtree with the path of /vol/vol1/users and the user name whose directory you want to create is mydomain\user1, you would create a directory with the following path: /vol/vol1/users/mydomain/user1.

If you created a volume named "home1" mounted at /home1, you would create a directory with the following path: /home1/mydomain/user1.

9. Verify that a user can successfully connect to the home share either by mapping a drive or connecting using the UNC path.

For example, if user mydomain\user1 wants to connect to the directory created in Step 8 that is located on SVM vs1, user1 would connect using the UNC path \\vs1\user1.

Example

The commands in the following example create a home directory configuration with the following settings:

- The share name is %w.
- The relative home directory path is %d/%w.
- The search path that is used to contain the home directories, /home1, is a volume configured with NTFS security style.
- The configuration is created on SVM vs1.

You can use this type of home directory configuration when users access their home directories from Windows hosts. You can also use this type of configuration when users access their home directories from Windows and UNIX hosts and the file system administrator uses Windows-based users and groups to control access to the file system.

cluster::> vserver cifs share create -vserver vs1 -share-name %w -path %d/%w -share-properties oplocks,browsable,changenotify,homedirectory cluster::> vserver cifs share show -vserver vs1 -share-name %w Vserver: vsl Share: %w CIFS Server NetBIOS Name: VS1 Path: %d/%w Share Properties: oplocks browsable changenotify homedirectory Symlink Properties: enable File Mode Creation Mask: -Directory Mode Creation Mask: -Share Comment: -Share ACL: Everyone / Full Control File Attribute Cache Lifetime: -Volume Name: -Offline Files: manual Vscan File-Operations Profile: standard cluster::> vserver cifs home-directory search-path add -vserver vs1 -path /home1 cluster::> vserver cifs home-directory search-path show Vserver Position Path -----1 vs1 /home1

Related information

Configuring home directories using the %u variable

Additional home directory configurations

Displaying information about an SMB user's home directory path

Configure home directories using the %u variable

You can create a home directory configuration where you designate the share name using the w variable but you use the u variable to designate the relative path to the home directory share. Users can then connect to their home share using dynamically shares created using their Windows user name without being aware of the actual name or path of the home directory.

Steps

- 1. Create a qtree to contain user's home directories: volume qtree create -vserver vserver_name -qtree-path qtree path
- 2. Verify that the qtree is using the correct security style: volume qtree show
- 3. If the qtree is not using the desired security style, change the security style using the volume qtree security command.
- 4. Add a home directory share: vserver cifs share create -vserver vserver -share-name %w -path %u -share-properties homedirectory ,...]

-vserver vserver specifies the CIFS-enabled storage virtual machine (SVM) on which to add the search path.

-share-name %w specifies the home directory share name. The share name is dynamically created as each user connects to their home directory and is of the form *windows_user_name*.



You can also use the %u variable for the -share-name option. This creates a relative share path that uses the mapped UNIX user name.

-path %u specifies the relative path to the home directory. The relative path is created dynamically as each user connects to their home directory and is of the form *mapped_UNIX_user_name*.



The value for this option can contain static elements as well. For example, eng/8u.

-share-properties homedirectory\[,...\] specifies the share properties for that share. You must specify the homedirectory value. You can specify additional share properties using a comma delimited list.

- 5. Verify that the share has the desired configuration using the vserver cifs share show command.
- 6. Add a home directory search path: vserver cifs home-directory search-path add -vserver vserver -path path

-vserver vserver specifies the CIFS-enabled SVM on which to add the search path.

-path path specifies the absolute directory path to the search path.

- 7. Verify that you successfully added the search path using the vserver cifs home-directory search-path show command.
- 8. If the UNIX user does not exist, create the UNIX user using the vserver services unix-user create command.



The UNIX user name to which you map the Windows user name must exist before mapping the user.

9. Create a name mapping for the Windows user to the UNIX user using the following command: vserver name-mapping create -vserver vserver_name -direction win-unix -priority integer -pattern windows_user_name -replacement unix_user_name



If name mappings already exist that map Windows users to UNIX users, you do not have to perform the mapping step.

The Windows user name is mapped to the corresponding UNIX user name. When the Windows user connects to their home directory share, they connect to a dynamically created home directory with a share name that corresponds to their Windows user name without being aware that the directory name corresponds to the UNIX user name.

10. For users with a home directory, create a corresponding directory in the qtree or volume designated to contain home directories.

For example, if you created a qtree with the path of /vol/vol1/users and the mapped UNIX user name of the user whose directory you want to create is "unixuser1", you would create a directory with the following path: /vol/vol1/users/unixuser1.

If you created a volume named "home1" mounted at /home1, you would create a directory with the following path: /home1/unixuser1.

11. Verify that a user can successfully connect to the home share either by mapping a drive or connecting using the UNC path.

For example, if user mydomain\user1 maps to UNIX user unixuser1 and wants to connect to the directory created in Step 10 that is located on SVM vs1, user1 would connect using the UNC path \\vs1\user1.

Example

The commands in the following example create a home directory configuration with the following settings:

- The share name is %w.
- The relative home directory path is %u.
- The search path that is used to contain the home directories, /home1, is a volume configured with UNIX security style.
- The configuration is created on SVM vs1.

You can use this type of home directory configuration when users access their home directories from both Windows hosts or Windows and UNIX hosts and the file system administrator uses UNIX-based users and groups to control access to the file system.

cluster::> vserver cifs share create -vserver vs1 -share-name %w -path %u -share-properties oplocks, browsable, changenotify, homedirectory cluster::> vserver cifs share show -vserver vs1 -share-name %u Vserver: vsl Share: %w CIFS Server NetBIOS Name: VS1 Path: %u Share Properties: oplocks browsable changenotify homedirectory Symlink Properties: enable File Mode Creation Mask: -Directory Mode Creation Mask: -Share Comment: -Share ACL: Everyone / Full Control File Attribute Cache Lifetime: -Volume Name: -Offline Files: manual Vscan File-Operations Profile: standard cluster::> vserver cifs home-directory search-path add -vserver vs1 -path /home1 cluster::> vserver cifs home-directory search-path show -vserver vs1 Vserver Position Path ----vs1 1 /home1 cluster::> vserver name-mapping create -vserver vs1 -direction win-unix -position 5 -pattern user1 -replacement unixuser1 cluster::> vserver name-mapping show -pattern user1 Direction Position Vserver ----- ----win-unix 5 Pattern: user1 vs1 Replacement: unixuser1

Related information

Creating a home directory configuration using the %w and %d variables

Additional home directory configurations

Displaying information about an SMB user's home directory path

Additional home directory configurations

You can create additional home directory configurations using the w, d, and u variables, which enables you to customize the home directory configuration to meet your needs.

You can create a number of home directory configurations using a combination of variables and static strings in the share names and search paths. The following table provides some examples illustrating how to create different home directory configurations:

Paths created when /vol1/user contains home directories	Share command
To create a share path \\vs1\~win_username that directs the user to /vol1/user/win_username	<pre>vserver cifs share create -share-name ~%w -path %w -share-properties oplocks,browsable,changenotify,homedire ctory</pre>
To create a share path \\vs1\win_username that directs the user to /vol1/user/domain/win_username	<pre>vserver cifs share create -share-name %w -path %d/%w -share-properties oplocks,browsable,changenotify,homedire ctory</pre>
To create a share path \\vs1\win_username that directs the user to /vol1/user/unix_username	<pre>vserver cifs share create -share-name %w -path %u -share-properties oplocks,browsable,changenotify,homedire ctory</pre>
To create a share path \\vs1\unix_username that directs the user to /vol1/user/unix_username	<pre>vserver cifs share create -share-name %u -path %u -share-properties oplocks,browsable,changenotify,homedire ctory</pre>

Commands for managing search paths

There are specific ONTAP commands for managing search paths for SMB home directory configurations. For example, there are commands for adding, removing, and displaying information about search paths. There is also a command for changing the search path order.

If you want to	Use this command
Add a search path	vserver cifs home-directory search-path add
Display search paths	vserver cifs home-directory search-path show

If you want to	Use this command
Change the search path order	vserver cifs home-directory search-path reorder
Remove a search path	vserver cifs home-directory search-path remove

See the man page for each command for more information.

Display information about an SMB user's home directory path

You can display an SMB user's home directory path on the storage virtual machine (SVM), which can be used if you have multiple CIFS home directory paths configured and you want to see which path holds the user's home directory.

Step

1. Display the home directory path by using the vserver cifs home-directory show-user command.

vserver cifs home-directory show-user -vserver vs1 -username user1

Vserver	User	Home Dir Path
vs1	user1	/home/user1

Related information

Managing accessibility to users' home directories

Manage accessibility to users' home directories

By default, a user's home directory can be accessed only by that user. For shares where the dynamic name of the share is preceded with a tilde (~), you can enable or disable access to users' home directories by Windows administrators or by any other user (public access).

Before you begin

Home directory shares on the storage virtual machine (SVM) must be configured with dynamic share names that are preceded with a tilde (~). The following cases illustrate share naming requirements:

Home directory share name	Example of command to connect to the share
~%d~%w	net use * \\IPaddress\~domain~user/u:credentials

Home directory share name	Example of command to connect to the share
~%w	net use * \\IPaddress\~user/u:credentials
~abc~%w	net use * \\IPaddress\abc~user/u:credentials

Step

1. Perform the appropriate action:

If you want to enable or disable access to users' home directories to	Enter the following
Windows administrators	<pre>vserver cifs home-directory modify -vserver vserver_name -is-home-dirs -access-for-admin-enabled {true false} The default is true.</pre>
Any user (public access)	a. Set the privilege level to advanced: set -privilege advanced
	b. Enable or disable access: vserver cifs home-directory modify -vserver vserver_name -is-home-dirs-access -for-public-enabled {true false} The default is false.
	c. Return to the admin privilege level: set -privilege admin

The following example enables public access to users' home directories:

```
set -privilege advanced
vserver cifs home-directory modify -vserver vs1 -is-home-dirs-access-for-public
-enabled true
set -privilege admin
```

Related information

Displaying information about an SMB user's home directory path

Configure SMB client access to UNIX symbolic links

How ONTAP enables you to provide SMB client access to UNIX symbolic links

A symbolic link is a file that is created in a UNIX environment that contains a reference to another file or directory. If a client accesses a symbolic link, the client is redirected to the target file or directory to which the symbolic link refers. ONTAP supports relative and absolute symbolic links, including widelinks (absolute links with targets outside the local file system).

ONTAP provides SMB clients the ability to follow UNIX symbolic links that are configured on the SVM. This feature is optional, and you can configure it on a per-share basis, using the <code>-symlink-properties</code> option of the <code>vserver cifs</code> share create command, with one of the following settings:

- · Enabled with read/write access
- · Enabled with read-only access
- Disabled by hiding symbolic links from SMB clients
- · Disabled with no access to symbolic links from SMB clients

If you enable symbolic links on a share, relative symbolic links work without further configuration.

If you enable symbolic links on a share, absolute symbolic links do not work right away. You must first create a mapping between the UNIX path of the symbolic link to the destination SMB path. When creating absolute symbolic link mappings, you can specify whether it is a local link or a *widelink*; widelinks can be links to file systems on other storage devices or links to file systems hosted in separate SVMs on the same ONTAP system. When you create a widelink, it must include the information for the client to follow; that is, you create a reparse point for the client to discover the directory junction point. If you create an absolute symbolic link to a file or directory outside of the local share but set the locality to local, ONTAP disallows access to the target.



If a client attempts to delete a local symbolic link (absolute or relative), only the symbolic link is deleted, not the target file or directory. However, if a client attempts to delete a widelink, it might delete the actual target file or directory to which the widelink refers. ONTAP does not have control over this because the client can explicitly open the target file or directory outside the SVM and delete it.

Reparse points and ONTAP file system services

A *reparse point* is an NTFS file system object that can be optionally stored on volumes along with a file. Reparse points provide SMB clients the ability to receive enhanced or extended file system services when working with NTFS style volumes. Reparse points consist of standard tags that identify the type of reparse point, and the content of the reparse point that can be retrieved by SMB clients for further processing by the client. Of the object types available for extended file system functionality, ONTAP implements support for NTFS symbolic links and directory junction points using reparse point tags. SMB clients that cannot understand the contents of a reparse point simply ignore it and don't provide the extended file system service that the reparse point might enable.

Directory junction points and ONTAP support for symbolic links

Directory junction points are locations within a file system directory structure that can refer to alternate locations where files are stored, either on a different path (symbolic links) or a separate storage device (widelinks). ONTAP SMB servers expose directory junction points to Windows clients as reparse points, allowing capable clients to obtain reparse point contents from ONTAP when a directory junction point is traversed. They can thereby navigate and connect to different paths or storage devices as though they were part of the same file system.

Enabling widelink support using reparse point options

The -is-use-junctions-as-reparse-points-enabled option is enabled by default in ONTAP 9. Not all SMB clients support widelinks, so the option to enable the information is configurable on a perprotocol version basis, allowing administrators to accommodate both supported and non-supported SMB clients. In ONTAP 9.2 and later releases, you must enable the option -widelink-as-reparse-point -versions for each client protocol that accesses the share using widelinks; the default is SMB1. In earlier releases, only widelinks accessed using the default SMB1 were reported, and systems using SMB2 or SMB3 were unable to access the widelinks.

Related information

- Windows backup applications and Unix-style symlinks
- Microsoft Documentation: Reparse Points

Limits when configuring UNIX symbolic links for SMB access

You need to be aware of certain limits when configuring UNIX symbolic links for SMB access.

Limit	Description
45	Maximum length of the CIFS server name that you can specify when using an FQDN for the CIFS server name.Image: Constant of the cife
80	Maximum length of the share name.
256	Maximum length of the UNIX path that you can specify when creating a symbolic link or when modifying an existing symbolic link's UNIX path.The UNIX path must start with a "/" (slash) and end with a "/". Both the beginning and ending slashes count as part of the 256-character limit.
256	Maximum length of the CIFS path that you can specify when creating a symbolic link or when modifying an existing symbolic link's CIFS path.The CIFS path must start with a "/" (slash) and end with a "/". Both the beginning and ending slashes count as part of the 256-character limit.

Related information

Creating symbolic link mappings for SMB shares

Control automatic DFS advertisements in ONTAP with a CIFS server option

A CIFS server option controls how DFS capabilities are advertised to SMB clients when connecting to shares. Because ONTAP uses DFS referrals when clients access symbolic links over SMB, you should be aware of what the impact is when disabling or enabling this option.

A CIFS server option determines whether the CIFS servers automatically advertise that they are DFS capable to SMB clients. By default, this option is enabled and the CIFS server always advertises that it is DFS capable

to SMB clients (even when connecting to shares where access to symbolic links is disabled). If you want the CIFS server to advertise that it is DFS capable to clients only when they are connecting to shares where access to symbolic links is enabled, you can disable this option.

You should be aware of what happens when this option is disabled:

- The share configurations for symbolic links is unchanged.
- If the share parameter is set to allow symbolic link access (either read-write access or read-only access), the CIFS server advertises DFS capabilities to clients connecting to that share.

Client connections and access to symbolic links continue without interruption.

• If the share parameter is set to not allow symbolic link access (either by disabling access or if the value for the share parameter is null), the CIFS server does not advertise DFS capabilities to clients connecting to that share.

Because clients have cached information that the CIFS server is DFS capable and it is no longer advertising that it is, clients that are connected to shares where symbolic link access is disabled might not be able to access these shares after the CIFS server option is disabled. After the option is disabled, you might need to reboot clients that are connected to these shares, thus clearing the cached information.

These changes do not apply to SMB 1.0 connections.

Configure UNIX symbolic link support on SMB shares

You can configure UNIX symbolic link support on SMB shares by specifying a symbolic link share-property setting when you create SMB shares or at any time by modifying existing SMB shares. UNIX symbolic link support is enabled by default. You can also disable UNIX symbolic link support on a share.

About this task

When configuring UNIX symbolic link support for SMB shares, you can choose one of the following settings:

Setting	Description
enable (DEPRECATED*)	Specifies that symbolic links are enabled for read- write access.
<pre>read_only (DEPRECATED*)</pre>	Specifies that symlinks are enabled for read-only access. This setting does not apply to widelinks. Widelink access is always read-write.
hide (DEPRECATED*)	Specifies that SMB clients are prevented from seeing symlinks.
no-strict-security	Specifies that clients follow symlinks outside of share boundaries.

Setting	Description
symlinks	Specifies that symlinks are enabled locally for read- write access. The DFS advertisements are not generated even if the CIFS option is-advertise- dfs-enabled is set to true. This is the default setting.
symlinks-and-widelinks	Specifies that both local symlinks and widelinks for read-write access. The DFS advertisements are generated for both local symlink and widelinks even if the CIFS option is-advertise-dfs-enabled is set to false.
disable	Specifies that symlinks and widelinks are disabled. The DFS advertisements are not generated even if the CIFS option is-advertise-dfs-enabled is set to true.
"" (null, not set)	Disables symbolic links on the share.
- (not set)	Disables symbolic links on the share.



*The *enable*, *hide*, and *read-only* parameters are deprecated and may be removed in a future release of ONTAP.

Steps

1. Configure or disable symbolic link support:

If it is	Enter
A new SMB share	<pre>vserver cifs share create -vserver vserver_name -share-name share_name -path path -symlink-properties {enable hide read-only "" - symlinks symlinks-and- widelinks disable},]</pre>
An existing SMB share	<pre>vserver cifs share modify -vserver vserver_name -share-name share_name -symlink-properties {enable hide read- only "" - symlinks symlinks-and- widelinks disable},]</pre>

2. Verify that the SMB share configuration is correct: vserver cifs share show -vserver vserver_name -share-name share_name -instance

Example

The following command creates an SMB share named "data1" with the UNIX symbolic link configuration set to enable:

```
cluster1::> vserver cifs share create -vserver vs1 -share-name data1 -path
/data1 -symlink-properties enable
cluster1::> vserver cifs share show -vserver vs1 -share-name data1
-instance
                          Vserver: vsl
                            Share: data1
         CIFS Server NetBIOS Name: VS1
                             Path: /data1
                 Share Properties: oplocks
                                   browsable
                                   changenotify
               Symlink Properties: enable
          File Mode Creation Mask: -
     Directory Mode Creation Mask: -
                    Share Comment: -
                        Share ACL: Everyone / Full Control
    File Attribute Cache Lifetime: -
                      Volume Name: -
                    Offline Files: manual
    Vscan File-Operations Profile: standard
Maximum Tree Connections on Share: 4294967295
       UNIX Group for File Create: -
```

Related information

Creating symbolic link mappings for SMB shares

Create symbolic link mappings for SMB shares

You can create mappings of UNIX symbolic links for SMB shares. You can either create a relative symbolic link, which refers to the file or folder relative to its parent folder, or you can create an absolute symbolic link, which refers to the file or folder using an absolute path.

About this task

Widelinks are not accessible from Mac OS X clients if you use SMB 2.x. When a user attempts to connect to a share using widelinks from a Mac OS X client, the attempt fails. However, you can use widelinks with Mac OS X clients if you use SMB 1.

Steps

1. To create symbolic link mappings for SMB shares: vserver cifs symlink create -vserver virtual_server_name -unix-path path -share-name share_name -cifs-path path [- cifs-server server_name] [-locality {local|free|widelink}] [-home-directory

-vserver virtual server name specifies the storage virtual machine (SVM) name.

-unix-path path specifies the UNIX path. The UNIX path must begin with a slash (/) and must end with a slash (/).

-share-name share name specifies the name of the SMB share to map.

-cifs-path path specifies the CIFS path. The CIFS path must begin with a slash (/) and must end with a slash (/).

-cifs-server server_name specifies the CIFS server name. The CIFS server name can be specified as a DNS name (for example, mynetwork.cifs.server.com), IP address, or NetBIOS name. The NetBIOS name can be determined by using the vserver cifs show command. If this optional parameter is not specified, the default value is the NetBIOS name of the local CIFS server.

-locality {local|free|widelink} specifies whether to create a local link, a free link or a wide symbolic link. A local symbolic link maps to the local SMB share. A free symbolic link can map anywhere on the local SMB server. A wide symbolic link maps to any SMB share on the network. If you do not specify this optional parameter, the default value is local.

-home-directory {true|false} specifies whether the target share is a home directory. Even though this parameter is optional, you must set this parameter to true when the target share is configured as a home directory. The default is false.

Example

The following command creates a symbolic link mapping on the SVM named vs1. It has the UNIX path /src/, the SMB share name "SOURCE", the CIFS path /mycompany/source/, and the CIFS server IP address 123.123.123.123, and it is a widelink.

```
cluster1::> vserver cifs symlink create -vserver vs1 -unix-path /src/
-share-name SOURCE -cifs-path "/mycompany/source/" -cifs-server
123.123.123.123 -locality widelink
```

Related information

Configuring UNIX symbolic link support on SMB shares

Commands for managing symbolic link mappings

There are specific ONTAP commands for managing symbolic link mappings.

If you want to	Use this command
Create a symbolic link mapping	vserver cifs symlink create
Display information about symbolic link mappings	vserver cifs symlink show

If you want to	Use this command
Modify a symbolic link mapping	vserver cifs symlink modify
Delete a symbolic link mapping	vserver cifs symlink delete

See the man page for each command for more information.

Windows backup applications and Unix-style symlinks

When a backup application running on Windows encounters a Unix-style symbolic link (symlink), the link is followed and the data is backed up. Beginning with ONTAP 9.15.1, you have the option of backing up the symlinks instead of the data. This feature is fully supported with ONTAP FlexGroups and FlexVols.

Overview

Before you change how ONTAP handles symlinks during a Windows backup operation, you should be familiar with the benefits, key concepts, and configuration options.

Benefits

When this feature is disabled or unavailable, each symlink is traversed and the data it links to is backed up. Because of this, unnecessary data can sometimes be backed up and in certain situations the application might end up in a loop. Backing up the symlinks instead avoids these issues. And because the symlink files are very small compared to the data in most cases, the backups take less time to complete. The overall performance of the cluster can also improve because of the reduced IO operations.

Windows server environment

This feature is supported for backup applications running on Windows. You should understand the relevant technical aspects of the environment before using it.

Extended attributes

Windows supports extended attributes (EA) which collectively form additional metadata optionally associated with the files. These attributes are used by various applications, such as the Windows Subsystem for Linux as described at File Permissions for WSL. Applications can request extended attributes for each file when reading data from ONTAP.

The symlinks are returned in the extended attributes when the feature is enabled. Therefore a backup application must provide standard EA support which is used to store the metadata. Some Windows utilities support and preserve the extended attributes. However, if the backup software does not support backing up and restoring the extended attributes, it will not preserve the metadata associated with each file and fail to process the symlinks properly.

Windows configuration

Backup applications running on a Microsoft Windows server can be granted a special privilege allowing them to bypass normal file security. This is typically done by adding the applications to the Backup Operators group. The apps can then back up and restore files as needed as well as perform other related system operations. There are subtle changes to the SMB protocol used by the backup applications which can be detected by ONTAP as the data is read and written.

Requirements

The symlink backup feature has several requirements including:

- Your cluster is running ONTAP 9.15.1 or later.
- A Windows backup application that has been granted special backup privileges.
- The backup application must also support extended attributes and request them during the backup operations.
- The ONTAP symlink backup feature is enabled for the applicable data SVM.

Configuration options

In addition to the ONTAP CLI, you can also manage this feature using the REST API. See What's new with the ONTAP REST API and automation for more information. The configuration determining how ONTAP processes the Unix-style symlinks must be performed separately for each SVM.

Enable the symlink backup feature in ONTAP

A configuration option has been introduced to an existing CLI command with ONTAP 9.15.1. You can use this option to enable or disable Unix-style symlink processing.

Before you begin

Review the basic Requirements. In addition:

- Be able to elevate your CLI privilege to the advanced level.
- Determine the data SVM you want to modify. The SVM vs1 is used in the example command.

Steps

1. Set the advanced privilege level.

set privilege advanced

2. Enable symlink file backup.

vserver cifs options modify -vserver vs1 -is-backup-symlink-enabled true

Use BranchCache to cache SMB share content at a branch office

Use BranchCache to cache SMB share content at a branch office overview

BranchCache was developed by Microsoft to enable caching of content on computers local to requesting clients. ONTAP implementation of BranchCache can reduce wide-area network (WAN) utilization and provide improved access response time when users in a branch office access content stored on storage virtual machines (SVMs) using SMB.

If you configure BranchCache, Windows BranchCache clients first retrieve content from the SVM and then cache the content on a computer within the branch office. If another BranchCache-enabled client in the branch office requests the same content, the SVM first authenticates and authorizes the requesting user. The SVM then determines whether the cached content is still up-to-date and, if it is, sends the client metadata about the cached content. The client then uses the metadata to retrieve content directly from the locally based cache.

Related information

Using offline files to allow caching of files for offline use

Requirements and guidelines

BranchCache version support

You should be aware of which BranchCache versions ONTAP supports.

ONTAP supports BranchCache 1 and the enhanced BranchCache 2:

• When you configure BranchCache on the SMB server for the storage virtual machine (SVM), you can enable BranchCache 1, BranchCache 2, or all versions.

By default, all versions are enabled.

• If you enable only BranchCache 2, the remote office Windows client machines must support BranchCache 2.

Only SMB 3.0 or later clients support BranchCache 2.

For more information about BranchCache versions, see the Microsoft TechNet Library.

Related information

Microsoft TechNet Library: technet.microsoft.com/en-us/library/

Network protocol support requirements

You must be aware of the network protocol requirements for implementing ONTAP BranchCache.

You can implement the ONTAP BranchCache feature over IPv4 and IPv6 networks using SMB 2.1 or later.

All CIFS servers and branch office machines participating in the BranchCache implementation must have the SMB 2.1 or later protocol enabled. SMB 2.1 has protocol extensions that allow a client to participate in a BranchCache environment. This is the minimum SMB protocol version that offers BranchCache support. SMB 2.1 supports version BranchCache version 1.

If you want to use BranchCache version 2, SMB 3.0 is the minimum supported version. All CIFS servers and branch office machines participating in a BranchCache 2 implementation must have SMB 3.0 or later enabled.

If you have remote offices where some of the clients support only SMB 2.1 and some of the clients support SMB 3.0, you can implement a BranchCache configuration on the CIFS server that provides caching support over both BranchCache 1 and BranchCache 2.



Even though the Microsoft BranchCache feature supports using both the HTTP/HTTPS and SMB protocols as file access protocols, ONTAP BranchCache only supports the use of SMB.

ONTAP and Windows hosts version requirements

ONTAP and branch office Windows hosts must meet certain version requirements before you can configure BranchCache.

Before configuring BranchCache, you must ensure that the version of ONTAP on the cluster and participating branch office clients support SMB 2.1 or later and support the BranchCache feature. If you configure Hosted Cache mode, you must also ensure that you use a supported host for the cache server.

BranchCache 1 is supported on the following ONTAP versions and Windows hosts:

- · Content server: storage virtual machine (SVM) with ONTAP
- · Cache server: Windows Server 2008 R2 or Windows Server 2012 or later
- Peer or client: Windows 7 Enterprise, Windows 7 Ultimate, Windows 8, Windows Server 2008 R2 or Windows Server 2012 or later

BranchCache 2 is supported on the following ONTAP versions and Windows hosts:

- Content server: SVM with ONTAP
- · Cache server: Windows Server 2012 or later
- Peer or client: Windows 8 or Windows Server 2012 or later

Reasons ONTAP invalidates BranchCache hashes

Understanding the reasons why ONTAP invalidates hashes can be helpful as you plan your BranchCache configuration. It can help you decide which operating mode you should configure and can help you choose on which shares to enable BranchCache.

ONTAP must manage BranchCache hashes to ensure that hashes are valid. If a hash is not valid, ONTAP invalidates the hash and computes a new hash the next time that content is requested, assuming that BranchCache is still enabled.

ONTAP invalidates hashes for the following reasons:

• The server key is modified.

If the server key is modified, ONTAP invalidates all hashes in the hash store.

• A hash is flushed from the cache because the BranchCache hash store maximum size has been reached.

This is a tunable parameter and can be modified to meet your business requirements.

- A file is modified either through SMB or NFS access.
- A file for which there are computed hashes is restored using the snap restore command.
- A volume that contains SMB shares that are BranchCache-enabled is restored using the snap restore command.

Guidelines for choosing the hash store location

When configuring BranchCache, you choose where to store hashes and what size the hash store should be. Understanding the guidelines when choosing the hash store

location and size can help you plan your BranchCache configuration on a CIFS-enabled SVM.

• You should locate the hash store on a volume where atime updates are permitted.

The access time on a hash file is used to keep frequently accessed files in the hash store. If atime updates are disabled, the creation time is used for this purpose. It is preferable to use atime to track frequently used files.

- You cannot store hashes on read-only file systems such as SnapMirror destinations and SnapLock volumes.
- If the maximum size of the hash store is reached, older hashes are flushed to make room for new hashes.

You can increase the maximum size of the hash store to reduce the amount of hashes that are flushed from the cache.

• If the volume on which you store hashes is unavailable or full, or if there is an issue with intra-cluster communication where the BranchCache service cannot retrieve hash information, BranchCache services are not available.

The volume might be unavailable because it is offline or because the storage administrator specified a new location for the hash store.

This does not cause issues with file access. If access to the hash store is impeded, ONTAP returns a Microsoft-defined error to the client, which causes the client to request the file using the normal SMB read request.

Related information

Configure BranchCache on the SMB server

Modify the BranchCache configuration

BranchCache recommendations

Before you configure BranchCache, there are certain recommendations you should keep in mind when deciding on which SMB shares you want to enable BranchCache caching.

You should keep the following recommendations in mind when deciding on which operating mode to use and on which SMB shares to enable BranchCache:

- The benefits of BranchCache are reduced when the data to be remotely cached changes frequently.
- BranchCache services are beneficial for shares containing file content that is reused by multiple remote office clients or by file content that is repeatedly accessed by a single remote user.
- Consider enabling caching for read-only content such as data in Snapshot copies and SnapMirror destinations.

Configure BranchCache

Configure BranchCache overview

You configure BranchCache on your SMB server using ONTAP commands. To implement

BranchCache, you must also configure your clients, and optionally your hosted cache servers at the branch offices where you want to cache content.

If you configure BranchCache to enable caching on a share-by-share basis, you must enable BranchCache on the SMB shares for which you want to provide BranchCache caching services.

Requirements for configuring BranchCache

After meeting some prerequisites, you can set up BranchCache.

The following requirements must be met before configuring BranchCache on the CIFS server for your SVM:

- ONTAP must be installed on all nodes in the cluster.
- CIFS must be licensed and a SMB server must be configured. The SMB license is included with ONTAP One. If you don't have ONTAP One and the license is not installed, contact your sales representative.
- IPv4 or IPv6 network connectivity must be configured.
- For BranchCache 1, SMB 2.1 or later must be enabled.
- For BranchCache 2, SMB 3.0 must be enabled and the remote Windows clients must support BranchCache 2.

Configure BranchCache on the SMB server

You can configure BranchCache to provide BranchCache services on a per-share basis. Alternatively, you can configure BranchCache to automatically enable caching on all SMB shares.

About this task

You can configure BranchCache on SVMs.

- You can create an all-shares BranchCache configuration if want to offer caching services for all content contained within all SMB shares on the CIFS server.
- You can create a per-share BranchCache configuration if you want to offer caching services for content contained within selected SMB shares on the CIFS server.

You must specify the following parameters when configuring BranchCache:

Required parameters	Description
SVM name	BranchCache is configured on a per SVM basis. You must specify on which CIFS-enabled SVM you want to configure the BranchCache service.

Required parameters	Description
Path to hash store	BranchCache hashes are stored in regular files on the SVM volume. You must specify the path to an existing directory where you want ONTAP to store the hash data.The BranchCache hash path must be read- writable. Read-only paths, such as Snapshot directories are not allowed. You can store hash data in a volume that contains other data or you can create a separate volume to store hash data. If the SVM is an SVM disaster recovery source, the hash path cannot be on the root volume. This is because the root volume is not replicated to the disaster recovery destination. The hash path can contain blanks and any valid file name characters.

You can optionally specify the following parameters:

Optional parameters	Description
Supported Versions	ONTAP support BranchCache 1 and 2. You can enable version 1, version 2, or both versions. The default is to enable both versions.
Maximum size of hash store	You can specify the size to use for the hash data store. If the hash data exceeds this value, ONTAP deletes older hashes to make room for newer hashes. The default size for the hash store is 1 GB. BranchCache performs more efficiently if hashes are not discarded in an overly aggressive manner. If you determine that hashes are discarded frequently because the hash store is full, you can increase the hash store size by modifying the BranchCache configuration.
Server key	You can specify a server key that the BranchCache service uses to prevent clients from impersonating the BranchCache server. If you do not specify a server key, one is randomly generated when you create the BranchCache configuration. You can set the server key to a specific value so that if multiple servers are providing BranchCache data for the same files, clients can use hashes from any server using that same server key. If the server key contains any spaces, you must enclose the server key in quotation marks.

Optional parameters	Description
Operating mode	The default is to enable BranchCache on a per-share basis.
	• To create a BranchCache configuration where you enable BranchCache on a per-share basis, you can either not specify this optional parameter or you can specify per-share.
	 To automatically enable BranchCache on all shares, you must set the operating mode to all- shares.

Steps

- 1. Enable SMB 2.1 and 3.0 as needed:
 - a. Set the privilege level to advanced: set -privilege advanced
 - b. Check the configured SVM SMB settings to determine whether all needed versions of SMB are enabled: vserver cifs options show -vserver vserver name
 - c. If necessary, enable SMB 2.1: vserver cifs options modify -vserver vserver_name -smb2-enabled true

The command enables both SMB 2.0 and SMB 2.1.

- d. If necessary, enable SMB 3.0: vserver cifs options modify -vserver vserver_name -smb3-enabled true
- e. Return to the admin privilege level: set -privilege admin
- 2. Configure BranchCache: vserver cifs branchcache create -vserver vserver_name -hash -store-path path [-hash-store-max-size {integer[KB|MB|GB|TB|PB]}] [-versions {v1-enable|v2-enable|enable-all] [-server-key text] -operating-mode {pershare|all-shares}

The specified hash storage path must exist and must reside on a volume managed by the SVM. The path must also be located on a read-writable volume. The command fails if the path is read-only or does not exist.

If you want to use the same server key for additional SVM BranchCache configurations, record the value you enter for the server key. The server key does not appear when you display information about the BranchCache configuration.

3. Verify that the BranchCache configuration is correct: vserver cifs branchcache show -vserver vserver_name

Examples

The following commands verify that both SMB 2.1 and 3.0 are enabled and configure BranchCache to automatically enable caching on all SMB shares on SVM vs1:

cluster1::> set -privilege advanced Warning: These advanced commands are potentially dangerous; use them only when directed to do so by technical support personnel. Do you wish to continue? (y or n): y cluster1::*> vserver cifs options show -vserver vs1 -fields smb2enabled, smb3-enabled vserver smb2-enabled smb3-enabled _____ ____ vs1 true true cluster1::*> set -privilege admin cluster1::> vserver cifs branchcache create -vserver vs1 -hash-store-path /hash data -hash-store-max-size 20GB -versions enable-all -server-key "my server key" -operating-mode all-shares cluster1::> vserver cifs branchcache show -vserver vs1 Vserver: vsl Supported BranchCache Versions: enable all Path to Hash Store: /hash data Maximum Size of the Hash Store: 20GB Encryption Key Used to Secure the Hashes: -CIFS BranchCache Operating Modes: all shares

The following commands verify that both SMB 2.1 and 3.0 are enabled, configure BranchCache to enable caching on a per-share basis on SVM vs1, and verify the BranchCache configuration:

cluster1::> set -privilege advanced Warning: These advanced commands are potentially dangerous; use them only when directed to do so by technical support personnel. Do you wish to continue? (y or n): y cluster1::*> vserver cifs options show -vserver vs1 -fields smb2enabled, smb3-enabled vserver smb2-enabled smb3-enabled _____ ____ vs1 true true cluster1::*> set -privilege admin cluster1::> vserver cifs branchcache create -vserver vs1 -hash-store-path /hash data -hash-store-max-size 20GB -versions enable-all -server-key "my server key" cluster1::> vserver cifs branchcache show -vserver vs1 Vserver: vs1 Supported BranchCache Versions: enable all Path to Hash Store: /hash data Maximum Size of the Hash Store: 20GB Encryption Key Used to Secure the Hashes: -CIFS BranchCache Operating Modes: per share

Related information

Requirements and guidelines: BranchCache version support

Where to find information about configuring BranchCache at the remote office

Create a BranchCache-enabled SMB share

Enable BranchCache on an existing SMB share

Modify the BranchCache configuration

Disable BranchCache on SMB shares overview

Delete the BranchCache configuration on SVMs

Where to find information about configuring BranchCache at the remote office

After configuring BranchCache on the SMB server, you must install and configure BranchCache on client computers and, optionally, on caching servers at your remote office. Microsoft provides instructions for configuring BranchCache at the remote office.

Instructions for configuring branch office clients and, optionally, caching servers to use BranchCache are on

the Microsoft BranchCache web site.

Microsoft BranchCache Docs: What's New

Configure BranchCache-enabled SMB shares

Configure BranchCache-enabled SMB shares overview

After you configure BranchCache on the SMB server and at the branch office, you can enable BranchCache on SMB shares that contain content that you want to allow clients at branch offices to cache.

BranchCache caching can be enabled on all SMB shares on the SMB server or on a share-by-share basis.

• If you enable BranchCache on a share-by-share basis, you can enable BranchCache as you create the share or by modifying existing shares.

If you enable caching on an existing SMB share, ONTAP begins computing hashes and sending metadata to clients requesting content as soon as you enable BranchCache on that share.

• Any clients that have an existing SMB connection to a share do not get BranchCache support if BranchCache is subsequently enabled on that share.

ONTAP advertises BranchCache support for a share at the time the SMB session is set up. Clients that already have established sessions when BranchCache is enabled need to disconnect and reconnect to use cached content for this share.



If BranchCache on a SMB share is subsequently disabled, ONTAP stops sending metadata to the requesting client. A client that needs data retrieves it directly from the content server (SMB server).

Create a BranchCache-enabled SMB share

You can enable BranchCache on an SMB share when you create the share by setting the branchcache share property.

About this task

• If BranchCache is enabled on the SMB share, the share must have the offline files configuration set to manual caching.

This is the default setting when you create a share.

- You can also specify additional optional share parameters when you create the BranchCache-enabled share.
- You can set the branchcache property on a share even if BranchCache is not configured and enabled on the storage virtual machine (SVM).

However, if you want the share to offer cached content, you must configure and enable BranchCache on the SVM.

• Since there are no default share properties applied to the share when you use the -share-properties parameter, you must specify all other share properties that you want applied to the share in addition to the

branchcache share property by using a comma-delimited list.

• For more information, see the man page for the vserver cifs share create command.

Step

- Create a BranchCache-enabled SMB share: vserver cifs share create -vserver vserver_name -share-name share_name -path path -share-properties branchcache[,...]
- 2. Verify that the BranchCache share property is set on the SMB share by using the vserver cifs share show command.

Example

The following command creates a BranchCache-enabled SMB share named "data" with a path of /data on SVM vs1. By default, the offline files setting is set to manual:

```
cluster1::> vserver cifs share create -vserver vs1 -share-name data -path
/data -share-properties branchcache, oplocks, browsable, changenotify
cluster1::> vserver cifs share show -vserver vs1 -share-name data
                      Vserver: vsl
                        Share: data
     CIFS Server NetBIOS Name: VS1
                         Path: /data
             Share Properties: branchcache
                                oplocks
                               browsable
                                changenotify
           Symlink Properties: enable
      File Mode Creation Mask: -
 Directory Mode Creation Mask: -
                Share Comment: -
                    Share ACL: Everyone / Full Control
File Attribute Cache Lifetime: -
                  Volume Name: data
                Offline Files: manual
Vscan File-Operations Profile: standard
```

Related information

Disabling BranchCache on a single SMB share

Enable BranchCache on an existing SMB share

You can enable BranchCache on an existing SMB share by adding the branchcache share property to the existing list of share properties.

About this task

• If BranchCache is enabled on the SMB share, the share must have the offline files configuration set to

manual caching.

If the existing share's offline files setting is not set to manual caching, you must configure it by modifying the share.

• You can set the branchcache property on a share even if BranchCache is not configured and enabled on the storage virtual machine (SVM).

However, if you want the share to offer cached content, you must configure and enable BranchCache on the SVM.

• When you add the branchcache share property to the share, existing share settings and share properties are preserved.

The BranchCache share property is added to the existing list of share properties. For more information about using the vserver cifs share properties add command, see the man pages.

Steps

- 1. If necessary, configure the offline files share setting for manual caching:
 - a. Determine what the offline files share setting is by using the vserver cifs share show command.
 - b. If the offline files share setting is not set to manual, change it to the required value: vserver cifs share modify -vserver vserver_name -share-name share_name -offline-files manual
- 2. Enable BranchCache on an existing SMB share: vserver cifs share properties add -vserver vserver_name -share-name share_name -share-properties branchcache
- 3. Verify that the BranchCache share property is set on the SMB share: vserver cifs share show -vserver vserver_name -share-name share_name

Example

The following command enables BranchCache on an existing SMB share named "data2" with a path of /data2 on SVM vs1:

cluster1::> vserver cifs share show -vserver vs1 -share-name data2 Vserver: vsl Share: data2 CIFS Server NetBIOS Name: VS1 Path: /data2 Share Properties: oplocks browsable changenotify showsnapshot Symlink Properties: -File Mode Creation Mask: -Directory Mode Creation Mask: -Share Comment: -Share ACL: Everyone / Full Control File Attribute Cache Lifetime: 10s Volume Name: -Offline Files: manual Vscan File-Operations Profile: standard cluster1::> vserver cifs share properties add -vserver vs1 -share-name data2 -share-properties branchcache cluster1::> vserver cifs share show -vserver vs1 -share-name data2 Vserver: vsl Share: data2 CIFS Server NetBIOS Name: VS1 Path: /data2 Share Properties: oplocks browsable showsnapshot changenotify branchcache Symlink Properties: -File Mode Creation Mask: -Directory Mode Creation Mask: -Share Comment: -Share ACL: Everyone / Full Control File Attribute Cache Lifetime: 10s Volume Name: -Offline Files: manual Vscan File-Operations Profile: standard

Related information
Disabling BranchCache on a single SMB share

Manage and monitor the BranchCache configuration

Modify BranchCache configurations

You can modify the configuration of the BranchCache service on SVMs, including changing the hash store directory path, the hash store maximum directory size, the operating mode, and which BranchCache versions are supported. You can also increase the size of the volume that contains the hash store.

Steps

1. Perform the appropriate action:

If you want to	Enter the following
Modify the hash store directory size	<pre>vserver cifs branchcache modify -vserver vserver_name -hash-store-max -size {integer[KB MB GB TB PB]}</pre>
Increase the size of the volume that contains the hash store	<pre>volume size -vserver vserver_name -volume volume_name -new-size new_size[k m g t] If the volume containing the hash store fills up, you might be able to increase the size of the volume. You can specify the new volume size as a number followed by a unit designation. Learn more about managing FlexVol volumes</pre>

If you want to	Enter the following
Modify the hash store directory path	vserver cifs branchcache modify -vserver vserver_name -hash-store-path path -flush-hashes {true false} If the SVM is an SVM disaster recovery source, the hash path cannot be on the root volume. This is because the root volume is not replicated to the disaster recovery destination.
	The BranchCache hash path can contain blanks and any valid file name characters.
	If you modify the hash path, -flush-hashes is a required parameter that specifies whether you want ONTAP to flush the hashes from the original hash store location. You can set the following values for the -flush-hashes parameter:
	 If you specify true, ONTAP deletes the hashes in the original location and creates new hashes in the new location as new requests are made by BranchCache-enabled clients.
	 If you specify false, the hashes are not flushed.
	In this case, you can choose to reuse the existing hashes later by changing the hash store path back to the original location.
Change the operating mode	<pre>vserver cifs branchcache modify -vserver vserver_name -operating-mode {per-share all-shares disable}</pre>
	You should be aware of the following when modifying the operating mode:
	 ONTAP advertises BranchCache support for a share when the SMB session is set up.
	 Clients that already have established sessions when BranchCache is enabled need to disconnect and reconnect to use cached content for this share.
Change the BranchCache version support	<pre>vserver cifs branchcache modify -vserver vserver_name -versions {v1- enable v2-enable enable-all}</pre>

2. Verify the configuration changes by using the vserver cifs branchcache show command.

Display information about BranchCache configurations

You can display information about BranchCache configurations on storage virtual machines (SVMs), which can be used when verifying a configuration or when determining current settings before modifying a configuration.

Step

1. Perform one of the following actions:

If you want to display	Enter this command
Summary information about BranchCache configurations on all SVMs	vserver cifs branchcache show
Detailed information about the configuration on a specific SVM	<pre>vserver cifs branchcache show -vserver vserver_name</pre>

Example

The following example displays information about the BranchCache configuration on SVM vs1:

```
cluster1::> vserver cifs branchcache show -vserver vs1
Vserver: vs1
Supported BranchCache Versions: enable_all
Path to Hash Store: /hash_data
Maximum Size of the Hash Store: 20GB
Encryption Key Used to Secure the Hashes: -
CIFS BranchCache Operating Modes: per_share
```

Change the BranchCache server key

You can change the BranchCache server key by modifying the BranchCache configuration on the storage virtual machine (SVM) and specifying a different server key.

About this task

You can set the server key to a specific value so that if multiple servers are providing BranchCache data for the same files, clients can use hashes from any server using that same server key.

When you change the server key, you must also flush the hash cache. After flushing the hashes, ONTAP creates new hashes as new requests are made by BranchCache-enabled clients.

Steps

 Change the server key by using the following command: vserver cifs branchcache modify -vserver vserver_name -server-key text -flush-hashes true

When configuring a new server key, you must also specify -flush-hashes and set the value to true.

2. Verify that the BranchCache configuration is correct by using the vserver cifs branchcache show command.

Example

The following example sets a new server key that contains spaces and flushes the hash cache on SVM vs1:

Related information

Reasons ONTAP invalidates BranchCache hashes

Pre-compute BranchCache hashes on specified paths

You can configure the BranchCache service to pre-compute hashes for a single file, for a directory, or for all files in a directory structure. This can be helpful if you want to compute hashes on data in a BranchCache-enabled share during off, non-peak hours.

About this task

If you want to collect a data sample before you display hash statistics, you must use the statistics start and optional statistics stop commands.

- You must specify the storage virtual machine (SVM) and path on which you want to pre-compute hashes.
- · You must also specify whether you want hashes computed recursively.
- If you want hashes computed recursively, the BranchCache service traverses the entire directory tree under the specified path, and computes hashes for each eligible object.

Steps

1. Pre-compute hashes as desired:

If you want to pre-compute hashes on	Enter the command
A single file or directory	vserver cifs branchcache hash-create -vserver vserver_name -path path -recurse false

If you want to pre-compute hashes on	Enter the command
Recursively on all files in a directory structure	vserver cifs branchcache hash-create -vserver vserver_name -path absolute_path -recurse true

- 2. Verify that hashes are being computed by using the statistics command:
 - a. Display statistics for the hashd object on the desired SVM instance: statistics show -object hashd -instance vserver_name
 - b. Verify that the number of hashes created is increasing by repeating the command.

Examples

The following example creates hashes on the path $/{\tt data}$ and on all contained files and subdirectories on SVM vs1:

cluster1::> vserver cifs branchcache hash-create -vserver vs1 -path /data -recurse true cluster1::> statistics show -object hashd -instance vs1 Object: hashd Instance: vs1 Start-time: 9/6/2012 19:09:54 End-time: 9/6/2012 19:11:15 Cluster: cluster1 Counter Value _____ branchcache hash created 85 branchcache hash files replaced 0 branchcache hash rejected 0 branchcache hash store bytes 0 branchcache hash store size 0 instance name vs1 node name node1 node uuid process name cluster1::> statistics show -object hashd -instance vs1 Object: hashd Instance: vsl Start-time: 9/6/2012 19:09:54 End-time: 9/6/2012 19:11:15 Cluster: cluster1 Counter Value _____ ____ branchcache hash created 92 branchcache hash files replaced 0 branchcache hash rejected 0 branchcache hash store bytes 0 branchcache hash store size 0 instance name vs1 node name node1 node uuid process_name

Related information

Performance monitoring setup

Flush hashes from the SVM BranchCache hash store

You can flush all cached hashes from the BranchCache hash store on the storage virtual machine (SVM). This can be useful if you have changed the branch office BranchCache configuration. For example, if you recently reconfigured the caching mode from distributed caching to hosted caching mode, you would want to flush the hash store.

About this task

After flushing the hashes, ONTAP creates new hashes as new requests are made by BranchCache-enabled clients.

Step

1. Flush the hashes from the BranchCache hash store: vserver cifs branchcache hash-flush -vserver vserver_name

vserver cifs branchcache hash-flush -vserver vs1

Display BranchCache statistics

You can display BranchCache statistics to, among other things, identify how well caching is performing, determine whether your configuration is providing cached content to clients, and determine whether hash files were deleted to make room for more recent hash data.

About this task

The hashd statistic object contains counters that provide statistical information about BranchCache hashes. The cifs statistic object contains counters that provide statistical information about BranchCache-related activity. You can collect and display information about these objects at the advanced-privilege level.

Steps

1. Set the privilege level to advanced: set -privilege advanced

```
cluster1::> set -privilege advanced
Warning: These advanced commands are potentially dangerous; use them
only when directed to do so by support personnel.
Do you want to continue? {y|n}: y
```

2. Display the BranchCache-related counters by using the statistics catalog counter show command.

For more information about statistics counters, see the man page for this command.

```
cluster1::*> statistics catalog counter show -object hashd
Object: hashd
Counter Description
```

_____ Number of times a request to generate branchcache hash created BranchCache hash for a file succeeded. branchcache hash files replaced Number of times a BranchCache hash file was deleted to make room for more recent hash data. This happens if the hash store size is exceeded. branchcache_hash_rejected Number of times a request to generate BranchCache hash data failed. branchcache hash store bytes Total number of bytes used to store hash data. branchcache hash store size Total space used to store BranchCache hash data for the Vserver. instance name Instance Name instance uuid Instance UUID node name System node name node uuid System node id 9 entries were displayed. cluster1::*> statistics catalog counter show -object cifs Object: cifs Description Counter -----Number of active searches over SMB and active_searches SMB2 Authentication refused after too many auth reject too many requests were made in rapid succession avg directory depth Average number of directories crossed by SMB and SMB2 path-based commands avg junction depth Average number of junctions crossed by SMB and SMB2 path-based commands branchcache hash fetch fail Total number of times a request to fetch hash data failed. These are failures when attempting to read existing hash data.

```
Ιt
                                does not include attempts to fetch hash
data
                                that has not yet been generated.
    branchcache hash fetch ok
                                Total number of times a request to fetch
hash
                                data succeeded.
    branchcache hash sent bytes Total number of bytes sent to clients
                                requesting hashes.
    branchcache_missing_hash_bytes
                                Total number of bytes of data that had
to be
                                read by the client because the hash for
that
                                content was not available on the server.
   ....Output truncated....
```

3. Collect BranchCache-related statistics by using the statistics start and statistics stop commands.

```
cluster1::*> statistics start -object cifs -vserver vs1 -sample-id 11
Statistics collection is being started for Sample-id: 11
cluster1::*> statistics stop -sample-id 11
Statistics collection is being stopped for Sample-id: 11
```

4. Display the collected BranchCache statistics by using the statistics show command.

```
cluster1::*> statistics show -object cifs -counter
branchcache hash sent bytes -sample-id 11
Object: cifs
Instance: vs1
Start-time: 12/26/2012 19:50:24
End-time: 12/26/2012 19:51:01
Cluster: cluster1
   Counter
                                                            Value
    _____
   branchcache hash sent bytes
                                                                0
   branchcache hash sent bytes
                                                                0
   branchcache hash sent bytes
                                                                0
   branchcache hash sent bytes
                                                                0
cluster1::*> statistics show -object cifs -counter
branchcache missing hash bytes -sample-id 11
Object: cifs
Instance: vs1
Start-time: 12/26/2012 19:50:24
End-time: 12/26/2012 19:51:01
Cluster: cluster1
   Counter
                                                            Value
                      _____ ___
   branchcache missing hash bytes
                                                                0
   branchcache missing hash bytes
                                                                0
   branchcache missing hash bytes
                                                                0
   branchcache missing hash bytes
                                                                0
```

5. Return to the admin privilege level: set -privilege admin

```
cluster1::*> set -privilege admin
```

Related information

Displaying statistics

Performance monitoring setup

Support for BranchCache Group Policy Objects

ONTAP BranchCache provides support for BranchCache Group Policy Objects (GPOs),

which allow centralized management for certain BranchCache configuration parameters. There are two GPOs used for BranchCache, the Hash Publication for BranchCache GPO and the Hash Version Support for BranchCache GPO.

Hash Publication for BranchCache GPO

The Hash Publication for BranchCache GPO corresponds to the <code>-operating-mode</code> parameter. When GPO updates occur, this value is applied to storage virtual machine (SVM) objects contained within the organizational unit (OU) to which the group policy applies.

Hash Version Support for BranchCache GPO

The Hash Version Support for BranchCache GPO corresponds to the -versions parameter. When GPO updates occur, this value is applied to SVM objects contained within the organizational unit to which the group policy applies.

Related information

Applying Group Policy Objects to CIFS servers

Display information about BranchCache Group Policy Objects

You can display information about the CIFS server's Group Policy Object (GPO) configuration to determine whether BranchCache GPOs are defined for the domain to which the CIFS server belongs and, if so, what the allowed settings are. You can also determine whether BranchCache GPO settings are applied to the CIFS server.

About this task

Even though a GPO setting is defined within the domain to which the CIFS server belongs, it is not necessarily applied to the organizational unit (OU) containing the CIFS-enabled storage virtual machine (SVM). Applied GPO setting are the subset of all defined GPOs that are applied to the CIFS-enabled SVM. BranchCache settings applied through GPOs override settings applied through the CLI.

Steps

1. Display the defined BranchCache GPO setting for the Active Directory domain by using the vserver cifs group-policy show-defined command.



This example does not display all of the available output fields for the command. Output is truncated.

```
cluster1::> vserver cifs group-policy show-defined -vserver vs1
Vserver: vsl
_____
      GPO Name: Default Domain Policy
      Level: Domain
      Status: enabled
  Advanced Audit Settings:
     Object Access:
          Central Access Policy Staging: failure
  Registry Settings:
     Refresh Time Interval: 22
     Refresh Random Offset: 8
     Hash Publication Mode for BranchCache: per-share
     Hash Version Support for BranchCache: version1
  [...]
    GPO Name: Resultant Set of Policy
      Status: enabled
  Advanced Audit Settings:
      Object Access:
          Central Access Policy Staging: failure
  Registry Settings:
     Refresh Time Interval: 22
     Refresh Random Offset: 8
     Hash Publication for Mode BranchCache: per-share
     Hash Version Support for BranchCache: version1
  [...]
```

2. Display the BranchCache GPO setting applied to the CIFS server by using the vserver cifs grouppolicy show-applied command. ``



This example does not display all of the available output fields for the command. Output is truncated.

```
cluster1::> vserver cifs group-policy show-applied -vserver vs1
Vserver: vsl
_____
    GPO Name: Default Domain Policy
      Level: Domain
      Status: enabled
  Advanced Audit Settings:
     Object Access:
          Central Access Policy Staging: failure
  Registry Settings:
      Refresh Time Interval: 22
      Refresh Random Offset: 8
      Hash Publication Mode for BranchCache: per-share
      Hash Version Support for BranchCache: version1
  [...]
    GPO Name: Resultant Set of Policy
      Level: RSOP
  Advanced Audit Settings:
      Object Access:
          Central Access Policy Staging: failure
  Registry Settings:
     Refresh Time Interval: 22
     Refresh Random Offset: 8
     Hash Publication Mode for BranchCache: per-share
     Hash Version Support for BranchCache: version1
 [...]
```

Related information

Enabling or disabling GPO support on a CIFS server

Disable BranchCache on SMB shares

Disable BranchCache on SMB shares overview

If you do not want to provide BranchCache caching services on certain SMB shares but you might want to provide caching services on those shares later, you can disable BranchCache on a share-by-share basis. If you have BranchCache configured to offer caching on all shares but you want to temporarily disable all caching services, you can modify the BranchCache configuration to stop automatic caching on all shares.

If BranchCache on an SMB share is subsequently disabled after first being enabled, ONTAP stops sending metadata to the requesting client. A client that needs data retrieves it directly from the content server (CIFS server on the storage virtual machine (SVM)).

Related information

Configuring BranchCache-enabled SMB shares

Disable BranchCache on a single SMB share

If you do not want to offer caching services on certain shares that previously offered cached content, you can disable BranchCache on an existing SMB share.

Step

1. Enter the following command: vserver cifs share properties remove -vserver vserver_name -share-name share_name -share-properties branchcache

The BranchCache share property is removed. Other applied share properties remain in effect.

Example

The following command disables BranchCache on an existing SMB share named "data2":

cluster1::> vserver cifs share show -vserver vs1 -share-name data2 Vserver: vsl Share: data2 CIFS Server NetBIOS Name: VS1 Path: /data2 Share Properties: oplocks browsable changenotify attributecache branchcache Symlink Properties: -File Mode Creation Mask: -Directory Mode Creation Mask: -Share Comment: -Share ACL: Everyone / Full Control File Attribute Cache Lifetime: 10s Volume Name: -Offline Files: manual Vscan File-Operations Profile: standard cluster1::> vserver cifs share properties remove -vserver vs1 -share-name data2 -share-properties branchcache cluster1::> vserver cifs share show -vserver vs1 -share-name data2 Vserver: vsl Share: data2 CIFS Server NetBIOS Name: VS1 Path: /data2 Share Properties: oplocks browsable changenotify attributecache Symlink Properties: -File Mode Creation Mask: -Directory Mode Creation Mask: -Share Comment: -Share ACL: Everyone / Full Control File Attribute Cache Lifetime: 10s Volume Name: -Offline Files: manual Vscan File-Operations Profile: standard

Stop automatic caching on all SMB shares

If your BranchCache configuration automatically enables caching on all SMB shares on each storage virtual machine (SVM), you can modify the BranchCache configuration to stop automatically caching content for all SMB shares.

About this task

To stop automatic caching on all SMB shares, you change the BranchCache operating mode to per-share caching.

Steps

- 1. Configure BranchCache to stop automatic caching on all SMB shares: vserver cifs branchcache modify -vserver vserver_name -operating-mode per-share
- 2. Verify that the BranchCache configuration is correct: vserver cifs branchcache show -vserver vserver_name

Example

The following command changes the BranchCache configuration on storage virtual machine (SVM, formerly known as Vserver) vs1 to stop automatic caching on all SMB shares:

Disable or enable BranchCache on the SVM

What happens when you disable or reenable BranchCache on the CIFS server

If you previously configured BranchCache but do not want the branch office clients to use cached content, you can disable caching on the CIFS server. You must be aware of what happens when you disable BranchCache.

When you disable BranchCache, ONTAP no longer computes hashes or sends the metadata to the requesting client. However, there is no interruption to file access. Thereafter, when BranchCache-enabled clients request metadata information for content they want to access, ONTAP responds with a Microsoft-defined error, which causes the client to send a second request, requesting the actual content. In response to the request for content, the CIFS server sends the actual content that is stored on the storage virtual machine (SVM).

After BranchCache is disabled on the CIFS server, SMB shares do not advertise BranchCache capabilities. To access data on new SMB connections, clients make normal read SMB requests.

You can reenable BranchCache on the CIFS server at any time.

- Because the hash store is not deleted when you disable BranchCache, ONTAP can use the stored hashes when replying to hash requests after you reenable BranchCache, provided that the requested hash is still valid.
- Any clients that have made SMB connections to BranchCache-enabled shares during the time when BranchCache was disabled do not get BranchCache support if BranchCache is subsequently reenabled.

This is because ONTAP advertises BranchCache support for a share at the time the SMB session is set up. Clients that established sessions to BranchCache-enabled shares while BranchCache was disabled need to disconnect and reconnect to use cached content for this share.



If you do not want to save the hash store after you disable BranchCache on a CIFS server, you can manually delete it. If you reenable BranchCache, you must ensure that the hash store directory exists. After BranchCache is reenabled, BranchCache-enabled shares advertise BranchCache capabilities. ONTAP creates new hashes as new requests are made by BranchCache-enabled clients.

Disable or enable BranchCache

You can disable BranchCache on the storage virtual machine (SVM) by changing the BranchCache operating mode to disabled. You can enable BranchCache at any time by changing the operating mode to either offer BranchCache services per-share or automatically for all shares.

Steps

1. Run the appropriate command:

If you want to	Then enter the following
Disable BranchCache	vserver cifs branchcache modify -vserver vserver_name -operating-mode disable
Enable BranchCache per share	vserver cifs branchcache modify -vserver vserver_name -operating-mode per-share
Enable BranchCache for all shares	<pre>vserver cifs branchcache modify -vserver vserver_name -operating-mode all-shares</pre>

2. Verify that the BranchCache operating mode is configured with the desired setting: vserver cifs branchcache show -vserver vserver_name

Example

The following example disables BranchCache on SVM vs1:

cluster1::> vserver cifs branchcache modify -vserver vsl -operating-mode disable cluster1::> vserver cifs branchcache show -vserver vsl Vserver: vsl Supported BranchCache Versions: enable_all Path to Hash Store: /hash_data Maximum Size of the Hash Store: 20GB Encryption Key Used to Secure the Hashes: -CIFS BranchCache Operating Modes: disable

Delete the BranchCache configuration on SVMs

What happens when you delete the BranchCache configuration

If you previously configured BranchCache but do not want the storage virtual machine (SVM) to continue providing cached content, you can delete the BranchCache configuration on the CIFS server. You must be aware of what happens when you delete the configuration.

When you delete the configuration, ONTAP removes the configuration information for that SVM from the cluster and stops the BranchCache service. You can choose whether ONTAP should delete the hash store on the SVM.

Deleting the BranchCache configuration does not disrupt access by BranchCache-enabled clients. Thereafter, when BranchCache-enabled clients request metadata information on existing SMB connections for content that is already cached, ONTAP responds with a Microsoft defined error, which causes the client to send a second request, requesting the actual content. In response to the request for content, the CIFS server sends the actual content that is stored on the SVM

After the BranchCache configuration is deleted, SMB shares do not advertise BranchCache capabilities. To access content that has not previously been cached using new SMB connections, clients make normal read SMB requests.

Delete the BranchCache configuration

The command you use for deleting the BranchCache service on your storage virtual machine (SVM) differs depending on whether you want to delete or keep existing hashes.

Step

1. Run the appropriate command:

If you want to	Then enter the following
Delete the BranchCache configuration and delete existing hashes	vserver cifs branchcache delete -vserver <i>vserver_name</i> -flush-hashes true

If you want to	Then enter the following
Delete the BranchCache configuration but keep existing hashes	vserver cifs branchcache delete -vserver <i>vserver_name</i> -flush-hashes false

Example

The following example deletes the BranchCache configuration on SVM vs1 and deletes all existing hashes:

```
cluster1::> vserver cifs branchcache delete -vserver vs1 -flush-hashes
true
```

What happens to BranchCache when reverting

It is important to understand what happens when you revert ONTAP to a release that does not support BranchCache.

• When you revert to a version of ONTAP that does not support BranchCache, the SMB shares do not advertise BranchCache capabilities to BranchCache-enabled clients; therefore, the clients do not request hash information.

Instead, they request the actual content using normal SMB read requests. In response to the request for content, the SMB server sends the actual content that is stored on the storage virtual machine (SVM).

• When a node hosting a hash store is reverted to a release that does not support BranchCache, the storage administrator needs to manually revert the BranchCache configuration using a command that is printed out during the revert.

This command deletes the BranchCache configuration and hashes.

After the revert completes, the storage administrator can manually delete the directory that contained the hash store if desired.

Related information

Deleting the BranchCache configuration on SVMs

Improve Microsoft remote copy performance

Improve Microsoft remote copy performance overview

Microsoft Offloaded Data Transfer (ODX), also known as *copy offload*, enables direct data transfers within or between compatible storage devices without transferring the data through the host computer.

ONTAP supports ODX for both the SMB and SAN protocols. The source can be either a CIFS server or LUN, and the destination can be either a CIFS server or LUN.

In non-ODX file transfers, the data is read from the source and is transferred across the network to the client

computer. The client computer transfers the data back over the network to the destination. In summary, the client computer reads the data from the source and writes it to the destination. With ODX file transfers, data is copied directly from the source to the destination.

Because ODX offloaded copies are performed directly between the source and destination storage, there are significant performance benefits. The performance benefits realized include faster copy time between source and destination, reduced resource utilization (CPU, memory) on the client, and reduced network I/O bandwidth utilization.

For SMB environments, this functionality is only available when both the client and the storage server support SMB 3.0 and the ODX feature. For SAN environments, this functionality is only available when both the client and the storage server support the ODX feature. Client computers that support ODX and have ODX enabled automatically and transparently use offloaded file transfer when moving or copying files. ODX is used irrespective of whether you drag-and-drop files through Windows Explorer or use command-line file copy commands, or whether a client application initiates file copy requests.

Related information

Improving client response time by providing SMB automatic node referrals with Auto Location

SMB configuration for Microsoft Hyper-V and SQL Server

How ODX works

ODX copy offload uses a token-based mechanism for reading and writing data within or between ODX-enabled CIFS servers. Instead of routing the data through the host, the CIFS server sends a small token, which represents the data, to the client. The ODX client presents that token to the destination server, which then can transfer the data represented by that token from the source to the destination.

When an ODX client learns that the CIFS server is ODX-capable, it opens the source file and requests a token from the CIFS server. After opening the destination file, the client uses the token to instruct the server to copy the data directly from the source to the destination.



The source and destination can be on the same storage virtual machine (SVM) or on different SVMs, depending on the scope of the copy operation.

The token serves as a point-in-time representation of the data. As an example, when you copy data between storage locations, a token representing a data segment is returned to the requesting client, which the client copies to the destination, thereby removing the need to copy the underlying data through the client.

ONTAP supports tokens that represent 8 MB of data. ODX copies of greater than 8 MB are performed by using multiple tokens, with each token representing 8 MB of data.

The following figure explains the steps that are involved with an ODX copy operation:



- 1. A user copies or moves a file by using Windows Explorer, a command-line interface, or as part of a virtual machine migration, or an application initiates file copies or moves.
- 2. The ODX-capable client automatically translates this transfer request into an ODX request.

The ODX request that is sent to the CIFS server contains a request for a token.

- 3. If ODX is enabled on the CIFS server and the connection is over SMB 3.0, the CIFS server generates a token, which is a logical representation of the data on the source.
- 4. The client receives a token that represents the data and sends it with the write request to the destination CIFS server.

This is the only data that is copied over the network from the source to the client and then from the client to the destination.

- 5. The token is delivered to the storage subsystem.
- 6. The SVM internally performs the copy or move.

If the file that is copied or moved is larger than 8 MB, multiple tokens are needed to perform the copy. Steps 2 through 6 as performed as needed to complete the copy.



If there is a failure with the ODX offloaded copy, the copy or move operation falls back to traditional reads and writes for the copy or move operation. Similarly, if the destination CIFS server does not support ODX or ODX is disabled, the copy or move operation falls back to traditional reads and writes for the copy or move operation.

Requirements for using ODX

Before you can use ODX for copy offloads with your storage virtual machine (SVM), you need to be aware of certain requirements.

ONTAP version requirements

ONTAP releases support ODX for copy offloads.

SMB version requirements

- ONTAP supports ODX with SMB 3.0 and later.
- SMB 3.0 must be enabled on the CIFS server before ODX can be enabled:
 - Enabling ODX also enables SMB 3.0, if it is not already enabled.
 - Disabling SMB 3.0 also disables ODX.

Windows server and client requirements

Before you can use ODX for copy offloads, the Windows client must support the feature.

The NetApp Interoperability Matrixcontains the latest information about supported Windows clients.

Volume requirements

- Source volumes must be a minimum of 1.25 GB.
- If you use compressed volumes, the compression type must be adaptive and only compression group size 8K is supported.

Secondary compression type is not supported.

Guidelines for using ODX

Before you can use ODX for copy offload, you need to be aware of the guidelines. For example, you need to know on which types of volumes you can use ODX and you need to understand the intra-cluster and inter-cluster ODX considerations.

Volume guidelines

- You cannot use ODX for copy offload with the following volume configurations:
 - $\,\circ\,$ Source volume size is less than 1.25 GB

The volume size must be 1.25 GB or larger to use ODX.

Read-only volumes

ODX is not used for files and folders residing in load-sharing mirrors or in SnapMirror or SnapVault destination volumes.

- · If the source volume is not deduplicated
- ODX copies are supported only for intra-cluster copies.

You cannot use ODX to copy files or folders to a volume in another cluster.

Other guidelines

• In SMB environments, to use ODX for copy offload, the files must be 256 kb or larger.

Smaller files are transferred using a traditional copy operation.

• ODX copy offload uses deduplication as part of the copy process.

If you do not want deduplication to occur on SVM volumes when copying or moving data, you should disable ODX copy offload on that SVM.

• The application that performs the data transfer must be written to support ODX.

Application operations that support ODX include the following:

- Hyper-V management operations, such as creating and converting virtual hard disks (VHDs), managing Snapshot copies, and copying files between virtual machines
- Windows Explorer operations
- · Windows PowerShell copy commands
- · Windows command prompt copy commands

Robocopy at the Windows command prompt supports ODX.



The applications must be running on Windows servers or clients that support ODX.

For more information about supported ODX applications on Windows servers and clients, consult the Microsoft TechNet Library.

Related information

Microsoft TechNet Library: technet.microsoft.com/en-us/library/

Use cases for ODX

You should be aware of the use cases for using ODX on SVMs so that you can determine under what circumstances ODX provides you with performance benefits.

Windows servers and clients that support ODX use copy offload as the default way of copying data across remote servers. If the Windows server or client does not support ODX or the ODX copy offload fails at any point, the copy or move operation falls back to traditional reads and writes for the copy or move operation.

The following use cases support using ODX copies and moves:

Intra-volume

The source and destination files or LUNs are within the same volume.

• Inter-volume, same node, same SVM

The source and destination files or LUNs are on different volumes that are located on the same node. The data is owned by the same SVM.

• Inter-volume, different nodes, same SVM

The source and destination files or LUNs are on different volumes that are located on different nodes. The data is owned by the same SVM.

• Inter-SVM, same node

The source and destination file or LUNs are on different volumes that are located on the same node. The data is owned by different SVMs.

• Inter-SVM, different nodes

The source and destination file or LUNs are on different volumes that are located on different nodes. The data is owned by different SVMs.

Inter-cluster

The source and destination LUNs are on different volumes that are located on different nodes across clusters. This is only supported for SAN and does not work for CIFS.

There are some additional special use cases:

• With the ONTAP ODX implementation, you can use ODX to copy files between SMB shares and FC or iSCSI attached virtual drives.

You can use Windows Explorer, the Windows CLI or PowerShell, Hyper-V, or other applications that support ODX to copy or move files seamlessly using ODX copy offload between SMB shares and connected LUNs, provided that the SMB shares and LUNs are on the same cluster.

- Hyper-V provides some additional use cases for ODX copy offload:
 - You can use ODX copy offload pass-through with Hyper-V to copy data within or across virtual hard disk (VHD) files or to copy data between mapped SMB shares and connected iSCSI LUNs within the same cluster.

This allows copies from guest operating systems to pass through to the underlying storage.

- When creating fixed-sized VHDs, ODX is used for initializing the disk with zeros, using a well-known zeroed token.
- ODX copy offload is used for virtual machine storage migration if the source and destination storage is on the same cluster.



To take advantage of the use cases for ODX copy offload pass-through with Hyper-V, the guest operating system must support ODX and the guest operating system's disks must be SCSI disks backed by storage (either SMB or SAN) that supports ODX. IDE disks on the guest operating system do not support ODX pass-through.

Enable or disable ODX

You can enable or disable ODX on storage virtual machines (SVMs). The default is to enable support for ODX copy offload if SMB 3.0 is also enabled.

Before you begin

SMB 3.0 must be enabled.

About this task

If you disable SMB 3.0, ONTAP also disables SMB ODX. If you reenable SMB 3.0, you must manually reenable SMB ODX.

Steps

- 1. Set the privilege level to advanced: set -privilege advanced
- 2. Perform one of the following actions:

If you want ODX copy offload to be	Enter the command
Enabled	vserver cifs options modify -vserver vserver_name -copy-offload-enabled true
Disabled	vserver cifs options modify -vserver vserver_name -copy-offload-enabled false

3. Return to the admin privilege level: set -privilege admin

Example

The following example enables ODX copy offload on SVM vs1:

```
cluster1::> set -privilege advanced
Warning: These advanced commands are potentially dangerous; use them
only when directed to do so by technical support personnel.
Do you wish to continue? (y or n): y
cluster1::*> vserver cifs options modify -vserver vs1 -copy-offload
-enabled true
cluster1::*> set -privilege admin
```

Related information

Available SMB server options

Improve client response time by providing SMB automatic node referrals with Auto Location

Improve client response time by providing SMB automatic node referrals with Auto Location overview

Auto Location uses SMB automatic node referrals to increase SMB client performance on storage virtual machines (SVMs). Automatic node referrals automatically redirect the requesting client to a LIF on the node SVM that is hosting the volume in which the data resides, which can lead to improved client response times.

When an SMB client connects to an SMB share hosted on the SVM, it might connect using a LIF that is on a node that does not own the requested data. The node to which the client is connected accesses data owned by another node by using the cluster network. The client can experience faster response times if the SMB connection uses a LIF located on the node containing the requested data:

• ONTAP provides this functionality by using Microsoft DFS referrals to inform SMB clients that a requested file or folder in the namespace is hosted somewhere else.

A node makes a referral when it determines that there is anSVM LIF on the node containing the data.

- Automatic node referrals are supported for IPv4 and IPv6 LIF IP addresses.
- Referrals are made based on the location of the root of the share through which the client is connected.
- The referral occurs during SMB negotiation.

The referral is made before the connection is established. After ONTAP refers the SMB client to the target node, the connection is made, and the client accesses data through the referred LIF path from that point on. This allows the clients faster access to the data and avoids extra cluster communication.

(i)

If a share spans multiple junction points and some of the junctions are to volumes contained on other nodes, data within the share is spread across multiple nodes. Because ONTAP provides referrals that are local to the root of the share, ONTAP must use the cluster network to retrieve the data contained within these non-local volumes. With this type of namespace architecture, automatic node referrals might not provide significant performance benefits.

If the node hosting the data does not have an available LIF, ONTAP establishes the connection using the LIF chosen by the client. After a file is opened by an SMB client, it continues to access the file through the same referred connection.

If, for any reason, the CIFS server cannot make a referral, there is no disruption to SMB service. The SMB connection is established as if automatic node referrals were not enabled.

Related information

Improving Microsoft remote copy performance

Requirements and guidelines for using automatic node referrals

Before you can use SMB automatic node referrals, also known as *autolocation*, you need to be aware of certain requirements, including which versions of ONTAP support the

feature. You also need to know about supported SMB protocol versions and certain other special guidelines.

ONTAP version and license requirements

- All nodes in the cluster must be running a version of ONTAP that supports automatic node referrals.
- Widelinks must be enabled on a SMB share to use autolocation.
- CIFS must be licensed, and an SMB server must exist on the SVMs. The SMB license is included with ONTAP One. If you don't have ONTAP One and the license is not installed, contact your sales representative.

SMB protocol version requirements

• For SVMs, ONTAP supports automatic node referrals on all versions of SMB.

SMB client requirements

All Microsoft clients supported by ONTAP support SMB automatic node referrals.

The Interoperability Matrix contains the latest information about which Windows clients ONTAP supports.

NetApp Interoperability Matrix Tool

Data LIF requirements

If you want to use a data LIF as a potential referral for SMB clients, you must create data LIFs with both NFS and CIFS enabled.

Automatic node referrals can fail to work if the target node contains data LIFs that are enabled only for the NFS protocol, or enabled only for the SMB protocol.

If this requirement is not met, data access is not affected. The SMB client maps the share using the original LIF that the client used to connect to the SVM.

NTLM authentication requirements when making a referred SMB connection

NTLM authentication must be allowed on the domain containing the CIFS server and on the domains containing clients that want to use automatic node referrals.

When making a referral, the SMB server refers an IP address to the Windows client. Because NTLM authentication is used when making a connection using an IP address, Kerberos authentication is not performed for referred connections.

This happens because the Windows client cannot craft the service principal name used by Kerberos (which is of the form service/NetBIOS name and service/FQDN), which means that the client cannot request a Kerberos ticket to the service.

Guidelines for using automatic node referrals with the home directory feature

When shares are configured with the home directory share property enabled, there can be one or more home directory search paths configured for a home directory configuration. The search paths can point to volumes contained on each node containing SVM volumes. Clients receive a referral and, if an active, local data LIF is available, connect through a referred LIF that is local to the home user's home directory.

There are guidelines when SMB 1.0 clients access dynamic home directories with automatic node referrals enabled. This is because SMB 1.0 clients require the automatic node referral before they have authenticated, which is before the SMB server has the user's name. However, SMB home directory access works correctly for SMB 1.0 clients if the following statements are true:

- SMB home directories are configured to use simple names, such as "%w" (Windows user name) or "%u" (mapped UNIX user name), and not domain-name style names, such as "`%d\%w `" (domain-name\user-name).
- When creating home directory shares, the CIFS home directory shares names are configured with variables ("%w" or "%u"), and not with static names, such as "HOME".

For SMB 2.x and SMB 3.0 clients, there are no special guidelines when accessing home directories using automatic node referrals.

Guidelines for disabling automatic node referrals on CIFS servers with existing referred connections

If you disable automatic node referrals after the option has been enabled, clients currently connected to a referred LIF keep the referred connection. Because ONTAP uses DFS referrals as the mechanism for SMB automatic node referrals, clients can even reconnect to the referred LIF after you disable the option until the client's cached DFS referral for the referred connection times out. This is true even in the case of a revert to a version of ONTAP that does not support automatic node referrals. Clients continue to use referrals until the DFS referral times out from the client's cache.

Autolocation uses SMB automatic node referrals to increase SMB client performance by referring clients to the LIF on the node that owns the data volume of an SVM. When an SMB client connects to an SMB share hosted on an SVM, it might connect using a LIF on a node that does not own the requested data and uses cluster interconnect network to retrieve data. The client can experience faster response times if the SMB connection uses a LIF located on the node containing the requested data.

ONTAP provides this functionality by using Microsoft Distributed File System (DFS) referrals to inform SMB clients that a requested file or folder in the namespace is hosted somewhere else. A node makes a referral when it determines that there is an SVM LIF on the node containing the data. Referrals are made based on the location of the root of the share through which the client is connected.

The referral occurs during SMB negotiation. The referral is made before the connection is established. After ONTAP refers the SMB client to the target node, the connection is made, and the client accesses data through the referred LIF path from that point on. This allows the clients faster access to the data and avoids extra cluster communication.

Guidelines for using automatic node referrals with Mac OS clients

Mac OS X clients do not support SMB automatic node referrals, even though the Mac OS supports Microsoft's Distributed File System (DFS). Windows clients make a DFS referral request before connecting to an SMB share. ONTAP provides a referral to a data LIF found on the same node that hosts the requested data, which leads to improved client response times. Although the Mac OS supports DFS, Mac OS clients do not behave exactly like Windows clients in this area.

Related information

How ONTAP enables dynamic home directories

Network management

NetApp Interoperability Matrix Tool

Support for SMB automatic node referrals

Before you enable SMB automatic node referrals, you should be aware that certain ONTAP functionality does not support referrals.

- The following types of volumes do not support SMB automatic node referrals:
 - · Read-only members of a load-sharing mirror
 - · Destination volume of a data-protection mirror
- Node referrals do not move alongside a LIF move.

If a client is using a referred connection over an SMB 2.x or SMB 3.0 connection and a data LIF moves nondisruptively, the client continues to use the same referred connection, even if the LIF is no longer local to the data.

• Node referrals do not move alongside a volume move.

If a client is using a referred connection over any SMB connection and a volume move occurs, the client continues to use the same referred connection, even if the volume is no longer located on the same node as the data LIF.

Enable or disable SMB automatic node referrals

You can enable SMB automatic node referrals to increase SMB client access performance. You can disable automatic node referrals if you do not want ONTAP to make referrals to SMB clients.

Before you begin

A CIFS server must be configured and running on the storage virtual machine (SVM).

About this task

The SMB automatic node referrals functionality is disabled by default. You can enable or disable this functionality on each SVM as required.

This option is available at the advanced privilege level.

Steps

- 1. Set the privilege level to advanced: set -privilege advanced
- 2. Enable or disable SMB automatic node referrals as required:

If you want SMB automatic node referrals to be	Enter the following command
Enabled	<pre>vserver cifs options modify -vserver vserver_name -is-referral-enabled true</pre>
Disabled	vserver cifs options modify -vserver vserver_name -is-referral-enabled false

The option setting takes effect for new SMB sessions. Clients with existing connection can utilize node referral only when their existing cache timeout expires.

3. Switch to the admin privilege level: set -privilege admin

Related information

Available SMB server options

Use statistics to monitor automatic node referral activity

To determine how many SMB connections are referred, you can monitor automatic node referral activity by using the statistics command. By monitoring referrals you can determine the extent to which automatic referrals are locating connections on nodes that host the shares and whether you should redistribute your data LIFs to provide better local access to shares on the CIFS server.

About this task

The cifs object provides several counters at the advanced privilege level that are helpful when monitoring SMB automatic node referrals:

• node_referral_issued

Number of clients that have been issued a referral to the share root's node after the client connected using a LIF hosted by a node different from the share root's node.

• node referral local

Number of clients that connected using a LIF hosted by the same node that hosts the share root. Local access generally provides optimal performance.

• node_referral_not_possible

Number of clients that have not been issued a referral to the node hosting the share root after connecting using a LIF hosted by a node different from the share root's node. This is because an active data LIF for the share root's node was not found.

• node_referral_remote

Number of clients that connected using a LIF hosted by a node different from the node that hosts the share root. Remote access might result in degraded performance.

You can monitor automatic node referral statistics on your storage virtual machine (SVM) by collecting and viewing data for a specific time period (a sample). You can view data from the sample if you do not stop data collection. Stopping data collection gives you a fixed sample. Not stopping data collection gives you the ability to get updated data that you can use to compare against previous queries. The comparison can help you identify performance trends.



To evaluate and use the information you gather from the statistics command, you should understand the distribution of clients in your environments.

Steps

- 1. Set the privilege level to advanced: set -privilege advanced
- 2. View automatic node referral statistics by using the statistics command.

This example views automatic node referral statistics by collecting and viewing data for a sampled time period:

a. Start the collection: statistics start -object cifs -instance vs1 -sample-id sample1

Statistics collection is being started for Sample-id: sample1

- b. Wait for the desired collection time to elapse.
- c. Stop the collection: statistics stop -sample-id sample1

```
Statistics collection is being stopped for Sample-id: sample1
```

d. View the automatic node referral statistics: statistics show -sample-id sample1 -counter node

```
Object: cifs
Instance: vs1
Start-time: 2/4/2013 19:27:02
End-time: 2/4/2013 19:30:11
Cluster: cluster1
   Counter
                                                       Value
    _____
   node name
                                                       node1
   node referral issued
                                                            0
   node referral local
                                                            1
   node referral not possible
                                                            2
   node referral remote
                                                            2
    . . .
   node name
                                                       node2
   node referral issued
                                                            2
   node referral local
                                                            1
   node referral not possible
                                                            0
   node referral remote
                                                            2
    . . .
```

Output displays counters for all nodes participating in SVM vs1. For clarity, only output fields related to automatic node referral statistics are provided in the example.

3. Return to the admin privilege level: set -privilege admin

Related information

Displaying statistics

Performance monitoring setup

Monitor client-side SMB automatic node referral information using a Windows client

To determine what referrals are made from the client's perspective, you can use the Windows dfsutil.exe utility.

The Remote Server Administration Tools (RSAT) kit available with Windows 7 and later clients contains the dfsutil.exe utility. Using this utility, you can display information about the contents of the referral cache as well as view information about each referral that the client is currently using. You can also use the utility to clear the client's referral cache. For more information, consult the Microsoft TechNet Library.

Related information

Microsoft TechNet Library: technet.microsoft.com/en-us/library/

Provide folder security on shares with access-based enumeration

Provide folder security on shares with access-based enumeration overview

When access-based enumeration (ABE) is enabled on an SMB share, users who do not have permission to access a folder or file contained within the share (whether through individual or group permission restrictions) do not see that shared resource displayed in their environment, although the share itself remains visible.

Conventional share properties allow you to specify which users (individually or in groups) have permission to view or modify files or folders contained within the share. However, they do not allow you to control whether folders or files within the share are visible to users who do not have permission to access them. This could pose problems if the names of these folders or files within the share describe sensitive information, such as the names of customers or products under development.

Access-based enumeration (ABE) extends share properties to include the enumeration of files and folders within the share. ABE therefore enables you to filter the display of files and folders within the share based on user access rights. That is, the share itself would be visible to all users, but files and folders within the share could be displayed to or hidden from designated users. In addition to protecting sensitive information in your workplace, ABE enables you to simplify the display of large directory structures for the benefit of users who do not need access to your full range of content. For example, the share itself would be visible to all users, but files and folders within the share could be displayed or hidden.

Learn about Performance impact when using SMB/CIFS Access Based Enumeration.

Enable or disable access-based enumeration on SMB shares

You can enable or disable access-based enumeration (ABE) on SMB shares to allow or

prevent users from seeing shared resources that they do not have permission to access.

About this task

By default, ABE is disabled.

Steps

1. Perform one of the following actions:

If you want to	Enter the command
Enable ABE on a new share	vserver cifs share create -vserver vserver_name -share-name share_name -path path -share-properties access- based-enumeration You can specify additional optional share settings and additional share properties when you create an SMB share. For more information, see the man page for the vserver cifs share create command.
Enable ABE on an existing share	<pre>vserver cifs share properties add -vserver vserver_name -share-name share_name -share-properties access- based-enumeration Existing share properties are preserved. The ABE share property is added to the existing list of share properties.</pre>
Disable ABE on an existing share	vserver cifs share properties remove -vserver vserver_name -share-name share_name -share-properties access- based-enumeration Other share properties are preserved. Only the ABE share property is removed from the list of share properties.

2. Verify that the share configuration is correct by using the vserver cifs share show command.

Examples

The following example creates an ABE SMB share named "sales" with a path of /sales on SVM vs1. The share is created with access-based-enumeration as a share property:

cluster1::> vserver cifs share create -vserver vs1 -share-name sales -path /sales -share-properties access-basedenumeration, oplocks, browsable, changenotify cluster1::> vserver cifs share show -vserver vs1 -share-name sales Vserver: vsl Share: sales CIFS Server NetBIOS Name: VS1 Path: /sales Share Properties: access-based-enumeration oplocks browsable changenotify Symlink Properties: enable File Mode Creation Mask: -Directory Mode Creation Mask: -Share Comment: -Share ACL: Everyone / Full Control File Attribute Cache Lifetime: -Volume Name: -Offline Files: manual Vscan File-Operations Profile: standard

The following example adds the access-based-enumeration share property to an SMB share named "data2":

```
cluster1::> vserver cifs share properties add -vserver vs1 -share-name
data2 -share-properties access-based-enumeration
cluster1::> vserver cifs share show -vserver vs1 -share-name data2 -fields
share-name, share-properties
server share-name share-properties
------ vs1 data2 oplocks, browsable, changenotify, access-based-enumeration
```

Related information

Adding or removing share properties on an existing SMB share

Enable or disable access-based enumeration from a Windows client

You can enable or disable access-based enumeration (ABE) on SMB shares from a Windows client, which allows you to configure this share setting without needing to connect to the CIFS server.



The abecmd utility is not available in new versions of Windows Server and Windows clients. It was released as part of Windows Server 2008. Support ended for Windows Server 2008 on January 14, 2020.

Steps

1. From a Windows client that supports ABE, enter the following command: abecmd [/enable | /disable] [/server CIFS server name] {/all | share name}

For more information about the abecmd command, see your Windows client documentation.

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at http://www.netapp.com/TM are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.