



Secure NFS access using export policies

ONTAP 9

NetApp
September 19, 2024

Table of Contents

- Secure NFS access using export policies 1
 - How export policies control client access to volumes or qtrees 1
 - Default export policy for SVMs 1
 - How export rules work 1
 - Manage clients with an unlisted security type 3
 - How security types determine client access levels 5
 - Manage superuser access requests 7
 - How ONTAP uses export policy caches 9
 - How the access cache works 10
 - How access cache parameters work 10
 - Remove an export policy from a qtree 11
 - Validate qtree IDs for qtree file operations 12
 - Export policy restrictions and nested junctions for FlexVol volumes 12

Secure NFS access using export policies

How export policies control client access to volumes or qtrees

Export policies contain one or more *export rules* that process each client access request. The result of the process determines whether the client is denied or granted access and what level of access. An export policy with export rules must exist on the storage virtual machine (SVM) for clients to access data.

You associate exactly one export policy with each volume or qtree to configure client access to the volume or qtree. The SVM can contain multiple export policies. This enables you to do the following for SVMs with multiple volumes or qtrees:

- Assign different export policies to each volume or qtree of the SVM for individual client access control to each volume or qtree in the SVM.
- Assign the same export policy to multiple volumes or qtrees of the SVM for identical client access control without having to create a new export policy for each volume or qtree.

If a client makes an access request that is not permitted by the applicable export policy, the request fails with a permission-denied message. If a client does not match any rule in the export policy, then access is denied. If an export policy is empty, then all accesses are implicitly denied.

You can modify an export policy dynamically on a system running ONTAP.

Default export policy for SVMs

Each SVM has a default export policy that contains no rules. An export policy with rules must exist before clients can access data on the SVM. Each FlexVol volume contained in the SVM must be associated with an export policy.

When you create an SVM, the storage system automatically creates a default export policy called `default` for the root volume of the SVM. You must create one or more rules for the default export policy before clients can access data on the SVM. Alternatively, you can create a custom export policy with rules. You can modify and rename the default export policy, but you cannot delete the default export policy.

When you create a FlexVol volume in its containing SVM, the storage system creates the volume and associates the volume with the default export policy for the root volume of the SVM. By default, each volume created in the SVM is associated with the default export policy for the root volume. You can use the default export policy for all volumes contained in the SVM, or you can create a unique export policy for each volume. You can associate multiple volumes with the same export policy.

How export rules work

Export rules are the functional elements of an export policy. Export rules match client access requests to a volume against specific parameters you configure to determine how to handle the client access requests.

An export policy must contain at least one export rule to allow access to clients. If an export policy contains

more than one rule, the rules are processed in the order in which they appear in the export policy. The rule order is dictated by the rule index number. If a rule matches a client, the permissions of that rule are used and no further rules are processed. If no rules match, the client is denied access.

You can configure export rules to determine client access permissions using the following criteria:

- The file access protocol used by the client sending the request, for example, NFSv4 or SMB.
- A client identifier, for example, host name or IP address.

The maximum size for the `-clientmatch` field is 4096 characters.

- The security type used by the client to authenticate, for example, Kerberos v5, NTLM, or AUTH_SYS.

If a rule specifies multiple criteria, the client must match all of them for the rule to apply.



Beginning with ONTAP 9.3, you can enable export policy configuration checking as a background job that records any rules violations in an error rule list. The `vserver export-policy config-checker` commands invoke the checker and display results, which you can use to verify your configuration and delete erroneous rules from the policy.

The commands only validate export configuration for host names, netgroups, and anonymous users.

Example

The export policy contains an export rule with the following parameters:

- `-protocol nfs3`
- `-clientmatch 10.1.16.0/255.255.255.0`
- `-rorule any`
- `-rwrule any`

The client access request is sent using the NFSv3 protocol and the client has the IP address 10.1.17.37.

Even though the client access protocol matches, the IP address of the client is in a different subnet from the one specified in the export rule. Therefore, client matching fails and this rule does not apply to this client.

Example

The export policy contains an export rule with the following parameters:

- `-protocol nfs`
- `-clientmatch 10.1.16.0/255.255.255.0`
- `-rorule any`
- `-rwrule any`

The client access request is sent using the NFSv4 protocol and the client has the IP address 10.1.16.54.

The client access protocol matches and the IP address of the client is in the specified subnet. Therefore, client matching is successful and this rule applies to this client. The client gets read-write access regardless of its security type.

Example

The export policy contains an export rule with the following parameters:

- `-protocol nfs3`
- `-clientmatch 10.1.16.0/255.255.255.0`
- `-rorule any`
- `-rwrule krb5,ntlm`

Client #1 has the IP address 10.1.16.207, sends an access request using the NFSv3 protocol, and authenticated with Kerberos v5.

Client #2 has the IP address 10.1.16.211, sends an access request using the NFSv3 protocol, and authenticated with AUTH_SYS.

The client access protocol and IP address matches for both clients. The read-only parameter allows read-only access to all clients regardless of the security type they authenticated with. Therefore both clients get read-only access. However, only client #1 gets read-write access because it used the approved security type Kerberos v5 to authenticate. Client #2 does not get read-write access.

Manage clients with an unlisted security type

When a client presents itself with a security type that is not listed in an access parameter of an export rule, you have the choice of either denying access to the client or mapping it to the anonymous user ID instead by using the option `none` in the access parameter.

A client might present itself with a security type that is not listed in an access parameter because it was authenticated with a different security type or was not authenticated at all (security type AUTH_NONE). By default, the client is automatically denied access to that level. However, you can add the option `none` to the access parameter. As a result, clients with an unlisted security style are mapped to the anonymous user ID instead. The `-anon` parameter determines what user ID is assigned to those clients. The user ID specified for the `-anon` parameter must be a valid user that is configured with permissions you deem appropriate for the anonymous user.

Valid values for the `-anon` parameter range from 0 to 65535.

User ID assigned to <code>-anon</code>	Resulting handling of client access requests
0 - 65533	The client access request is mapped to the anonymous user ID and gets access depending on the permissions configured for this user.
65534	The client access request is mapped to the user <code>nobody</code> and gets access depending on the permissions configured for this user. This is the default.

User ID assigned to <code>-anon</code>	Resulting handling of client access requests
65535	The access request from any client is denied when mapped to this ID and the client presents itself with security type <code>AUTH_NONE</code> . The access request from clients with user ID 0 is denied when mapped to this ID and the client presents itself with any other security type.

When using the option `none`, it is important to remember that the read-only parameter is processed first. Consider the following guidelines when configuring export rules for clients with unlisted security types:

Read-only includes <code>none</code>	Read-write includes <code>none</code>	Resulting access for clients with unlisted security types
No	No	Denied
No	Yes	Denied because read-only is processed first
Yes	No	Read-only as anonymous
Yes	Yes	Read-write as anonymous

Example

The export policy contains an export rule with the following parameters:

- `-protocol nfs3`
- `-clientmatch 10.1.16.0/255.255.255.0`
- `-rorule sys,none`
- `-rwrule any`
- `-anon 70`

Client #1 has the IP address 10.1.16.207, sends an access request using the NFSv3 protocol, and authenticated with Kerberos v5.

Client #2 has the IP address 10.1.16.211, sends an access request using the NFSv3 protocol, and authenticated with `AUTH_SYS`.

Client #3 has the IP address 10.1.16.234, sends an access request using the NFSv3 protocol, and did not authenticate (meaning security type `AUTH_NONE`).

The client access protocol and IP address matches for all three clients. The read-only parameter allows read-only access to clients with their own user ID that authenticated with `AUTH_SYS`. The read-only parameter allows read-only access as the anonymous user with user ID 70 to clients that authenticated using any other security type. The read-write parameter allows read-write access to any security type, but in this case only applies to clients already filtered by the read-only rule.

Therefore, clients #1 and #3 get read-write access only as the anonymous user with user ID 70. Client #2 gets

read-write access with its own user ID.

Example

The export policy contains an export rule with the following parameters:

- `-protocol nfs3`
- `-clientmatch 10.1.16.0/255.255.255.0`
- `-rorule sys,none`
- `-rwrule none`
- `-anon 70`

Client #1 has the IP address 10.1.16.207, sends an access request using the NFSv3 protocol, and authenticated with Kerberos v5.

Client #2 has the IP address 10.1.16.211, sends an access request using the NFSv3 protocol, and authenticated with AUTH_SYS.

Client #3 has the IP address 10.1.16.234, sends an access request using the NFSv3 protocol, and did not authenticate (meaning security type AUTH_NONE).

The client access protocol and IP address matches for all three clients. The read-only parameter allows read-only access to clients with their own user ID that authenticated with AUTH_SYS. The read-only parameter allows read-only access as the anonymous user with user ID 70 to clients that authenticated using any other security type. The read-write parameter allows read-write access only as the anonymous user.

Therefore, client #1 and client #3 get read-write access only as the anonymous user with user ID 70. Client #2 gets read-only access with its own user ID but is denied read-write access.

How security types determine client access levels

The security type that the client authenticated with plays a special role in export rules. You must understand how the security type determines the levels of access the client gets to a volume or qtree.

The three possible access levels are as follows:

1. Read-only
2. Read-write
3. Superuser (for clients with user ID 0)

Because the access level by security type is evaluated in this order, you must observe the following rules when constructing access level parameters in export rules:

For a client to get access level...	These access parameters must match the client's security type...
Normal user read-only	Read-only (<code>-rorule</code>)

For a client to get access level...	These access parameters must match the client's security type...
Normal user read-write	Read-only (<code>-rorule</code>) and read-write (<code>-rwrule</code>)
Superuser read-only	Read-only (<code>-rorule</code>) and <code>-superuser</code>
Superuser read-write	Read-only (<code>-rorule</code>) and read-write (<code>-rwrule</code>) and <code>-superuser</code>

The following are valid security types for each of these three access parameters:

- `any`
- `none`
- `never`

This security type is not valid for use with the `-superuser` parameter.

- `krb5`
- `krb5i`
- `krb5p`
- `ntlm`
- `sys`

When matching a client's security type against each of the three access parameters, there are three possible outcomes:

If the client's security type...	Then the client...
Matches the one specified in the access parameter.	Gets access for that level with its own user ID.
Does not match the one specified, but the access parameter includes the option <code>none</code> .	Gets access for that level but as the anonymous user with the user ID specified by the <code>-anon</code> parameter.
Does not match the one specified and the access parameter does not include the option <code>none</code> .	Does not get any access for that level. This does not apply to the <code>-superuser</code> parameter because it always includes <code>none</code> even when not specified.

Example

The export policy contains an export rule with the following parameters:

- `-protocol nfs3`
- `-clientmatch 10.1.16.0/255.255.255.0`
- `-rorule any`

- `-rwrule sys,krb5`
- `-superuser krb5`

Client #1 has the IP address 10.1.16.207, has user ID 0, sends an access request using the NFSv3 protocol, and authenticated with Kerberos v5.

Client #2 has the IP address 10.1.16.211, has user ID 0, sends an access request using the NFSv3 protocol, and authenticated with AUTH_SYS.

Client #3 has the IP address 10.1.16.234, has user ID 0, sends an access request using the NFSv3 protocol, and did not authenticate (AUTH_NONE).

The client access protocol and IP address matches all three clients. The read-only parameter allows read-only access to all clients regardless of security type. The read-write parameter allows read-write access to clients with their own user ID that authenticated with AUTH_SYS or Kerberos v5. The superuser parameter allows superuser access to clients with user ID 0 that authenticated with Kerberos v5.

Therefore, client #1 gets superuser read-write access because it matches all three access parameters. Client #2 gets read-write access but not superuser access. Client #3 gets read-only access but not superuser access.

Manage superuser access requests

When you configure export policies, you need to consider what you want to happen if the storage system receives a client access request with user ID 0, meaning as a superuser, and set up your export rules accordingly.

In the UNIX world, a user with the user ID 0 is known as the superuser, typically called root, who has unlimited access rights on a system. Using superuser privileges can be dangerous for several reasons, including breach of system and data security.

By default, ONTAP maps clients presenting with user ID 0 to the anonymous user. However, you can specify the `-superuser` parameter in export rules to determine how to handle clients presenting with user ID 0 depending on their security type. The following are valid options for the `-superuser` parameter:

- `any`
- `none`

This is the default setting if you do not specify the `-superuser` parameter.

- `krb5`
- `ntlm`
- `sys`

There are two different ways how clients presenting with user ID 0 are handled, depending on the `-superuser` parameter configuration:

If the <code>-superuser</code> parameter and the client's security type...	Then the client...
Match	Gets superuser access with user ID 0.
Do not match	Gets access as the anonymous user with the user ID specified by the <code>-anon</code> parameter and its assigned permissions. This is regardless of whether the read-only or read-write parameter specifies the option <code>none</code> .

If a client presents with user ID 0 to access a volume with NTFS security style and the `-superuser` parameter is set to `none`, ONTAP uses the name mapping for the anonymous user to obtain the proper credentials.

Example

The export policy contains an export rule with the following parameters:

- `-protocol nfs3`
- `-clientmatch 10.1.16.0/255.255.255.0`
- `-rorule any`
- `-rwrule krb5,ntlm`
- `-anon 127`

Client #1 has the IP address 10.1.16.207, has user ID 746, sends an access request using the NFSv3 protocol, and authenticated with Kerberos v5.

Client #2 has the IP address 10.1.16.211, has user ID 0, sends an access request using the NFSv3 protocol, and authenticated with AUTH_SYS.

The client access protocol and IP address matches for both clients. The read-only parameter allows read-only access to all clients regardless of the security type they authenticated with. However, only client #1 gets read-write access because it used the approved security type Kerberos v5 to authenticate.

Client #2 does not get superuser access. Instead, it gets mapped to anonymous because the `-superuser` parameter is not specified. This means it defaults to `none` and automatically maps user ID 0 to anonymous. Client #2 also only gets read-only access because its security type did not match the read-write parameter.

Example

The export policy contains an export rule with the following parameters:

- `-protocol nfs3`
- `-clientmatch 10.1.16.0/255.255.255.0`
- `-rorule any`
- `-rwrule krb5,ntlm`
- `-superuser krb5`
- `-anon 0`

Client #1 has the IP address 10.1.16.207, has user ID 0, sends an access request using the NFSv3 protocol, and authenticated with Kerberos v5.

Client #2 has the IP address 10.1.16.211, has user ID 0, sends an access request using the NFSv3 protocol, and authenticated with AUTH_SYS.

The client access protocol and IP address matches for both clients. The read-only parameter allows read-only access to all clients regardless of the security type they authenticated with. However, only client #1 gets read-write access because it used the approved security type Kerberos v5 to authenticate. Client #2 does not get read-write access.

The export rule allows superuser access for clients with user ID 0. Client #1 gets superuser access because it matches the user ID and security type for the read-only and `-superuser` parameters. Client #2 does not get read-write or superuser access because its security type does not match the read-write parameter or the `-superuser` parameter. Instead, client #2 is mapped to the anonymous user, which in this case has the user ID 0.

How ONTAP uses export policy caches

To improve system performance, ONTAP uses local caches to store information such as host names and netgroups. This enables ONTAP to process export policy rules more quickly than retrieving the information from external sources. Understanding what the caches are and what they do can help you troubleshoot client access issues.

You configure export policies to control client access to NFS exports. Each export policy contains rules, and each rule contains parameters to match the rule to clients requesting access. Some of these parameters require ONTAP to contact an external source, such as DNS or NIS servers, to resolve objects such as domain names, host names, or netgroups.

These communications with external sources take a small amount of time. To increase performance, ONTAP reduces the amount of time it takes to resolve export policy rule objects by storing information locally on each node in several caches.

Cache name	Type of information stored
Access	Mappings of clients to corresponding export policies
Name	Mappings of UNIX user names to corresponding UNIX user IDs
ID	Mappings of UNIX user IDs to corresponding UNIX user IDs and extended UNIX group IDs
Host	Mappings of host names to corresponding IP addresses
Netgroup	Mappings of netgroups to corresponding IP addresses of members
Showmount	List of exported directories from SVM namespace

If you change information on the external name servers in your environment after ONTAP retrieved and stored it locally, the caches might now contain outdated information. Although ONTAP refreshes caches automatically after certain time periods, different caches have different expiration and refresh times and algorithms.

Another possible reason for caches to contain outdated information is when ONTAP attempts to refresh cached information but encounters a failure when attempting to communicate with name servers. If this happens, ONTAP continues to use the information currently stored in the local caches to prevent client disruption.

As a result, client access requests that are supposed to succeed might fail, and client access requests that are supposed to fail might succeed. You can view and manually flush some of the export policy caches when troubleshooting such client access issues.

How the access cache works

ONTAP uses an access cache to store the results of export policy rule evaluation for client access operations to a volume or qtree. This results in performance improvements because the information can be retrieved much faster from the access cache than going through the export policy rule evaluation process every time a client sends an I/O request.

Whenever an NFS client sends an I/O request to access data on a volume or qtree, ONTAP must evaluate each I/O request to determine whether to grant or deny the I/O request. This evaluation involves checking every export policy rule of the export policy associated with the volume or qtree. If the path to the volume or qtree involves crossing one or more junction points, this might require performing this check for multiple export policies along the path.

Note that this evaluation occurs for every I/O request sent from an NFS client, such as read, write, list, copy and other operations; not just for initial mount requests.

After ONTAP has identified the applicable export policy rules and decided whether to allow or deny the request, ONTAP then creates an entry in the access cache to store this information.

When an NFS client sends an I/O request, ONTAP notes the IP address of the client, the ID of the SVM, and the export policy associated with the target volume or qtree, and first checks the access cache for a matching entry. If a matching entry exists in the access cache, ONTAP uses the stored information to allow or deny the I/O request. If a matching entry does not exist, ONTAP then goes through the normal process of evaluating all applicable policy rules as explained above.

Access cache entries that are not actively used are not refreshed. This reduces unnecessary and wasteful communication with external name serves.

Retrieving the information from the access cache is much faster than going through the entire export policy rule evaluation process for every I/O request. Therefore, using the access cache greatly improves performance by reducing the overhead of client access checks.

How access cache parameters work

Several parameters control the refresh periods for entries in the access cache. Understanding how these parameters work enables you to modify them to tune the access cache and balance performance with how recent the stored information is.

The access cache stores entries consisting of one or more export rules that apply to clients attempting to

access volumes or qtrees. These entries are stored for a certain amount of time before they are refreshed. The refresh time is determined by access cache parameters and depends on the type of access cache entry.

You can specify access cache parameters for individual SVMs. This allows the parameters to differ according to SVM access requirements. Access cache entries that are not actively used are not refreshed, which reduces unnecessary and wasteful communication with external name serves.

Access cache entry type	Description	Refresh period in seconds
Positive entries	Access cache entries that have not resulted in access denial to clients.	Minimum: 300 Maximum: 86,400 Default: 3,600
Negative entries	Access cache entries that have resulted in access denial to clients.	Minimum: 60 Maximum: 86,400 Default: 3,600

Example

An NFS client attempts to access a volume on a cluster. ONTAP matches the client to an export policy rule and determines that the client gets access based on the export policy rule configuration. ONTAP stores the export policy rule in the access cache as a positive entry. By default, ONTAP keeps the positive entry in the access cache for one hour (3,600 seconds), and then automatically refreshes the entry to keep the information current.

To prevent the access cache from filling up unnecessarily, there is an additional parameter to clear existing access cache entries that have not been used for a certain time period to decide client access. This `-harvest -timeout` parameter has an allowed range of 60 through 2,592,000 seconds and a default setting of 86,400 seconds.

Remove an export policy from a qtree

If you decide you do not want a specific export policy assigned to a qtree any longer, you can remove the export policy by modifying the qtree to inherit the export policy of the containing volume instead. You can do this by using the `volume qtree modify` command with the `-export-policy` parameter and an empty name string (`""`).

Steps

1. To remove an export policy from a qtree, enter the following command:

```
volume qtree modify -vserver vservers_name -qtree-path  
/vol/volume_name/qtree_name -export-policy ""
```

2. Verify that the qtree was modified accordingly:

```
volume qtree show -qtree qtree_name -fields export-policy
```

Validate qtree IDs for qtree file operations

ONTAP can perform an optional additional validation of qtree IDs. This validation ensures that client file operation requests use a valid qtree ID and that clients can only move files within the same qtree. You can enable or disable this validation by modifying the `-validate-qtrees-export` parameter. This parameter is enabled by default.

About this task

This parameter is only effective when you have assigned an export policy directly to one or more qtrees on the storage virtual machine (SVM).

Steps

1. Set the privilege level to advanced:

```
set -privilege advanced
```

2. Perform one of the following actions:

If you want qtree ID validation to be...	Enter the following command...
Enabled	<pre>vserver nfs modify -vserver vserver_name -validate-qtrees-export enabled</pre>
Disabled	<pre>vserver nfs modify -vserver vserver_name -validate-qtrees-export disabled</pre>

3. Return to the admin privilege level:

```
set -privilege admin
```

Export policy restrictions and nested junctions for FlexVol volumes

If you configured export policies to set a less restrictive policy on a nested junction but a more restrictive policy on a higher level junction, access to the lower level junction might fail.

You should ensure that higher level junctions have less restrictive export policies than lower level junctions.

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.