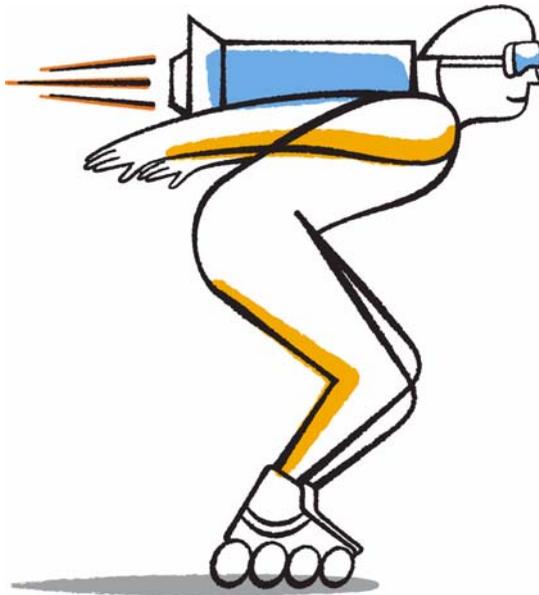




Data ONTAP[®] 8.1

NFSv3 Express Guide

For 7-Mode Administrators Learning Cluster-Mode



NetApp, Inc.
495 East Java Drive
Sunnyvale, CA 94089
U.S.

Telephone: +1 (408) 822-6000
Fax: +1 (408) 822-4501
Support telephone: +1 (888) 463-8277
Web: www.netapp.com
Feedback: doccomments@netapp.com

Part number: 215-07231_A0
Updated for Data ONTAP 8.1.1
September 2012

Contents

Deciding whether to use this guide	4
Mode differences in NFS setup	5
How exports differ in Cluster-Mode	5
File path mount points in Cluster-Mode	7
Export policy and rule concepts	9
What the default export policy is	9
How export rules are defined in System Manager	10
Examples of export policies in Cluster-Mode	14
Worksheet: Cluster and network information	18
Setting up NFS	20
Creating a Vserver that supports NFS	21
Creating an export policy in System Manager	23
Creating FlexVol volumes	26
Applying export policies to volumes	27
Setting up a test client and verifying system operation	28
Configuring user authentication	29
Configuring an NIS domain	29
Configuring local users and groups	30
Verifying authentication	30
Next steps	32
Where to go next	33
Copyright information	34
Trademark information	35
How to send your comments	36
Index	37

Deciding whether to use this guide

This guide describes how to set up NFSv3 access for clients with Data ONTAP operating in Cluster-Mode. It includes information to help administrators who are familiar with Data ONTAP operating in 7-Mode.

You should use this guide if you want a standard configuration following NetApp best practices, and you do not want information about all the available options or a lot of conceptual background for the tasks.

- This guide assumes that your storage system has been successfully installed and a cluster has been created.
- This guide assumes you download and run OnCommand System Manager 2.0.2 or later for all applicable tasks.
It does not include procedures using the Data ONTAP CLI except when the CLI is the only way to complete a task.
- This guide describes how to configure a system with FlexVol volumes only.
It does not include information about configuring Infinite Volumes.

If these assumptions are not correct for your installation, or if you want more conceptual background information, you should see the following documentation instead:

- *Data ONTAP Software Setup Guide for Cluster-Mode* (for new systems)
- *Data ONTAP System Administration Guide for Cluster-Mode* (for Vserver creation)
- *Data ONTAP File Access and Protocols Management Guide for Cluster-Mode* (for NFS and CIFS)
- *OnCommand System Manager Help* (available both from within the product and as a PDF download)

This documentation is available from the Product Documentation section of the NetApp Support Site.

Related information

[Documentation on the NetApp Support Site: support.netapp.com](http://support.netapp.com)

Mode differences in NFS setup

The overall process of setting up NFS version 3 access to a storage system in a Cluster-Mode environment is similar to the same process in a 7-Mode environment. However, how you perform the tasks differs.

In a Cluster-Mode environment, after the storage system has been set up and a cluster and aggregates have been created, you perform the following tasks:

1. Create an NFS server.
2. Define exports to control client access.
3. Configure user authentication for clients.
4. Set up a test client, and verify system operation.

This list of tasks is similar to what you do to set up NFS access in a 7-Mode environment.

Of these tasks, the one that differs the most is defining exports. In Cluster-Mode there is no `/etc/exports` file and no `exportfs` command. Instead, you must define an export policy. Export policies permit you to control client access in much the same way as you did in 7-Mode, but give you additional functionality such as the ability to reuse the same export policy for multiple volumes.

Another difference is that the functionality provided by the client-based command `showmount` is not available. One way of seeing the mount points that are available in a Cluster-Mode environment is to use System Manager. Select the Vserver, then go to **Storage > Namespace**.

How exports differ in Cluster-Mode

Cluster-Mode exports are defined and used differently than they are in 7-Mode environments.

Table 1: Overview of differences

Areas of difference	7-Mode	Cluster-Mode
How exports are defined	Exports are defined in the <code>/etc/exports</code> file.	Exports are defined by creating an export policy within a Vserver. A Vserver can include more than one export policy.

Areas of difference	7-Mode	Cluster-Mode
Scope of export	<ul style="list-style-type: none"> • Exports apply to a specified file path or qtree. • You must create a separate entry in <code>/etc/exports</code> for each file path or qtree. • Exports are persistent only if they are defined in the <code>/etc/exports</code> file. 	<ul style="list-style-type: none"> • Export policies apply to an entire volume, including all of the file paths and qtrees contained in the volume. • Export policies can be applied to more than one volume if you want. • All export policies are persistent across system restarts.
Fencing (specifying different access for specific clients to the same resources)	To provide specific clients different access to a single exported resource, you have to list each client and its permitted access in the <code>/etc/exports</code> file.	<p>Export policies are composed of a number of individual export rules. Each export rule defines the access that a specified client has to a resource.</p> <p>To specify different access for specific clients, you have to create an export rule for each client, and then add the rules to the export policy.</p>
Name aliasing	When you define an export, you can choose to make the name of the export different than the name of the file path. You should use the <code>-actual</code> parameter when defining the export in the <code>/etc/exports</code> file.	<p>You can choose to make the name of the exported volume different than the actual volume name.</p> <p>To do this, you must mount the volume with a custom junction path name within the Vserver's namespace.</p> <p>Note: By default, volumes are mounted with their volume name. To customize a volume's junction path name you need to unmount it, rename it, and then remount it.</p>

Areas of difference	7-Mode	Cluster-Mode
How root users are treated		When creating an export policy, you have to explicitly state which clients permit superuser access for users with UNIX user ID = 0.

References

For other important information about export policies, see the following references:

- *Data ONTAP File Access and Protocols Management Guide for Cluster-Mode.*
- *Technical Report TR 4067: NFSv3/v4 in Data ONTAP 8.1 Operating in Cluster-Mode Implementation Guide*
- *Knowledge Base article 3011272: How do export-policies work in Data ONTAP GX, 8.0 and 8.1 Cluster-Mode?*

This documentation is available from the NetApp Support Site.

Related concepts

[Export policy and rule concepts](#) on page 9

Related information

[Documentation on the NetApp Support Site: support.netapp.com](http://support.netapp.com)

[Knowledge Base on the NetApp Support Site: kb.netapp.com/support/index?page=home](http://kb.netapp.com/support/index?page=home)

File path mount points in Cluster-Mode

In 7-Mode, you can export any part of a file path, and can also create an alias for its name. It is possible to re-create these features of 7-Mode file path mount points in Cluster-Mode by applying an export policy to nested volume junctions.

How to re-create 7-Mode file path mount points

In a 7-Mode environment, you can export any element of a file path. For example, to specify read/write access to `/dir1` for one client while providing another client read-only access to `/dir1/subdir`, you can create separate entries for each file path in the `/etc/exports` file.

In a Cluster-Mode environment, you can achieve the same effect using export policies and nested junctions:

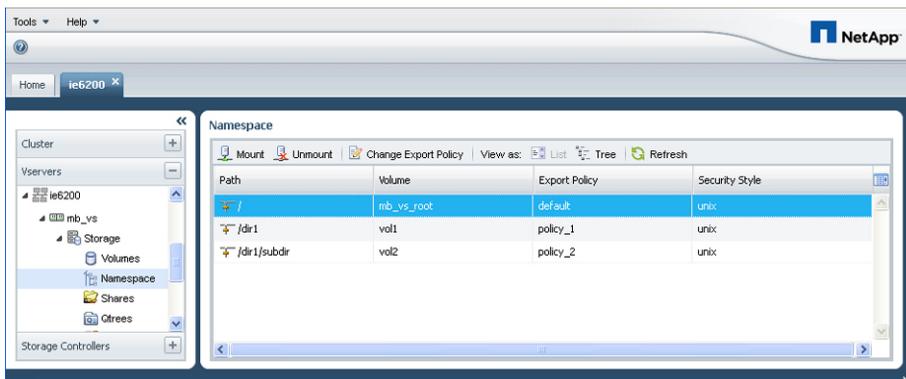
1. Create a separate volume for each set of data that requires a distinct export policy. That is, create two volumes in your Vserver, `vol1` and `vol2`.
2. Junction these volumes within the namespace of the Vserver such that the volumes are mounted in a hierarchy, and have the name that you require. That is, mount `vol1` with the junction name

`dir1` and a junction path of `/`. Mount `vol2` with the junction name `subdir` and a junction path of `dir1`.

3. Apply separate export policies to each volume. That is, apply a read/write export policy to `vol1` and a read-only export policy to `vol2`.

The net effect is the same as it was in the 7-Mode example: one client can mount `/dir1` read/write, while the second client can mount `/dir1/subdir` read-only.

The following image of the **Storage > Namespace** screen in System Manager shows these junctioned volumes with the appropriate export policies applied:



Export policies and nested junctions in Cluster-Mode

When using nested junctions, you should be aware that applying a more restrictive export policy to a higher-level junction might affect a client's access to volumes that are junctioned lower in the hierarchy. Do not restrict access for a client to a junction if that client requires access to any of its child junctions.

For example, if you create an the export policy `policy1` that permits a client read/write access to `subdir` but does not permit any access to `dir1`, the client must mount `/dir1/subdir` directly to access it. If the client mounts `/`, it can see `dir1` but cannot see what is within it, and therefore cannot change directories to `/dir1/subdir`.

References

For more information about export policies and junctioned volumes, see *Knowledge Base article 1013380: How to give a client full control to a junctioned volume over NFS, but no access (not even read access) to the root volume of a Vserver*

Related information

Knowledge Base on the NetApp Support Site: kb.netapp.com/support/index?page=home

Export policy and rule concepts

Export policies enable you to restrict access to volumes to clients that match specific IP addresses and specific authentication types. An export policy with export rules must exist on a Vserver for clients to access data.

You can associate exactly one export policy with each volume. However, a Vserver can contain multiple export policies. This enables you to assign different export policies to each volume in a Vserver, or to assign a given export policy to more than one volume. A Data ONTAP cluster can contain up to 1,024 export policies.

Export policies consist of individual *export rules*. An export policy can contain a large number of rules (approximately 4,000).

Important things to know about export rules:

- Each export rule specifies access permissions to volumes for one or more clients. The clients can be specified by host name, IP address, or netgroup.
- Export rules can use host entries from a netgroup.
- Export rules are processed in the order in which they appear in the export policy. The rule order is dictated by the rule index number.
- Export rules specify the authentication types that are required for both read-only and read/write operations.
- To have any access to a volume, matching clients must authenticate with the authentication type specified by the read-only rule.
- To have write access to the volume, matching clients must authenticate with the authentication type specified by the read/write rule.

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 volume's 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 Data ONTAP.

What the default export policy is

Each Vserver has a default export policy. When a volume is created within a Vserver, the default export policy of that Vserver is automatically assigned to the newly created volume.

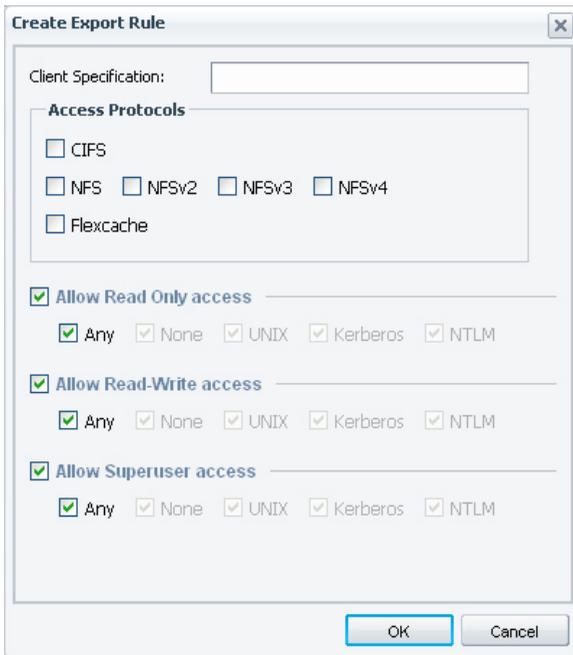
You can modify the default export policy by adding one or more export rules that grant the access permissions that you require, or you can create custom export policies. Creating and applying custom export policies helps ensure that you review and apply exactly the permissions that you intend for every volume.

How export rules are defined in System Manager

You should understand how export rules are defined in System Manager to help you create or modify the rules that you use to control access to volumes.

How export rules are defined

The Create Export Rule dialog box is used to add an export rule to an export policy:



The following table explains the options available in the dialog box:

Options in the Create Export Rule dialog box	Description
<p>Client Specification</p>	<p>You can specify the client in any of the following ways:</p> <ul style="list-style-type: none"> • As a host name, for instance, host1 • As an IPv4 address; for example, 10.1.12.24 • As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.10/24 • As a netgroup, with the netgroup name preceded by the @ character; for instance, @netgroup <p>Note: Netgroups can be specified even if the Vserver is not part of an NIS or LDAP domain.</p> <ul style="list-style-type: none"> • As a domain name preceded by the "." character; for instance, .example.com <p>Note: Entering an IP address range, such as 10.1.12.10-10.1.12.70, is not allowed. Entries in this format are interpreted as a text string and are treated as a host name.</p> <p>Note: You can enter the IPv4 address 0.0.0.0/0 to provide access to all hosts.</p>
<p>Access Protocols</p>	<p>If you do not select any access protocols, the export rule defaults to “any”. This means that the specified client is permitted access using any protocol.</p> <p>If you select a specific protocol, the client is granted access only using that protocol. For example, if you select NFSv3, the export rule does not permit access using NFSv4, NFSv2, or CIFS.</p>

Options in the Create Export Rule dialog box	Description
<p>Allow Read Only access</p> <p>Allow Read-Write access</p> <p>Allow Superuser access</p>	<p>If you do not select an access level, that type of access is not allowed to any client.</p> <p>For each selected access level, you must also select the security type that clients must authenticate with in order to be granted that access.</p> <ul style="list-style-type: none"> • Choosing Any grants access to a user from a matching client, regardless of the security type used. • Choosing UNIX, Kerberos, or NTLM means that only users that are authenticated using the specified security type or types are granted access. • Selecting None means that if users are authenticated using any other authentication method, they are granted access as anonymous users.

How to control the processing order of export rules

Rules within an export policy are evaluated “top down.”

In the Create Export Policy dialog box, you can control the processing order of the rules by selecting a rule and clicking **Up** or **Down**.

Policy Name:

Copy Rules from

VServer:

Source Policy:

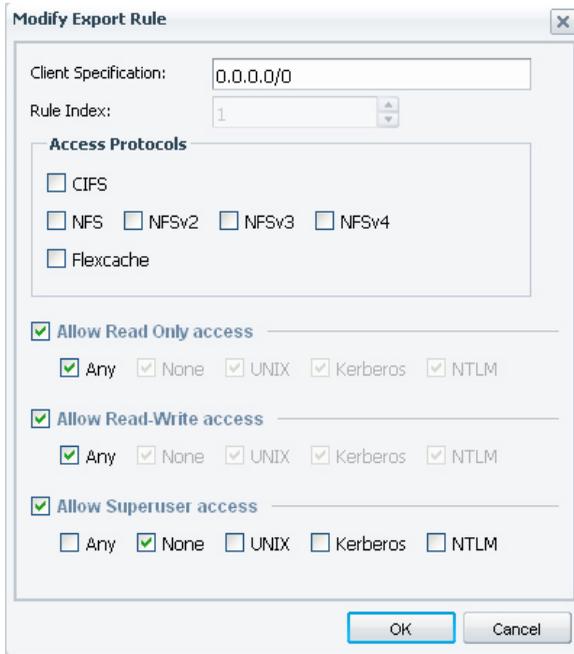
Client Specification	Access Proto...	Read-Only Rule	Read-Write R...	Superuser
@netgroup_1	NFSv3	UNIX	Never	UNIX
0.0.0.0/0	NFSv3	Never	Any	Never

Buttons: Up, Down, Add, Edit, Delete, Create, Cancel

You can also specify the order of an export rule by explicitly setting the export rule's index number when adding a rule to an existing policy in System Manager, or when creating an export rule at the command line.

About the default export policy created in System Manager

When you create a Vserver using System Manager, a default export policy with a single rule is created and applied to the root volume of that Vserver. The default export rule is shown in the following image.



This export rule permits any client to access the root volume of the Vserver, using any access protocol. Users from the matching clients are granted read-only and read/write access to the root volume when they are authenticated by any method, but no one is permitted root access.

Examples of export policies in Cluster-Mode

You can review example export policies to better understand how export policies work in Cluster-Mode.

Sample 7-Mode export policy

The following is an example of a 7-Mode export as it appears in the `/etc/export` file:

```
/vol/vol1 -sec=sys,ro=@readonly_netgroup,rw=@readwrite_netgroup1:
@readwrite_netgroup2:@rootaccess_netgroup,root=@rootaccess_netgroup
```

To reproduce this export as a Cluster-Mode export policy, you have to create an export policy with four export rules, and then apply the export policy to the volume `vol1`.

Rule	Element	Value
Rule 1	Client Specification	@readonly_netgroup1
	Rule Index (or position of export rule in the list of rules)	1
	Access Protocols	NFS
	Allow Read Only access	Selected, with UNIX selected for the authentication method
Rule 2	Client Specification	@rootaccess_netgroup
	Rule Index	2
	Access Protocols	NFS
	Allow Superuser access	Selected, with UNIX selected for the authentication method
Rule 3	Client Specification	@readwrite_netgroup1
	Rule Index	3
	Access Protocols	NFS
	Allow Read Write access	Selected, with UNIX selected for the authentication method
Rule 4	Client Specification	@readwrite_netgroup2
	Rule Index	4
	Access Protocols	NFS
	Allow Read Write access	Selected, with UNIX selected for the authentication method

Export policy that implements fencing

The following example shows how to provide read/write access to some clients and read-only access to others.

- Read/write access is granted to all users from the host name `host1`
- Clients from the domain `example.com` are provided read-only access

This export policy requires two rules:

Rule	Element	Value
Rule 1	Client Specification	host1
	Rule Index (or position of export rule in the list of rules)	1
	Access Protocols	NFSv3
	Allow Read-Write access	Selected, with Any selected for the authentication method.
Rule 2	Client Specification	.example.com
	Rule Index	2
	Access Protocols	NFSv3
	Allow Read Only access	Selected, with Any selected for the authentication method

Export policy where the rule index value is important

This example shows how to define a set of restrictions that permit the following access to a volume:

- Read-only access to clients in the netgroup “netgroup_1” for users that are authenticated with system authentication
- Superuser access for "root" users in "netgroup_1" for users that use system authentication
- Read/write access to all the other clients

Implementing these restrictions requires the creation of an export policy that contains two rules:

Rule	Element	Value
Rule 1	Client Specification	@netgroup_1
	Rule Index (or position of export rule in the list of rules)	1
	Access Protocols	NFSv3
	Allow Read Only access	Selected, with UNIX selected
	Allow Superuser access	Selected, with UNIX selected
Rule 2	Client Specification	0.0.0.0/0
	Rule Index	2
	Access Protocols	NFSv3
	Allow Read-Write access	Any

Note: The order of the rules is important because rules in an export policy are processed in numerical order, and processing stops after a rule is satisfied for a client. Therefore, if you swapped the two rules such that Rule 2 had a rule index value of 1, then clients in `netgroup_1` will have read/write access to the volume.

Also, this export policy enables superuser access only for users of clients in `netgroup_1`. Users with root access to clients that are not in `netgroup_1` are mapped to the UNIX user ID=65534, or "Nobody."

Related concepts

[Export policy and rule concepts](#) on page 9

Related tasks

[Creating an export policy in System Manager](#) on page 23

Related references

[How export rules are defined in System Manager](#) on page 10

Worksheet: Cluster and network information

You can use this worksheet to collect the information that you require to set up NFS access for clients. You must obtain the IP addresses and other information about your cluster and network from your storage and network administrators before you begin.

Information for Vserver configuration

Domain Name Service (DNS)

		Host name	IP address
	Primary domain Server		
	Search domains (optional: up to 5)		
1			
2			
3			
4			
5			
	Name servers (at least 1; up to 3)		
1			
2			
3			

Data interfaces

In Cluster-Mode, data interfaces are hosted on logical interfaces, known as data LIFs. A LIF has a home port and a home node, but is not tied to a single physical interface. You must create at least one data LIF on the Vserver for use by the NFS protocol.

Protocol	Interface name	Home node:port	IP address	Netmask	Gateway
NFS					

Name service and authentication information

Network Information Service (NIS)

	Host name or IP address
Primary NIS domain	
NIS server	
Secondary domains (optional)	

Local users and groups

	Group name	Group ID
1		
2		
3		
...		
n		

	User name	User ID	Group ID
1			
2			
3			
...			
n			

Setting up NFS

You can review the tasks involved in setting up NFS, the prerequisites to the setup, and the optional tasks that you might choose to perform after the setup.

Before you begin

- Cluster must be created
- Data aggregate must be created
- Time zone and NTP must be configured



Setting up NFS

Create a Vserver that supports NFS.



Create an export policy.



Create a volume.



Apply the export policy to the volume.

Test



Configure NIS or local users and groups.

Test



Optional next tasks

- Update the export rules for root.
- Configure routing.
- Configure failover groups.

Creating a Vserver that supports NFS

The Vserver is the basic building block of all Cluster-Mode configurations. The Vserver contains the volume, NFS server, and logical interfaces used to access data using NFS.

Before you begin

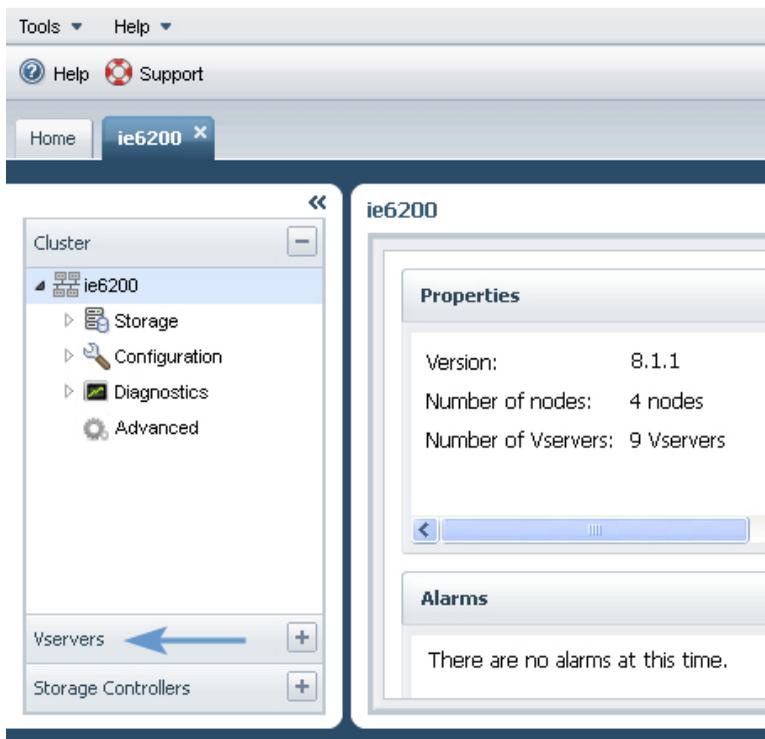
- The cluster must have at least one data aggregate available that you can use when creating the Vserver.
- You must have set the time zone and synchronized the time across the cluster by configuring NTP.

You can verify this in System Manager by selecting the cluster and clicking **Configuration > System Tools > DateTime**. This prevents authentication errors, and ensures that time stamps in log files are consistent across the cluster.

- The cluster must have an NFS license.
You can verify this in System Manager by selecting the cluster and clicking **Configuration > System Tools > Licenses**.

Steps

1. Start System Manager, and on the **Home** tab, double-click the appropriate storage system.
2. In the left navigation pane, click **Vservers**.



3. In the left navigation pane, select the cluster and click **Create**.
4. Follow the instructions in the **Create Vserver** wizard:
 - a) When prompted to select a protocol, select **NFS**.
 - b) Configure DNS for the Vserver by entering the name of the server that provides primary domain name service for the storage system and information about at least one additional name server.
 - c) Create one or more data interfaces.

Data interfaces are used by NFS and CIFS clients to access the storage system. For a given interface, you can enable one protocol or both protocols.
 - d) When prompted, select **NFS version 3**.
 - e) Select **Local User** and create a test local user and group to use when verifying your system configuration.

To simplify the set up process, do not select NIS or LDAP at this time.
5. After completing the wizard, select the new Vserver in the left navigation pane and view its storage and configuration.

Result

Your Vserver is created with the following:

- An NFS server that supports NFS version 3
- A root volume that has the UNIX security style
- DNS enabled
- One or more logical data interfaces for client access to data
- A test local user and group to use when verifying system operation

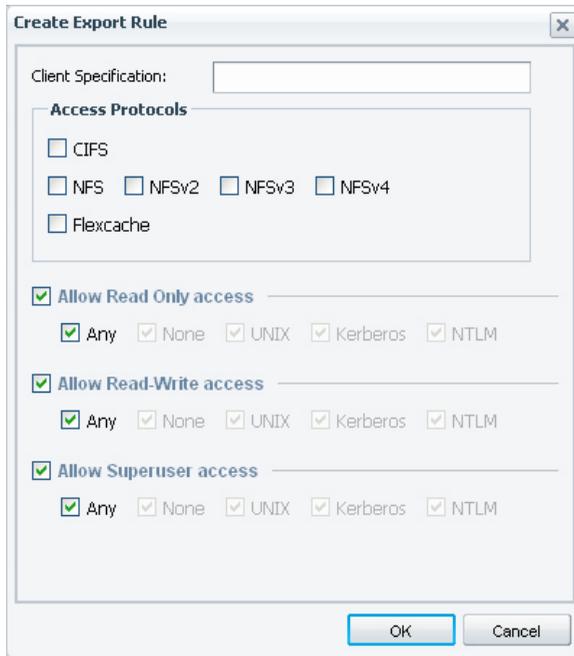
The Vserver is started automatically. It contains a root volume under which you can junction additional volumes for data storage.

Creating an export policy in System Manager

Export policies contain a set of rules to specify the access that clients have to volumes in a Vserver. You must create an export policy for your Vserver.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Policies > Export Policies**.
4. Click **Create Policy** and specify a policy name.
5. Click **Add** to add an export rule to the policy.
6. In the **Create Export Rule** dialog box, perform the following steps:
 - a) Specify the client that requires access to the data.



- b) Select **NFSv3**.
- c) Select one or more access types and its security type.
- d) Click **OK**.

The first rule is added to the export policy.

- 7. Click **Add** and complete the dialog box to add another export rule to the policy.

The export policy now has two rules that are processed in the order that they are shown.

Policy Name:

Copy Rules from

VServer:

Source Policy:

Client Specification	Access Proto...	Read-Only Rule	Read-Write R...	Superuser
@netgroup_1	NFSv3	UNIX	Never	UNIX
0.0.0.0/0	NFSv3	Never	Any	Never

Up
Down

Add Edit Delete

Create Cancel

8. To change which rule is evaluated first, select a rule and click **Up** or **Down** to change its position in the list.
9. Click **Create** to create the export policy.

Policy	Client Specification	Access Protocols	Read-Only Rule	Read-Write Rule	Superuser
default					
rule_index_order_k					
1	@netgroup_1	NFSv3	UNIX	Never	UNIX
2	0.0.0.0/0	NFSv3	Never	Any	Never

When a client attempts to connect, the storage system checks the client against the first export rule. If the client matches the specification, the first rule is applied and the other rules are not checked. Rule processing stops after a match is found. Therefore, if you swap the order of the two rules shown in this export policy, then all clients, including those in `netgroup_1`, are given read/write access to the volume.

Related concepts

[Export policy and rule concepts](#) on page 9

[What the default export policy is](#) on page 9

[How exports differ in Cluster-Mode](#) on page 5

[Examples of export policies in Cluster-Mode](#) on page 14

Related references

[How export rules are defined in System Manager](#) on page 10

Creating FlexVol volumes

You must create a FlexVol volume for your data.

Before you begin

The cluster must contain a non-root aggregate and a Vserver.

About this task

You should always create a volume for your data, rather than storing data in the root volume of a Vserver.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > Volumes**.
4. Click **Create**.
5. If you want to change the default name, specify a new name.
6. Select the containing aggregate for the volume.
7. Select the type of storage for which you are creating this volume.
8. Specify the size of the volume, and accept the default value for the snapshot reserve.

The default space reserved for Snapshot copies is zero percent for SAN and VMware volumes. For NAS volumes, the default is five percent on storage systems running Data ONTAP 8.1.

9. Click **Create**.
10. Verify that the volume you created is included in the list of volumes in the **Volume** window.

The volume is created with UNIX style security and UNIX 700 "read write execute" permissions for the Owner.

After you finish

To change the security style or UNIX permissions for the volume, go to **Storage > Volumes**. Select the volume and click **Edit**.

Applying export policies to volumes

When a volume is created, it automatically inherits the default export policy of the root volume of the Vserver. This procedure describes how to apply your customized export policy.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > Namespace**.
4. Select the volume and click **Change Export Policy**.
5. Select the export policy and click **Change**.

6. Verify that the Export Policy column in the **Namespace** window displays the export policy that you applied to the volume.

Setting up a test client and verifying system operation

You should verify that you have configured NFS correctly by connecting a test client to your storage system and writing and reading data to it.

About this task

You should test your configuration with the test local user that you created when you created the Vserver.

Steps

1. Log in to a client system that you have configured for NFS access.

Example

If you configured the example export policy where rule order is important, log in to a client system that is **not** in `netgroup_1`. Client systems in `netgroup_1` are restricted from writing data.

2. Change the directory to the mount folder:

```
cd /mnt
```

3. Create a mount folder that is named `vsNFS-v3`:

```
mkdir /mnt/vsNFS-v3
```

4. Mount the volume `vol1` at this new directory:

```
mount -t nfs -o nfsvers=3,hard IP_Address:/vol1 /mnt/vsNFS-v3
```

You should use the IP address of the data interface that you configured when enabling NFS for the Vserver.

5. Change the directory to `vol1`:

```
cd vol1
```

6. Create a test file:

```
touch testfile
```

If your configuration is correct, the file is created successfully.

7. List the directory's contents to confirm that the file was created:

```
ls -l
```

8. Remove the test file:

```
rm -r testfile
```

9. Continue to test the other rules in your test export policy.

Example

If you are using the example export policy where rule order is important, log into a client system in `netgroup_1` and attempt to create a file. Your attempt should fail because clients in `netgroup_1` have read-only access.

10. After verifying the export policy, use System Manager to delete the test user:

- a) Select the Vserver, and then click **Configuration > Local Users and Groups > UNIX**.
- b) Select the test user, and click **Delete**.

Result

You have confirmed that you have enabled NFS access to your storage system, and you have correctly implemented an export policy.

Related concepts

[Examples of export policies in Cluster-Mode](#) on page 14

Configuring user authentication

You can edit the Vserver configuration to enable the use of NIS name services or to add local users and groups for system authentication.

Configuring an NIS domain

If you are using NIS, configure an NIS domain for the Vserver. Only one NIS domain can be active on a Vserver at any given time.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Services > NIS**.
4. Click **Create**.
5. Type the NIS domain name and add one or more NIS servers.
6. Click **Create**.

Configuring local users and groups

If you are using system authentication, you should add local users and groups to the Vserver to enable those users to access the storage system.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Local Users and Groups > UNIX**.
4. Add the groups and users that you require.

Verifying authentication

After you have configured user authentication, you should repeat the test of reading and writing data to verify that that you have configured authentication correctly.

Before you begin

- You must have ensured that you can read and write data with a test local user.
- You must have created an export policy that permits read/write access for the client whose access you are testing.

About this task

This procedure repeats the verification steps that you performed earlier for a test user for your production clients. This enables you to ensure that your clients that use system authentication or NIS authentication are configured correctly.

Steps

1. Log in to a client system that you have configured for NFS access as a user authenticated by one of the methods you want to verify.

2. Change the directory to the mount folder that you created in your initial test:

```
cd /mnt/vsNFS-v3
```

3. Mount the volume vol1 at this new directory:

```
mount -t nfs -o nfsvers=3,hard IPAddress:/vol1 /mnt/vsNFS-v3
```

Use the IP address of the data interface that you configured when enabling NFS for the Vserver.

4. Change the directory to vol1:

```
cd vol1
```

5. Create a test file:

```
touch testfile
```

If your configuration is correct, the file is created successfully.

6. List the directory's contents to confirm that the file is created:

```
ls -l
```

7. Remove the test file:

```
rm -r testfile
```

8. Continue to test the other rules in your export policy.

Result

You have confirmed that you have enabled NFS access to your storage system, including configuring your user authentication.

Related tasks

[Setting up a test client and verifying system operation](#) on page 28

Next steps

After access is configured, you might want to further configure your storage system for use in your environment.

Optional tasks

You can complete the following optional tasks after NFS access has been configured and tested:

Task	Reference
Update export rules for root users, which can be done at the command line. Note: By default, export rules map the root user to "nobody", or User ID = 65534. That is, by default export rules implement root squashing.	<i>Data ONTAP File Access and Protocols Management Guide for Cluster-Mode</i> <i>Knowledge Base article 3011272: How do export-policies work in Data ONTAP GX, 8.0 and 8.1 Cluster-Mode?</i>
Configure routing groups and routes for the Vserver's data interfaces.	<i>Data ONTAP Network Management Guide for Cluster-Mode</i>
Configure failover groups for data interfaces.	<i>Data ONTAP Network Management Guide for Cluster-Mode</i>

Related information

[Documentation on the NetApp Support Site: support.netapp.com](http://support.netapp.com)

[Knowledge Base on the NetApp Support Site: kb.netapp.com/support/index?page=home](http://kb.netapp.com/support/index?page=home)

Where to find additional information

There are additional documents and tools to help you learn about the additional setup and configuration steps your cluster might require.

Documentation references

<i>OnCommand System Manager Help</i>	Describes how to use OnCommand System Manager to complete typical tasks. Available both from within the product and as a PDF download.
<i>Data ONTAP 7-Mode to Cluster-Mode Command Map</i>	Provides a mapping of 7-Mode commands to Cluster-Mode commands.
<i>Data ONTAP File Access and Protocols Management Guide for Cluster-Mode</i>	Describes how to manage file access on NetApp systems with CIFS and NFS protocols.
<i>Data ONTAP Logical Storage Management Guide for Cluster-Mode</i>	Describes how to efficiently manage your logical storage resources on NetApp systems running Data ONTAP operating in Cluster-Mode, using volumes, FlexClone volumes, files and LUNs, FlexCache volumes, deduplication, compression, qtrees, and quotas.
<i>Data ONTAP Network Management Guide for Cluster-Mode</i>	Describes how to connect your cluster to your Ethernet networks and how to manage logical interfaces (LIFs).
<i>Data ONTAP System Administration Guide for Cluster-Mode</i>	Describes general system administration for NetApp systems running Data ONTAP operating in Cluster-Mode.

This documentation is available from the Product Documentation section of the NetApp Support Site.

Tool reference

The following tool can help you manage your storage system. It is available from the NetApp Support Site.

Interoperability Matrix Tool (IMT)	Lists supported combinations of hardware components, software versions, firmware, and drivers.
---	--

Related information

[Documentation on the NetApp Support Site: support.netapp.com](http://support.netapp.com)

Copyright information

Copyright © 1994–2012 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.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

NetApp, the NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Cryptainer, CryptoShred, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, FAServer, FastStak, FilerView, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Manage ONTAP, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP, ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, SnapDirector, SnapDrive, SnapFilter, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the web at www.ibm.com/legal/copytrade.shtml.

Apple is a registered trademark and QuickTime is a trademark of Apple, Inc. in the United States and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the United States and/or other countries. RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the United States and/or other countries.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks.

NetApp, Inc. NetCache is certified RealSystem compatible.

How to send your comments

You can help us to improve the quality of our documentation by sending us your feedback.

Your feedback is important in helping us to provide the most accurate and high-quality information. If you have suggestions for improving this document, send us your comments by email to doccomments@netapp.com. To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

Index

A

- about this guide 4
- additional information
 - finding 33
 - optional tasks 32
- audience 4
- authentication
 - configuring different types 29
 - local users and groups 30
 - system 30
 - verifying 30

C

- Cluster-Mode
 - how exports differ 5, 7
 - how NFS setup differs 5
 - default export policy 9
 - export policies 5, 7
 - export policies examples 14–16
- configuring
 - local users and groups 30
 - a NIS domain 29
- creating
 - export policies 23
 - FlexVol volumes 26

D

- domains
 - NIS, configuring 29

E

- examples
 - default export policy 10, 12, 13
 - export policies 14–16
- export policies
 - about default 9
 - and 7-Mode file path exports 7, 8
 - applying them to volumes 27
 - concepts 9
 - creating 23
 - differences in Cluster-Mode 5, 7

- examples 14–16
 - nested junctions 7, 8
- export rules
 - defining in System Manager 10, 12, 13
 - specifying order of 10, 12, 13

F

- FlexVol volumes
 - creating 26

N

- next steps
 - optional tasks 32
- NFS
 - setup overview 20
 - creating Vserver to support 21
 - difference in setup for Cluster-Mode 5
 - optional tasks 32
 - verifying access with authentication 30
 - verifying with a test client 28
 - worksheet to gather setup information 18
- NIS
 - domain, configuring 29

S

- System Manager
 - export rules reference 10, 12, 13

T

- test client
 - verifying NFS access 28

V

- volumes
 - applying export policies to 27
- Vservers
 - creating 21
 - gathering configuration information 18