

IBM System Storage N series



Data ONTAP 8.1 MultiStore Management Guide For 7-Mode

Contents

Preface	9
About this guide	9
Supported features	9
Websites	9
Getting information, help, and service	10
Before you call	10
Using the documentation	10
Hardware service and support	11
Firmware updates	11
How to send your comments	11
What MultiStore is	12
Benefits of using MultiStore	12
MultiStore for consolidating servers	13
MultiStore for service providers and enterprises	14
MultiStore for disaster recovery and data migration	14
The default vFiler unit	15
Number of vFiler units allowed	15
Data access from the hosting storage system	16
Hosting storage system tasks	16
MultiStore management	17
Types of vFiler unit administrators	17
Enabling MultiStore	17
Disabling MultiStore	18
Prerequisites for creating vFiler units	19
Storage guidelines	20
Naming guidelines	20
Language guidelines	21
Quota guidelines	21
HA pair guidelines	21
SAN guidelines	22
The vFiler commands	22
Ensuring that the network interface is ready	23

Creating a vFiler unit	24
Setting up a vFiler unit	25
vFiler unit storage management from the hosting storage system	27
Effects of adding, removing, and moving vFiler unit resources	27
Adding resources to a vFiler unit	27
Requirements for moving and removing resources	28
Removing resources from a vFiler unit	28
Moving resources between vFiler units	29
Default limits on the number of vFiler units	30
Maximum vFiler units allowed	30
Viewing the current limit on the number of vFiler units	30
Increasing the vFiler unit limit	31
Decreasing the vFiler unit limit	31
What the vfiler rename command does	32
Renaming a vFiler unit	32
Stopping a vFiler unit	33
Destroying a vFiler unit	34
Restoring a vFiler unit	35
Starting a vFiler unit	36
Protocols supported by a vFiler unit	37
Allowing a protocol on a vFiler unit	37
Effects of disallowing protocols on a vFiler unit	38
Disallowing a protocol on a vFiler unit	38
Displaying the vFiler unit status	39
Identifying commands that you can execute from a vFiler unit	39
Logging in to a vFiler unit	40
Limits on the number of interactive SSH sessions	41
Executing commands from a vFiler unit	41
Executing commands from the hosting storage system	42
Executing RSH commands for a vFiler unit	42
Executing SSH commands for a vFiler unit	43
List of RSH and SSH commands	44
Effects of storage system reboot on a vFiler unit	46
Volumes and qtrees on a vFiler unit	46
Effects of taking a vFiler unit volume offline	46
Changes required after volumes are renamed	47

Who can change volume or qtree security styles and oplock settings	47
Differences in qtree command output	47
Viewing all qtrees and the owner vFiler units	47
Backup of vFiler units	48
NDMP support	48
Available NDMP options	49
Support for the ndmpcopy command	49
NDMP command support	49
NDMP password support	49
LUNs on a vFiler unit	49
Guidelines for managing iSCSI LUNs and igrroups on a vFiler unit	50
The iSCSI service on a vFiler unit	51
LUN and igrup limitations on vFiler units	52
Networking guidelines	52
The routed daemon on vFiler units	52
Command for changing the routing table in the default IPspace	52
The /etc/dgateways file	53
SnapMirror technology on the hosting storage system	53
Guidelines for using SnapMirror technology	53
Determining the status of SnapMirror relationships	54
Deduplication with vFiler units	54
Running deduplication commands on a vFiler unit	55
Data compression on vFiler units	56
Running data compression commands on a vFiler unit	56
How vFiler units work with FlexClone files and FlexClone LUNs	57
Backup of vFiler units using SnapVault	57
Where to enter SnapVault commands	58
Determining the status of SnapVault relationships	58
SNMP support on vFiler units	58
vFiler unit data from MIBs	58
Monitoring performance and statistics	59
Viewing storage system statistics	59
Viewing uptime statistics	59
Viewing NFS statistics	59
Viewing CIFS statistics	60
What an IPspace is	61

Guidelines for vFiler unit participation in an IPspace	61
IPspace application scenario	61
Interface participation in an IPspace	63
Routing in an IPspace	63
Advantages of using VLAN tagging for IPspaces	64
HA pair and IPspaces	64
IPspace naming requirement	65
IPspace assignment requirement	65
Asymmetric HA pair setup	65
Specifying partners in an asymmetric HA pair setup	65
Creating an IPspace	66
IPspace and the routed daemon	67
Listing IPspaces on a storage system	67
Removing an IP address from an interface	67
Assigning an interface to an IPspace	68
Destroying IPspaces	69
Creating a vFiler unit in a nondefault IPspace	69
File system access using NFS and CIFS	72
Path name specification for NFS exports or CIFS shares	72
vFiler unit preparation for NFS	72
Starting the NFS protocol	73
Exporting all file systems in /etc/exports	73
vFiler unit preparation for CIFS	74
Commands run from the hosting storage system	74
Local user accounts for vFiler units	75
Virus protection for CIFS	75
Virus scanner registration	75
Virus scanning on vFiler units	76
Effect of virus scanner availability on CIFS access	76
Configuring virus scanning for a vFiler unit	76
Disaster recovery using MultiStore	77
Checking and preparing the storage system	77
Storage checklist	79
Checking the network	80
Network checklist	83
Secure communication for disaster recovery	84

Creating a disaster recovery vFiler unit	84
Disaster recovery with SnapMirror IP address based verification and IPv6 addresses	86
Deleting the disaster recovery vFiler unit	87
The vfiler dr configure command	87
Activating the disaster recovery vFiler unit	88
What activating the disaster recovery vFiler unit does	90
Resynchronizing the vFiler unit	90
Handling resynchronization failures	93
Reactivating the original vFiler unit by using SnapMirror commands	94
Reactivating the original vFiler unit by using vfiler dr commands	96
Re-creating the vFiler unit on a replacement storage system	98
Data migration using MultiStore	100
Secure communication for data migration	100
How migrating a vFiler unit affects clients	101
Offline migration of vFiler units	101
The vfiler migrate commands	101
Migrating a vFiler unit by copying data	102
Adjusting client and network configurations if migrating to a different subnet	104
vFiler unit migration without copying data	105
Prerequisites for vFiler unit migration between the nodes of an HA pair ...	105
Guidelines for setting up volumes to support vFiler unit migration in an HA pair	106
vFiler unit migration in an HA pair	106
Migrating a vFiler unit by using the vfiler migrate -m nocopy command ..	107
What IBM N series Data Motion for vFiler is	107
Features supported by Data Motion for vFiler	108
Considerations for online migration of vFiler units	108
Option required for online migration	109
Stages of a vFiler unit migration	109
How to perform online migration of vFiler units	111
Viewing the status of a vFiler unit migration	111
Commands not allowed during the cutover phase of online migration	112
Target portal group management for online migration of vFiler units	114

Data migration implications for IP-based target portal group management	115
Enabling IP-based target portal group management	116
Creating IP-based target portal groups	118
Adding IP addresses to IP-based target portal groups	119
Removing IP addresses from IP-based target portal groups	119
Destroying IP-based target portal groups	120
Displaying IP-based target portal group information	120
Disk space management using quotas	122
Allowing or disallowing quotas for a volume	122
Quota specification management	123
Turning on or turning off quotas from a vFiler unit	123
When quota thresholds and soft quotas are exceeded	124
How you can resize quotas	125
How the quotas file works	125
Displaying the quota status	125
Displaying a quota report	126
Copyright information	127
Trademark information	128
Index	133

Preface

About this guide

This document applies to IBM N series systems running Data ONTAP, including systems with gateway functionality. If the term *7-Mode* is used in the document, it refers to Data ONTAP operating in 7-Mode, which has the same features and functionality found in the prior Data ONTAP 7.1, 7.2, and 7.3 release families.

Note: In this document, the term *gateway* describes IBM N series storage systems that have been ordered with gateway functionality. Gateways support various types of storage, and they are used with third-party disk storage systems—for example, disk storage systems from IBM, HP®, Hitachi Data Systems®, and EMC®. In this case, disk storage for customer data and the RAID controller functionality is provided by the back-end disk storage system. A gateway might also be used with disk storage expansion units specifically designed for the IBM N series models.

The term *filer* describes IBM N series storage systems that either contain internal disk storage or attach to disk storage expansion units specifically designed for the IBM N series storage systems. Filer storage systems do not support using third-party disk storage systems.

Supported features

IBM System Storage N series storage systems are driven by NetApp Data ONTAP software. Some features described in the product software documentation are neither offered nor supported by IBM. Please contact your local IBM representative or reseller for further details.

Information about supported features can also be found on the N series support website (accessed and navigated as described in *Websites* on page 9).

Websites

IBM maintains pages on the World Wide Web where you can get the latest technical information and download device drivers and updates. The following web pages provide N series information:

- A listing of currently available N series products and features can be found at the following web page:
www.ibm.com/storage/nas/
- The IBM System Storage N series support website requires users to register in order to obtain access to N series support content on the web. To understand how the N series support web

content is organized and navigated, and to access the N series support website, refer to the following publicly accessible web page:

www.ibm.com/storage/support/nseries/

This web page also provides links to AutoSupport information as well as other important N series product resources.

- IBM System Storage N series products attach to a variety of servers and operating systems. To determine the latest supported attachments, go to the IBM N series interoperability matrix at the following web page:
www.ibm.com/systems/storage/network/interophome.html
- For the latest N series hardware product documentation, including planning, installation and setup, and hardware monitoring, service and diagnostics, see the IBM N series Information Center at the following web page:
publib.boulder.ibm.com/infocenter/nasinfo/nseries/index.jsp

Getting information, help, and service

If you need help, service, or technical assistance or just want more information about IBM products, you will find a wide variety of sources available from IBM to assist you. This section contains information about where to go for additional information about IBM and IBM products, what to do if you experience a problem with your IBM N series product, and whom to call for service, if it is necessary.

Before you call

Before you call, make sure you have taken these steps to try to solve the problem yourself:

- Check all cables to make sure they are connected.
- Check the power switches to make sure the system is turned on.
- Use the troubleshooting information in your system documentation and use the diagnostic tools that come with your system.
- Refer to the N series support website (accessed and navigated as described in [Websites](#) on page 9) for information on known problems and limitations.

Using the documentation

The latest versions of N series software documentation, including Data ONTAP and other software products, are available on the N series support website (accessed and navigated as described in [Websites](#) on page 9).

Current N series hardware product documentation is shipped with your hardware product in printed documents or as PDF files on a documentation CD. For the latest N series hardware product documentation PDFs, go to the N series support website.

Hardware documentation, including planning, installation and setup, and hardware monitoring, service, and diagnostics, is also provided in an IBM N series Information Center at the following web page:

publib.boulder.ibm.com/infocenter/nasinfo/nseries/index.jsp

Hardware service and support

You can receive hardware service through IBM Integrated Technology Services. Visit the following web page for support telephone numbers:

www.ibm.com/planetwide/

Firmware updates

IBM N series product firmware is embedded in Data ONTAP. As with all devices, ensure that you run the latest level of firmware. Any firmware updates are posted to the N series support website (accessed and navigated as described in [Websites](#) on page 9).

Note: If you do not see new firmware updates on the N series support website, you are running the latest level of firmware.

Verify that the latest level of firmware is installed on your machine before contacting IBM for technical support.

How to send your comments

Your feedback helps us to provide the most accurate and high-quality information. If you have comments or suggestions for improving this document, please send them by email to starpubs@us.ibm.com.

Be sure to include the following:

- Exact publication title
- Publication form number (for example, GC26-1234-02)
- Page, table, or illustration numbers
- A detailed description of any information that should be changed

What MultiStore is

MultiStore is a feature for Data ONTAP software that enables you to partition the resources of a single storage system so that it appears as multiple "vFiler unit" storage systems on your network.

Each storage system created as a result of the partitioning is called a *vFiler unit*. A vFiler unit, using the resources assigned, delivers file access and block access services to its clients the same way that a storage system does.

The storage system on which you create vFiler units is called the *hosting storage system*. The storage and network resources used by the vFiler units exist on the hosting storage system.

The storage resource assigned to a vFiler unit can be one or more qtrees or volumes. The network resource assigned can be one or more base IP addresses or IP aliases associated with network interfaces. You can add storage and network resources to a vFiler unit at any time. You can also remove these resources from a vFiler unit at any time.

You can use IPv4 and IPv6 addresses as network resources that can be assigned to the vFiler units. To use IPv6 addresses on the vFiler units, you must enable the IPv6 protocol on the hosting storage system. You must not disable the IPv6 protocol on a hosting storage system that has vFiler units. If you do so, you see a warning message similar to the following: `vfilers are configured with IPv6 addresses. This option cannot be disabled.`

For more information about enabling or disabling IPv6 on the hosting storage system, see the *Data ONTAP Network Management Guide for 7-Mode*.

Benefits of using MultiStore

You can use MultiStore features such as virtualization, consolidation and management of storage requirements, security, disaster recovery, and data migration for your provisioning needs.

- **Virtualization**
MultiStore enables you to manage tasks such as storage administration, provisioning, and management.
- **Consolidation and ease of management**
Application service providers can consolidate the storage requirements of their customers. You can reduce management costs while offering independent, domain-specific storage management.
- **Security**
Security is one of the key concerns when storage is consolidated either within an organization or by an application service provider. Using vFiler units enables you to have different security domains within the same storage system.
- **Delegation of management**

Administrators of vFiler unit can manage all vFiler units that they are authorized to access. However, vFiler unit administrators have access rights different from those of storage system administrators.

- Disaster recovery and data migration

MultiStore enables you to migrate or back up data from one storage system to another without extensive reconfiguration on the destination storage system.

Related concepts


[MultiStore for disaster recovery and data migration](#) on page 14

Multistore for consolidating servers

If you manage multiple application or operating system servers, you can store all of the data on one storage system for easier administration. You can consolidate the servers by partitioning the storage system into vFiler units and then copying the data from the servers to the vFiler units.

If the vFiler units are for CIFS users, you can set up the vFiler units to use the same computer names as the servers. This enables CIFS clients to share resources without having to remap their drives or search for the new server in Network Neighborhood. If the vFiler units are for NFS users, the NFS clients might need to remount the file systems, or they can use the automounter to automatically mount the file systems from the new locations.

The following illustration shows how multiple servers can be consolidated and can share the same resources:



A storage system that does not have MultiStore enabled can participate in only one security domain. Therefore, if your environment requires that different groups of CIFS users be in different domains,

you must use multiple storage systems. MultiStore enables you to install each vFiler unit in the appropriate domain while keeping all of the data on the same physical storage system.

Because you can set up NIS and DNS servers for individual vFiler units, after you consolidate the servers on one storage system, network clients of the vFiler units can continue to use the same NIS and DNS servers as before.

MultiStore for service providers and enterprises

Service providers, such as ISPs and SSPs, can partition the resources of a storage system to create many vFiler units for client companies. Similarly, the information technology (IT) department of an enterprise can create vFiler units for various organizations, or customers, within the enterprise.

The administrator for each customer can manage and view files only on the assigned vFiler unit, not on other vFiler units that reside on the same storage system. In addition, there is no data flow between vFiler units. A customer using a vFiler unit is assured that no sensitive information is exposed to other customers that store data on the same storage system.

For example, an SSP can create the following vFiler units on a storage system:

- A vFiler unit named vFilerA
It uses the `/vol/vol1` volume and the `e0` interface on the storage system.
It is leased to CompanyA.
- A vFiler unit named vFilerB
It uses the `/vol/vol2` volume and the `e1` interface on the storage system.
It is leased to CompanyB.

Although both CompanyA and CompanyB store data on the same storage system, network traffic for each company is restricted to the specified interface. The administrator at CompanyA (that uses NFS to access data) cannot use the `showmount` command on a UNIX client to view directories on the storage system that are outside the `/vol/vol1` volume. Similarly, the administrator at CompanyB (that uses CIFS to access data) cannot browse any shared directories that are outside the `/vol/vol2` volume.

MultiStore for disaster recovery and data migration

MultiStore enables easier migration and mirroring of data because all the information about users, CIFS shares, NFS exports, LUNs, igroups, ACLs, and so on, is encapsulated in the vFiler unit. If a disaster occurs, you can activate this vFiler unit on the destination storage systems with minimum reconfiguration.

You do not have to edit the files' ACLs, local user group definitions, user mapping information, and so on, before users can access the data.

Note: The static routing information is not carried to the destination storage system.

NFS and iSCSI users experience minimum disruption in service when the vFiler unit on the destination storage system starts serving data instead of the vFiler unit on the source storage system.

The default vFiler unit

When you enable MultiStore, Data ONTAP automatically creates a default vFiler unit on the hosting storage system that is named vfiler0. The vfiler0 unit owns all the resources of the storage system.

When you create vFiler units and assign resources to them, the resources are assigned from vfiler0. Therefore, vfiler0 owns all resources that are not owned by non-default vFiler units.

The default vFiler unit exists as long as MultiStore is enabled. On a storage system with MultiStore enabled, you cannot rename or destroy vfiler0.

All information provided about the vFiler units is applicable to vfiler0, unless noted otherwise.

Number of vFiler units allowed

There are limits on the number of vFiler units allowed in a storage system that has MultiStore enabled. You can have a maximum of 65 vFiler units on a storage system.

You can create 64 vFiler units on a storage system. The 65th vFiler unit is vfiler0, which is created automatically when MultiStore is enabled on the storage system. The default vFiler unit exists as long as MultiStore is enabled.

In a High Availability (HA) pair, you can create up to 64 vFiler units on each node of the HA pair, for a maximum of 130 vFiler units in the HA pair.

Note:

These limits can be exceeded only during a takeover scenario, when one storage system takes over the resources of a vFiler unit in another storage system.

You can create a maximum of 16 vFiler units in N3400 systems.

Related tasks

[Viewing the current limit on the number of vFiler units](#) on page 30

Related references

[Maximum vFiler units allowed](#) on page 30

Data access from the hosting storage system

As the hosting storage system administrator, you can access all the data contained in a vFiler unit by using the `vfiler` context or the `vfiler run` commands. However, after you assign a qtree or volume to a vFiler unit, you no longer have access to the data in that qtree or volume.

For example, if you create a vFiler unit with the `/vol/vol1` volume, you can configure the `/etc(exports` file to mount the `/vol/vol1` volume. However, after you create the vFiler unit, an attempt to mount the `/vol/vol1` volume from the hosting storage system results in the following error message:

```
".../vol/vol1 belongs to vFiler unit A, cannot mount from vfiler0."
```

Related concepts

[*vFiler unit storage management from the hosting storage system*](#) on page 27

Related tasks

[*Executing commands from the hosting storage system*](#) on page 42

Hosting storage system tasks

You can perform tasks related to managing the resources on the hosting storage system in the same way that you perform them on a storage system without MultiStore.

You can use either the command line or System Manager to perform the following tasks:

- Managing volumes, disks, and RAID groups
- Increasing data availability through Snapshot management, SnapMirror management, and volume copy
- Backing up and recover data

These tasks are covered in detail in the *Data ONTAP SAN Administration Guide for 7-Mode*, the *Data ONTAP Storage Management Guide for 7-Mode*, and the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode*.

MultiStore management

You can manage MultiStore from the hosting storage system by using the command line or System Manager. You can perform tasks such as creating, starting or stopping, and destroying vFiler units. You can also manage resources and monitor the status of vFiler units.

Types of vFiler unit administrators

Two types of administrators can manage the vFiler units on a storage system: a vfiler0 administrator and a vFiler administrator.

A vfiler0 administrator can manage the hosting storage system and the other vFiler units on the storage system. A vfiler0 administrator has all the privileges to manage the vFiler unit and the storage system. A vfiler0 administrator can switch to any vFiler unit's context and execute the commands from that vFiler unit. However, when the vfiler0 administrator switches the context to a vFiler unit, the privileges are the same as those of a vFiler administrator.

A vFiler administrator can manage a specific non-default vFiler unit on a storage system. A vFiler administrator does not have all the privileges of a vfiler0 administrator. The vFiler administrator role and the credentials are specific to a non-default vFiler unit. After logging in to a vFiler unit, a vFiler administrator can identify the list of commands to execute on a vFiler unit.

Related tasks

[Logging in to a vFiler unit](#) on page 40

[Identifying commands that you can execute from a vFiler unit](#) on page 39

Enabling MultiStore

To begin using MultiStore, you must enable MultiStore. After you enable MultiStore, you can create vFiler units on the storage system.

About this task

Depending on the hardware platforms, you can enable MultiStore either by adding the license key or by turning the option on.

You must enable the MultiStore license for the following hardware platforms:

- N5300 and N5600
- N6000 series
- N7000 series

For all other supported hardware platforms, you must enable the MultiStore option.

Step

1. Depending on your hardware platform, enter the appropriate command to enable MultiStore:

If you want to...	Enter the following command...
Enable the MultiStore option	<code>options licensed_feature.multistore.enable on</code>
Enable the MultiStore license	<code>license add license_key</code>

Result

Enabling MultiStore has the following effects on the storage system:

- You can use the `vfiler` and `ipspace` commands.
- Data ONTAP starts logging the status of vFiler units and sends the information to technical support by using the AutoSupport feature.
- The routed daemon is enabled, but only in `vfiler0`.
- The `ip.match_any_ifaddr` option is set to `off`.

Note: You can turn the `ip.match_any_ifaddr` option to `on` if this option was enabled before enabling MultiStore.

- The vFiler unit limit (the number of vFiler units you can create on this storage system, including `vfiler0`) is set to a default value between three and 11, depending on the memory capacity of the hosting storage system.

Related tasks

[Viewing the current limit on the number of vFiler units](#) on page 30

Related references

[Maximum vFiler units allowed](#) on page 30

Disabling MultiStore

If you are not using vFiler units, you can disable MultiStore.

Before you begin

- You must have destroyed all the vFiler units on the storage system.
- You must have destroyed all the IPspaces.

About this task

Depending on the hardware platforms, you can disable MultiStore either by deleting the license key or by turning the option off.

You must disable the MultiStore license for the following hardware platforms:

- N5300 and N5600
- N6000 series
- N7000 series

For all other supported hardware platforms, you must disable the MultiStore option.

Step

1. Depending on your hardware platform, enter the appropriate command to disable MultiStore:

If you want to...	Enter the following command...
Disable the MultiStore option	<code>options licensed_feature.multistore.enable off</code>
Disable the MultiStore license	<code>license delete multistore</code>

Result

Disabling MultiStore has the following effects:

- MultiStore becomes unavailable immediately.
- You can no longer use the `vfiler` and `ipspace` commands.

Related tasks

[Destroying a vFiler unit](#) on page 34

Prerequisites for creating vFiler units

Before you create vFiler units, you must ensure that you have created at least one unit of storage (qtree, traditional volume, or FlexVol volume). You must also ensure that there is a free IP address in the network, which is required when you create a vFiler unit.

The storage unit that contains information about configuring the vFiler unit must be writable. It must not be a read-only file system, such as the destination volume or qtree in a SnapMirror relationship.

Storage guidelines

When you assign storage units to vFiler units, the first storage unit (qtree, traditional volume or FlexVol volume) assigned to the vFiler unit should not be removed as long as the vFiler unit exists. You cannot assign aggregates to a vFiler unit.

The first storage unit assigned to a vFiler unit is called the primary storage unit. The primary storage unit contains information about configuring the vFiler unit. Although you can remove storage units from a vFiler unit at any time after the vFiler unit is created, the primary storage unit must remain for as long as the vFiler unit exists. The primary storage unit has the same security characteristics it had before it was transferred to the vFiler unit.

When you create a new vFiler unit, C\$ share-level permissions are restricted to administrators only, but file-level security is not modified. The vFiler unit administrator can set more restrictive file-level permissions.

If the qtree or volume to be used as the primary storage unit contains a `/etc` directory, the data in the directory is lost after you add the qtree or volume to a vFiler unit. Data in qtrees and volumes that are used as non-primary storage units are preserved. A volume assigned to a vFiler unit must not be the storage system's root volume. However, you can assign qtrees in the root volume to a vFiler unit. A volume assigned to a vFiler unit can be a traditional volume or a FlexVol volume.

For information about traditional volumes, FlexVol volumes, and aggregates, see the *Data ONTAP Storage Management Guide for 7-Mode*.

FlexCache volumes can be created on the default vFiler unit, `vfiler0`, but cannot be assigned or moved to any other vFiler unit.

A qtree is assigned to a vFiler unit, owned by a vFiler unit, or associated with a vFiler unit only if that qtree is added as a resource to a vFiler unit. If a volume containing a qtree is added as a resource to a vFiler unit, then the qtree implicitly becomes a resource of that vFiler unit. If the vFiler unit administrator needs to create qtrees on the vFiler unit, the administrator assigns volumes instead of qtrees to the vFiler unit when creating the vFiler unit. This is because qtrees can be created only at the root of a volume.

If you anticipate that you might have to move the disks that are used for the vFiler unit's storage from one storage system to another, you should assign volumes, instead of qtrees, to the vFiler unit.

When managing NFS exports, CIFS shares, quotas, and options, vFiler unit administrators need to enter the complete path names of the storage resources used by the vFiler units in commands and configuration files. Therefore, the storage system administrator should choose volume and qtree names appropriately so that the complete path names beginning with `filer_name:/vol/vol_name` can be shared with the vFiler unit administrators.

Naming guidelines

You must follow certain naming guidelines to create or rename a vFiler unit successfully.

- The name can contain up to 31 alphanumeric ASCII characters.
- The name can contain the dash (-) and the underscore (_) characters.

However, the name should not begin with a dash.

- The name is case-insensitive.
- The name must be unique.

You can include the name of the hosting storage system as part of the vFiler unit name so that you can easily determine the storage system that contains the vFiler unit—for example, mycompanyss1_vfiler1.

- The name vFiler0 should not be used because it is the name of the default vFiler unit.

Language guidelines

Administrators of vFiler units need to edit the `/etc/quotas` and `/etc/usermap.cfg` files for their vFiler units. These files support Unicode and root volume UNIX encoding. To ensure that vFiler unit administrators can edit these files without requiring Unicode-capable editors, you should create vFiler units on a storage system whose root volume language can be used for editing.

Quota guidelines

When you create a vFiler unit, the ownership of a volume or qtree is changed from the hosting storage system to the vFiler unit that is created. This change requires that quotas be turned off for the affected volume before you create the vFiler unit. You can turn on the quotas for the volume after the vFiler unit is created.

HA pair guidelines

When you set up vFiler units in an HA pair, both nodes must have MultiStore enabled to take over the partner. IPspaces created in one node in an HA pair must also be created in the partner system for failover to function correctly.

- You can create up to 64 vFiler units on each member of an HA pair, depending on the memory capacity of the hosting storage systems.
- The vFiler units hosted by the storage systems of the HA pair are created and configured independently.

This means each storage system can host a different number of vFiler units, and the vFiler unit configurations on the storage systems can be different from each other.

- In takeover mode, the functioning storage system takes over all the vFiler units created on the failed storage system.

These vFiler units include the vFiler units you have created and the default unit vfiler0.

Therefore, for vFiler units on the failed storage system to work correctly after takeover, each network interface used by a vFiler unit in an HA pair must have a partner interface.

Related tasks

[Viewing the current limit on the number of vFiler units](#) on page 30

SAN guidelines

When you create vFiler units in a SAN environment, note that only iSCSI is supported and FCP is not supported.

You must remember the following points when you create vFiler units in a SAN environment:

- SCSI LUNs and igroups are supported on all vFiler units managed separately for each vFiler unit.
- FC LUNs and igroups are not supported on vFiler unit.
FC LUNs and igroups are supported only on the hosting storage system.
- When you create a vFiler unit on a storage system on which iSCSI is licensed, the iSCSI service is automatically started on the vFiler unit.

Note: If iSCSI is licensed on the hosting storage system and the storage to be allocated to the vFiler unit contains LUNs, you should unmap the LUNs.

For more information about LUNs, see the *Data ONTAP SAN Administration Guide for 7-Mode*.

- Starting a vFiler unit starts iSCSI packet processing for that vFiler unit.
- Stopping a vFiler unit stops iSCSI packet processing for that vFiler unit.

The vFiler commands

The `vfiler` commands, which are supported only on the hosting storage system, enable the hosting storage system administrator to set up vFiler units and manage vFiler unit resources and Data ONTAP features on individual vFiler units. Each `vfiler` command has a different syntax.

The general `vfiler` command syntax is as follows:

`vfiler command vfilertemplate options...`

Some `vfiler` commands support the `vfilertemplate` option. `vfilertemplate` can be any of the following:

- A vFiler unit name
- A comma-separated list of vFiler unit names
- An IPspace declaration
- An asterisk (*) used as a wildcard

For example, you can run the `setup` command for the vFiler unit `vfiler1`:

`vfiler run vfiler1 setup`

If you use the asterisk, the command takes effect on all vFiler units, including `vfiler0` (the hosting storage system), unless the command cannot be applied to `vfiler0`. See the `na_vfiler(1)` man page for more information.

Some `vfiler` commands include a complete path name for the qtree or volume that is assigned to the specified vFiler unit.

Related concepts

[LUNs on a vFiler unit](#) on page 49

[File system access using NFS and CIFS](#) on page 72

[Disk space management using quotas](#) on page 122

[Virus protection for CIFS](#) on page 75

Ensuring that the network interface is ready

Before you create a vFiler unit, you must ensure that the network interface is ready.

Steps

1. If the IP address for the vFiler unit is a base IP address for an interface, enter the following command to change the state of the interface for the IP address to down:

ifconfig interface down

Example

The following command changes the state of the e0 interface to down:

ifconfig e0 down

2. If the IP address for the vFiler unit is an IP alias for an interface, enter the following command to remove the alias:

ifconfig interface -alias address

Example

The following command removes the IP alias from the e0 interface:

ifconfig e0 -alias 123.123.123.123

3. If the IP alias is currently assigned to an interface, enter the following command to remove the alias:

ifconfig interface -alias address

If the IP alias is currently unassigned, the network interface is ready.

4. If the base IP address for the vFiler unit is assigned to an interface in the `up` state, enter the following command to change the state of the interface to down:

ifconfig interface down

If the base IP address for the vFiler unit is assigned to an interface in the `down` state, the network interface is ready.

Creating a vFiler unit

You can create a vFiler unit when you want to partition the storage and network resources of a single storage system. You can create a vFiler unit by using the `vfiler create` command from the command-line interface. You can also use System Manager to create a vFiler unit.

About this task

When you create and associate a vFiler unit with the IPspace that is used by a stopped vFiler unit belonging to the same subnet, the default or static routes are migrated from the stopped vFiler unit to the newly created vFiler unit.

Steps

1. From the CLI, enter the following command:

```
vfiler create vfiler_name [-s ipspace] -i ip_address [ -i
ip_address ] ... path [ path ] ...
```

vfiler_name is the name of the vFiler unit.

ipspace is the IPspace the vFiler unit must belong to.

ip_address is an IP address of the vFiler unit.

path is the complete path name to an existing volume or qtree.

The first path name is the storage unit that contains the `/etc` directory. The `/etc` directory contains the configuration information about the vFiler unit.

Example

The following command creates a vFiler unit with two IP addresses, one volume, and one qtree:

```
vfiler create vfiler1 -i 123.123.123.123 -i 123.123.123.124 /vol/
vol1 /vol/vol2/qtree2
```

For more information about the `vfiler create` command, see the `na_vfiler(1)` man page.

2. Respond to the prompts to set up the storage system, and to set up CIFS if necessary.
3. To set up CIFS on a vFiler unit with IPv6 address, perform the following steps:
 - a) Skip the cifs setup by entering `n` at the cifs setup prompt.
 - b) Enable cifs on the vFiler unit by entering the following command:

```
vfiler run vfiler1 options cifs.ipv6.enable on
```

- c) Set up cifs on the vFiler unit by entering the following command:

```
vfiler run vfiler1 cifs setup
```

Result

The setup process does the following on the new vFiler unit:

- Starts NFS (if NFS is licensed on the hosting storage system) and configures the vFiler unit's primary storage (root volume) to be exported to the vFiler unit's administration host (using an entry in the `/etc/exports` file)
- Configures the vFiler unit's IP addresses and adds the appropriate entries to `/etc/rc`
- Creates a “pseudo-root” that allows CIFS clients to see all the storage that has been assigned to the vFiler unit as a single tree
- Starts the iSCSI service (if iSCSI is licensed on the hosting storage system)

As the vFiler unit administrator, you can now mount the root volume of the vFiler unit, edit `/etc/exports` to suit your needs, and rerun the `exportfs` command.

Related concepts

[File system access using NFS and CIFS](#) on page 72

[Naming guidelines](#) on page 20

Related tasks

[Setting up a vFiler unit](#) on page 25

Setting up a vFiler unit

You can set or modify the network configuration of a vFiler unit by using the `setup` command.

About this task

If you have used `vfiler create` command with the `-n` option to create a vFiler unit, then the Setup wizard is not prompted to configure the vFiler unit's network. In this case, you have to configure the vFiler unit's network by using the `setup` command.

The `setup` command does not prompt you for the time zone. All vFiler units are in the same time zone as the hosting storage system.

Unlike the `setup` command for the storage system, the `setup` command for a vFiler unit does not cause NFS to start running, as it runs automatically when you create a vFiler unit.

The IPv6 address options are supported for the `setup` command.

Steps

- To run setup from the default vFiler unit, enter the following command:

```
vfiler run vfiler_name setup
```

The general `setup` command syntax with IPv6 address, when used in the vFiler context, is as follows:

```
setup [-e ifname:ipv4_address| [ipv6_address]:netmask|prefixlen ...] [-d DNS_domain_name:DNS_server_ipv4_address| [DNS_server_ipv6_address] ...]
[-n NIS_domain_name:NIS_server_ipv4_address| [NIS_server_ipv6_address] ...] [-a ipv4_address| [ipv6_address]]
name:ipv4_address| [ipv6_address] [-p root_password]
```

Note: Ensure that you enter the IPv6 address in [] brackets.

The general setup command syntax, when used in the vFiler context, is as follows:

```
setup [-e ifname:ipv4_address:netmask ...] [-d DNS domain
name:DNS_server_ipv4_address...] [-n NIS_domain_name:NIS_server_ipv4_address ...] [-a ipv4_address|
name:ipv4_address| [-p root_password]]
```

The `setup` command displays prompts for you to configure the vFiler unit. After you respond to all the prompts, configuration files, such as the `/etc/exports` file, are created in the `/etc` directory for the vFiler unit.

You can use any of these options with the `setup` command:

- The `-e` option creates the bindings of the vFiler unit and provides netmask with IPv4 address and prefixlen with IPv6 address.
Netmask is of the form a.b.c.d, where a, b, c, and d should be between 0 and 255. Prefixlen is an integer between 0 and 127.
- The `-e` option creates the bindings of the vFiler unit's IP addresses.
- The `-d` option specifies a DNS domain name and the IP addresses of one or more DNS servers.
- The `-n` option specifies a NIS domain name and the IP addresses of one or more NIS servers.
- The `-a` option specifies the administrator host name and IP address.
- The `-p` option sets the password of the vFiler unit's root user.

Note: If the vFiler unit runs the CIFS protocol, go to the next step. Otherwise, `setup` is complete.

2. To configure CIFS protocol, enter the following command:

```
vfiler run vfiler_name cifs setup
```

The `cifs setup` command displays prompts for you to configure CIFS on the vFiler unit. After you respond to all the prompts, CIFS starts running.

Related concepts

[File system access using NFS and CIFS](#) on page 72

vFiler unit storage management from the hosting storage system

As the physical storage system administrator, if you need to manage storage resources that belong to a vFiler unit but you do not have administrative access to the vFiler unit, you can temporarily move the vFiler unit's resources, or temporarily destroy the vFiler unit.

Note: Before taking either of the following actions, you should unmap any LUNs that have been created in the affected storage resources. For instructions, see the *Data ONTAP SAN Administration Guide for 7-Mode*.

- Temporarily move the resources to the hosting storage system.
However, you cannot move the vFiler unit's primary /etc volume.
- Temporarily destroy the vFiler unit.
This returns ownership of all resources to the hosting storage system. No user data is modified when you destroy a vFiler unit.

Depending on which action you take, you can later move the storage resources back to the vFiler unit, or restore the vFiler unit.

Effects of adding, removing, and moving vFiler unit resources

Adding, removing, or moving vFiler unit resources affect only the association between the vFiler unit and the resources. It does not have any effect on user data in the vFiler unit.

- After you add storage resources to a vFiler unit, the resources are moved from the hosting storage system to a vFiler unit.
- After you remove storage resources from a vFiler unit, the resources are removed from the vFiler unit to the hosting storage system.
- After you add an IP address to a vFiler unit, you can assign the address as an IP alias of an interface or assign the address to a network interface that has not been configured.
- After you remove an IP address from a vFiler unit, the IP address becomes an unassigned IP address.
- After you move resources from one vFiler unit to another, the resources are moved from the resource vFiler unit and added to the destination vFiler unit.

Adding resources to a vFiler unit

To partition the storage and network resources of a single storage system so that it appears as multiple storage systems on the network, you must add resources to a vFiler unit.

Step

1. Enter the following command:

```
vfiler add vfiler_name [ -f ] [ -i ip_address [ -i ip_address ] ... ]
[ path [ path ... ] ]
```

You can use the `-f` option to skip warning messages.

Example

The following command adds an IP address and a volume to an existing vFiler unit:

```
vfiler add vfiler1 -i 123.123.123.125 /vol/vol3
```

Example

The following example adds the IPv4 address 69.2.2.2 and the IPv6 address 3000::B9:23 in vfiler2.

```
vfiler add vfiler2 -i 3000::B9:23 -i 69.2.2.2
```

Requirements for moving and removing resources

When you move or remove vFiler unit resources, both the source and destination vFiler units must be in the same IPspace. If the resource that is being moved or removed is a storage unit, then the storage unit must not contain the vFiler unit's `/etc` directory.

- If a storage unit is to be moved or removed, and it contains any CIFS shares, home directories, or open files and directories, you must remove the CIFS shares, remove the home directories from the list of home directories, or close open files and directories.
- If the IP address is an IP alias, the alias must be removed. If the IP address is not an IP alias, the network interface associated with the address must not be configured.

Removing resources from a vFiler unit

You can remove resources you have added to the vFiler unit. For example, you might have to remove resources when an SSP wants to reduce the amount of storage used, to reduce operating costs.

Before you begin

- If a storage unit you want to remove contains LUNs, you must have unmapped the LUNs. For more information about LUNs, see the *Data ONTAP SAN Administration Guide for 7-Mode*.
- You must have stopped the FlexClone file and FlexClone LUN operations running on the storage unit from the non-default vFiler context.

Step

1. To remove the resources from a vFiler unit, enter the following command:

```
vfiler remove vfiler_name [ -f ] [ -i ip_address [ -i ip_address ] ... ]
[ path [ path ... ] ]
```

You can use the `-f` option to skip warning messages.

Example

The following command removes an IP address and a volume from an existing vFiler unit:

```
vfiler remove vfiler1 -i 123.123.123.125 /vol/vol3
```

Moving resources between vFiler units

You can move storage resources such as volumes and qtrees from one vFiler unit to another. When you move the storage resources to another vFiler unit, the ownership of the volume, qtree, and the data changes.

Before you begin

- If a storage unit you want to move from one vFiler unit to another contains LUNs, you must have unmapped the LUNs.
For more information about unmapping the LUNs, see the *Data ONTAP SAN Administration Guide for 7-Mode*.
- You must have stopped the FlexClone file and FlexClone LUN operations on the storage unit from the nondefault vFiler context.
For more information about FlexClone file and FlexClone LUN operations, see the *Data ONTAP Storage Management Guide for 7-Mode*.
- You must have disabled the SnapVault technology on the source vFiler unit.
For more information about the SnapVault technology, see the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode*.

About this task

- After you move a storage unit from one vFiler unit to another, the security information associated with the files in the storage unit is retained.
As a result, users might be unable to access files properly.
- If you reassign a volume from one vFiler unit to another, Data ONTAP turns off quotas for the volume.
After the volume is moved, you can turn quotas on again for the volume from the destination vFiler unit.
- If you reassign a qtree from one vFiler unit to another, Data ONTAP turns off quotas for the volume containing the qtree on both the source vFiler unit and the destination vFiler unit.
After the qtree is moved, you can turn on the quotas again for the volume.
- When resources are being moved, all the network connections to those resources are terminated.

Step

1. To move the resources between vFiler units, enter the following command:

```
vfiler move source_vfiler destination_vfiler [ -f ] [ -i ip_address [ -i ip_address ] ... ] [ path [ path ... ] ]
```

You can use the `-f` option to skip warning messages.

Example

The following command moves an IP address and a volume from one vFiler unit to another:

```
vfiler move vfiler1 vfiler2 -i 123.123.123.125 /vol/vol3
```

Default limits on the number of vFiler units

By default, there are limits on the number of vFiler units in storage systems that have MultiStore enabled.

The default limit on a storage system with the memory capacity of 4 GB or more is 11 vFiler units.

Note: The default limit on N3400 systems is five vFiler units.

This limit includes vfiler0. Therefore, a limit of 11 vFiler units means that you can create a maximum of 10 vFiler units on each node of an HA pair.

Maximum vFiler units allowed

There are limits on the number of vFiler units that you can create, depending on the available memory of the storage system.

All the supported systems in Data ONTAP 8.1 can have a maximum of 65 vFiler units.

Note: You can create a maximum of 16 vFiler units in N3400 systems.

You can use the `sysconfig -v` command to verify the memory size of your storage system.

Viewing the current limit on the number of vFiler units

To determine whether you want to increase or decrease the current limit on the number of vFiler units that you can have on a hosting storage system, you can view the current limit.

About this task

The vFiler unit limit specifies the maximum number of vFiler units that can exist on the hosting storage system. Because the limit includes the hosting storage system, vfiler0 (which always exists if MultiStore is enabled), the number of vFiler units you can create on a storage system is one less than the vFiler unit limit set on a storage system.

In an HA pair during a takeover, the limit on the number of vFiler units on each storage system is applicable within the context of that storage system.

Step

1. To view the current limit on the number of vFiler units, enter the following command:

```
vfiler limit
```

Example

The following command displays the current limit of the vFiler units:

```
Storage system> vfiler limit
Current limit: 11
```

Current in use: 1
Platform hard limit: 65

Increasing the vFiler unit limit

If you need more partitions on your hosting storage system, you can increase the vFiler unit limit. The maximum number of vFiler units you can have on a storage system depends on the memory capacity of the hosting storage system.

Step

1. To increase the vFiler unit limit, enter the following command:

```
vfiler limit number
```

Example

To increase the number of vFiler units that you can create to 15, enter the following command on the hosting storage system:

```
vfiler limit 16
```

The limit is set to a number that must be one more than the number you create, because one vFiler unit is created automatically when you enable MultiStore.

Result

In an HA pair, the `vfiler limit` command sets the vFiler units limit on each of the nodes of the HA pair.

After you finish

For the change to take effect, you must reboot the storage system (or each storage system in an HA pair).

Decreasing the vFiler unit limit

You can decrease the vFiler unit limit when you need fewer partitions on your hosting storage system.

About this task

When you decrease the limit, the change is effective immediately and does not require a reboot of the storage system.

Step

1. To decrease the vFiler unit limit, enter the following command:

```
vfiler limit number
```

Example

To reduce the number of vFiler units that you can create from 15 to 10, enter the following command on the hosting storage system:

```
vfiler limit 11
```

The limit is set to a number that must be one more than the number you create, because one vFiler unit is created automatically when you enable MultiStore.

What the `vfiler rename` command does

The `vfiler rename` command renames the vFiler unit. You can rename the vFiler unit when you want the vFiler unit to have a unique name. The command changes the name of the vFiler unit only within Data ONTAP.

The `vfiler rename` command does not re-broadcast the new name to the CIFS domain controllers or the NetBIOS name servers because these protocols might be using a different name for the vFiler unit from the name that Data ONTAP uses. To change the name mapping in the CIFS domain controllers, you should run `CIFS setup` for each of these protocols.

Renaming a vFiler unit

You can use the `vfiler rename` command to rename the vFiler unit, for example, when you want the vFiler unit to have a unique name.

About this task

You should not rename a vFiler unit while it is being migrated. If you rename a vFiler unit that is being migrated, the `migrate` command on the remote system fails.

The new name for the vFiler unit should not exist on the storage system or on the partner storage system in an HA pair.

Although Data ONTAP allows the storage system and its partner to have vFiler units with identical names, it is easier to administer the storage systems if each vFiler unit has a unique name.

Step

1. Enter the following command to rename the vFiler unit:

```
vfiler rename old_vfiler_name new_vfiler_name
```

Example

The following command renames the vFiler unit `vfiler1` as `vfiler2`:

```
vfiler rename vfiler1 vfiler2
```

Stopping a vFiler unit

You can stop a vFiler unit if you need to troubleshoot or destroy a vFiler unit.

About this task

After you stop a vFiler unit, the vFiler unit can no longer receive packets from clients.

Note: You cannot stop vfiler0.

The stopped state is not persistent across reboots. When you reboot the storage system, the vFiler unit that was stopped before the reboot operation resumes automatically.

Note: If the vFiler unit has an active interactive SSH session, you must terminate the session before stopping the vFiler unit.

When you stop a vFiler unit that has static or default routes associated with it, then these routes are migrated to a running vFiler unit in the same IPspace and belonging to the same subnet.

Step

1. Enter the following command to stop the vFiler unit:

```
vfiler stop vfilertemplate
```

vfilertemplate is the name of the vFiler unit that you want to stop.

Example

Assume the storage system supports two vFiler units: vfiler1 and vfiler2. The following command stops all vFiler units, except vfiler0:

```
vfiler stop *
```

The following message appears after you enter the command:

```
vfiler stop *
vfiler1                      stopped
vfiler2                      stopped
```

Related concepts

[The vFiler commands](#) on page 22

Destroying a vFiler unit

If you want to return storage resources back to the hosting storage system (and the storage administrator's domain), you should destroy the vFiler unit that owns the storage resources.

Before you begin

- You must have unmapped any LUNs that are mapped to the vFiler unit's storage. For information about how to unmap any LUNs, see the *Data ONTAP SAN Administration Guide for 7-Mode*.
- You must have stopped the vFiler unit by using the `vfiler stop` command.
- You must have stopped the FlexClone file and FlexClone LUN operations running on the storage unit from the nondefault vFiler context.

Step

1. To destroy the vFiler unit, enter the following command:

```
vfiler destroy [ -f ] vfiler_name
```

The `-f` (force) option in the `vfiler destroy` command destroys the vFiler unit immediately. Without the `-f` option, the command displays a confirmation prompt.

Attention: When there are multiple vFiler units in an IPspace, you must ensure that no routes used by other vFiler units are associated with the vFiler unit that you want to destroy. Otherwise, destroying this vFiler unit will render the other vFiler units in the IPspace inaccessible.

Result

Destroying a vFiler unit has the following effects:

- Clients using LUNs experience an interruption in service.
- All resources associated with the vFiler unit are released to the hosting storage system.
- There is no loss of data as the data can be accessed from the hosting storage system.
- All the vFiler unit's IP addresses are not configured and the corresponding entries in the storage system's `/etc/rc` file are removed.
- If the vFiler unit that is destroyed was not in the same IPspace as the hosting storage system, the IP addresses previously owned by the vFiler unit are not available for use after you destroy the vFiler unit.
- The effects on the quotas specified in the destroyed vFiler unit are the same as the effects on the quotas when you move resources from a vFiler unit to the hosting storage system.

Related tasks

[Moving resources between vFiler units](#) on page 29

Restoring a vFiler unit

You can restore a vFiler unit by using the `vfiler create -r` command. You restore a vFiler unit in the following scenarios: during disaster recovery, data migration (if the migration fails and the vFiler unit on the source storage system is destroyed), or if a vFiler unit is accidentally deleted.

Before you begin

The `/etc` configuration directory must exist in the source or destination storage system. If the `/etc` configuration directory does not exist in the source or destination storage system, you can copy the `/etc` configuration directory from the destination or source storage system by using the SnapMirror feature, `ndmpcopy`, or from tape.

Step

1. Enter the appropriate command to restore the destroyed vFiler unit:

If you want to...	Enter the following command...
Restore the destroyed vFiler unit with the old name	<code>vfiler create oldname -r original_rootvolume</code>
Restore the destroyed vFiler unit with a new name	<code>vfiler create newname -r original_rootvolume -b oldname</code>

oldname is the name of the original vFiler unit.

original_rootvolume is the path of the root volume that was specified when creating the vFiler unit.

newname is the name of the new vFiler unit.

For more information about the `vfiler create` command, see the `na_vfiler(1)` man page.

Example

Example

The following command restores `vfiler2` that was destroyed:

```
vfiler create vfiler2 -r /vol/vol1
```

The following command restores `vfiler2` that was destroyed, with a new name `vfiler3`:

```
vfiler create vfiler3 -r /vol/vol1 -b vfiler2
```

Related tasks

[Setting up a vFiler unit](#) on page 25

Starting a vFiler unit

You can start a vFiler unit that is in the stopped state by using the `vfiler start` command.

About this task

After a vFiler unit starts, it is in a running state and can receive packets from clients. For example, if iSCSI is licensed on the storage system, starting or stopping a vFiler unit starts or stops iSCSI packet processing for that vFiler unit.

Note: You cannot start `vfiler0`.

When you start a vFiler unit that is associated with the IPspace that is used by a stopped vFiler unit belonging to the same subnet, the default or static routes are migrated from the stopped vFiler unit to the started vFiler unit.

Step

1. Enter the following command to start a vFiler unit:

```
vfiler start vfilertemplate
```

Example

Assume the storage system supports two vFiler units: `vfiler1` and `vfiler2`. The following command starts all vFiler units, except `vfiler0`:

```
vfiler start *
```

The following message appears after you enter the `vfiler start` command:

```
vfiler start *
The default vfiler cannot be stopped or started.
Vfiler vfiler1 is already running.
Vfiler vfiler2 is already running.
vfiler0 running
vfiler1 running
vfiler2 running
```

Related concepts

[The vFiler commands](#) on page 22

Protocols supported by a vFiler unit

All protocols that are supported by the hosting storage system are supported by vFiler units. However, you can select the protocols that you want to allow on each of the vFiler units by using the `vfiler allow` command.

The following protocols are supported by vFiler units:

- CIFS
- NFS
- RSH
- SSH
- iSCSI
- FTP
- HTTP

The maximum number of FTP connections to a vFiler unit is determined by the `ftpd.max_connections` option set on the hosting storage system. The value set for this option is shared among the vFiler units on a storage system.

Allowing a protocol on a vFiler unit

You can select the protocols that you want to allow on the vFiler units by using the `vfiler allow` command.

Before you begin

To allow CIFS, NFS, FTP, or HTTP on a vFiler unit, each of the protocols must have an active license on the hosting storage system.

Step

1. Enter the following command to allow a protocol on a vFiler unit:

```
vfiler allow vfilertemplate proto=protocol ...
```

`protocol` can be any one of the following supported protocols: nfs, cifs, iscsi, rsh, ssh, ftp, or http.

Example

The following command allows the NFS and RSH protocols on the vFiler unit named vfiler1:

```
vfiler allow vfiler1 proto=nfs proto=rsh
```

Related concepts

[The vFiler commands](#) on page 22

Effects of disallowing protocols on a vFiler unit

When you disallow protocols on a vFiler unit, it has certain effects, depending on the protocol that is disabled. For example, if you disable HTTP or FTP, no new FTP or HTTP connections are allowed on the vFiler unit.

Effects of disallowing iSCSI: After iSCSI is disallowed on a vFiler unit, the following conditions apply on that vFiler unit:

- You cannot start iSCSI.
- No new iSCSI sessions are allowed.
- iSCSI commands on existing sessions are rejected.

Effects of disallowing FTP: After FTP is disallowed on a vFiler unit, no new FTP connections are allowed on the vFiler unit. However, transfers that started before FTP was disallowed are completed.

Effects of disallowing HTTP: After HTTP is disallowed on a vFiler unit, no new HTTP connections are allowed on the vFiler unit. Each new request receives a 503 `HTTP server is disabled` message. Any existing connections remain active.

Disallowing RSH and SSH: Although the Data ONTAP `rsh.enable` and `ssh.enable` option values (On or Off) determine whether the RSH or SSH server is enabled or disabled on a storage system, disallowing RSH or SSH on a vFiler unit is independent of the value for that option. A vFiler unit can be configured to restrict RSH or SSH even when the corresponding `enable` option is set to On.

Note: To allow RSH on a vFiler unit, you must have the `rsh.enable` option set to On. To allow SSH on a vFiler unit, you must have the `ssh.enable` option set to On.

Disallowing a protocol on a vFiler unit

You can disallow protocols on a vFiler unit by using the `vfiler disallow` command. For example, for security reasons you might want to disallow access to the vFiler unit by SSH or RSH.

About this task

If the CIFS, NFS, iSCSI, FTP, or HTTP protocol is running and you disallow it, the protocol continues to run until the storage system reboots. However, packets destined for the vFiler unit are ignored.

Step

1. Enter the following command to disallow a protocol for a vFiler unit:

```
vfiler disallow vfilertemplate proto=protocol ...
```

protocol can be any one of the following supported protocols: nfs, cifs, iscsi, rsh, ssh, ftp, or http.

Example

The following command disallows the NFS and RSH protocols on the vFiler unit named vfiler1:

```
vfiler disallow vfiler1 proto=nfs proto=rsh
```

Related concepts

[The vFiler commands](#) on page 22

Displaying the vFiler unit status

You can use the `vfiler status` command to check the status of a vFiler unit. For example, you might want to check if a vFiler unit is in the stopped or running state.

Step

1. Enter the following command:

```
vfiler status
```

Result

The `vfiler status` command shows the following information about the vFiler unit:

- The state of the vFiler unit, whether stopped or running.
- Whether the IP addresses that have been assigned are configured and the interfaces that they are bound to.
- Which protocols are allowed or disallowed.

See the `na_vfiler(1)` man page for more information.

Identifying commands that you can execute from a vFiler unit

You can execute only some of the Data ONTAP commands from a vFiler unit. As a vfiler0 administrator or any other vFiler unit administrator, you can identify the commands that you can execute from a vFiler unit either from the vfiler0 context or any other vFiler unit context.

Before you begin

You must have understood the consideration and restrictions of the `vfiler` commands contained in the vFiler considerations section of each command's man page.

Step

1. Based on whether you are a vfiler0 administrator or vFiler administrator, perform the appropriate action to identify the commands that you can execute from a vFiler unit:

If you are a...	Then...
vfiler0 administrator	<p>Perform one of the following actions:</p> <ul style="list-style-type: none"> • Enter the following command from the vfiler0 context: <p>vfiler run vfilertemplate ? <i>vfilertemplate</i> is the name of the vFiler unit.</p> <ul style="list-style-type: none"> • After changing the context to the vFiler unit, enter the following command from the vFiler context: <p>?</p>
Any nondefault vFiler unit's administrator	<p>Enter the following command:</p> <p>?</p>

Logging in to a vFiler unit

Starting with Data ONTAP 8.1, as a vFiler unit administrator, you can log in to a vFiler unit from an appropriate Secure Shell client application, such as PuTTY for Windows hosts or OpenSSH for UNIX hosts. You can execute commands directly on a vFiler unit through an interactive SSH session.

About this task

- You can also use IPv6 addresses to establish an interactive SSH session.
- You can establish only one interactive SSH session with a vFiler unit at a time.
- You can start an interactive session using a password or without a password.

Step

1. To log in to a vFiler unit, perform the appropriate action based on the operating system:

If your host is running...	Then...
Windows operating system	<ol style="list-style-type: none"> a. Enter the IP address of the vFiler unit in the client application. b. If prompted, enter the user name and password at the login prompt.

If your host is running...	Then...
UNIX or Linux operating system	<p>Enter the following command from the client application:</p> <pre>ssh user_name@vfiler_IP</pre> <p><i>user_name</i> is the name of the user.</p> <p><i>vfiler_IP</i> is the IP address of the vFiler unit.</p>

Result

After you log in, you can execute the vfiler commands on the vFiler unit.

Related tasks

[Identifying commands that you can execute from a vFiler unit](#) on page 39

Limits on the number of interactive SSH sessions

There are limits on the number of concurrent interactive SSH sessions that you can run on a storage system, depending on the number of vFiler units allowed on that storage system.

A vFiler unit can have only one active interactive SSH session established at a time. For all the supported platforms in Data ONTAP 8.1, there can be a maximum of 32 concurrent active interactive SSH sessions.

Note: There can be a maximum of eight concurrent active interactive SSH sessions on an N3400 system.

After you have established the maximum number of interactive SSH sessions, you cannot start a new interactive SSH session until an active interactive SSH session is closed.

For example, a storage system with 65 vFiler units can have 32 concurrent active interactive SSH sessions. You cannot establish an interactive SSH session until one of the active interactive SSH sessions is closed.

Executing commands from a vFiler unit

To work with the data on a vFiler unit, you can execute commands directly from the vFiler unit. These commands are run in the same way that they are run from the vFiler unit's console. These commands can be run only on storage resources owned by the vFiler unit.

Steps

1. Enter the following command to switch to a particular vFiler context:

```
vfiler context vfiler_name
```

Example

```
vfiler context vfiler1
```

You are now in the context of vfiler1.

2. Enter the command you want to run from that particular vFiler context.

Example

The following command shows all the commands that are available from the context of the vFiler unit:

```
?
```

3. To return to the context of the hosting storage system, enter the following command:

```
vfiler context vfiler0
```

Executing commands from the hosting storage system

The storage system administrator can execute commands for any vFiler unit from the hosting storage system. However, a vFiler unit administrator cannot execute any commands from the hosting storage system.

Step

1. Enter the following command:

```
vfiler run vfilertemplate command
```

Example

```
vfiler run vfiler1 setup
```

You can now run the `setup` command for vfiler1 from the vfiler0 context.

Related tasks

[Executing commands from a vFiler unit](#) on page 41

Executing RSH commands for a vFiler unit

You can execute commands for a vFiler unit through RSH. The RSH protocol over IPv6 is supported on vFiler units when they are configured to use IPv6 addresses. You can also execute commands from a default vFiler unit to a nondefault vFiler unit by using the `vfiler run` command.

Before you begin

- The RSH protocol must be allowed for the vFiler unit.
By default, RSH is disabled.

- To enable the RSH protocol for the vFiler unit, the `rsh.enable` option for the vFiler unit must be set to `on`.
- You must enter the command on a client of the vFiler unit that is permitted to have RSH access to the vFiler unit.
The client must be one of the hosts specified by the `rsh.access` option for the vFiler unit.

About this task

Note: You cannot launch RSH as an interactive shell or issue a vFiler command that requires user interaction through RSH.

Step

1. Enter the following command to execute commands for a vFiler unit by using RSH:

```
rsh -l user:password vfiler_IP_address command
```

`user` and `password` is your user name and password.

`vfiler_IP_address` is the IP address of the vFiler unit.

`command` is any command that you want to execute.

Example

The following command displays all options on the vFiler unit with the IP address 123.123.123.1:

```
rsh 123.123.123.1 options
```

Executing SSH commands for a vFiler unit

You can execute commands for a vFiler unit by using SSH. You can also execute commands from the default vFiler unit to a nondefault vFiler unit by using the `vfiler run` command.

Before you begin

- The SSH protocol must be enabled.
You can enable the SSH protocol by setting the `ssh.enable` option to `on`. By default, SSH is allowed on a vFiler unit.
- You must enter the command on a client of the vFiler unit that is permitted to have SSH access to the vFiler unit.
The client must be one of the hosts specified by the `ssh.access` option for the vFiler unit.

About this task

Starting with Data ONTAP 8.1, you can start SSH as an interactive shell, and issue `vfiler` commands that require user interaction through SSH.

Step

1. Enter the following command to execute commands for a vFiler unit by using SSH:

```
ssh vfiler_IP_address command
```

- If you specify *command*, you start a non-interactive SSH session.
- If you do not specify *command*, you start an interactive SSH session. You can enter *command* when the storage system prompt appears.

Example

The following command launches a non-interactive SSH session to display all options on the vFiler unit with the IP address 123.123.123.1:

```
ssh user_name@123.123.123.1 options
```

Example

The following command launches an interactive SSH session to vFiler unit with the IP address 123.123.123.1:

```
ssh user_name@123.123.123.1
```

Note: The SSH options that are available through a non-interactive SSH session are also available through the interactive SSH session.

List of RSH and SSH commands

When you use RSH or SSH, you can only execute some of the commands that are normally usable on a vFiler unit.

The following list shows the RSH and SSH commands that you can execute on a vFiler unit.

- ?
- cifs
- clone
- config
- df
- dns
- echo
- exportfs
- fpolicy
- fsecurity
- help
- hostname *
- igrpoup
- ipsec

- iscsi
- keymgr
- lock
- lun
- nbtstat
- ndmpcopy
- ndmpd
- nfs
- nfsstat
- nis
- options
- passwd
- qtree
- quota
- route
- sectrace
- secureadmin
- sftp
- sis
- snap
- snapmirror
- snapvault
- useradmin
- vol
- vscan
- wcc
- ypcat
- ypgroup
- ypmatch
- ypwhich

Note: The `hostname` command is only for displaying, and not changing the name of the vFiler unit.

Effects of storage system reboot on a vFiler unit

Rebooting a storage system can affect the state of a protocol or a vFiler unit located in the storage system. If you allow or disallow a protocol on a vFiler unit, the allowed or disallowed protocol persists across reboots.

If you want to allow a protocol again after a reboot, you must reenable the protocol. For example, if you disallow NFS for a vFiler unit and then reboot the storage system, NFS remains disallowed for the vFiler unit after the reboot.

When you stop a vFiler unit and then reboot the storage system, the stopped vFiler unit starts running again after the reboot.

Volumes and qtrees on a vFiler unit

A volume assigned to a vFiler unit can be a traditional volume or a FlexVol volume. You cannot assign an aggregate to a vFiler unit. You can create qtrees in a volume only if the vFiler unit owns that volume.

For information about traditional volumes, FlexVol volumes, and aggregates, see the *Data ONTAP Storage Management Guide for 7-Mode*.

Example: Creating a qtree in the default vFiler unit

Assume the /vol/vol0 volume is owned by vfiler0. You can use the following command to create a qtree in the /vol/vol0 volume:

```
qtree create /vol/vol0/qtree1
```

Example: Creating a qtree in a nondefault vFiler unit

Assume the /vol/vol1 volume is owned by vfiler1. The administrator for vfiler1 can use the following command to create a qtree in the /vol/vol1 volume:

```
rsh vfiler1 qtree create /vol/vol1/qtree2
```

Effects of taking a vFiler unit volume offline

Taking a volume that is used for vFiler unit storage offline affects protocols used by the vFiler unit and LUNs in the volume.

The effects are as follows:

- All CIFS shares and NFS exports in the volume are deactivated.
- If the volume contains the /etc directory for a vFiler unit, the vFiler unit stops running. The vFiler unit starts running again only after you bring the volume back online.

- All LUNs become unavailable.

Changes required after volumes are renamed

After you rename a volume that is used for vFiler unit storage, you should change the path names used in the vFiler unit's `/etc/exports` file accordingly. You should also verify that CIFS shares and quotas are configured properly.

Who can change volume or qtree security styles and oplock settings

You can change the security style and oplock setting for a volume or qtree only if you are the owner of that volume or qtree.

- If a vFiler unit owns a volume, you can change the security styles or oplock settings for the volume and all qtrees on the volume from the vFiler unit.

If you want to apply a security policy to a volume of the vFiler unit, then you must add the `security.conf` file in the `/etc` folder of the vFiler unit.
- If a vFiler unit owns qtrees on a volume owned by the hosting storage system, you can change the security styles or oplock settings from the vFiler unit only for the qtrees the vFiler unit owns.
- If the hosting storage system owns a volume that contains qtrees assigned to vFiler units, you can change the security styles or oplock settings from the hosting storage system only for the qtrees the hosting storage system owns.

Differences in qtree command output

The `qtree` command output changes, depending on whether you enter the command from the vFiler unit or the hosting storage system.

- If you enter the `qtree` command from the hosting storage system, the command displays information about all qtrees on the storage system, irrespective of whether the qtrees are owned by the hosting storage system or vFiler units.
- If you enter the `qtree` command from a vFiler unit, the command displays information about qtrees on that vFiler unit only.
- If you enter the `qtree` command without arguments from a vFiler unit, a qtree that is the destination qtree for SnapMirror is shown as `read_only` in the Status column.

Viewing all qtrees and the owner vFiler units

To view a list of qtrees grouped by the vFiler units that own the qtrees, you must run the `vfiler run` command from the hosting storage system.

Step

1. Enter the following command to view all qtrees and the owner vFiler units:

```
vfiler run * qtree status
```

Backup of vFiler units

You can back up vFiler unit data from the hosting storage system if you want to back up all vFiler units at the same time. If you want to back up individual vFiler unit's data separately, you should back up from a vFiler unit's client.

You should keep in mind the following points when you plan vFiler unit backups:

- From the hosting storage system, you can back up storage units owned by vFiler units—for example, by using the `dump` command.
You can back up all vFiler units at the same time. This method does not separate the data by vFiler unit; therefore, it is not suitable if each vFiler unit's data must be backed up separately.
- From a client of a vFiler unit, you can back up only that vFiler unit's data, but not any other vFiler unit's data.
 - A CIFS client can mount “/” from a vFiler unit and see a virtual tree comprising all of that vFiler unit's storage units.
 - A CIFS client can back up the entire data, including both CIFS and NFS data of a vFiler unit.
 - An NFS client cannot see a virtual tree for the vFiler unit.
 - An NFS client can back up all of the vFiler unit's NFS data, but not its CIFS data.

If you want to back up an individual vFiler unit's data separately, you can back up from a client (particularly a CIFS client). This backup method does not allow you to back up all vFiler units at the same time.

NDMP support

NDMP supports vFiler units. Because each vFiler unit has its own NDMP server, NDMP enables you to back up or restore each vFiler unit independently, and you can set NDMP options on each vFiler unit.

NDMP support for a vFiler unit is identical to NDMP support for a storage system, except in the following areas:

- Local tape backup and restore commands are not supported in individual vFiler units. Commands that access physical tape drive resources must be executed in the default vFiler (vfiler0) context.
- NDMP SAN management commands are not supported in the individual vFiler unit context. These commands must be executed in the default vFiler (vfiler0) context.
- VERITAS NDMP management commands are not supported in the individual vFiler unit context. These commands must be executed on the storage system.
- There is a hard limit of 160 concurrent NDMP sessions per storage system. Therefore, an NDMP server running on a vFiler unit might return an `All sessions used up` message even when there are no active sessions running on the vFiler unit.

- The vFiler unit must have a volume as the root storage resource for NDMP operations.

Available NDMP options

All the NDMP options from the `options` command are available on the default vFiler unit (`vfiler0`). They are `ndmpd.access`, `ndmpd.authtype`, `ndmpd.connectlog.enabled`, `ndmpd.enable`, `ndmpd.ignore_ctime.enabled`, `ndmpd.offset_map.enable`, `ndmpd.password.length` and `ndmpd.preferred_interface`.

These NDMP options are also available on nondefault vFiler units, except the `ndmpd.preferred_interface` option.

Support for the `ndmpcopy` command

The `ndmpcopy` command uses NDMP over external IP interfaces. Therefore, you must first ensure that you have network connectivity, name resolution, and NDMP services configured properly at the source and destination locations before attempting to use the `ndmpcopy` command.

You can use the `ndmpcopy` command to copy data from one vFiler unit to another vFiler unit, or between different locations on the same vFiler unit.

NDMP command support

You can use the `ndmpd` command to manage and monitor the NDMP service for individual vFiler units. When you enable the NDMP service on a vFiler unit, the service is enabled only for that vFiler unit and not for all vFiler units.

Also, when you use the `ndmpd` command to monitor NDMP services and sessions in an individual nondefault vFiler unit context, it displays information only about the vFiler unit you are currently monitoring.

NDMP password support

When you use the NDMP commands on the storage system, you should use the storage system's root user's password in the `ndmpcopy` command. For enhanced security, the NDMP root user for individual nondefault vFiler units has a separate user name and password.

To view a nondefault vFiler unit's root user or nonroot password on any vFiler unit, you should use the `ndmpd password` command with that user name. This command lists the NDMP user password required by the `ndmpcopy` command.

LUNs on a vFiler unit

LUNs are portions in a storage system that, when exported, look and act like local disks on the importing host. Data on a LUN can be managed at the block level (for example, by a database manager) as well as at the file level. A LUN is the basic unit of storage in a SAN.

Guidelines for managing iSCSI LUNs and igroups on a vFiler unit

Data ONTAP allows you to create and manage a separate set of iSCSI LUNs and igroups on each vFiler unit. As the vFiler unit or hosting storage system administrator, you need to be aware of certain considerations.

For a host importing LUNs, a vFiler unit looks and behaves as a storage system. Administrators of those hosts need not be aware that the LUN resides on a storage unit owned by a vFiler unit.

You must remember the following points when you manage LUNs on a storage system on which MultiStore is enabled:

- You must create and manage LUNs from the vFiler unit that owns the storage containing the LUNs.
- A vFiler unit is aware only of those LUNs that are in the storage units owned by the vFiler unit. When executed on a vFiler unit, the `lun show` command displays only that vFiler unit's LUNs.
- Ownership of LUNs changes with the ownership of the storage unit that contains the LUNs.
- LUNs must be unmapped before you can move the storage unit containing the LUNs. Therefore, you must unmap all affected LUNs before performing any of the following tasks:
 - Assigning storage that contains LUNs to a vFiler unit, either when you create the vFiler unit or later
 - Destroying a vFiler unit that owns storage containing LUNs
 - Moving storage that contains LUNs from one vFiler unit to another, or between a vFiler unit and the hosting storage system

If you try to move a storage unit without unmapping the LUNs it contains, the operation fails. For instructions on unmapping LUNs, see the *Data ONTAP SAN Administration Guide for 7-Mode*.

Note: You do not have to unmap LUNs when you migrate a vFiler unit or replace it for disaster recovery purposes.

- igroups are owned by the vFiler unit on which they are created.
- Like LUNs, a vFiler unit is aware only of those igroups that it owns. When executed on a vFiler unit, the `igroup show` command displays only that vFiler unit's igroups.
- LUNs must be mapped to igroups owned by the vFiler unit that owns the LUNs.
- Each vFiler unit has its own namespace for LUNs and igroups:
 - igroups on different vFiler units can have the same initiator.
 - LUNs on different vFiler units can have the same LUN ID.
- When you migrate a vFiler unit or replicate it for disaster recovery purposes, LUNs owned by the vFiler unit are also migrated or replicated, along with their maps, igroups, and iSCSI configuration (the node names and the state of the iSCSI service). However, iSCSI authentication is not migrated or replicated.

Related concepts

[Disaster recovery using MultiStore](#) on page 77

Related tasks

[Executing commands from a vFiler unit](#) on page 41

The iSCSI service on a vFiler unit

In general, the iSCSI service operates on individual vFiler units, treating each of the vFiler units like a physical storage system. But the iSCSI software adapter (iSCSI software target) and the commands that manage and report on it, and the underlying NICs, operate on the hosting storage system.

Because an iSCSI adapter on a storage system has only one identity (there are no vFiler unit-specific adapter names), there is only one set of iSCSI sessions and statistics.

For more information about the iSCSI service, see the *Data ONTAP SAN Administration Guide for 7-Mode*.

You must keep in mind the following considerations about the iSCSI service on the hosting storage system:

- The `iswt` command, which you use to manage the iSCSI service on the storage system's NICs, operates on the hosting storage system, not on individual vFiler units.
- The iSCSI software target driver allows the storage system to be accessed as an iSCSI target device over the storage system's standard network interfaces.
- If iSCSI is licensed on the hosting storage system, the iSCSI service is available by default during the vFiler unit setup.
- The `iscsi stats` command displays statistics by hosting storage system and iSCSI adapter.

You must keep in mind the following considerations about the iSCSI service on individual vFiler unit:

- The iSCSI protocol can be allowed or disallowed, and the iSCSI service can be started or stopped, for each vFiler unit.
- The iSCSI software adapter is online or offline for each vFiler unit, depending on whether the iSCSI service is running or stopped on that vFiler unit.
- Each vFiler unit has its own iSCSI node name, which includes the vFiler unit's UUID.
- Portal groups are defined for each vFiler unit.
- iSCSI subcommands operate specifically on each vFiler unit on which they are executed, except for the `iscsi stats` command.
- You should configure iSCSI security separately for each vFiler unit.

This includes setting the default authentication mode: none, deny, or CHAP. For more information about CHAP, see the *Data ONTAP SAN Administration Guide for 7-Mode*.

- In the case of CHAP, there is a separate list of initiators and passwords for each vFiler unit.
- You can configure iSNS separately for each vFiler unit.

You can use `iscsi isns` subcommands on each vFiler unit to do the following:

- Configuring which iSNS server to use
- Turning on or turn off iSNS registration

When created, a vFiler unit's iSNS configuration is in the "not configured" state, regardless of its state on the hosting storage system.

For more information, see the `na_iscsi(1)` man page.

LUN and igrup limitations on vFiler units

FC LUNs are supported only on the hosting storage system, not on a vFiler unit. You can create only iSCSI igrups on vFiler units. You cannot create FC igrups on vFiler units.

You must keep in mind the following limitations when you create LUNs:

- You can create FC igrups only on the hosting storage system.
- FC-connected hosts can access only those LUNs that are owned by the hosting storage system.
- The `fcp` command does not recognize vFiler units.

For detailed information about FC LUNs, see the *Data ONTAP SAN Administration Guide for 7-Mode*.

Networking guidelines

To understand how vFiler units function, you must know how vFiler units operate with routing tables and gateways.

The routed daemon on vFiler units

When you enable MultiStore, the `routed` daemon is enabled by default only on the default vFiler unit `vfiler0`.

When vFiler units are licensed and the `routed` daemon is on, the console displays the following message:

```
routed on Fri Nov 4 22:42:10 GMT
[ip.drd.vfiler.info:info]:Although vFiler units are licensed, the
routing daemon runs in the default IP space only.
```

Command for changing the routing table in the default IPspace

As vFiler units in the same IPspace share one routing table, you can change the routing table by entering the `route` command from the hosting storage system.

The `route` command has the following syntax:

```
route [-fn] add|delete [host|net] destination [gateway metric]
```

For more information about the `route` command and options, see the `na_route(1)` man page.

You can include the `route` command in the storage system `/etc/rc` file so that the routes are added automatically each time the storage system is rebooted.

Related concepts

[What an IPspace is](#) on page 61

The /etc/dgateways file

Only the hosting storage system contains the /etc/dgateways file. vFiler units do not maintain their own /etc/dgateways file.

SnapMirror technology on the hosting storage system

The SnapMirror feature for mirroring volumes and qtrees works with MultiStore after the SnapMirror feature is licensed on the source and destination storage systems. You can enter SnapMirror commands from the default vFiler unit vfiler0 or nondefault vFiler units.

SnapMirror commands entered from the default vFiler unit can be used to make changes on and display information about all the nondefault vFiler units on the hosting storage system. SnapMirror commands entered from a nondefault vFiler unit makes changes on or displays information only about that specific vFiler unit.

For backward compatibility, the default vFiler unit can operate SnapMirror commands on all volumes and qtrees, even if they are owned by other vFiler units.

If vFiler unit storage volumes and qtrees are mirrored by vfiler0, the SnapMirror relationship is reflected only on vfiler0.

Guidelines for using SnapMirror technology

You can use the `snapmirror` command on a nondefault vFiler unit in the same way as you do on a storage system. You can perform SnapMirror operations on a vFiler unit only when a vFiler unit has a volume as root storage resource. However, there are some exceptions.

The exceptions are as follows:

- Qtree SnapMirror is only supported for qtrees inside volumes owned by a vFiler unit.
- Qtree SnapMirror is only supported if a vFiler unit is rooted on a volume.
- Tape devices are not supported.
- SnapMirror sources and destinations cannot be qtrees in shared volumes.
- Synchronous SnapMirror is not supported.

For more information, see the `na_snapmirror(1)` man page.

Additionally, SnapMirror in a MultiStore context has the following features:

- The feature can be turned on and off independently on each vFiler unit.
- The `snapmirror.access`, `snapmirror.checkip.enable`, and `snapmirror.enable` options can be managed independently on each vFiler unit.
- Each vFiler unit has its own `snapmirror.conf` file in the `/etc` directory.
- A nondefault vFiler unit can only operate on the volumes and qtrees the vFiler unit owns.

- vFiler units do not require additional SnapMirror licenses and they use the same license as the storage system.
- SnapMirror relationships established between vFiler units are maintained during vFiler unit migration.
- SnapMirror destination updates are supported on both the hosting storage system and the vFiler unit.

Note: SnapMirror relationships between vFiler units and all the Snapshot copies in vFiler units are destroyed before a revert operation.

When specifying a path name in the `/etc/snapmirror.conf` file, ensure that you use the storage system name, and not the vFiler unit name. For more information, see the `na_snapmirror.conf(5)` man page.

Determining the status of SnapMirror relationships

On a vFiler unit, you can display active transfer entries related only to that vFiler unit. On the physical storage unit, you can display active transfer entries from all vFiler units. Inactive transfers are displayed only on the relevant vFiler unit.

Step

1. To display a comprehensive and readable list of SnapMirror transfers, enter the following command:

```
vfiler run * snapmirror status
```

This command cascades through all vFiler units and lists their transfers.

Deduplication with vFiler units

Deduplication enables you to save space in vFiler units by eliminating redundant data blocks within the vFiler units. You can run deduplication commands from both default and nondefault vFiler unit contexts.

Deduplication is used on redundant data only on the FlexVol volumes owned by a vFiler unit (this is called FlexVol volume granularity). Currently, deduplication is supported at the FlexVol volume level. It is not supported at the qtree level, although a vFiler unit can own volumes and qtrees.

Storage owned by a vFiler unit is not accessible by any other vFiler unit by using deduplication commands. From the default vFiler unit, you can execute deduplication commands on volumes owned by any vFiler unit in the storage system. From a nondefault vFiler unit's context, you cannot execute any deduplication command on a FlexVol volume that is not owned by the current vFiler unit.

You can configure the maximum deduplication session limit per vFiler unit.

Note: You can use the `sis.max_vfiler_active_ops` option to limit the number of active deduplication instances on a vFiler unit.

You can have a maximum of eight deduplication sessions, which is also the default limit. The minimum deduplication session limit is one. The hosting storage system allows a maximum of eight concurrent deduplication operations and they are shared among all hosted vFiler units. However, on a 32-bit platform, the maximum number of concurrent deduplication operations per storage system is five.

During a vFiler unit offline or online migration, deduplicated volumes in the vFiler unit are also migrated. The FlexVol volumes on the destination vFiler unit inherit the deduplication attributes of the source vFiler unit.

Deduplicated volumes in a vFiler unit can be recovered during disaster recovery. All FlexVol volumes on the destination vFiler unit inherit the deduplication attributes of the source vFiler unit.

Note:

- You must enable Deduplication on the source storage system.
- You need not enable Deduplication on the destination storage system.

However, if there is a situation in which the source storage system is down and the destination storage system becomes the new source storage system, you must enable deduplication to continue. The best practice is to have deduplication enabled on both locations.

For more information about deduplication, see the *Data ONTAP Storage Management Guide for 7-Mode*.

Related information

N series support website: www.ibm.com/storage/support/nseries/

Running deduplication commands on a vFiler unit

You can run deduplication commands from interfaces such as CLI and SSH.

Steps

1. Enter the following command to switch to a particular vFiler context:

```
vfiler context vfiler_name
```

Example

```
vfiler context vfiler1
```

2. Run sis commands from the particular vFiler context:

```
sis on vol_name
```

Example

```
sis on /vol/vola
```

Note: All sis commands are supported.

For more information about `sis` commands, see the *Data ONTAP Storage Management Guide for 7-Mode*.

You can run deduplication commands on a nondefault vFiler unit by using the `vfiler run` command from the `vfiler0` context. For example, `vfiler run [-q] vfilertemplate sis command`.

The output of these commands is specific to the vFiler context, which means that it shows information about all the volumes in the current vFiler unit only.

The `vfiler run` command runs on the vFiler unit or units specified in `vfilertemplate`. If more than one vFiler unit is specified, the command runs separately for each vFiler unit.

Data compression on vFiler units

Data compression enables you to reduce space required to store data in vFiler units by compressing data blocks within a FlexVol volume. It is not supported at the qtree level. To use data compression, you must enable it on the required volumes of the vFiler unit.

Starting with Data ONTAP 8.1, data compression is supported on vFiler units.

As the hosting storage system (`vfiler0`) administrator, you can execute data compression commands on volumes owned by any vFiler unit in the storage system within the respective vFiler context. As a vFiler unit administrator, you can execute data compression commands only on the volumes owned by the vFiler unit.

Compressed data in volumes are also migrated during online migration. The FlexVol volumes on the destination vFiler unit inherit the data compression attributes of the source vFiler unit. Online migration of vFiler units with compressed data requires less bandwidth, and takes less time to complete.

For more information about data compression, see the *Data ONTAP Storage Management Guide for 7-Mode*.

Running data compression commands on a vFiler unit

You can run data compression commands from interfaces such as CLI and SSH.

Steps

- From the vFiler unit context, enter the one of the following commands to enable data compression:

```
sis config -C true -I false vol_name
```

`vol_name` is the name of the volume.

When the option `C` is set to true, data compression is enabled. By default, it is disabled.

When the option `I` is set to true, inline compression is enabled. To enable inline compression, data compression must be enabled.

2. Enter the required data compression commands.

Note: All `sis` commands are supported.

For more information about `sis` commands, see the *Data ONTAP Storage Management Guide for 7-Mode*.

Example

```
vfilerA::> sis config -C true /vol/volA
```

How vFiler units work with FlexClone files and FlexClone LUNs

You can manage FlexClone files and FlexClone LUNs of parent files and parent LUNs in both the default and nondefault vFiler contexts.

If you are a default vFiler administrator, you can manage the FlexClone files and FlexClone LUNs of all the FlexClone volumes in both the default and the nondefault vFiler units. However, if you are a nondefault vFiler administrator, you can only manage the FlexClone files and FlexClone LUNs of the FlexClone volumes in that nondefault vFiler unit.

While creating FlexClone files and FlexClone LUNs of FlexClone volumes in a nondefault vFiler unit, you must ensure that the source file path and the destination file path are in the same vFiler unit.

For more information about FlexClone files and FlexClone LUNs, see the *Data ONTAP Storage Management Guide for 7-Mode*.

Backup of vFiler units using SnapVault

You can back up volumes and qtrees either from the default vFiler unit (`vfiler0`) or from nondefault vFiler units.

You must be aware of the following considerations when managing SnapVault volumes on the source or destination storage system:

- If the SnapVault volumes are owned by the default vFiler unit, you can manage them only from the default vFiler unit.
- If the SnapVault volumes are owned by a nondefault vFiler unit, you can manage them either from that nondefault vFiler unit or from the default vFiler unit.

Note: When performing SnapVault operations on a nondefault vFiler unit, you must ensure that the nondefault vFiler unit has a volume as the root storage resource. You must not have a qtree as the root storage resource. For more information, see the `na_snapvault(1)` man page.

You must license SnapVault on the source and destination storage systems before you use the SnapVault feature. Additional SnapVault licenses are not required for vFiler units.

Where to enter SnapVault commands

Depending on your storage system, you can enter SnapVault commands either from the default storage system (`vfiler0`) or from any nondefault vFiler unit.

Any command entered on the default vFiler unit makes changes on or displays information about all the vFiler units on the hosting storage system. Some commands entered on a nondefault vFiler unit make changes on or display information only about that specific vFiler unit—for example, `snapvault status`, and `snapvault snap sched`.

Note: The `snapvault.access` and `snapvault.enable` options can be managed independently on each vFiler unit without affecting the SnapVault options at the default vFiler unit and other nondefault vFiler units.

Determining the status of SnapVault relationships

On a vFiler unit, the `status` command displays active transfer entries related only to that vFiler unit. On the storage system, the `status` command displays active transfer entries from all vFiler units. Inactive transfers are displayed only on the relevant vFiler unit.

Step

1. To display a comprehensive and readable list of SnapVault transfers, enter the following command:

```
vfiler run * snapvault status
```

This command cascades through all vFiler units and lists their transfers.

SNMP support on vFiler units

SNMP is supported only on the hosting storage system and is not supported on individual vFiler units. You can enable SNMP on the hosting storage system to collect data about vFiler units.

vFiler unit data from MIBs

Data about vFiler units can be collected from the standard Management Information Base (MIB) and from the Data ONTAP custom Management Information Base.

In the standard MIB, all vFiler unit data is global. It pertains to the sum of data from all vFiler units on the storage system, with the following exceptions:

- Statistics related to network interfaces are for the interfaces in the default IPspace.
- TCP statistics include data only from the connections and listen sockets in the default vFiler unit.
- UDP statistics include data only from sockets in the default vFiler unit.
- Quota information is gathered for each volume.

If the hosting storage system or a vFiler unit owns a volume with quotas, quota information is provided for the hosting storage system or the vFiler unit owning the volume. If a vFiler unit owns qtrees in a volume that it does not own, no quota information is provided for the vFiler unit.

In the Data ONTAP custom MIB, a group named vFiler is included. It provides information about each vFiler unit, such as the MultiStore license, IP address, protocols allowed, and so on.

Monitoring performance and statistics

You can view storage system statistics, NFS statistics, and CIFS statistics to determine how well your vFiler units are performing.

Viewing storage system statistics

You can view the storage system statistics only for the sum of statistics generated by all vFiler units, including `vfiler0`. You cannot view the statistics of a particular vFiler unit.

Step

1. Enter the following command to view the storage system statistics:

`sysstat`

Viewing uptime statistics

You can view the uptime statistics only for the storage system. You cannot view the uptime statistics for specific vFiler units.

Step

1. Enter the following command to view the uptime statistics:

`uptime`

Viewing NFS statistics

You can view the NFS statistics for the entire storage system or for specified vFiler units by using the `nfsstat` command.

Step

1. Enter the `nfsstat` command required for the statistics you want to view:

If you want to view the statistics of... Enter the following command from the hosting storage system...

All vFiler units together **nfsstat**

Specified vFiler units **vfiler run vfilertemplate nfsstat**

The hosting storage system **vfiler run vfiler0 nfsstat**

Viewing CIFS statistics

You can view the CIFS statistics for the entire storage system or for specified vFiler units by using the `cifs stat` command.

Step

1. Enter the `cifs stat` command required for the statistics you want to view:

If you want to view statistics of... Enter the following command from the hosting storage system...

All vFiler units together **cifs stat**

Specified vFiler units **vfiler run vfilertemplate cifs stat**

The hosting storage system **vfiler run vfiler0 cifs stat**

What an IPspace is

An IPspace defines a distinct IP address space in which vFiler units can participate. IP addresses defined for an IPspace are applicable only within that IPspace. A distinct routing table is maintained for each IPspace. No cross-IPspace traffic routing happens.

Note: IPspaces support IPv4 and IPv6 addresses on their routing domains.

Each IPspace has a unique loopback interface assigned to it. The loopback traffic of each IPspace is completely isolated from the loopback traffic on other IPspaces.

Guidelines for vFiler unit participation in an IPspace

When you assign an IPspace to a vFiler unit, you must ensure that the vFiler unit has a unique IP address within that IPspace. After you assign an IPspace to a vFiler unit, you cannot change the IPspace without destroying the vFiler unit.

An IPspace can contain multiple vFiler units. However, a vFiler unit can belong only to one IPspace. A vFiler unit in one IPspace can have the same IP address as a vFiler unit in a different IPspace.

Each vFiler unit must have one IP address on the interface that leads to the default gateway of the assigned IPspace. This requirement ensures that the vFiler unit is reachable from within the IPspace.

When you configure a new IP address on a vFiler unit, the static or default routes belonging to a stopped vFiler unit are automatically migrated to a running vFiler unit belonging to the same subnet and in the same IPspace. Also, the new IP address is associated with the running vFiler unit.


Note: This is supported only for IPv4 addresses.

IPspace application scenario

A typical application of an IPspace is when an SSP needs to connect customers of company A and company B to a storage system on the SSP's premises.

The SSP creates two vFiler units on the physical storage system—one per customer—and provides a dedicated network path from one vFiler unit to company A's network and one from the other vFiler unit to company B's network.

This deployment should work if both companies are using non-private IP address ranges. However, the following illustration shows both companies using the same private address ranges:




Both companies use the private IP address subnet 10.0.0.0, causing the following problems:

- The two vFiler units on the storage system at the SSP location have conflicting IP addresses if both companies decide to use the same IP address for their respective vFiler units.
- Even when the two companies agree on using different IP addresses for their vFiler units, problems arise: if any client in Company A's network has the same IP address as a client in Company B's network, packets destined for a client in A's address space might get routed to a client in B's address space, and vice versa.
- Assume the two companies decide to use mutually exclusive address spaces (for example, Company A uses 10.0.0.0 with a network mask of 255.128.0.0 and Company B uses 10.128.0.0 with a network mask of 255.128.0.0). The SSP needs to configure static routes on the storage system to route traffic appropriately to A's and B's networks. This solution is neither scalable (because of static routes) nor secure (broadcast traffic is sent to all interfaces of the storage system).

To overcome these problems, two IPspaces are defined on the storage system—one per vFiler unit. Because a distinct routing table is maintained for each IPspace and no cross-IPspace traffic is routed, the data for each company is securely routed to its respective network even if the two vFiler units are

configured in the 10.0.0.0 address space, as shown in the following illustration:



Additionally, the IP addresses referred to by the various configuration files, such as the `/etc/hosts` file, the `/etc/hosts.equiv` file, and the `/etc/rc` file, are relative to that IPspace. Therefore, the IPspaces allow the SSP to configure the same IP address for the configuration and authentication data for both vFiler units, without conflict.

Interface participation in an IPspace

If MultiStore is enabled on the storage system, all its IP-addressable interfaces, including interfaces such as interface groups and VLAN, belong to the default IPspace. The default IPspace exists automatically and cannot be renamed or destroyed.

When you create a new IPspace, you assign interfaces to the new IPspace from the default IPspace. An interface can belong only to one IPspace.

Routing in an IPspace

A distinct routing table is maintained for each IPspace. All vFiler units participating in an IPspace share its routing table. The IP address of the interface and the IPspace identifier are used to identify the vFiler unit for which the packet is intended.

All packets coming in through an interface are tagged with the IPspace identifier of the IPspace to which the interface belongs.

All outgoing traffic uses the IPspace identifier of the vFiler unit that is generating the traffic to determine the routing table to use. Data ONTAP ensures that packets generated by the vFiler units of an IPspace are transmitted through the interfaces that belong to that IPspace.

Note: Broadcast packets are restricted to the vFiler units within the destination IPspace.

Advantages of using VLAN tagging for IPspaces

You can use VLAN tagging for IPspaces to provide traffic separation for customers, to set up more IPspaces, and to securely deliver packets to a vFiler unit in an IPspace.

Traffic separation

VLAN tagging for IPspaces provides traffic separation from each customer to the storage system without incurring the cost of additional network devices, such as switches.

Without VLANs, you must provide physically separate network connections to ensure that the traffic from each customer is forwarded securely to and from the storage system. This solution is neither cost-effective nor scalable.

With VLAN tagging, you can set up distinct VLANs for each customer on a single switch. Thus, VLAN tagging provides an alternative to physically separate networks.

More IPspaces with VLAN tagging

Dedicating at least one physical interface per IPspace limits the number of IPspaces that can be set up on a storage system to the number of physical interfaces available on the storage system. VLAN tagging enables you to overcome this limitation.

VLAN tags can be used to forward traffic to appropriate IPspaces in cases where more than one IPspace shares the same physical interface.

Secure delivery of packets to a vFiler unit in an IPspace

VLANs inherently confine the broadcast domains. Therefore, only vFiler units belonging to a VLAN receive broadcasts intended for that VLAN, even if multiple vFiler units share a physical network interface.

HA pair and IPspaces

You can use IPspaces in an HA pair. However, you should keep in mind the requirements for naming and assigning IPspaces.

IPspace naming requirement

The names of IPspaces to which the partner interfaces are assigned must be the same on both storage systems.

For example, in an HA pair of storage system A and storage system B, if IPspaceA is created on storage system A, an IPspaceA must also exist on storage system B.

IPspace assignment requirement

The partner interfaces on both partners must be assigned to IPspaces with the same name on their respective storage systems.

For example, in an HA pair of storage system A and storage system B, interface e4 of storage system B is the takeover interface of interface e0 of storage system A, and vice versa. If interface e0 belongs to IPspaceA on storage system A, interface e4 must belong to IPspaceA on storage system B.

Asymmetric HA pair setup

In an asymmetric HA pair setup, the vFiler unit-IPspace configuration in one HA pair is different from that of its partner. For example, each partner might have a different number of vFiler units configured in a specific IPspace, or one partner might have no vFiler units.

A “standby” HA pair is an example of an asymmetric HA pair setup in which one of the hosts is connected to minimal storage of its own. This host takes over its partner’s storage and vFiler units if the partner fails. In such a configuration, the standby host might not have enough storage to support the number of vFiler units that the primary host has.

Specifying partners in an asymmetric HA pair setup

You can use the interface name of the partner instead of the IP address to specify the partner interface when setting up an HA pair.

About this task

In this example, two storage systems, storage system1 and storage system2, are configured as shown in the following table:

Storage system	vFiler unit (associated IP address)	IPspace	Interface name
storage system1	vfiler0 (1.1.1.1)	default	e0
	vFiler1 (2.1.1.1)	ips1	e1
	vFiler2 (3.1.1.1)	ips2	e2

Storage system	vFiler unit (associated IP address)	IPspace	Interface name
storage system2	vfiler0 (1.1.1.2)	default	e0
			e4a
			e4b

Steps

1. Specify partner interfaces on storage system1 by creating a host-partner relationship using the following commands:

```
ifconfig e0 1.1.1.1 netmask 255.255.255.0 partner e0
ifconfig e1 2.1.1.1 netmask 255.255.255.0 partner e4a
ifconfig e2 3.1.1.1 netmask 255.255.255.0 partner e4b
```

2. Specify partner interfaces on storage system2 by completing the following steps:

- a) Create two IPspaces: ips1 and ips2.
- b) Assign interface e4a to ips1 and interface e4b to ips2.
- c) Create a host-partner relationship on storage system2 by entering the following commands:

```
ifconfig e0 1.1.1.2 netmask 255.255.255.0 partner e0
ifconfig e4a partner e1
ifconfig e4b partner e2
```

Creating an IPspace

IPspaces are distinct IP address spaces in which vFiler units reside. You create IPspaces when you need your vFiler units to have their own secure storage, administration, and routing.

About this task

You can have a maximum of 101 IPspaces per storage system. Of the 101 IPspaces, one is created by default when you enable MultiStore on your storage system. You can create the remaining 100 IPspaces on the storage system.

You can use an alphanumeric string, 1 to 31 characters long, as the IPspace name.

All the IPspace names you create on a storage system must be unique. However, the interfaces that are mapped to each other in an HA pair must belong to the same IPspace.

Step

1. To create an IPspace, enter the following command:

```
ipspace create ipspacename
```

ipspacename is the IPspace name that you want to create.

Example

To create IPspace1 on a storage system, enter the following command:

```
ipspace create ipspace1
```

After you finish

You can now assign an interface to the IPspace.

IPspace and the routed daemon

When MultiStore is enabled, the `routed` daemon runs only in the default IPspace. The `routed` daemon is disabled when you create a new IPspace. It is enabled automatically after the next reboot. Optionally, you can enable the `routed` daemon by using the `route on` command, without rebooting the storage system.

Listing IPspaces on a storage system

When you want to see which interfaces are assigned to each IPspace, you can use the `ipspace list` command.

Step

1. Enter the following command:

```
ipspace list
```

Example

If you enter the `ipspace list` command on a storage system that has three nondefault IPspaces, you see the following output:

```
Number of ipspaces configured: 4 default-ipspace (e3) ipspace1 (e2d)
ipspace2 (e2c) ipspace3 (e10 e2b sf_vif)
```

Removing an IP address from an interface

You must remove all the IP addresses from an interface before you can assign that interface to an IPspace.

Step

1. To remove an IP address from an interface, perform the appropriate action:

If the IP address is...	Then...
An alias	<ul style="list-style-type: none"> Enter the following command to remove an IPv4 address: <code>ifconfig interface -alias address</code> <code>interface</code> is the name of the interface. <code>address</code> is the IP address configured for the alias. Enter the following command to remove an IPv6 address: <code>ifconfig interface inet6 -alias IPv6_address</code> <code>interface</code> is the name of the interface. <code>IPv6_address</code> is the IPv6 address configured on the interface.
The primary IP address	<p>Enter the following command: <code>ifconfig interface 0.0.0.0</code> <code>interface</code> is the name of the interface.</p> <p>Note: To remove the primary IP address of an interface, you must remove all the aliases from that interface.</p>

Result

You can now assign the interface to an IPspace.

Related concepts

[Effects of adding, removing, and moving vFiler unit resources](#) on page 27

Assigning an interface to an IPspace

To assign an interface to an IPspace, you must ensure that the interface does not have a configured IP address. You can assign only one interface at a time. If IPv6 is enabled on your storage system and you bring up an interface, an IPv6 address is automatically configured on the interface.

About this task

Note: You cannot assign the e0M interface to a nondefault IPspace.

Step

- To assign an interface to an IPspace, enter the following command:

`ipspace assign ipspacename interface_name`

`ipspacename` is the IPspace name to which the interface is assigned.

interface_name is the name of the interface to be assigned.

Destroying IPspaces

If you no longer need an IPspace, you can destroy it.

Before you begin

You must have ensured that there are no network interfaces or vFiler units associated with the IPspace you want to destroy.

Step

1. To destroy an IPspace, enter the following command:

```
ipspace destroy ipspace_name
```

ipspace_name is the IPspace name that you want to destroy.

Related tasks

[Listing IPspaces on a storage system](#) on page 67

[Assigning an interface to an IPspace](#) on page 68

[Displaying the vFiler unit status](#) on page 39

[Destroying a vFiler unit](#) on page 34

Creating a vFiler unit in a nondefault IPspace

If you are a Storage Service Provider (SSP), you might have different clients that require complete network isolation. Creating vFiler units in separate, or nondefault IPspaces ensures that the data on these vFiler units is not routed to other vFiler units or the hosting storage system in other IPspaces.

Before you begin

- You must have created at least one unit of storage (qtrees, traditional or FlexVol volumes).
- The storage unit that contains information about configuring the vFiler unit must be writable. It must not be a read-only file system, such as the destination volume or qtree in a SnapMirror relationship.
- The IP address to be used by the vFiler unit must not have been configured when you created the vFiler unit.
- You must have created an IPspace.
- You must have verified that each interface to be used by the vFiler unit is ready for configuration.

Steps

1. Assign an interface to be used by the vFiler unit to the newly created IPspace by entering the following command:

```
ipspace assign ipspacename interface_name
```

ipspacename is the newly created IPspace.

2. Create the vFiler unit by entering the following command:

```
vfiler create vfiler_name -n -s ipspace -i ip_address [ -i ip_address ] ... path [ path ] ...
```

ipspace is the IPspace in which the vFiler unit's IP addresses reside.

ip_address is an IP address.

path is the complete path name to an existing volume or qtree.

The first path name is the storage unit that contains the `/etc` directory, which contains the configuration information about the vFiler unit.

Note: You must use the `-n` option of the `vfiler create` command to specify IP addresses for interfaces assigned to different IPspaces. The `setup` command (which runs automatically after the `vfiler create` command unless you use the `-n` option) does not allow duplicate IP addresses even if they are for interfaces in the same IPspace.

Example

The following command shows how to create a vFiler unit named vFiler1 in the IPspace named `ipspace1`, using the IP address `123.123.123.123` and the `/vol/vol1` volume as resources:

```
vfiler create vFiler1 -n -s ipspace1 -i 123.123.123.123 /vol/vol1
```

3. Depending on the type of IP address, perform the appropriate action:

If you want to create... Enter the following command...

An alias IP address **ifconfig interface alias *address* up**

A primary IP address **ifconfig interface *address***

Note: When you run the `ifconfig interface address` command, the existing primary IP address is replaced with the new IP address.

Each interface can have one IP and one IPv6 primary address. You can add or delete additional addresses (aliases) after the primary IP address is created. You cannot delete the primary IP address if any aliases exist, but you can modify it to change the address.

4. Modify the routing table for the vFiler unit by entering the following command:

```
vfiler run vfiler_name route add [inet] default gateway_metric
```

```
vfiler run vfiler_name route add [inet6] default gateway_metric
```

Example

The following example shows how to add a route to the routing table used by the vFiler unit named vFiler1:

```
vfiler run vFiler1 route add default 1.2.3.4 1
```

5. For each interface used by the vFiler unit, add the following information to the `/etc/rc` file on the hosting storage system:

- The same `ifconfig` command that was used to configure the interface as Up.
- The same `route` command that was used to add a route to the routing table.

This configures the interfaces as Up and enforces the `route` commands across reboots.

6. Optional: If the hosting storage system is part of an HA pair, edit the `/etc/rc` file in each partner of the HA pair to define a partner interface for each interface that the vFiler unit uses.

Example

The following command shows how to ensure that interface `e10` has a partner `e10` interface:

```
ifconfig e10 partner e10
```

Alternatively, you can run the `setup` command from the storage system and enter the partner interface for each interface.

Attention: If you use the storage system `setup` command to automatically define a partner interface, it clears all existing information about vFiler unit configuration in `/etc/rc` file.

7. Optional: If the IPspace does not have a default gateway, establish a route to the default gateway in the IPspace that any vFiler in the IPspace can use by entering the following command:

```
vfiler run vfiler_name route add default gateway_metric
```

Example

The following command shows how to add a route to the default gateway for the IPspace used by the vFiler unit named vFiler1:

```
vfiler run vFiler1 route add default 1.2.3.1 1
```

```
vfiler run vFiler1 route add inet6 default fc20::abcd 1
```

Related concepts

[Prerequisites for creating vFiler units](#) on page 19

Related tasks

[Creating an IPspace](#) on page 66

[Assigning an interface to an IPspace](#) on page 68

[Ensuring that the network interface is ready](#) on page 23

File system access using NFS and CIFS

To access a vFiler unit's file system using NFS and CIFS, you must prepare the vFiler unit by using the NFS and CIFS protocols, respectively.

For CIFS clients, the root of the primary storage qtree is the root (“/”) of a vFiler unit's file system. If “/” is shared, a CIFS client mapping to it can browse all of the vFiler unit's storage in a single tree. This mechanism is called the vFiler unit's *pseudo-root*.

NFS clients must import discrete storage units as they are defined on the hosting storage systems. Pseudo-root directories are not available to NFS clients.

Path name specification for NFS exports or CIFS shares

When you specify a path to export to NFS clients or to share with CIFS clients, you should use the complete path name.

Example of a path for an NFS export

Assume a vFiler unit named vfiler1 uses the /vol/vol1 volume for storage. To export the home directory at the root of this volume to the clients of vfiler1, you should use /vol/vol1/home in the /etc/exports file or in the exportfs command.

Example of a path for a CIFS share

Assume a vFiler unit named vfiler1 uses the hosting storage system's /vol/vol1 volume as its primary storage. To share the entire volume, and all other storage owned by the vFiler unit, in a single tree, you should specify / as the share path. To offer the home directory at the root of this volume as the home share, specify /home as the path name for the home share. The vFiler unit mechanism that makes this possible is known as *pseudo-root*.

vFiler unit preparation for NFS

To prepare the vFiler unit for NFS, you must start the NFS protocol and export discrete storage units to NFS clients in the same way they are defined on the hosting storage system.

However, you must do this only if you want to start NFS manually.

Note: When you use an IPv6 address, you should enter the following command: `vfiler run vfiler_name options nfs.ipv6.enable on`.

Related tasks

[Executing commands from a vFiler unit](#) on page 41

Starting the NFS protocol

To start the NFS protocol on the vFiler unit, you should use the `nfs on` command.

Step

1. To start the NFS protocol, perform appropriate action based on your role:

If you are a...	Then...
vfiler0 administrator	<ul style="list-style-type: none"> From the hosting storage system, enter the following command: <code>vfiler run vfiler_name nfs on</code> <p>Note: When you use an IPv6 address, you must enter the following command: <code>vfiler run vfiler_name options nfs.ipv6.enable on</code></p>
Any other vFiler unit administrator	From the vFiler unit, enter the following command: <code>nfs on</code>

Result

The NFS protocol server starts running on the vFiler unit.

Exporting all file systems in /etc/exports

As pseudo-root directories are not available to NFS clients, to start using NFS on a vFiler unit, NFS clients must import discrete storage units in the same way they are defined on the hosting storage system.

Step

1. Complete one of the following steps:

- From the vFiler unit, enter the following command:
`exportfs -a`
- From the hosting storage system, enter the following command:
`vfiler run vfiler_name exportfs -a`
- From a vFiler unit client that is allowed to connect to the vFiler unit through RSH, enter the following command:
`rsh vfiler_name exportfs -a`

vFiler unit preparation for CIFS

You can run the `cifs setup` command to configure CIFS on the vFiler unit. You can also configure CIFS while creating the vFiler unit by using the `vfiler create` command.

From the hosting storage system, you can use the `vfiler run` command to issue CIFS commands for vFiler units. You can use System Manager to manage user accounts and shares.

Note: The `cifs setup` command is not available on RSH.

Note: When you use an IPv6 address, you should enter the following command: `vfiler run vfiler_name options cifs.ipv6.enable on`.

Data ONTAP does not limit the number of users, shares, open files, and locked files on a per vFiler unit basis.

Related tasks

[Executing commands from a vFiler unit](#) on page 41

Commands run from the hosting storage system

Server Manager does not perform all the functions of the `cifs shares -add` and `cifs shares -change` commands. You can execute the commands that Server Manager cannot run from the vFiler unit's CLI (through the `vfiler context` command) or from the hosting storage system (through the `vfiler run` command).

The following commands cannot be run from Server Manager:

- `cifs shares -add-forcegroup group_name`
- `cifs shares -add share_name pathname -nosymlink_strict_security`
- `cifs shares -add- widelink`
- `cifs shares -add -novscan`
- `cifs shares -add -novscanread`
- `cifs shares -change share_name { -forcegroup group_name | -noforcegroup }`
- `cifs shares -change share_name { -symlink_strict_security | -nosymlink_strict_security }`
- `cifs shares -change share_name { -widelink | -nowidelink }`
- `cifs shares -change share_name { -vscan | -novscan }`
- `cifs shares -change share_name { -vscanread | -novscanread }`

Local user accounts for vFiler units

From the hosting storage system, you can use the `useradmin` command to create local accounts for CIFS users of each vFiler unit. Each vFiler unit supports up to 96 local user accounts.

The maximum number of vFiler unit user accounts per storage system is 96 times the maximum number of vFiler units for that storage system.

For additional information about using local accounts for CIFS authentication, see the *Data ONTAP File Access and Protocols Management Guide for 7-Mode*.

For detailed information about managing local user accounts by using the `useradmin` command, see the *Data ONTAP Storage Management Guide for 7-Mode*.

Related tasks

[Viewing the current limit on the number of vFiler units](#) on page 30

Virus protection for CIFS

You can perform virus scanning for vFiler units that run the CIFS protocol. The hosting storage system administrator can configure virus scanning on files owned by the hosting storage system and files owned by vFiler units. Administrators of vFiler unit can configure virus scanning only for vFiler units they administer.

Virus scanner registration

Virus scanners can be registered with the hosting storage system or any vFiler unit. A virus scanner that is registered with a vFiler unit always takes precedence over the virus scanner registered with the hosting storage system.

You can determine whether files on a vFiler unit are scanned by virus scanners registered with the hosting storage system or the vFiler unit. If a virus scanner is registered with a vFiler unit and is functional, the files on the vFiler unit are scanned by the scanner that is registered with the vFiler unit.

If the scanner that is registered with a vFiler unit becomes unavailable, then the hosting storage system's virus scanner scans the files on the vFiler unit on the following conditions:

- The `vscan options use_host_scanners` command is set to On
- A scanner is registered with the hosting storage system

When the scanner local to the vFiler unit becomes available, it takes over from the hosting storage system's scanner.

Virus scanning on vFiler units

To scan files on nondefault vFiler units, virus scanning must be enabled on these vFiler units. Files on the default vFiler unit can be scanned only by a virus scanner registered with the hosting storage system. Also, only hosting storage system administrators can run virus scanning on vfiler0.

Before you run virus scanning on nondefault vFiler units, ensure the following requirements are met:

- Virus scanning must be enabled for a nondefault vFiler unit.
- A virus scanner must be registered either with the vFiler unit or hosting storage system, and the vFiler unit must be allowed to use it.

Effect of virus scanner availability on CIFS access

Although virus scanning is enabled and the `mandatory_scan` option for the `vscan` command is set to `On`, CIFS clients of the vFiler unit are not allowed to open any files on the vFiler unit if no virus scanner is available.

Configuring virus scanning for a vFiler unit

You can configure virus scanning either to use or not use the virus scanner registered with the hosting storage system, if the virus scanner registered with the vFiler unit is unavailable.

Steps

1. Enter the following command to specify the virus scanner for scanning files on a vFiler unit:

```
vfiler run vfilertemplate vscan options use_host_scanners on | off
```

When you set the `use_host_scanners` option to `On`, the vFiler unit uses the virus scanner registered with the hosting storage system. When this option is set to `On`, the hosting storage system and its vFiler units share the virus scanner. However, the vFiler unit uses the virus scanner registered with the hosting storage system only when the virus scanner registered with the vFiler unit is unavailable.

You can set the `use_host_scanners` option to `Off` if you do not want to allow the vFiler unit to use the virus scanner registered with the hosting storage system.

Note: The `use_host_scanners` option is applicable only to a vFiler unit you created. You cannot set it on `vfiler0` or a storage system.

2. Enter the following command to enable or disable virus scanning:

```
vfiler run vfilertemplate vscan on | off
```

Disaster recovery using MultiStore

MultiStore supports disaster recovery, and you can prepare for disaster recovery by creating a backup vFiler unit to prevent loss of data if a disaster occurs.

You can safeguard information by creating vFiler units on the destination storage system, which remain inactive unless a disaster occurs. You should perform checks to ensure that the storage system and network are ready for disaster recovery.

Related concepts

What an IPspace is on page 61

Data migration using MultiStore on page 100

Checking and preparing the storage system

You must ensure that the destination storage system can support the disaster recovery vFiler unit.

Steps

1. Verify that the destination storage system has enough storage space to hold the source vFiler unit's volumes.
 - a) On the source storage system, enter the `vfiler status -r` command to see the volumes that the vFiler unit is using.
 - b) Enter the `df` command on each of those volumes to check the disk space being used. The destination volumes must have at least the same amount of space that are being used on the source volumes. You can run the `df` command on the destination storage system to check this.

Note: If the source and destination storage systems use different-sized disks and have different block sizes, you should adjust the `df` numbers accordingly.

2. Perform the appropriate action based on what you learned about space in Step 1:

Note: The destination volumes must have the same amount of space as the source volumes.

If the destination volumes... Then...

Have enough space	Go to the Step 3.
Do not have enough space	<ul style="list-style-type: none">a. Install new disk shelves.b. Use the <code>aggr add</code> command to add new disks to the destination volumes.

3. Ensure that the destination storage system has the same volume structure as the source and the volumes to be used by the destination vFiler unit are not used by any other vFiler unit.

The volumes to be used by the destination vFiler unit must have the same path names as those used by the source vFiler unit.

If the destination storage system...	Then...
Has a volume whose path name matches the path name used by the source vFiler unit and the volumes to be used by the destination vFiler unit are not used by any other vFiler unit	Go to Step 4.
Has a volume whose path name matches the path name of the source vFiler unit, but the volume is used by another vFiler unit	<p>Perform one of the following steps:</p> <ul style="list-style-type: none"> • If the volume is the root volume of the vFiler unit, use the <code>vfiler destroy</code> command to destroy the vFiler unit. • If the volume can be removed, use the <code>vfiler remove</code> command to disassociate the volume from that vFiler unit. • If the volume cannot be destroyed or removed, use the <code>vol rename</code> command to rename the volume. Then, create a new volume with the old name of the volume you just renamed.
Does not have any path name that matches the name used by the source vFiler unit or by volumes to be used by the destination	<p>Perform one of the following steps:</p> <ul style="list-style-type: none"> • For traditional and FlexVol volumes, use the <code>vol create</code> command on the destination storage system to create volumes whose names match those being used by the source vFiler unit. • Use the <code>vol rename</code> to rename a volume. <p>For more information, see the <i>Data ONTAP Storage Management Guide for 7-Mode</i>.</p>

4. You must ensure that the destination volumes do not contain any qtree with the same name as that used by a qtree in the source volumes.

For disaster recovery, SnapMirror is used, and SnapMirror replicates qtree names from the source to the destination volume. Therefore, ensure that the qtree names on the source do not exist on the destination.

If there are...	Then...
Qtrees in the destination volumes that have names matching the names of qtrees in the source volumes	<p>Rename the matching qtrees in the destination volumes.</p> <p>To rename a qtree, move it from a client the same way you move a directory or folder. For more information, see the <i>Data ONTAP Storage Management Guide for 7-Mode</i>.</p>

If there are...	Then...
No matching qtree names in the destination volumes	Go to Step 5.

5. Verify whether quotas are being enforced from the hosting storage system.

To verify where quotas are being enforced from, enter the following command from the hosting storage system:

quota report

Quotas enforced from the vFiler unit are copied to the new vFiler unit, but quotas enforced from the hosting storage system are not copied.

6. Depending on the result of Step 5, perform the appropriate action:

If quotas for qtrees used by the vFiler unit...	Then...
Are being enforced	Go to the Step 7.Go to the Step 8.
Are not being enforced	You have completed checking the storage system.

7. Keep a record of the storage system's /vol/vol0/etc/quotas file for future reference.

8. Copy the relevant entries into the destination storage system's /vol/vol0/etc/quotas file.

Storage checklist

You can use the storage checklist to record storage system information and to ensure that your systems are ready to use for disaster recovery.

- How much disk space is used on the source storage system's volumes?
df of volumes of the source storage system:
- How much disk space is free on the destination storage system's volumes?
df of volumes of the destination storage system:
- Have you added enough disks to the destination volumes, if required? _____
- Do the path names of the source and destination volumes match? _____
- If you are managing qtree-based vFiler units, do any destination volume qtree names match those on the source volume? _____
- Have you copied storage system-based quota information from the source to the destination storage system's /etc/quotas file? _____

Checking the network

Before setting up a disaster recovery vFiler unit, you must check whether the source and destination storage systems are on the same subnet, the IPspace used by the vFiler unit, and whether the destination vFiler unit can access the same NIS and DNS servers as the source vFiler unit.

Steps

1. Check whether the destination vFiler unit can take over the source vFiler unit's IP addresses by displaying information about all the network interfaces of the destination vFiler unit.

You must use the `ifconfig -a` command on the source vFiler unit and the destination storage system to display information about all the network interfaces.

You can reuse the source IP addresses and aliases on the destination vFiler unit if the destination vFiler unit is on the same subnet as the source vFiler unit.

2. Depending on the result of Step 1, perform one of the following actions:

If the source and destination storage systems...	Then...
Are on the same subnet	<p>Go to Step 3.</p> <p>Note: This is the default for the <code>vfiler migrate</code> command.</p>
Are on different subnets	<ol style="list-style-type: none"> Obtain all new IP addresses that are in use on the source vFiler unit for the destination vFiler unit. <p>Note: You might need to replicate subnet-separation arrangements that exist on the source vFiler unit. For example, the source vFiler unit might use one IP address for a service network and another for an administration network.</p> Make a note of the new IP addresses on a worksheet. <p>The <code>vfiler dr configure</code> command prompts you for these addresses when you create the disaster recovery vFiler unit. The <code>vfiler migrate</code> command also prompts you for these addresses, but you might then need to run the <code>setup</code> command on the destination vFiler unit.</p>

3. Check whether the source vFiler unit is using the default IPspace.

To display information about IPspaces and the interfaces assigned to them, you must use the `ipspace list` command on the source vFiler unit.

4. Depending on the result of Step 3, perform one of the following actions:

If the ipspace list command reports...	Then...
Default-ipspace	Go to Step 5.
Something other than default-ipspace	<ul style="list-style-type: none"> a. Use the <code>ipspace create</code> command to create a corresponding IPspace with the same name on the destination storage system. b. Use the <code>ipspace assign</code> command to assign physical interfaces to the IPspace. <p>These interfaces should be attached to the same physical network.</p>

5. Check whether the destination vFiler unit has access to the same NIS servers as the source.

Note: You can skip this check if the source and destination vFiler units are on the same subnet.

To see the NIS servers that are available to the source vFiler unit, use the `nis info` command.

Note: The `ypwhich` command shows the server to which the storage system is currently bound.

6. Depending on the result of Step 5, perform one of the following actions:

If...	Then...
The destination vFiler unit can use the same NIS servers as the source vFiler unit	<p>Go to Step 7.</p> <p>Note: This is the default for the <code>vfiler migrate</code> command.</p>
The destination vFiler unit cannot use the same NIS servers	<ul style="list-style-type: none"> a. Find the NIS servers that are available for the destination storage system. b. Make a note of the IP addresses of those servers on the network checklist. <p>The <code>vfiler dr configure</code> command prompts you for these addresses when you create the disaster recovery vFiler unit.</p> <p>The <code>vfiler migrate</code> command does not prompt you for these addresses. If you move a vFiler unit to a different subnet, you might need to run the <code>setup</code> command on the destination vFiler unit.</p>

7. Check whether the destination vFiler unit has access to the same DNS servers as the source.

Note: You can skip this check if the source and destination vFiler units are on the same subnet.

To see what DNS servers are available to the source vFiler unit, use the `dns info` on command.

8. Depending on the result of Step 7, perform one of the following actions:

If...	Then...
The destination vFiler unit can use the same DNS servers as the source vFiler unit	<p>Go to Step 9.</p> <p>Note: This is the default for the <code>vfiler migrate</code> command.</p>
The destination vFiler unit cannot use the same DNS servers	<p>a. Find the DNS servers that are available for the destination storage system.</p> <p>b. Make a note of the IP addresses of those servers on a worksheet.</p> <p>The <code>vfiler dr configure</code> command prompts you for these addresses.</p> <p>The <code>vfiler migrate</code> command does not prompt you for these addresses. If you move a vFiler unit to a different subnet, you might need to run the <code>setup</code> command on the destination vFiler unit.</p>

9. Check whether the destination vFiler unit has access to the same WINS servers and the same Windows security network as the source.

10. Depending on the result of Step 9, perform one of the following actions:

If the destination vFiler unit...	Then...
Can use the same WINS servers and Windows security network as the source vFiler unit	<p>Go to Step 11.</p>
Cannot use the same WINS servers and Windows security network	<p>a. Find the name and type (Windows NT 4 or Windows 2000) of the domain the destination vFiler unit is in.</p> <p>b. Note this information on the network checklist.</p> <p>When you activate the disaster recovery vFiler unit, you have to configure it into the new domain.</p> <p>If you move a vFiler unit into a different domain, you have to configure it into the new domain.</p>

11. Check whether the destination vFiler unit can use the same trusted host for vFiler unit administration as the source vFiler unit.

12. Depending on the result of Step 11, perform one of the following actions:

If the destination vFiler unit...	Then...
Can use the same trusted host as the source vFiler unit	You have completed this task.
Cannot use the same trusted host	<p>a. Find the name of the new trusted host.</p> <p>b. Note this information on the network checklist.</p> <p>You must configure the new trusted-host information after configuring the disaster recovery vFiler unit, or after moving the vFiler unit.</p>

Related tasks

[Creating a disaster recovery vFiler unit](#) on page 84

[Adjusting client and network configurations if migrating to a different subnet](#) on page 104

[Activating the disaster recovery vFiler unit](#) on page 88

Related references

[Storage checklist](#) on page 79

Network checklist

You can use the network checklist to record network information and to ensure that your systems are ready to use for disaster recovery.

- Are there enough IP addresses available for the vFiler unit on the destination network?

old interface: _____ new interface: _____
 old interface: _____ new interface: _____
 old interface: _____ new interface: _____
 old interface: _____ new interface: _____

Note: Check syntax carefully. Interface names are case-sensitive.

- Have you created the number of nondefault IPspaces, if any are required?
- Have you gathered all the authority servers?

old NIS domain: _____ new NIS domain: _____
 old NIS IP address: _____ new NIS IP address: _____
 old NIS IP address : _____ new NIS IP address: _____
 old DNS domain: _____ new DNS domain: _____
 old DNS IP address : _____ new DNS IP address: _____
 old DNS IP address : _____ new DNS IP address: _____
 old DNS IP address : _____ new DNS IP address: _____
 old WINS IP address : _____ new WINS IP address: _____
 old WINS IP address : _____ new WINS IP address: _____

old NT domain type: NT4 W2K

old domain name (FQDN and NetBIOS): _____

new NT domain type: NT4 W2K

new domain name (FQDN and NetBIOS): _____

- Can you use the same trusted host for vFiler unit administration?

old trusted host name: _____ new trusted host name: _____

Related tasks

[Checking the network](#) on page 80

Secure communication for disaster recovery

Configuring disaster recovery with MultiStore requires communication between the source and destination storage systems. During disaster recovery configuration, commands are sent from the destination storage system to the source storage system, and are authenticated by an administrative user name and password.

By default, the commands and the authentication user name and password are sent in cleartext by using the RSH protocol.

To enable secure communication, you must enable SSL by using the `secureadmin` command on the default vFiler unit (vfiler0) context. In addition, you must use the `-c secure` option of the `vfiler dr configure` command to send the commands and the authentication user name and password by using the encrypted SSL protocol.

Note: When you use encrypted SSL protocol, you must ensure the following requirements are met:

- SSL is enabled on the source vFiler unit.
- The `httpd.enable`, `httpd.admin.enable`, and `httpd.admin.ssl.enable` options are turned on at the source vFiler unit.

SSL is not supported from the nondefault vFiler unit context.

Creating a disaster recovery vFiler unit

You can create a disaster recovery vFiler unit on a destination storage system that has the storage capacity, characteristics, and the network connectivity to host an identical copy of the vFiler unit on the source storage system, which can serve data if the original vFiler unit fails to serve data.

Before you begin

- You must have prepared the destination storage system.

- SnapMirror technology must have been licensed and enabled on both the source and the destination storage systems.
- The source and destination storage system can communicate with each other over the network.
- For a multipath asynchronous SnapMirror relationship, you must have specified the mode of connection 'multi' in the `snapmirror.conf` file.

Note: The source address and the destination address in the `snapmirror.conf` file must match the host names of the source and destination storage systems.
- The destination volumes must be online.
- You must know the source storage system's administrative ID and password.

About this task

Attention: On the disaster recovery storage system, you must protect any volumes that have the same names as the volumes on the original vFiler unit. Otherwise, data in those volumes is lost.

Steps

1. On the destination storage system, enter the following command:

```
vfiler dr configure source_vfiler@source_filer
```

Note:

- You must use the `-c` secure option to enable secure communication.

Note: You must ensure that when you are using IPv6 addresses, the `httpd.ipv6.enable` option is set to on.

For more information about the `-c` option, see the `na_vfiler(1)` man page.

- If you want to set up synchronous SnapMirror between the source and destination storage systems, you should use the `-s` option of the `vfiler dr configure` command.

For more information about the `-s` option, see the `na_vfiler(1)` man page.

- If you want to set up multiple path for synchronous SnapMirror relationship from the source to the destination storage system, you should use the `-a` option of the `vfiler dr configure` command.

For more information, see the section on using SnapMirror over multiple paths in the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode*.

- If you do not want to initialize an existing SnapMirror relationship between the source and the destination storage system after the baseline transfer, you should use the `-u` option of the `vfiler dr configure` command.

2. Optional: For multipath asynchronous SnapMirror relationship, edit the SnapMirror connection names in the `snapmirror.conf` file.

Example

```
Multipath_connection_name: volume1 dst_filer:volume1 - 0-59/3 * * *
```

3. Respond to the login prompt with a valid administrative ID and password for the source storage system.
4. Respond to the IP address and binding prompts.
5. Respond to the NIS and DNS server prompts.
6. Optional: Monitor the progress of the disaster recovery by using the following command:

```
vfiler dr status source_vfiler@source_filer
```

When the `vfiler dr status` command output shows that all the storage units of the source vFiler unit are mirrored, the disaster recovery vFiler unit has been created. However, the disaster recovery vFiler unit has not been started.

Note: The `vfiler dr configure` command might take some time to complete while the volumes are being replicated.

Related concepts

[Disaster recovery using MultiStore](#) on page 77

[The vfiler dr configure command](#) on page 87

Related references

[Storage checklist](#) on page 79

[Network checklist](#) on page 83

Disaster recovery with SnapMirror IP address based verification and IPv6 addresses

During disaster recovery configuration, if the option `snapmirror.checkip.enable` is set to `On` on the source storage system, and the IPv6 address of the source storage system is used for communication, SnapMirror requires the IPv6 address of the destination storage system to set up a SnapMirror relationship.

You must specify the destination storage system's IPv6 address in the list of allowed addresses for the SnapMirror relationship on the source storage system. You can add the destination storage system's IPv6 address either to the `snapmirror.access` option or to the `/etc/snapmirror.allow` file if the `snapmirror.access` option is set to `legacy`.

For more information about the `snapmirror.checkip.enable` option, see the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode*.

Related information

[N series support website: \[www.ibm.com/storage/support/nseries/\]\(http://www.ibm.com/storage/support/nseries/\)](#)

Deleting the disaster recovery vFiler unit

You can delete the disaster recovery vFiler unit at any time after setting it up.

Step

1. On the destination storage system, enter the following command:

```
vfiler dr delete source_vfiler@source_filer
```

Example

To remove a disaster recovery vFiler unit on the destination storage system "StorageSystem 2" created for a vFiler unit "vfiler1" on the source storage system "StorageSystem1", even if SnapMirror errors exists, enter the following command on "StorageSystem 2":

```
vfiler dr delete -f vfiler1@StorageSystem1
```

Result

Before removing the disaster recovery vFiler unit, the `vfiler dr delete` command removes all SnapMirror relationships, and any other configuration information related to the disaster recovery vFiler unit, from the source vFiler unit.

If any errors are detected in the SnapMirror relationships, the deletion of the vFiler unit is canceled. To ignore SnapMirror errors and remove the disaster recovery vFiler unit, you can use the `-f` option available in the `vfiler dr delete` command.

The `vfiler dr configure` command

With the `vfiler dr configure` command you can set up a disaster recovery vFiler unit. This command uses the Data ONTAP SnapMirror feature as its underlying technology. The `-a` option of the `vfiler dr configure` command enables you to set multiple paths for the SnapMirror configuration.

The `vfiler dr configure` command performs the following actions:

- Checks whether the destination storage system can receive the source data.
- Configures and runs SnapMirror to copy the data from the source to the destination vFiler unit.
 - iSCSI LUNs (including the LUN maps) are copied from the source vFiler unit to the destination vFiler unit.
 - igroups and the iSCSI configuration, including node names and the iSCSI service state, are copied to the destination vFiler unit.
 - iSCSI authentication is not copied to the destination vFiler unit.

- Saves the IP configuration and binding information you supplied when you created the disaster recovery vFiler unit.
- Saves the NIS and DNS server information you supply.
- Saves the quota information from the source vFiler unit's `/etc/quotas` file.
- Causes a baseline transfer to occur from the source to the destination.
- Sets the incremental update interval from the source to the destination to be once every three minutes.
 - If you want to change the default setting, you should edit the `etc/snapmirror.conf` file as described in the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode*.
 - The `vfiler dr configure` command automatically configures everything that SnapMirror requires for regular updates. No other SnapMirror configuration is necessary.
- Overwrites all data on the volumes of the destination vFiler unit.
You must protect any volumes on the destination storage system that have the same name as the volumes on the source vFiler unit. Otherwise, data in the volumes of the destination vFiler unit is lost.
- Creates a vFiler unit on the destination storage system as the disaster recovery backup vFiler unit. This vFiler unit is stopped and cannot be started except when a disaster occurs. Before activation, the vFiler unit responds only to the `vfiler dr delete`, `vfiler dr status`, and `vfiler dr resync` commands. You should not use `ifconfig` command to configure its addresses.

Note: With the `-u` option, the `vfiler dr configure` command does not initialize the SnapMirror relationship between the source and the destination storage system.

For more information about the options of the `vfiler dr configure` command, see the `na_vfiler(1)` man page.

Related tasks

[Creating a disaster recovery vFiler unit](#) on page 84

[Activating the disaster recovery vFiler unit](#) on page 88

Activating the disaster recovery vFiler unit

After a disaster, you can keep serving data by switching to the disaster recovery vFiler unit while trying to recover the original vFiler unit that was damaged during the disaster.

Steps

1. On the destination storage system, enter the following command:

```
vfiler dr activate source_vfiler@source_filer
```

2. Configure the DNS or NIS servers:

If you specified...	Then...
A different set of DNS or NIS servers	Copy the <code>/etc/hosts.equiv.bak</code> file to the <code>/etc/hosts.equiv</code> file.
The same set of DNS or NIS servers	Go to Step 3.

When you activate the vFiler unit in the event of a disaster, the `/etc/hosts.equiv` file can be overwritten. If you specified a different set of DNS or NIS servers for the disaster recovery location when you created the disaster recovery vFiler unit, the existing `/etc/hosts.equiv` file is overwritten and the old file is copied to an `/etc/hosts.equiv.bak` file.

3. To change the name of the Windows domain controller, use the `cifs prefdc` command.

4. To change the Windows WINS server, use the `cifs setup` command.

Note: If the Windows domain has changed, you might have to change the permissions on the Windows data files to allow your users the same access they had in the old domain.

5. Make adjustments on the clients, such as remounting volumes and qtrees.

6. Add static route entry if required, because static routing information is not carried to the destination storage system.

After you finish

- If you copied quota information to the destination storage system's `/etc/quotas` file, activate the quotas on that storage system. For activating quota on each volume, use the following command:

`quota on volume_name`

- Edit the disaster recovery vFiler unit's `/etc/hosts.equiv` file by adding the name of the trusted host for administering the disaster recovery vFiler unit.

Note: If the trusted host is either a Windows or a UNIX system and the trusted user is not the root user, you need to add the user name as well. For example:

`adminhost joe_smith`

- Add the path to the root volume and the name of the trusted host to the disaster recovery vFiler unit's `/etc/exports` file.

`/vol/vf1_root access=adminhost, root=adminhost`

- If the vFiler unit's storage units contain iSCSI LUNs, reconfigure iSCSI authentication. For instructions, see the *Data ONTAP SAN Administration Guide for 7-Mode*.

What activating the disaster recovery vFiler unit does

When you activate a disaster recovery vFiler unit, the original vFiler unit is replaced with the disaster recovery vFiler unit.

The `vfiler dr activate` command performs the following actions:

- Breaks the SnapMirror relationships between the source and destination storage systems
- Activates the disaster recovery vFiler unit which then starts and responds to all commands that vFiler units support
- Brings LUNs online
- Configures IP bindings on the destination vFiler unit according to the information you provided in the `vfiler dr configure` command, adding the destination IP information to the destination `/etc/rc` file
Any IP information that pertains only to the source vFiler unit is removed from the destination `/etc/rc` file.
- Configures the NIS and DNS servers according to the information you provided to the `vfiler dr configure` command
- Configures any quota information saved by the `vfiler dr configure` command

Resynchronizing the vFiler unit

You can resynchronize the original vFiler unit with the currently activated disaster recovery vFiler unit before reactivating the original vFiler unit. If you resynchronize, you do not have to delete the original vFiler unit and create a new vFiler unit. A baseline transfer is also not required.

Before you begin

- You must have only volumes (traditional or FlexVol volumes) and not qtrees as the storage elements for the vFiler unit.
- The source and destination vFiler units must contain identical volumes.
- The size of the volumes on the source and destination vFiler units must be the same.
- The vFiler unit from which you are updating must have been activated.
- The original vFiler unit must not be in the process of migration.
- If new storage elements have been added to a disaster recovery activated vFiler unit, the newly added storage elements must exist on the original storage system as well.
- The original vFiler unit you are resynchronizing must be in a stopped state.

About this task

If you do not resynchronize the original vFiler unit, baseline transfer occurs between the new vFiler unit and the disaster recovery vFiler unit.

The `vfiler dr resync` command performs the following actions:

- Resynchronizes all storage elements that belong to the disaster recovery vFiler unit, including the volumes that were added to the disaster recovery vFiler unit after it was activated.
- Sets the incremental update interval from the source to the destination to be once every three minutes. Three minutes is the default setting. If you want to change the default setting, you should edit the `etc/snapmirror.conf` file as described in the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode*. If you edit the SnapMirror relationship schedules in the `etc/snapmirror.conf` file after executing `vfiler dr configure` command, the `vfiler dr resync` command does not modify the SnapMirror relationship schedules for the existing volumes.

Attention:

On the storage system on which you are resynchronizing the original vFiler unit, you must protect any volumes that have the same names as the volumes on the disaster recovery vFiler unit.

If a volume with the same name exists, the volume is automatically added and initialized for SnapMirror transfers from the disaster recovery vFiler unit. Any existing data on the newly added volume is lost.

If you want to delete old Snapshot copies after executing the `vfiler dr resync` command, you must first release them by using the `snapmirror release` command. If you fail to do this and Snapshot copies are deleted, then further SnapMirror updates fail to synchronize.

See the `na_snapmirror(1)` man page for more information about `snapmirror` commands.

Steps

1. On the original storage system, enter the following command:

```
vfiler dr resync [-l authinfo] [-a alt-remote, alt-local] [-c secure][-s] dr_vfilername@disaster_recovery_filer
```

`authinfo` is the authentication information specified in the `username:password` format, where `username` is the login name of the administration host on the disaster recovery storage system and `password` is the password for that user name. If you do not specify the authentication information in the `vfiler dr resync` command, you are prompted for it when you run the command.

`alt-remote` is the alternate host name or IP address of the source (the disaster recovery storage system, in this case).

`alt-local` is the alternate host name or IP address of the destination (the original storage system, in this case).

`-c` enables you to use the secure command channel.

This option is used only in the `vfiler dr` commands that interact with the remote storage system.

`-s` enables you to set up a synchronous SnapMirror relationship between the source and destination storage systems.

`dr_vfilername` is the name of the disaster recovery vFiler unit that is currently in the activated state.

`disaster_recovery_filer` is the name of the storage system on which the currently activated disaster recovery vFiler unit exists.

Note:

When you run the `vfiler dr resync` command on the disaster recovery vFiler unit to resynchronize it with the original vFiler unit, you must specify the same options that were used with the `vfiler dr configure` command for a baseline transfer.

2. After the resynchronization operation is complete, enter the following command on the storage system on which the original vFiler unit exists to verify the status of the vFiler unit that was resynchronized:

`vfiler status -r original_vfilername`

The original vFiler unit is now in a stopped, DR backup state. This is because the `vfiler dr resync` command does not activate the vFiler unit on resynchronizing. The vFiler unit continues to behave as a backup disaster recovery vFiler unit until you use the `vfiler dr activate` command to reactivate it.

3. Stop the disaster recovery vFiler unit by entering the following command:

`vfiler stop disaster_recover_filer`

`disaster_recover_filer` is the name of the disaster recovery vFiler unit.

4. Activate the original vFiler unit by entering the following command:

`vfiler dr activate dr_vfilername@disaster_recover_filer`

`dr_vfilername` is the name of the disaster recovery vFiler unit that is currently in the stopped state.

`disaster_recovery_filer` is the name of the storage system on which the currently stopped disaster recovery vFiler unit exists.

5. Reinstate the disaster recovery vFiler unit on the disaster recovery storage system:

`vfiler dr resync [-l authinfo] [-a alt-remote, alt-local] [-c secure] [-s] original_vfilername@original_filer`

Related tasks

[Activating the disaster recovery vFiler unit](#) on page 88

[Creating a disaster recovery vFiler unit](#) on page 84

Handling resynchronization failures

You can take some corrective action if the resynchronization of the original vFiler unit with the disaster recovery vFiler unit is interrupted or not completed.

Steps

1. Check whether the vFiler unit you were resynchronizing exists on the storage system on which you were running the `vfiler dr resync` command, enter the following command:

`vfiler status`

2. Depending on whether the vFiler unit exists, perform the appropriate action:

If...	Then...
The vFiler unit exists	Stop the vFiler unit by entering the following command: <code>vfiler stop vfilername</code> <code>vfilername</code> is the name of the vFiler unit.
The vFiler unit does not exist	Go to Step 3.

3. Create the vFiler unit by entering the following command:

`vfiler create -r vfilername pathname`

4. Resynchronize the vFiler unit by entering the following command:

`vfiler dr resync [-l authinfo] [-a alt-remote, alt-local] [-c secure] [-s] dr_vfilername@disaster_recovery_filer`

5. Depending on the error message you get, perform the appropriate action:

If you get...	Then...
A “volume offline or does not exist” error message	<ol style="list-style-type: none"> Make the volume online or create it. Resynchronize the vFiler unit.
Volume resync error messages	Reactivate the original vFiler unit by using SnapMirror commands.

Related tasks

[Reactivating the original vFiler unit by using SnapMirror commands](#) on page 94

[Resynchronizing the vFiler unit](#) on page 90

Reactivating the original vFiler unit by using SnapMirror commands

If the storage system is not physically damaged but has failed temporarily and if the storage element associated with the vFiler unit is a qtree, you can reactivate the vFiler unit by using SnapMirror commands. When you reactivate, the disaster recovery configuration is re-created.

Steps

1. Boot the original storage system and interrupt the boot process by pressing the Del or Esc key while the memory self-test is in progress.

Note: If you do not press the Del or Esc key in time, you can press Ctrl-c when prompted later during the boot, select option 5 (maintenance mode), and enter the `halt` command.

2. At the loader prompt, set the `no-vfiler-ips?` parameter as follows:

```
setenv no-vfiler-ips? true
```

This ensures that the storage system does not try to bind IP addresses already being used by the disaster recovery vFiler unit. When the storage system boots, the original vFiler unit starts running. However, it does not accept any read or write requests because its interfaces are not configured.

3. Verify that the `snapmirror.access` option on the disaster recovery storage system is set to `legacy` by entering the following command on the disaster recovery storage system:

```
options snapmirror
```

4. Depending on the result of the `options snapmirror` command, perform the appropriate action:

If the <code>options snapmirror</code> command returns...	Then...
<code>snapmirror.access legacy</code>	Add the host name of the original storage system to the <code>/etc/snapmirror.allow</code> file.
A list of host names that does not contain the name of the original storage system	Use the <code>options snapmirror</code> command to add the host name of the original storage system. For example: <pre>options snapmirror.access host=f3070-216-45,f6070-214-72</pre>

5. Resynchronize each volume and qtree owned by the original vFiler unit by entering the following command on the original storage system that you are trying to activate:

```
snapmirror resync -S disaster_recovery_filer:/pathname original_filer:/pathname
```

Example

```
snapmirror resync -s drfiler:/vol/vfiler1/qtree1 prfiler:/vol/vfiler1/qtree1
```

If the `snapmirror resync` command fails with a message that there are no matching Snapshot copies, you might have accidentally deleted the Snapshot copies that SnapMirror depends on. You must then initialize SnapMirror relationship between the original storage system and destination storage system by using the `snapmirror initialize` command.

For more information about the `snapmirror initialize` command, see the `na_snapmirror(1)` man page.

6. Stop the disaster recovery vFiler unit by entering the following command on the disaster recovery storage system:

```
vfiler stop vfilername
```

`vfilername` is the name of the disaster recovery vFiler unit.

7. Run the `setup` command on the disaster recovery vFiler unit and unconfigure its IP addresses by following prompts of the wizard.
8. Update the data on the original vFiler unit by entering the following command on the original storage system:

```
snapmirror update -s disaster_recovery_filer:/pathname original_filer:/pathname
```

Example

```
snapmirror update -s drfiler:/vol/vfiler1/qtree1 prfiler:/vol/vfiler1/qtree1
```

Note: You do not have to update the data on the original vFiler unit by using the `snapmirror update` command, if you have resynchronized each volume and qtree owned by the original vFiler unit.

9. Stop SnapMirror copy transfers to the disaster recovery vFiler unit.

For each volume and qtree owned by the original vFiler unit, enter the following command on the original storage system:

```
snapmirror quiesce pathname
```

Example

```
snapmirror quiesce /vol/vfiler1/qtree1
```

Note: This operation can take a long time. Use Ctrl-C to interrupt it, if required.

10. Verify that all the paths are quiesced by entering the following command:

```
snapmirror status
```

The `status` column in the output should show each path as Quiesced.

11. Break the SnapMirror relationship.

For each volume and qtree owned by the original vFiler unit, enter the following command on the original storage system:

```
snapmirror break pathname
```

Example

```
snapmirror break /vol/vfiler1/qtree1
```

12. Start the original vFiler unit by entering the following command:

```
vfiler start vfilername
```

vfilername is the name of the original vFiler unit.

13. Run the `setup` command on the original vFiler to configure the vFiler unit's IP addresses and the NIS and DNS servers.

You have completed reactivating the original storage system.

14. If the storage units that have been copied contain iSCSI LUNs, check that the iSCSI configuration on the original vFiler unit is not affected.

You might have to remap the LUNs and re-create the initiator groups (igroups) if you have created new LUNs on the destination vFiler unit.

15. If the storage units that are copied contain iSCSI LUNs, bring the LUNs back online on the original vFiler unit.

16. From the disaster recovery storage system, resynchronize the original vFiler unit and the disaster recovery vFiler unit.

Related tasks

[Resynchronizing the vFiler unit](#) on page 90

Reactivating the original vFiler unit by using `vfiler dr` commands

If the storage system is severely damaged, or you are not familiar with using SnapMirror commands, you can reactivate the vFiler unit by using `vfiler dr` commands.

Steps

1. Boot the original storage system and interrupt the boot process by pressing the Del or Esc key while the memory self-test is in progress.

Note: If you do not press the Del or Esc key in time, you can press Ctrl-C when prompted later during the boot, choose option 5 (maintenance mode), and enter `halt`.

- At the loader prompt, set the `no-vfiler-ips?` parameter as follows:

```
setenv no-vfiler-ips? true
```

This ensures that the storage system does not try to bind IP addresses already being used by the disaster recovery vFiler unit.

- At the loader prompt, enter the `boot` command.
- To destroy the original vFiler unit, enter the following command on the original storage system:

```
vfiler destroy vfilername
```

`vfilername` is the name of the original vFiler unit.

- Stop the disaster recovery vFiler unit by using the `vfiler stop` command.
- Create the disaster recovery vFiler unit.
- Update the data on the original storage system.

For each volume and qtree owned by the new vFiler unit, enter the following commands on the original storage system (or its replacement) to update the data:

```
snapmirror break name
```

```
snapmirror resync -S disaster_recovery_filer:name production_filer:name
```

`disaster_recovery_filer` is the name of the disaster recovery storage system.

`production_filer` is the name of the original storage system.

`name` is the volume name or path name of the qtree.

Example

For the volume `vol1`, enter the following commands:

```
snapmirror break drfiler:vol1
```

```
snapmirror resync -S drfiler:vol1 prfiler:vol1
```

Example

For the qtree `qtree1`, enter the following commands:

```
snapmirror break drfiler:/vol/vol2/qtree1
```

```
snapmirror resync -S drfiler:/vol/vol2/qtree1 prfiler:/vol/vol2/qtree1
```

- Create a new vFiler unit on the original storage system (or its replacement).

The original vFiler unit is now reactivated on the original storage system or its replacement.

- Reinstate the disaster recovery vFiler unit on the disaster recovery storage system:

If the storage elements for the vFiler unit are...

Volumes	On the disaster recovery storage system, enter the following command: <code>vfiler dr resync [-l authinfo] [-a alt-remote, alt-local] [-c secure] [-s] original_vfilername@original_filer</code>
Qtrees	<p>a. On the disaster recovery storage system, destroy the disaster recovery vFiler unit by entering the following command on the disaster recovery hosting storage system:</p> <p style="padding-left: 40px;"><code>vfiler destroy vfilername</code></p> <p><code>vfilername</code> is the name of the original vFiler unit.</p> <p>b. On the disaster recovery storage system, re-create the disaster recovery vFiler unit.</p>

Related concepts

[Disaster recovery using MultiStore](#) on page 77

Related tasks

[Creating a disaster recovery vFiler unit](#) on page 84

[Re-creating the vFiler unit on a replacement storage system](#) on page 98

[Checking and preparing the storage system](#) on page 77

[Checking the network](#) on page 80

Re-creating the vFiler unit on a replacement storage system

You can re-create the original vFiler unit on a replacement storage system if the original storage system is damaged beyond repair.

Steps

1. Boot the replacement storage system.
2. Stop the disaster recovery vFiler unit by using the `vfiler stop` command.
3. Prepare the new vFiler unit on the original storage system (or its replacement).
4. To update the data on the original storage system for each volume and qtree owned by the new vFiler unit, enter the following commands on the original storage system (or its replacement):

`snapmirror break name`

`snapmirror resync -s disaster_recovery_filer:name production_filer:name`

`disaster_recovery_filer` is the name of the disaster recovery storage system.

production_filer is the name of the original storage system.

name is the volume name or path name of the qtree.

Example

For the volume *vol1*, enter the following commands:

```
snapmirror break drfiler:vol1
snapmirror resync -S drfiler:vol1 prfiler:vol1
```

Example

For the qtree *qtree1*, enter the following commands:

```
snapmirror break drfiler:/vol/vol2/qtree1
snapmirror resync -S drfiler:/vol/vol2/qtree1 prfiler:/vol/vol2/qtree1
```

5. Create the new vFiler unit on the original storage system (or its replacement).

The original vFiler unit is now reactivated on the original storage system or its replacement.

6. Reinstate the disaster recovery vFiler unit on the disaster recovery storage system:

If the storage Then...
 elements for
 the vFiler unit
 are...

Volumes On the disaster recovery storage system, enter the following command:

```
vfiler dr resync [-l authinfo] [-a alt-remote, alt-local] [-c secure] [-s] original_vfilername@original_filer
```

Qtrees

- On the disaster recovery storage system, destroy the disaster recovery vFiler unit by entering the following command on the disaster recovery hosting storage system:


```
vfiler destroy vfilername
```

vfilername is the name of the original vFiler unit.
- On the disaster recovery storage system, re-create the disaster recovery vFiler unit.

Related concepts

[Disaster recovery using MultiStore](#) on page 77

Related tasks

[Reactivating the original vFiler unit by using vfiler dr commands](#) on page 96

[Creating a disaster recovery vFiler unit](#) on page 84

[Resynchronizing the vFiler unit](#) on page 90

Data migration using MultiStore

MultiStore enables you to migrate data from one storage system to another without extensive reconfiguration on the destination storage system. Migration moves a specified vFiler unit from a remote storage system to a local one. Migration is initiated on the destination storage system that hosts the vFiler unit after the migration.

Migrating data across storage systems enables you to manage the workload efficiently. Migration automatically destroys the source vFiler unit and activates the destination vFiler unit. The destination then starts serving data to its clients automatically. Only the vFiler unit configuration is destroyed on the source, not the data contained in the vFiler unit.

When performing offline migration, static route entry can be added, if required, because static routing information is not carried to the destination storage system.

You can add static route entry, if required, because static routing information is not carried to the destination storage system.

However, when performing online migration, static routing information is carried to the destination storage system.

MultiStore supports both online migration and offline migration. Offline migration is the default method of migration.

Related tasks

[Checking and preparing the storage system](#) on page 77

[Checking the network](#) on page 80

Secure communication for data migration

Configuring data migration with MultiStore requires communication between the source and the destination storage systems. During data migration, commands are sent from the destination storage system to the source storage system, and are authenticated by an administrative user name and password.

By default, the commands, and the authentication user name and password are sent in cleartext by using the RSH protocol.

To enable secure communication, you must enable SSL by using the `secureadmin` command on the default vFiler unit (vfiler0) context. In addition, you must use the `-c secure` option of the `vfiler migrate` command to send the commands, and the authentication user name and password by using the encrypted SSL protocol.

Note: When you use encrypted SSL protocol, ensure the following requirements are met:

- SSL is enabled on the source vFiler unit.
- The `httpd.enable`, `httpd.admin.enable`, and `httpd.admin.ssl.enable` options are turned on at the source vFiler unit.

SSL is not supported from the nondefault vFiler unit context.

How migrating a vFiler unit affects clients

When you migrate a vFiler unit to another storage system on the same subnet, you must reconnect CIFS clients. If the vFiler unit owns qtree resources and not the volumes in which qtree resides, you must remount all NFS exports on these qtrees.

When the vFiler unit owns the full volumes, NFS mounts persist the migration.

You can continue to access LUNs without any interruption. Although an iSCSI host is briefly disconnected from the source vFiler unit, an initiator hides this brief disruption from applications accessing the LUNs.

Offline migration of vFiler units

Data ONTAP supports two methods of offline migration. Data from a source vFiler unit is copied to a destination vFiler unit by using SnapMirror. Alternatively, vFiler units can be migrated between the storage systems in an HA pair by using the `vfiler migrate -m nocopy` command. During offline migration, users cannot access data from the vFiler units.

The period of time from when the source vFiler unit becomes inaccessible and the destination vFiler unit becomes available is called the cutover period.

The cutover period duration increases with the number of volumes owned by the vFiler unit. Because users and other applications cannot access the vFiler unit that is being migrated during the cutover period, offline migration is disruptive.

Note: Offline migration is not allowed if the source vFiler unit has an active interactive SSH session. You must terminate the active interactive SSH session before initiating offline migration.

Related concepts

[Considerations for online migration of vFiler units](#) on page 108

The `vfiler migrate` commands

You use two commands to migrate a vFiler unit by using the SnapMirror feature: `vfiler migrate start` and `vfiler migrate complete`.

The `vfiler migrate start` command does the following:

- Checks if the destination storage system can receive the source data.

- Configures and runs SnapMirror to copy the data from the source to the destination vFiler unit.
- Saves the quota information from the `/etc/quotas` file of the source vFiler unit.

The `vfiler migrate complete` command does the following:

- Stops the source vFiler unit.
- Updates the data on the destination vFiler unit.
- Breaks the SnapMirror relationships.
- Configures IP bindings on the destination vFiler unit according to the information you provided when running the `vfiler migrate start` command, and adds the destination IP information to the destination `/etc/rc` file.

Any IP information that pertains only to the source vFiler unit is removed from the destination `/etc/rc` file.

- Configures any quota information saved by the `vfiler migrate start` command.
- Destroys the source vFiler unit.
- Brings LUNs online by using the migrated LUN maps and igroups.

Note: vFiler unit migration fails to start if there is an active interactive SSH session on the source vFiler unit. You must terminate the SSH session before starting vFiler unit migration.

Migrating a vFiler unit by copying data

You might want to copy data from one storage system to another for moving the workload from an older vFiler unit that is to be replaced.

Before you begin

- Your storage systems and network must be ready for migration.
- You must have stopped the FlexClone file and FlexClone LUN operations running on the storage unit from the nondefault vFiler unit context.
- You must have prepared the destination storage system.
- SnapMirror technology must have been licensed and enabled on both the source and the destination storage systems.
- The source and destination storage systems must communicate with each other over the network (for example, by means of DNS lookup or entries in the `/etc/hosts` file).
- The destination volumes must be online.
- To log in to the source storage system, you must know the administrative user ID and password.

About this task

Attention: This procedure destroys the original vFiler unit after replicating it on the destination storage system.

Steps

1. On the destination storage system, enter one of the following commands:

If the time by when the console is locked...	Then...
Is important (for example, if you want to lock the console for a minimum amount of time)	Enter the following command: vfiler migrate start source_vfiler@source_filer
Is not important (for example, if you want to migrate the vFiler unit overnight)	Enter the following command: vfiler migrate source_vfiler@source_filer Note: If you use this command, skip Step 4 and Step 5. For more information, see the <code>na_vfiler(1)</code> man page.

2. Respond to the login prompt with a valid administrative ID and password for the source storage system.
3. Respond to the IP address and binding prompts.
4. Monitor the progress of the migration by using the following command:
vfiler migrate status source_vfiler@source_filer
Note: The `vfiler migrate` command might take some time to complete, especially if a source qtree has many millions of inodes.
5. When the status command reports that SnapMirror has replicated all the storage units of the source vFiler unit, you can either complete the migration or cancel the migration:

If you want to... Then...	
Complete the migration	Enter the following command: vfiler migrate complete [-l user:passwd] [-c secure] source_vfiler@source_filer Note: You have to use the user name and password when the password is changed at the remote storage system after <code>vfiler migrate start</code> is started.
Cancel the migration	Enter the following command: vfiler migrate cancel source_vfiler@source_filer This destroys the destination vFiler unit and removes the SnapMirror and other migration-related configuration information from the source vFiler unit.

6. If you copied quota information to the destination storage system's `/etc/quotas` file when you prepared the destination storage system, activate the quotas on that storage system. For activating quotas on each of the volumes, use the following command:

quota on *volume_name*

7. Remount the qtrees without changing volume-level mounts.

If you have moved the vFiler unit to a different subnet, CIFS domain, or Windows domain, you must rerun the CIFS setup. Also, you have to make adjustments on the clients and modify data-file security attributes.

8. Reconfigure iSCSI authentication if the vFiler unit's storage units contain iSCSI LUNs.

For instructions, see the *Data ONTAP SAN Administration Guide for 7-Mode*.

Related tasks

[Checking and preparing the storage system](#) on page 77

[Checking the network](#) on page 80

[Adjusting client and network configurations if migrating to a different subnet](#) on page 104

[Creating a disaster recovery vFiler unit](#) on page 84

Related references

[Storage checklist](#) on page 79

[Network checklist](#) on page 83

Adjusting client and network configurations if migrating to a different subnet

If you have moved a vFiler unit to a different subnet, you might have to configure the network servers and make adjustments on the clients. If you have moved the vFiler unit to a different Windows domain, you might also have to modify data-file security attributes, and run CIFS setup.

About this task

You must perform this task on the destination vFiler unit.

Steps

1. To configure NIS and DNS servers, run `setup`.
2. To change the name of the Windows domain controller, use the `cifs prefdc` command.
3. To change the Windows WINS server, run the `cifs setup` command.

Note: If the Windows domain has changed, you might have to change the permissions on the Windows data files to allow your users the same access they had in the old domain.

4. To change the trusted host, perform the following steps:
 - a) Edit the vFiler unit's `/etc/hosts.equiv` file, adding the name of the trusted host for administering the vFiler unit.

- b) Add the path to the root volume and the name of the trusted host to the vFiler unit's /etc/exports file.

Example

```
/vol/vol0 access=adminhost, root=adminhost
```

5. Remount volumes and qtrees on the clients.

Related references

[Network checklist](#) on page 83

vFiler unit migration without copying data

vFiler unit migration in an HA pair is the no-copy transfer of a volume-level vFiler unit from the source node to the other. vFiler unit migration in an HA pair uses software-based disk ownership to transfer ownership of the aggregate that contains the vFiler unit from the original source node to the destination node in an HA pair.

The migration operation is quickly completed as the vFiler unit migration is performed through transfer of disk ownership rather than by copying data from one set of disks to another.

Prerequisites for vFiler unit migration between the nodes of an HA pair

Before you use the `vfiler migrate -m nocopy` command to migrate a vFiler unit between the nodes of an HA pair, you must ensure that the nodes in the HA pair meet certain requirements, and also ensure that the required licenses are enabled.

The prerequisites are as follows:

- MultiStore must be enabled on each node of the configuration.
- Any license that is enabled on the source node must also be enabled on the destination node. For example, if CIFS is licensed on the source node of the HA pair, CIFS must also be licensed on the destination node of the HA pair. Otherwise, moving the vFiler unit causes CIFS to be unavailable for that vFiler unit.
- Both the nodes of the HA pair must be connected correctly to the storage shelves in the HA pair. The disks must be visible to the nodes of the source and destination HA pair.
- The destination node must have an Ethernet connection to the same subnet that the source node uses to ensure that software-based disk ownership changes are transparent to NFS users.
- The volumes assigned to the vFiler unit can either be traditional volumes or FlexVol volumes. If the volumes are FlexVol volumes, the containing aggregate must contain only volumes belonging to the migrating vFiler unit. For information about traditional and FlexVol volumes, see the *Data ONTAP Storage Management Guide for 7-Mode*.
- The vFiler unit's storage units must all be composed of complete volumes; that is, the vFiler unit's paths must use the form `/vol/volname`. Migration of storage units that name specific volume subdirectories—for example, `/vol/volname/qtree`—is not supported.
- The volume containing the configuration information of the vFiler unit (`/vol/volname`) must be writable.

- The destination storage system must not contain aggregates with the same names as the aggregates in the source storage system.

For more information about HA pair, see the *Data ONTAP High-Availability and MetroCluster Configuration Guide for 7-Mode*.

Related information

N series support website: www.ibm.com/storage/support/nseries/

Guidelines for setting up volumes to support vFiler unit migration in an HA pair

When a vFiler unit is migrated, all volumes associated with that vFiler unit are moved. vFiler unit migration in an HA pair cannot migrate a subset of the volumes that are managed by the vFiler unit it is migrating.

When you create vFiler units, keep in mind the following points:

- The destination node of the HA pair must be able to accommodate the vFiler unit with all its associated volumes.
- The names of the vFiler units and volumes being moved from the source to the destination must be unique on the destination.

Although you can rename a volume at the destination, it is best not to do so. This is because if you have NFS clients, renaming of the volume is not transparent to the NFS clients. When the storage system uses NFS to export a file system, the volume name is part of the exported path name. NFS clients try to mount by using the old path name. Therefore, to access the data after the vFiler unit has been migrated, clients must remount by using the new path name.

vFiler unit migration in an HA pair

To perform a vFiler unit migration in an HA pair, you can use the `vfiler migrate -m nocopy` command.

The `vfiler migrate -m nocopy` command performs the following:

- Verifies that no vFiler unit with the same name exists on the destination node of the HA pair.
- Verifies that both the source and destination nodes in the HA pair run the same version of Data ONTAP.
- Saves the IP configuration and binding information that you supplied when you created the vFiler unit.
- Saves the quota information from the source vFiler unit's `/etc/quotas` file.
- Stops the source vFiler unit.
- Destroys the source vFiler unit.
- Rewrites the disk ownership information so that the ownership of the vFiler unit volumes is transferred from the source node in the HA pair to the destination node in the HA pair.
- Re-creates the vFiler unit on the destination node in the HA pair.

Migrating a vFiler unit by using the `vfiler migrate -m nocopy` command

If many clients are using the same vFiler unit, thereby affecting the performance of its hosting node, and the other node of the HA pair is lightly loaded, you can transfer ownership of that vFiler unit to the hosting node's HA pair partner to balance the load processing on the two nodes.

Before you begin

You must ensure that there are no FlexClone file and FlexClone LUN operations running on the storage unit from the nondefault vFiler unit context.

Steps

1. On the destination node in the HA pair, enter the following command:

```
vfiler migrate -m nocopy vfilername@source_c1_partner
```

vfilername is the name of the vFiler unit that you are migrating.

source_c1_partner is the HA pair from which you are moving the vFiler unit.

For detailed information, see the `na_vfiler(1)` man page.

2. Answer the prompts, including the following information:

- A valid administrative login ID and password
- The IP address and binding information for the destination vFiler unit

The vFiler unit is migrated from the source node HA pair to the destination node in the HA pair.

3. Verify that the vFiler unit was moved by entering the following command on the destination node in the HA pair:

```
vfiler status -r vfilername
```

What IBM N series Data Motion for vFiler is

IBM N series Data Motion for vFiler is a data migration solution that integrates virtual storage, mirroring, and provisioning software technologies so that you can perform migrations nondisruptively in both physical and virtual environments.

By using the N series Management Console provisioning capability interface, you can migrate data from one storage system to another, as long as the data is contained in vFiler units and associated with datasets. Migration operations are performed transparently, so users are unaware of the migration operation being performed. The operations are also nondisruptive, so users retain access to migrated data, and the hosts and applications that access the migrated data do not require reconfiguration.

The application interfaces and documentation commonly refer to the Data Motion for vFiler capability as "online migration," "online dataset migration," "transparent migration", or "online vFiler unit migration."

Features supported by Data Motion for vFiler

The Data ONTAP storage efficiency features, such as deduplication, data compression, and FlexClones are supported with Data Motion for vFiler. Data Motion for vFiler uses Sync SnapMirror and SnapMirror as underlying technologies.

Considerations for online migration of vFiler units

There are certain considerations when migrating data from a source vFiler unit to the destination vFiler unit by using online migration.

You must be aware of the following considerations when performing online migration:

- The source and the destination storage systems should be in the same subnet.
- You can continue to run applications on vFiler units during online migration.
- Online migration is not allowed if the source vFiler unit has an active, interactive SSH session. You must terminate the active, interactive SSH session before initiating online migration.
- Online migration is not supported on the vFiler units that are configured with IPv6 addresses.
- Online migration is nondisruptive only for vFiler units configured to use NFS or iSCSI protocols.
- If the vFiler unit is configured to use CIFS, then before starting online migration, the provisioning application terminates the CIFS connection.

CIFS clients must reconnect all terminated CIFS connections after the completion of online migration.

- Online migration is supported only at the volumes and FlexClone levels. You cannot perform online migration at a qtree level.
- Online migration is not supported between vFiler units containing 32-bit volumes and 64-bit volumes.
- Online migration supports replication of similar volumes by using SnapMirror as the underlying technology.

vFiler units can have a combination of volumes contained in 32-bit and 64-bit aggregates. However, the source and destination vFiler units must have the volumes in similar aggregate types for online migration.

For example, if the source vFiler unit has one volume on a 32-bit aggregate and one volume on a 64-bit aggregate, then the destination storage system must have one 32-bit aggregate and one 64-bit aggregate for online migration. If the source vFiler unit has only 32-bit aggregates, the destination vFiler unit's volumes must be contained in 32-bit aggregates.

Online migration is supported only when the vFiler unit's 32-bit volumes are contained in 32-bit aggregates and 64-bit volumes are contained in 64-bit aggregates.

- Online migration is not supported during aggregate expansion from 32-bit to 64-bit on the source storage system or destination storage system.
- The size of the migrating volumes should be equal to or greater than 10 GB.

- The static routing information is carried to the destination storage system.

For more information about online migration, see the *OnCommand Unified Manager Guide to Common Provisioning and Data Protection Workflows for 7-Mode*.

Related concepts

[Offline migration of vFiler units](#) on page 101
[Data migration using MultiStore](#) on page 100

Related tasks

[Checking and preparing the storage system](#) on page 77
[Checking the network](#) on page 80

Related references

[Storage checklist](#) on page 79
[Network checklist](#) on page 83

Related information

[N series support website: \[www.ibm.com/storage/support/nseries/\]\(http://www.ibm.com/storage/support/nseries/\)](#)

Option required for online migration

You must enable the `httpd.admin.enable` option in the source and the destination storage systems to perform online migration of vFiler units.

Stages of a vFiler unit migration

Online migration of vFiler units consists of three stages: vFiler unit migrate start, vFiler unit migrate prepare, and vFiler unit migrate complete. You can view the different states of the migrating vFiler unit by using the `vfiler status` command or by viewing messages on the console of the hosting storage system.

vFiler unit migrate start stage

A SnapMirror relationship is established and baseline transfer is initialized during the vFiler unit migrate start stage.

The vFiler unit in the source storage system is in the `running` state and is accessible by users connected to the vFiler unit.

Also, a temporary vFiler unit is created in the destination storage system. This temporary vFiler unit is in the `stopped, migrating` state.

vFiler unit migrate prepare stage

During the vFiler unit migrate prepare stage, the source and destination vFiler units are prepared for the final stage of migration. The vFiler unit in the source storage system is accessible by users connected to the vFiler unit and is in the `running, migrate_prepare` state.

The temporary vFiler unit that was created in the destination storage system during the previous stage is in the `stopped, migrating` state.

In the vFiler unit migrate prepare stage, MultiStore verifies the following:

- If there is any active interactive SSH session to the migrating vFiler unit.
- If the vFiler context of the migrating vFiler unit is active.
- The SnapMirror relationship between the source and the destination is synchronized.
- The aggregates that own the volumes assigned to the vFiler unit are not in degraded or reconstructing state at the source or at the destination storage system.
- The destination or source storage system is in takeover mode. If either of the storage systems is in takeover mode, then the migrate complete stage returns an error message.
- The destination or source storage system is in the process of being taken over.
- Volumes are added or removed from source storage system after the vFiler unit migrate start stage.
- The volumes owned by the vFiler unit being migrated are online.
- The source storage system has no other vFiler units in the cutover stage of the migration.

vFiler unit migrate complete stage

In the vFiler unit migrate complete stage, MultiStore verifies the following:

- The SnapMirror relationship between the source and the destination is synchronized.
- The aggregates that own the volumes assigned to the vFiler unit are not in degraded or reconstructing state at the source or at the destination storage system.
- The destination or source storage system is in takeover mode. If either of the storage systems is in takeover mode, then the migrate complete stage returns an error message.
- The destination or source storage system is in the process of being taken over.
- Volumes are added or removed from the source storage system after the vFiler unit migrate start stage.
- The volumes owned by the vFiler unit being migrated are online.
- The source storage system has no other vFiler units in the cutover stage of the migration.

The vFiler unit in the source storage system is now in the `stopped, migrating_source` state and is not accessible by users connected to the vFiler unit.

The temporary vFiler unit that was created in the destination storage system is now in the `running` state. This vFiler unit is now accessible by the users who were connected to the source vFiler unit. Users do not have to reconfigure or restart any applications they were running on the original vFiler unit.

How to perform online migration of vFiler units

You can perform online migration of vFiler units only from the N series Management Console provisioning capability. This capability performs premigration storage checks, postmigration tasks, and automates features such as destination volume provisioning, network provisioning, rapid rollback, and rapid restart of online migration.

When you start online migration from the N series Management Console provisioning capability, all SnapMirror, SnapVault, NDMP, dump and restore operations that are currently running on the vFiler unit that is being migrated are aborted before cutover is initiated by the N series Management Console provisioning capability. If there are any currently running deduplication, FlexClone file and LUN, or LUN clone split operations, online migration fails. You must retry online migration after these operations are complete. New SnapMirror and SnapVault operations are not allowed.

If online migration fails, error messages are displayed on the console of the hosting storage system. You can also view the log of error messages in the `/etc/messages` file of the hosting storage system.

Before you start online migration from the N series Management Console provisioning capability, it is best to configure static routes on the destination storage system, when the source storage system has more than 20 routes.

For more information about using the N series Management Console provisioning capability to perform online migration, premigration storage checks, and postmigration tasks, see the *OnCommand Unified Manager Guide to Common Provisioning and Data Protection Workflows for 7-Mode*.

Viewing the status of a vFiler unit migration

You can view the messages on the CLI of the hosting storage system or the default vFiler unit `vfiler0` to check the status of the migrating vFiler unit. You can also use the `vfiler status` command to view the status of the source or destination vFiler units.

Step

1. Enter the following command in the CLI of the default vFiler unit `vfiler0`:

```
vfiler status vfiler_name
```

`vfiler_name` is the name of the vFiler unit.

Example

Assume a vFiler unit `vfilerA` is being migrated. The following command shows the status of the vFiler unit `vfilerA`:

```
vfiler status vfilerA
```

The output of this command shows the migration stage of the vFiler unit `vfilerA` at that point in time. For example, if the vFiler unit is in the migrate prepare stage, the `vfiler status vfilerA` command shows the following output:

```
vfiler status -r vfilerA
vfilerA      running, migrate_prepare
  ipspace: default-ipspace IP address: 12.21.1.1 [e0a]
  Path: /vol/vol_test1 [/etc]
  Path: /vol/smvol4
  UUID: 679b86c0-be3a-11de-aee6-00a09807609b
```

Commands not allowed during the cutover phase of online migration

During the cutover phase of online migration of a vFiler unit, you must not run some of the Data ONTAP commands for the vFiler unit that is being migrated from the CLI of the hosting storage system or the default vFiler unit vfiler0.

MultiStore commands that are not allowed during the cutover phase of online migration

- vfiler add
- vfiler allow
- vfiler context
- vfiler disallow
- vfiler destroy
- vfiler dr resync
- vfiler move
- vfiler remove
- vfiler rename
- vfiler run
- vfiler start
- vfiler stop

Volume and aggregate commands that are not allowed during the cutover phase of online migration

- vol copy
- vol rename
- vol online
- vol offline
- vol restrict
- vol options
- vol size
- vol clone create
- vol clone split
- aggr split

For more information about volume and aggregate commands, see the *Data ONTAP Storage Management Guide for 7-Mode*.

Snapshot commands that are not allowed during the cutover phase of online migration

- snap create
- snap delete
- snap list
- snap delta
- snap rename
- snap sched
- snap restore
- snap reclaimable
- snap reserve
- snap autodelete

For more information about Snapshot commands, see the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode*.

SnapMirror commands that are not allowed during the cutover phase of online migration

- snapmirror resync
- snapmirror initialize

For more information about SnapMirror commands, see the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode*.

Deduplication and data compression commands that are not allowed during the cutover phase of online migration

- sis on
- sis off
- sis config -s
- sis start
- sis start -s

For more information about deduplication commands, see the *Data ONTAP Storage Management Guide for 7-Mode*.

Tape backup and NDMP commands that are not allowed during the cutover phase of online migration

- dump
- restore
- ndmpcopy

For more information about NDMP commands, see the *Data Protection Tape Backup and Recovery Guide for 7-Mode*.

Related concepts

[Data access from the hosting storage system](#) on page 16

Related tasks

[Moving resources between vFiler units](#) on page 29

[Adding resources to a vFiler unit](#) on page 27

[Destroying a vFiler unit](#) on page 34

[Renaming a vFiler unit](#) on page 32

[Starting a vFiler unit](#) on page 36

[Stopping a vFiler unit](#) on page 33

[Allowing a protocol on a vFiler unit](#) on page 37

[Disallowing a protocol on a vFiler unit](#) on page 38

[Resynchronizing the vFiler unit](#) on page 90

Target portal group management for online migration of vFiler units

Target portal groups enable you to efficiently manage iSCSI sessions between initiators and targets. Although Data ONTAP manages target portal groups by network interface by default, you can also use IP address, starting with Data ONTAP 7.3.3. This is required if you want to perform an online migration of vFiler units, which allows you to nondisruptively migrate data from one storage system to another.

Note: The N series Management Console provisioning capability is required for performing online migrations of vFiler units.

When you migrate data, the target portal group tag on the destination network interface must be identical to the target portal group tag on the source network interface. This is problematic in a MultiStore environment because the source and destination storage systems might be of different hardware platforms. Changing the target portal group tags after migration is not sufficient because some hosts, such as HP-UX and Solaris, do not support dynamic iSCSI target discovery, resulting in a disruption of service to those hosts in the process.

Note: If offline (disruptive) migrations are not problematic in your environment, or if all of your hosts support dynamic iSCSI target discovery, then IP-based target portal group management is unnecessary.

If you choose to implement IP-based target portal groups by enabling the `iscsi.ip_based_tpgroup` option, interface-based target portal groups are automatically converted to IP-based target portal groups, and any future target portal group assignments are IP-based as well. However, note that if you are migrating between a system with IP-based target portal

groups and a system with interface-based target portal groups, the target portal group information is lost and the iSCSI service might be disrupted.

Note: ALUA is not supported with IP-based target portal groups.

For more information on target portal groups, see the *Block Access Management Guide for iSCSI and FC*.

For more information about the N series Management Console provisioning capability, see the *Provisioning Manager and Protection Manager Guide to Common Workflows for Administrators*.

Related information

Data ONTAP documentation on the N series support website: www.ibm.com/storage/support/nseries/

Data migration implications for IP-based target portal group management

You must understand how target portal groups are affected before beginning the migration.

The following table describes the various online migration scenarios and how they affect your target portal group configurations:

Online migration scenario	Impact to target portal groups
Migration between two systems running Data ONTAP 7.3.2 and earlier.	No change Existing interface-based target portal groups are not converted to IP-based target portal groups.
Migration between a system running Data ONTAP 7.3.2 or earlier and a system running 7.3.3 or later.	No change Existing interface-based target portal groups are not converted to IP-based target portal groups.
Migration between two systems running Data ONTAP 7.3.3 or later with interface-based target portal groups.	No change Existing interface-based target portal groups are not converted to IP-based target portal groups.
Migration between two systems running Data ONTAP 7.3.3 or later with IP-based target portal groups.	IP-based target portal group assignments are preserved and there is no disruption in the iSCSI service.
Migration between a system running Data ONTAP 8.1 or later with interface-based target portal groups and a system running 7.3.2 or earlier with interface-based target portal groups.	No change Target portal group assignments are preserved.
Migration between a system running Data ONTAP 8.1 or later with IP-based target portal groups and any system with interface-based target portal groups.	Target portal group information is lost and the iSCSI service might be disrupted.

Related information

N series support website: www.ibm.com/storage/support/nseries/

Enabling IP-based target portal group management

If you want to perform online migrations of vFiler unit, you must enable IP-based target portal groups on your vFiler units.

About this task

When you enable IP-based target portal groups, the existing interface-based target portal groups are automatically converted to IP-based target portal groups. However, note that the interface-based target portal groups remain intact for the default vFiler unit.

Step

1. Enter the following command:

```
vfiler run vFiler_unit options iscsi.ip_based_tpgroup on
```

The existing interface-based target portal groups are converted to IP-based target portal groups with no disruption in service to the host.

Example

Before enabling IP-based target portal groups, the target port group information for vFiler unit 2 (vf2) looks like this:

```
system1>vfiler run vf2 iscsi tpgroup show
TPGTag  Name          Member Interfaces
32      user_defined32 (none)
1000    e0_default     e0
1002    e11b_default   e11b
1003    e11c_default   e11c
1004    e11d_default   e11d
1005    e9a_default    e9a
1006    e9b_default    e9b
1007    e10a_default   e10a
1008    e10b_default   e10b
2000    vif_e0-1_default vif_e0-1
2001    vif_e0-2_default vif_e0-2
2002    vif_e0-3_default vif_e0-3
2003    vif_e11a-1_default vif_e11a-1
2004    vif_e11a-2_default vif_e11a-2
2005    vif_e11a-3_default vif_e11a-3
```

Each interface is associated with various IP addresses, and some of those are assigned to vFiler unit vf2. For example:

```
system1> vfiler run vf2 iscsi portal show
Network portals:
IP address      TCP Port  TPGroup  Interface
```

10.60.155.104	3260	1000	e0
192.168.11.100	3260	2003	vif_ella-1
192.168.11.101	3260	2003	vif_ella-1
192.168.13.100	3260	2005	vif_ella-3
192.168.13.101	3260	2005	vif_ella-3

After enabling IP-based target portal groups for vf2, the relevant interface-based target portal groups for vf2 are nondisruptively converted to IP-based target portal groups.

```
system1> vfiler run vf2 options iscsi.ip_based_tpgroup on

system1> vfiler run -q vf2 iscsi ip_tpgroup show
TPGTag  Name                      Member IP Addresses
1000    e0_default                10.60.155.104
2003    vif_ella-1_default        192.168.11.100, 192.168.11.101
2005    vif_ella-3_default        192.168.13.100, 192.168.13.101

system1> vfiler run -q vf2 iscsi portal show
Network portals:
IP address      TCP Port  TPGroup  Interface
10.60.155.104  3260     1000     e0
192.168.11.100 3260     2003     vif_ella-1
192.168.11.101 3260     2003     vif_ella-1
192.168.13.100 3260     2005     vif_ella-3
192.168.13.101 3260     2005     vif_ella-3
```

If you configure another IP address for vf2, then a new default IP-based target portal group (4000) is automatically created. For example:

```
system1> vfiler add vf2 -i 192.168.13.102

system1> ifconfig vif_ella-3 alias 192.168.13.102

system1> vfiler run vf2 iscsi ip_tpgroup show
TPGTag  Name                      Member IP Addresses
1000    e0_default                10.60.155.104
2003    vif_ella-1_default        192.168.11.100, 192.168.11.101
2005    vif_ella-3_default        192.168.13.100, 192.168.13.101
4000    192.168.13.102_default   192.168.13.102

system1> vfiler run vf2 iscsi portal show
Network portals:
IP address      TCP Port  TPGroup  Interface
10.60.155.104  3260     1000     e0
192.168.11.100 3260     2003     vif_ella-1
192.168.11.101 3260     2003     vif_ella-1
192.168.13.100 3260     2005     vif_ella-3
```

192.168.13.101	3260	2005	vif_e11a-3
192.168.13.102	3260	4000	vif_e11a-3

After you finish

After you enable IP-based target portal group management, it is recommended to leave it enabled. However, if you must disable IP-based target portal groups for some reason, enter the following command:

```
options iscsi.ip_based_tpgroup off
```

As a result, any IP-based target portal group information is discarded, and the interface-based target portal group information is reenabled. Note that this process might disrupt the iSCSI service to the hosts.

Also note that if an IP address is unassigned from a vFiler unit or unconfigured from the network interface, that IP address is no longer a valid iSCSI portal. However, the IP-based target portal group to which that IP address belonged remains intact so that if you add the IP address back later, it is automatically assigned back to the original target portal group.

Creating IP-based target portal groups

You can create new IP-based target portal groups in which to add and remove existing IP addresses.

Before you begin

IP-based target portal group management must be enabled by entering the following command:

```
options iscsi.ip_based_tpgroup on
```

Step

1. Enter the following command:

```
vfiler run vFiler_unit ip_tpgroup create [-f] [-t | tag] tpgroup_name IP address...
```

-f forces the new group to be created, even if that terminates an existing session using one of the IP addresses being added to the group.

-t sets the target portal group tag to the specified value. In general, you should accept the default tag value.

tpgroup_name is the target portal group name.

IP address is the list of IP addresses to include in the group, separated by spaces.

Example

```
vfiler run vfiler2 iscsi ip_tpgroup create -t 233 vfiler2_tpg1
10.1.3.5
```

After you finish

You can add and remove IP addresses from the new group.

Adding IP addresses to IP-based target portal groups

You can use the `iscsi ip_tpgroup add` command to add an IP address to an existing IP-based target portal group.

Before you begin

- IP-based target portal group management must be enabled.
- There must be at least one existing IP-based target portal group.

Step

- Enter the following command:

```
vfiler run vFiler_unit iscsi ip_tpgroup add [-f] tpgroup_name IP address ...
```

`-f` forces the new group to be created, even if that terminates an existing session using one of the IP addresses being added to the group.

`tpgroup_name` is the target portal group name.

`IP address` is the list of IP addresses to include in the group, separated by spaces.

Example

```
vfiler run vfiler2 iscsi ip_tpgroup add vfiler2_tpg1 192.168.2.1 192.112.2.1
```

Removing IP addresses from IP-based target portal groups

In the course of reconfiguring your network, you might need to remove one or more IP addresses from an IP-based target portal group.

Step

- Enter the following command:

```
vfiler run vFiler_unit iscsi ip_tpgroup remove [-f] tpgroup_name IP address ...
```

`-f` forces the new group to be created, even if that terminates an existing session using one of the IP addresses being added to the group.

`tpgroup_name` is the target portal group name.

`IP address` is the list of IP addresses to remove from the group, separated by spaces.

Example

```
vfiler run vfiler2 iscsi ip_tpgroup remove vfiler2_tpg1 192.112.2.1
```

Destroying IP-based target portal groups

If necessary, you can destroy IP-based target portal groups.

Before you begin

No active sessions must be in progress.

Step

1. Enter the following command:

```
vfiler run vFiler unit iscsi ip_tpgroup destroy [-f] tpgroup_name
```

-f forces the group to be destroyed, even if that terminates an existing session using one of the IP addresses in the group.

tpgroup_name is the target portal group name.

The target portal group is destroyed, and if there are active iSCSI sessions, a warning message indicates that those connections are lost.

Example

```
vfiler run vfiler2 iscsi ip_tpgroup destroy vfiler2_tpg1
```

Displaying IP-based target portal group information

You can use the `iscsi ip_tpgroup show` command to display important information about your IP-based target portal groups, including target portal group tags, target portal group names, and the IP addresses that belong to each group.

Step

1. Enter the following command:

```
vfiler run vFiler_unit iscsi ip_tpgroup show
```

Example

```
system1> vfiler run vfiler2 iscsi ip_tpgroup show
TPGTag  Name          Member IP Addresses
      1  vfiler2_migrate_test0  (none)
      2  vfiler2_migrate_test1  (none)
      3  vfiler2_migrate_test3  (none)
    100  user_defined_tpg1    (none)
```

128	vfiler2_ui_review	1.1.1.1
1007	e10a_default	10.1.1.8
1008	e10b_default	1.1.1.2
4000	10.1.1.5_default	10.1.1.5
4001	10.60.155.104_default	10.60.155.104
4002	192.168.1.1_default	192.168.1.1

Disk space management using quotas

You can apply user, group, and qtree quotas on a vFiler unit in the same way that you apply on a storage system. When you create a vFiler unit, quotas are automatically turned off on both the hosting storage system and the new vFiler unit.

However, quotas are turned off for all the volumes that you assign to the new vFiler unit, and for all the volumes from which you assign qtrees to the new vFiler unit.

Quotas are also turned off for volumes that you move from one vFiler unit to another. To activate quotas again, you must allow them and turn them on.

On a hosting storage system licensed for vFiler units, the hosting storage system administrator must allow quotas for a volume before you can turn on quotas or turn off quotas for the volume. By default, quotas are allowed on all volumes. Only hosting storage system administrators can allow or disallow quotas for a volume.

Allowing or disallowing quotas for a volume

As a hosting storage system administrator, you must allow quotas for a volume on the hosting storage system before you can turn on or turn off quotas. By default, quotas are allowed on all volumes.

Step

1. To allow or disallow quotas for a volume, enter the following command:

```
quota allow | disallow volume
```

volume is the name of the volume for which you want to allow or disallow the quota.

After you enter the `quota allow` command, you can turn on quotas for the specified volume from a vFiler unit.

After you enter the `quota disallow` command, vFiler units are prevented from turning quotas on for the specified volume. If quotas are currently turned on for any volume in vFiler units, they are turned off immediately.

Result

If you disallow quotas on a volume, the following effects occur on all vFiler units that have storage units in the volume:

- If quotas are currently turned off, you or the vFiler unit administrator cannot turn on quotas for that volume.
- If quotas are currently turned on, they are turned off immediately and cannot be turned back on.

Quota specification management

The vFiler unit administrator specifies the size of each quota in the vFiler unit's `/etc/quotas` file. The vFiler unit administrator tracks and limits the amount of disk space and the number of files each user, group, or qtree uses.

If a qtree owned by the vFiler unit resides on a volume owned by the hosting storage system, then the hosting storage system administrator can also specify a quota for the qtree in the hosting storage system's `/etc/quotas` file. The following example shows how qtree quota on the hosting storage system affects a vFiler unit qtree:

How qtree quota on the hosting storage system affects a vFiler unit qtree

Assume that the `/vol/vol1/qtree1` qtree is a storage unit of the vFiler unit, and the `/vol/vol1` volume is owned by the hosting storage system.

In the `/etc/quotas` file of the vFiler unit, the vFiler unit administrator specifies that this qtree is limited to 20 GB of disk space. In the `/etc/quotas` file of the hosting storage system, the storage system administrator can specify the disk space limit for the qtree as 10 GB. Therefore, if quotas are turned on from the hosting storage system for the `/vol/vol1` volume, the qtree cannot exceed the limit in either of the `/etc/quotas` files, whichever is lower. In this example, the qtree cannot exceed 10 GB.

The hosting storage system administrator controls the usage of quotas on each volume that the storage system owns by using the `quota allow` and `quota disallow` commands. When the hosting storage system administrator allows quotas, the vFiler unit administrator can turn quotas on or off on the vFiler units by using the `quota on` and `quota off` commands, respectively.

For more information about quotas, see the *Data ONTAP Storage Management Guide for 7-Mode*.

Turning on or turning off quotas from a vFiler unit

You can turn on quotas by using the `quota on volume` command, and you can turn off quotas by using the `quota off volume` command. The `quota on volume` command activates quotas on the specified volume based on the contents of `/etc/quotas`.

Before you begin

You must have added a valid entry for a newly created volume in the `/etc/quotas` file of the vFiler unit before turning on quotas on that volume.

For more information about quotas, see the *Data ONTAP Storage Management Guide for 7-Mode*.

About this task

Changes made to /etc/quotas do not take effect the next time the quota on or quota resize command is executed. Turning quotas off by using the quota off volume command deactivates quotas on the specified volume. You can turn quotas on and off on a per-volume basis for a vFiler unit. After you turn on quotas for a particular volume, Data ONTAP initializes quotas for the storage units residing on the volume that is owned by the vFiler unit. The on or off states of quotas are persistent and stay set after reboots.

Step

1. To turn quotas on or off for a volume owned by a vFiler unit, follow the instructions appropriate to your situation:

If you manage the vFiler unit	Then...
-------------------------------	---------

The hosting storage system	Enter the following command: <code>vfiler run vfilertemplate quota on off volume</code>
The vFiler unit	Enter the following command through an RSH connection to the vFiler unit: <code>quota on off volume</code>

Note: Whenever a qtree is explicitly reassigned to a vFiler unit, you must reenable the quota manually if quotas are used. Qtrees are explicitly reassigned to vFiler units when you create vFiler units (using the `vfiler create` command) or when you move qtrees between vFiler units (using the `vfiler move`, `vfiler add`, or `vfiler remove` commands).

When quota thresholds and soft quotas are exceeded

When a threshold or soft quota defined on a vFiler unit is exceeded, a warning message is logged on the storage system console.

You see a warning message similar to the following:

```
[vfiler1@quota.softlimit.exceeded:notice]: Threshold exceeded for tree 3 on
volume vol1 for vfiler "vfiler1"
```

How you can resize quotas

When you use the `quota resize` command, Data ONTAP rereads the quotas file for the specified volume. You can resize quotas only for certain types of changes to the quotas file, otherwise, you have to reinitialize quotas.

For more information about resizing quotas, see the *Data ONTAP Storage Management Guide for 7-Mode*.

How the quotas file works

The quotas file, found in the `/etc` directory, contains one or more entries specifying limit or tracking quotas for qtrees, groups, and users. The file can contain default (general) and specific entries.

Displaying the quota status

You can display the quota status for any volume on which your vFiler unit owns storage space.

Step

1. To display the quota status, perform one of the following step:

If you manage the vFiler unit from...	Then...
The hosting storage system	Enter the following command: <code>vfiler run vfilertemplate quota</code>
The vFiler unit	Enter the following command through an RSH connection to the vFiler unit: <code>quota</code>

The command displays quota status information about all the volumes in which the vFiler unit owns storage space. The following is a sample message in the command output:

```
vol0: quotas are on.
vol1: quotas are off.
vol2: quotas are disabled.
```

For more information about quotas, see the *Data ONTAP Storage Management Guide for 7-Mode*.

Displaying a quota report

You display a quota report using the `quota report` command. You can display a quota report for all quotas or for a specific file, directory, qtree or volume by specifying a pathname.

Step

1. To display a quota report, enter the following command:

```
quota report [path]
```

You can display a quota report for all quotas or for a specific file, directory, qtree or volume by specifying a path.

You can control the format and fields displayed using the `quota report` command options. For more information on the available options, see the `na_quota(1)` man page.

Copyright and trademark information

Copyright ©1994 - 2012 NetApp, Inc. All rights reserved. Printed in the U.S.A.

Portions copyright © 2012 IBM Corporation. All rights reserved.

US Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

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.

References in this documentation to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any of IBM's or NetApp's intellectual property rights may be used instead of the IBM or NetApp product, program, or service. Evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM and NetApp, are the user's responsibility.

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.A. 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

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 <http://www.ibm.com/legal/copytrade.shtml>

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

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.

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.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe on any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, N.Y. 10504-1785
U.S.A.

For additional information, visit the web at:
<http://www.ibm.com/ibm/licensing/contact/>

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM web sites are provided for convenience only and do not in any manner serve as an endorsement of those web sites. The materials at those web sites are not part of the materials for this IBM product and use of those web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

Index

/etc/dgateways file 53

/etc/exports file

 exporting all file systems in 73

A

adding resources to a vFiler unit 27

administrators, types of 17

aggr add command

 adding new disks 77

B

backing up a vFiler unit 48

backup vFiler unit 84

C

CIFS

 support for a vFiler unit 37

 local user accounts 75

 path names for shares 72

 statistics 59

 virus scanning 76

cifs setup command 25

cifs stat command 59

clients, effect of vFiler unit move on 101

commands

 ipspace assign 81

 ipspace create 81

 ipspace list 80

 options snapmirror 94

 options snapmirror.access 94

 vfiler dr configure 80, 86

 vfiler migrate 82

 ypwhich 80

commands for a vFiler unit

 vfiler command, purpose of 22

 entering through rsh 44

 vfiler dr activate command 88

 vfiler migrate start command 87, 101

 vfiler move command 29

 vfiler remove command 28

 vfiler setup command 25

 vfiler start command 36

vfiler status command 39, 77

configuring

 virus scanning 75, 76

consolidating multiple servers 12

create a vFiler unit

 prerequisites 19

Creating

 IPspace 66

creating a vFiler unit

 in nondefault IPspaces 69

 required status of network interface 23

Creating a vFiler unit 24

creating volumes 78

cutover period 101, 108

D

daemon

 enabling the routed 17

data compression commands 56

data migration

 secure communication for 100

Data Motion

 supported features 108

Data Motion for vFiler

 online migration component 107

deduplication 54

default IPspace, interfaces 63

destination storage system, preparing for migration 77

destroying a vFiler unit 34

destroying IPspaces 69

disabling MultiStore 18

disaster recovery

 quotas 79

 secure communication for 84

 Secure Sockets Layer

See SSL

 vFiler unit 84

displaying

 quota status 79, 125

distinct IP address space 61

domains

 DNS servers and vFiler unit 13

 NIS servers 13

F

FCP LUNs 52
 FilerView 16
 FlexVol volume 46
 FTP support for a vFiler unit 37

H

HA pair
 guidelines 21
 hosting storage system
 access to data on a vFiler unit 16
 manage a vFiler unit 27
 quotas not copied to new vFiler unit 79
 HTTP support for a vFiler unit 37

I

IP addresses
 configuring for vFiler unit creation 23
 IPspace
 assigning an interface 68
 creating 66
 secure routing 61
 typical applications 61
 IPspaces
 guidelines 61
 interfaces 63
 naming requirement 65
 routing tables
 incoming packets for an IPspace, outgoing
 packets for an IPspace 63
 iSCSI and a vFiler unit 51
 iSCSI LUNs and igroups
 guidelines 50
 iSCSI support for a vFiler unit 37

L

language, guidelines for
 /etc/quotas file
 etc/usermap.cfg, language for encoding 21
 language for encoding 21
 local user accounts 75
 LUNs on a vFiler unit 49

M

Managing MultiStore 17

managing storage system
 volumes, disks, and RAID groups
 backups and data recovery 16
 mandatory_scan option 76
 maximum number of interactive SSH sessions 41
 maximum number of vFiler units 31
 migrate
 by copying data 101
 storage system data 14
 Migration of vFiler units 100
 moving a vFiler unit 79
 moving resources, about 28
 multiple security domains 13
 multiple server consolidation 12
 multistore
 overview 12
 MultiStore, enabling, or disabling 17

N

NDMP 48
 network checks for migration to destination storage
 system 80
 Network Data Management Protocol 48
 network interfaces
 configuring down 23
 for a vFiler unit 15
 network resources
 base IP address
 IP alias 23
 requirements for moving and removing 28
 network resources, storage resources 27
 NFS
 support for a vFiler unit 37
 path names for exporting 72
 starting the protocol 72, 73
 statistics 59
 nfsstat command 59
 NIS servers
 and migration, disaster recovery 81
 no-vfiler-ips? variable, setting 97
 not supported commands
 MultiStore commands
 ndmp commands 112, 113
 SnapMirror commands 112, 113

O

offline migration 101
 online migration

- option required 109
- options 109
- supporting Data Motion for vFiler 107
- online migration of vFiler units
 - considerations 108

P

- partitioning system resources 14
- path names, for NFS exports and CIFS shares 72
- performance, monitoring 59
- primary unit 20
- protocols
 - allowing or disallowing, CIFS or NFS 38
 - supported for a vFiler unit 37

Q

- qtree command output (how it differs for a vFiler unit and hosting storage system) 47
- qtrees
 - who can create on a vFiler unit 47
- quota report command 79
- quota reports
 - displaying 126
- quotas
 - effects of destroying a vFiler unit 34
 - guideline for 21
 - prerequisite for turning on and off 122, 123
 - resizing 125
 - types supported for a vFiler unit 122
 - who specifies 123

R

- rebooting storage system, effects on a vFiler unit 46
- removing IP addresses from an interface 67
- removing IP aliases from an interface IP alias 23
- resizing quotas 125
- resources
 - assigning 24
 - guidelines for assigning 20
- Restore 35
- restricting storage system traffic 14
- resynchronizing 90
- routed daemon
 - and IPspace 67
 - effects of disabling 52
 - enabling 17
- routing table
 - for the storage system 52
 - vFiler unit in default IPspace 52

rsh

- access to vFiler unit from clients 42
- enable option 42

RSH

- support for a vFiler unit 37

S

- SAN guidelines 22
- Server Manager 74
- setting up a vFiler unit 25
- setup command for a vFiler unit 25
- sis 54
- SnapMirror
 - options snapmirror.access command 94
 - using to reactivate vFiler unit 94
- snapmirror commands
 - snapmirror break command 94
 - snapmirror quiesce 94
 - snapmirror status 95
 - snapmirror update command 94
- snapmirror update command 98
- snapmirror.allow file 94
- SnapVault
 - backup of vFiler units using 57
- SSH 37, 43
- starting a vFiler unit 36
- storage checklist 79
- storage checks for migration 77
- storage checks for migration to destination storage system 83
- storage resources
 - assigning volumes 20
 - requirements for moving and removing 28

- storage system partitioning 12
- storage system reboot, effects on a vFiler unit 46
- storage system resources
 - partitioning 14
- subnet, moving vFiler unit 104
- System Manager
 - configuring resources 24

T

- target portal groups
 - adding IP addresses to IP-based groups 119
 - data migration implications for 115
 - deleting IP-based groups 120
 - displaying information about IP-based groups 120
 - enabling IP-based 116
- traditional volume for a vFiler unit 46

trusted host

changing after migration 104

U

uptime command 59

User Manager 74

V

vFiler

FlexClone files and FlexClone LUNs 57

vfiler allow 37

vfiler context command 39

vfiler dr resync command 90

vfiler rename 32

vFiler unit

increasing the limit 30

decreasing the limit 31

default 15

disaster recovery

checking storage space 83

disaster recovery: 90

displaying status 39

maximum number 30

migration stages 109, 110

moving

to different subnet 104

to different Windows domain 104

prerequisites for migration 105

protocols supported 15

reactivating using vfiler dr commands 96

reactivating via SnapMirror commands 94

rename 32

replacing on original storage system 98

resources,IP addresses,moving resources,adding

and removing vFiler unit resources 27

resynchronization (resync)

handling failures 93

setup 25

states, stopped, or running 33

viewing status 111

vFiler unit administrators

types 17

vFiler unit limit 30

vFiler unit migration

using the vfiler migrate command 106

by copying data 102

vFiler units

deduplication on 55

identifying commands 39

logging in to 40

online migration considerations 108

vfiler0

included in vFiler unit limit 30

vfilertemplate, defined 22

virus scanning

registering scanners 75

requirements for 76

VLAN

tagging for traffic separation

tagging for more IPspaces 64

vol create command 78

volumes in a vFiler unit

effects of renaming 47

taking offline 46

W

WINS server

and migration, disaster recovery 82

changing after migration 104



NA 210-05491_A0, Printed in USA

GA32-1041-02

