NetApp

Basics of ONTAP 9

Cluster Mode administration

LAB GUIDE

V3.3 (ONTAP 9.4)

Igor Juran

August 2018

Contents

Lab Objectives

This lab shows how to do basic tasks to quickly and easily set up a FAS controller to allow quick start with NetApp disk array to hosting a CIFS share, NFS export, and iSCSI LUN. The lab also demonstrates the unique advantage of NetApp disk arrays – snapshots and theirs integration with server operating system and/or applications.

This lab's environment includes Windows and Linux hosts suitable for establishing CIFS, NFS, and iSCSI client connections but this lab guide does not include any instructions for configuring those clients to work with the FAS disk array.

Usernames and passwords

System	Username	Password
Windows 2008 (DC)	Administrator	passW0rd
CentOS	root	passW0rd
NetApp Simbox	root	passW0rd
Domain user NetApp	netapp	Netap123

What is not possible to do on this environment.

Simulate ONTAP 9 supports most ONTAP functionality and the majority of its features. However, Simulate ONTAP has some functional limitations and does not support some Data ONTAP features because of the virtualization architecture.

Simulate ONTAP has the following limitations:

- Non-Volatile RAM (NVRAM) is simulated and is not persistent.
- Data loss might occur after power cycles. You must shut down the simulator properly to avoid data loss.
- You can have a maximum of four simulated disk shelves with 14 disk drives per shelf, for a total of 56 drives per simulator.
- Each simulated drive is limited to 9 GB. Note: The simulator image is pre-configured with 28 x 1 GB disks; 14 each on simulated disk shelves 0 and 1 and 14 simulated 520 MB SSD drives in shelve 2.
- You can create 64-bit aggregates, but they are limited to a maximum of 9 GB per simulated disk drive.
- Simulate ONTAP is not suitable for applications that require high performance or heavy I/O.

Simulate ONTAP does not support the following features:

- High Availability (CFO/SFO)
- Fibre channel and SAN connectivity
- CFE, BIOS, shelf FW, and so on

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B

IČO 357 034 66
IČ DPH SK2020341158

LAB 1: Overview

TASK1: Put a "P" or "L" beside each item in the following list of Data ONTAP cluster-mode concepts to indicate whether it is physical or logical.

____ node

____ disk

____ aggregate

____ virtual server (Vserver)

____ cluster

____ network port

____ flexible volume

____ Snapshot copy

____ SnapMirror copy

____ host bus adapter (HBA)

____ LIF

TASK2: Put an "N," "A," or "D" (or a combination of those letters) beside each item to indicate whether it has a node Vserver, administration Vserver, or data Vserver scope.

____ disk

____ namespace

____ data LIF

____ network port

____ cluster management LIF

____ flexible volume

____ aggregate

____ Snapshot copy

____ host bus adapter (HBA)

____ LUN

END OF EXERCISE

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B

IČO 357 034 66
IČ DPH SK2020341158

LAB 2: Installation and configuration

TASK 1: Creating a two-node cluster

Aim of this exercise is to create a cluster, configure the first node in the cluster, and then add a second node to the cluster.

Step 1.

From the Windows machine in your lab kit, use PuTTY to open a telnet session to “Console node1” to the first storage controller in your lab kit and verify that the preconfigured entry is correct. (We have not configured IP addresses yet so we cannot use IP address).

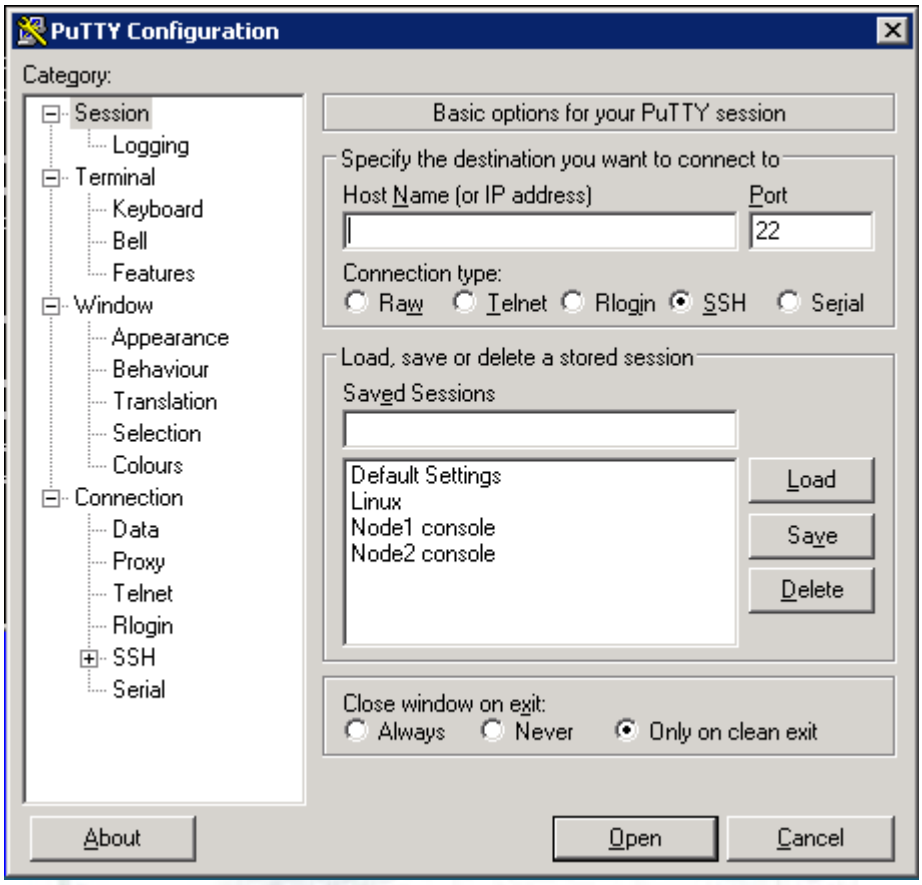

Step 2.

When node boots, then the Cluster Setup wizard starts on the console.

Welcome to the cluster setup wizard.
You can enter the following commands at any time:
"help" or "?" - if you want to have a question clarified,
"back" - if you want to change previously answered questions, and
"exit" or "quit" - if you want to quit the cluster setup wizard.

Any changes you made before quitting will be saved.

You can return to cluster setup at any time by typing "cluster setup".
To accept a default or omit a question, do not enter a value.

This system will send event messages and periodic reports to NetApp Technical Support. To disable this feature, enter
autosupport modify -support disable
within 24 hours.

Enabling AutoSupport can significantly speed problem determination and resolution should a problem occur on your system.
For further information on AutoSupport, see:
http://support.netapp.com/autosupport/

*Type yes to confirm and continue {yes}:**yes***

Enter the node management interface port [e0c]: ***e0c***
Enter the node management interface IP address: ***10.34.x1.31***
Enter the node management interface netmask: ***255.255.255.0***
Enter the node management interface default gateway: ***10.34.x1.1***
A node management interface on port e0c with IP address 10.34.x1.31 has been created.

Use your web browser to complete cluster setup by accessing https://10.34.x1.31

Otherwise, press Enter to complete cluster setup using the command line interface: **<Enter>**

Step 3.

Do you want to create a new cluster or join an existing cluster? {create, join}:
NOTE: If you make a mistake during cluster setup, press **Ctrl-C** to interrupt the wizard. Log on with admin username and restart the wizard by re-entering *cluster setup* at the cluster shell prompt.

Step 4.

At the prompt, enter create to create a new cluster.

Do you want to create a new cluster or join an existing cluster? {create, join}: **create**

Step 5.

When prompted about using the node as a single node cluster, reply no because this will be a multi-node cluster.

Do you intend for this node to be used as a single node cluster? {yes, no} [no]: **no**

Step 6.

Enter **yes** to accept the default values for cluster network configuration.

System Defaults:

Private cluster network ports [e0a, e0b].

Cluster port MTU values will be set to 1500.

Cluster interface IP addresses will be automatically generated.

Do you want to use these defaults? {yes, no} [yes]: **yes**

Step 7.

Enter the cluster administrator's (username "admin") password: **passW0rd**

Retype the password: **passW0rd**

It can take several minutes to create cluster interfaces...

Step 8.

Name the cluster **cluster1**.

Step 1 of 5: Create a Cluster

You can type "back", "exit", or "help" at any question.

Enter the cluster name: **cluster1**

Step 9.

Creating cluster cluster1

Starting cluster support services.....

Cluster cluster1 has been created.

Step 10.

Do not add additional license keys at this time but press **Enter** to continue.

Step 2 of 5: Add Feature License Keys

You can type "back", "exit", or "help" at any question.

Enter an additional license key []:

Step 11.

Configure the cluster management resources:

a. Type **e0c** as the port to host the cluster management interface.

b. Enter the IP address, netmask, and default gateway.

Insert the IP addresses depend on your POD number.

Step 3 of 5: Set Up a Vserver for Cluster Administration

You can type "back", "exit", or "help" at any question.

Enter the cluster management interface port [e0d]: **e0c**

Enter the cluster management interface IP address: **10.34.x1.21**

Enter the cluster management interface netmask: **255.255.255.0**

Enter the cluster management interface default gateway: **10.34.x1.1**

A cluster management interface on port e0c with IP address 10.34.x1.21 has been created. You can use this address to connect to and manage the cluster.

Step 12.

Enter the domain name and IP address of the DNS server.

Enter the DNS domain names: **ntap.lab1**

Enter the name server IP addresses: **10.34.x1.10**

DNS lookup for the admin Vserver will use the ntap.lab1 domain.

Step 4 of 5: Configure Storage Failover (SFO)

You can type "back", "exit", or "help" at any question.

SFO will not be enabled on a non-HA system.

Step 5 of 5: Set Up the Node

You can type "back", "exit", or "help" at any question.

Where is the controller located []: **Bratislava/Slovakia**

Step 13.

Cluster "cluster1" has been created.

To complete cluster setup, you must join each additional node to the cluster by running "system node show-discovered" and "cluster add-node" from a node in the cluster.

To complete system configuration, you can use either OnCommand System Manager or the Data ONTAP command-line interface.

To access OnCommand System Manager, point your web browser to the cluster management IP address (https://10.34.x1.21).

To access the command-line interface, connect to the cluster management IP address (for example, ssh admin@10.34.x1.21).

Step 14.

If the completion of the cluster setup wizard logs your management session off, log back in as *admin*.

Step 15.

To check the status of your new cluster, enter the **cluster show** command at the console.

cluster1::> cluster show

Node Health Eligibility

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B

IČO 357 034 66
IČ DPH SK2020341158

---------- ------- ------------

cluster1-01 *true* *true*

Step 16.

Disable autosupport

cluster1::> autosupport modify -support disable

Step 17.

Log off from terminal (telnet) session and log on trough SSH session to cluster management.

TASK 2: Add a second node to cluster

Step 1.

From the Windows machine in your lab kit, use PuTTY to open a telnet session to the second storage controller in your lab kit. It is preconfigured in PuTTY as Console node2.

Step 2.

Follow the cluster setup wizard.

Welcome to the cluster setup wizard.

You can enter the following commands at any time:
"help" or "?" - if you want to have a question clarified,
"back" - if you want to change previously answered questions, and
"exit" or "quit" - if you want to quit the cluster setup wizard.
Any changes you made before quitting will be saved.

You can return to cluster setup at any time by typing "cluster setup".
To accept a default or omit a question, do not enter a value.

This system will send event messages and periodic reports to NetApp Technical Support. To disable this feature, enter
autosupport modify -support disable
within 24 hours.

Enabling AutoSupport can significantly speed problem determination and resolution should a problem occur on your system.
For further information on AutoSupport, see:
http://support.netapp.com/autosupport/

Type yes to confirm and continue {yes}: **yes**

Step 3.

Configure node management:

a. Type **e0c** as the default port to host the node management interface.

b. Enter the IP address, netmask, and default gateway.

c. Press **Enter** to leave AutoSupport enabled.

Enter the node management interface port [e0c]: **e0c**

Enter the node management interface IP address: **10.34.x1.32**

Enter the node management interface netmask: **255.255.255.0**

Enter the node management interface default gateway: **10.34.x1.1**

A node management interface on port e0c with IP address 10.34.x1.32 has been created.

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B
IČO 357 034 66
IČ DPH SK2020341158

Use your web browser to complete cluster setup by accessing https://10.34.x1.32

Otherwise, press Enter to complete cluster setup using the command line interface: **<Enter>**

Step 4.

Do you want to create a new cluster or join an existing cluster? {create, join}: **join**

System Defaults:

Private cluster network ports [e0a,e0b].

Cluster port MTU values will be set to 1500.

Cluster interface IP addresses will be automatically generated.

Do you want to use these defaults? {yes, no} [yes]: **yes**

It can take several minutes to create cluster interfaces...

Step 1 of 3: Join an Existing Cluster

You can type "back", "exit", or "help" at any question.

Step 5.

Enter the IP address of an interface on the private cluster network from the cluster you want to join:

Step 6.

First we need to identify private cluster network from the cluster we created.

Switch to window with **SSH session to cluster management**. Enter command:

cluster1::> network interface show

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
Cluster						
	cluster1-01_clus1					
		up/up	169.254.yyy.yy/16	cluster1-01	e0a	true
	cluster1-01_clus2					
		up/up	169.254.zzz.zz/16	cluster1-01	e0b	true
cluster1						
	cluster1-01_mgmt1					
		up/up	10.34.x1.31/24	cluster1-01	e0c	true
	cluster_mgmt	up/up	10.34.x1.21/24	cluster1-01	e0c	true

4 entries were displayed.

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B

IČO 357 034 66
IČ DPH SK2020341158

Step 7.

Enter **cluster1-01_clus1 ip address** to join the existing cluster.

Enter the IP address of an interface on the private cluster network from the

cluster you want to join: 169.254.yyy.yy

Joining cluster at address 169.254.yyy.yy

Starting cluster support services

This node has joined the cluster cluster1.

Step 2 of 3: Configure Storage Failover (SFO)

You can type "back", "exit", or "help" at any question.

SFO will not be enabled on a non-HA system.

Step 3 of 3: Set Up the Node

You can type "back", "exit", or "help" at any question.

This node has been joined to cluster "cluster1".

To complete cluster setup, you must join each additional node to the cluster

by running "system node show-discovered" and "cluster add-node" from a node in the cluster.

To complete system configuration, you can use either OnCommand System Manager

or the Data ONTAP command-line interface.

To access OnCommand System Manager, point your web browser to the cluster

management IP address (https://10.34.x1.21).

To access the command-line interface, connect to the cluster management

IP address (for example, ssh admin@10.34.x1.21).

login:

Step 8.

Log off from terminal (telnet) session.

Step 9.

To check the status of your new cluster, enter the cluster show command at the CLI.

cluster1::> cluster show

Node	*Health*	*Eligibility*
cluster1-01	*true*	*true*
cluster1-02	*true*	*true*

2 entries were displayed.

Step 9.

Type **network interface show** and observe the cluster interfaces on both nodes and their associated IPs, noticing that the IP addresses are in a private network range.

cluster1::> network interface show

TASK 3: Preserve free space on vol0 volumes

Step 1.

This step is necessary only when running a virtualized cluster. The virtual disks attached to each node are small compared to real hard disks, resulting in unrealistic sizes for aggr0 and vol0.

Step 2.

Turn off snapshots on vol0 on both nodes.

cluster1::> system node run -node cluster1-01 vol options vol0 nosnap on

cluster1::> system node run -node cluster1-02 vol options vol0 nosnap on

Step 3.

Set snap reserve on vol0 to 0%.

cluster1::> system node run -node cluster1-01 snap reserve vol0 0

cluster1::> system node run -node cluster1-02 snap reserve vol0 0

Step 4.

From the command line, show the aggregates:

cluster1::> stor aggr show

You should see only two aggregates, an aggr0 owned by each of the two nodes in the cluster.

NOTE: The #Vols column shows that each aggregate contains one volume. Those are the vol0 volumes for each node. You cannot move them from their respective aggregates as you will be able to move the volumes that you will create later.

Step 5.

Rename the aggr0 aggregates so that they are more easily identifiable.

cluster1::> aggr rename -aggregate aggr0_cluster1_01 -newname aggr0_n1

cluster1::> aggr rename -aggregate aggr0_cluster1_02 -newname aggr0_n2

Step 6.

Verify the new names.

cluster1::> stor aggr show

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B
IČO 357 034 66
IČ DPH SK2020341158

The aggr0 aggregates have only one data disk each. The size of the data disk is 1GB.

Step 7.

Next steps till end of the exercise are specific only for the simulator environment. Add 3 data disks to each aggr0 aggregate.

cluster1::> aggr add-disk -aggregate aggr0_n1 -diskcount 3

cluster1::> aggr add-disk -aggregate aggr0_n2 -diskcount 3

Step 8.

Increase the size of vol0 to 3GB.

cluster1::> system node run -node cluster1-01 vol size vol0 +2g

cluster1::> system node run -node cluster1-02 vol size vol0 +2g

Step 9.

Verify the vol0 settings.

cluster1::> volume show -vserver cluster1-01 -volume vol0

cluster1::> volume show -vserver cluster1-02 -volume vol0

Step 10.

Assign all unassigned disk drives

cluster1::> storage disk show -container-type unassigned

…

cluster1::> storage disk assign -all -node cluster1-01

cluster1::> storage disk assign -all -node cluster1-02

cluster1::> storage disk show -container-type unassigned

There are no entries matching your query.

END OF EXERCISE

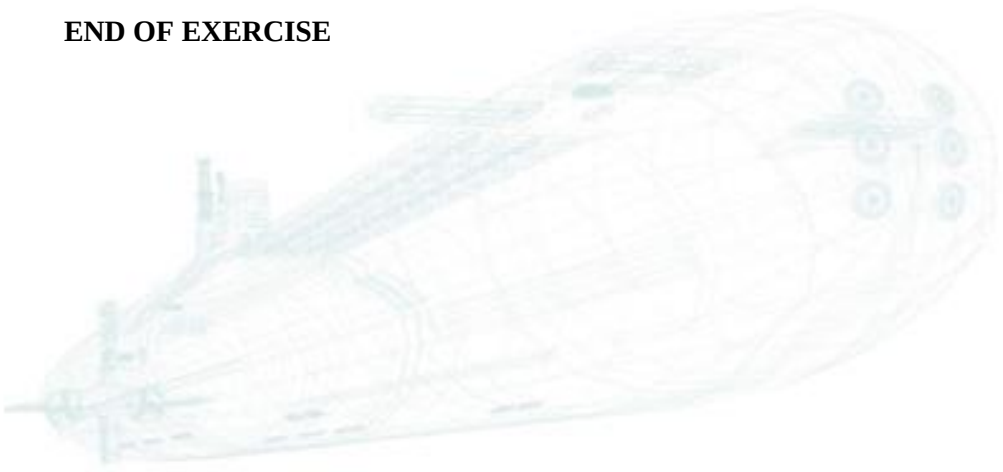

LAB 3: Cluster administration basics

In this exercise, you use the CLI and NetApp OnCommand GUI to explore how to manage a cluster. You also familiarize yourself with NetApp licensing.

Objectives of this exercise

By the end of this exercise, you should be able to:

- Connect to the command shell and explore the command hierarchy
- Review command options
- Compare privilege levels
- Use partial commands and complete commands with the Tab key
- Install NetApp OnCommand GUI
- Configure NetApp OnCommand GUI for your cluster
- Manage feature licenses

TASK 1. Connect to the command shell and explore the command hierarchy

Step 1.

From one of your node management sessions, enter *network interface show* and find the IP address of your cluster management interface.

cluster1::> network interface show

```
            Logical     Status     Network            Current       Current Is
Vserver     Interface   Admin/Oper Address/Mask       Node          Port    Home
----------- ----------- ---------- ------------------ ------------- ------- ----
cluster1
            cluster_mgmt up/up    10.34.x1.21/24     cluster1-01   e0c     true
cluster1-01
            clus1        up/up    169.254.12.222/16  cluster1-01   e0a     true
            clus2        up/up    169.254.131.174/16 cluster1-01   e0b     true
            mgmt1        up/up    10.34.x1.31/24     cluster1-01   e0d     true
cluster1-02
            clus1        up/up    169.254.245.33/16  cluster1-02   e0a     true
            clus2        up/up    169.254.95.73/16   cluster1-02   e0b     true
            mgmt1        up/up    10.34.x1.32/24     cluster1-02   e0d     true
7 entries were displayed.
```

Step 2.

In PuTTY, configure a new Secure Shell (SSH) session for the cluster management interface that you created in the previous exercise, and click Save to preserve the entry.

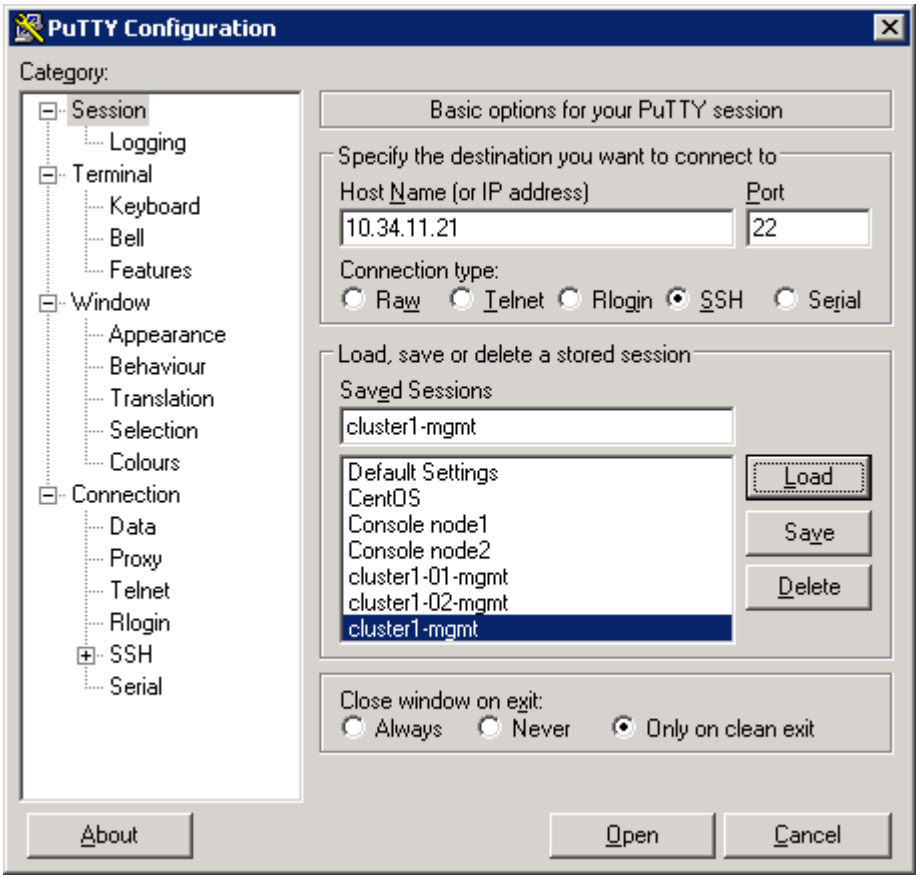

Step 3.

Using PuTTY, log in to the cluster management interface as admin, and compare the interface and the output of some command shell commands on both the cluster management session and a node management session. Are there any differences?

Step 4.

Close any open node management sessions by typing **exit** then pressing **Enter**, but leave the cluster management session open.

Step 5.

Type **?** to review the commands and command directories at the top level of the command hierarchy.

cluster1::>?

NOTE: You don't need to press Enter after typing a question mark.

Step 6.

Review the objects in the *storage* command directory.

cluster1::> storage ?

Step 7.

From the command shell, review the top-level command directories.

cluster1::>?

NOTE: An entry that ends with a "right angle bracket" (>) symbol is a command directory rather than a command. The structure resembles a UNIX or a DOS shell, in that you cannot execute command directory names as you do commands, but you can navigate to them. Command directories can contain subdirectories, commands, or both. Command directories provide contextual and hierarchical grouping of commands, so the command structure is not flat.

Step 8.

Go into the cluster directory.

cluster1::> cluster <Enter>

cluster1::cluster>

NOTE: The question mark can be used at any level of the command hierarchy to see what commands and directories are available within that context. Notice that the command shell prompt changes to indicate which context you're in.

Step 9.

Look at the available commands and directories at this level.

cluster1::cluster>?

Step 10.

Go into the **statistics** directory.

cluster1::cluster> statistics

cluster1::cluster statistics>

You're now in the cluster statistics context.

Step 11.

See what's available at this level.

cluster1::cluster statistics>?

Step 12.

Go back (up) one level by typing two periods and then pressing the **Enter** key.

cluster1::cluster statistics>..

cluster1::cluster>

Step 13.

Notice that you're back at the cluster level.

NOTE: From any level, you can enter *top* to go directly to the top of the entire command hierarchy.

Step 14.

Examine the manual page for the cluster command directory.

cluster1::cluster> man cluster

Step 15.

Enter **q** to exit the manual page.

Step 16.

Examine the manual page for the cluster modify directory and compare the output with the output of the man cluster command in the previous step.

cluster1::cluster> man cluster modify

TASK 2: Review command options

Step 1.

Go to the storage aggregate level within the command shell.

cluster1::> storage aggr

Step 2.

From the storage aggregate level, run this command:

cluster1::storage aggregate> modify ?

NOTE: Square brackets ([]) indicate optional command elements. The output of this command shows the parameter –aggregate with brackets around the parameter name, but not around the parameter value. This means that the parameter name is optional, but the value is required. You can enter the aggregate name as a positional parameter rather than a named parameter, to save keystrokes. All other parameters and values are optional, but brackets surround both parameter and value because, if you provide one, you must provide the other (the value cannot be specified based on positional).

In this case, the aggregate name is required to determine which aggregate will be modified. Although the other parameters are technically optional, at least one of them should be specified for this command to be meaningful (that is, to actually modify an attribute of this aggregate).

Step 3.

Review the options for the storage aggregate scrub command.

cluster1::storage aggregate> scrub ?

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B
IČO 357 034 66
IČ DPH SK2020341158

NOTE: As with the modify command, the aggregate name is required, but the parameter name is optional. In addition, the action value is required, but the parameter name (–action) is optional. Two possible forms of the command are:

- *storage aggregate scrub –aggregate aggr0 –action start*
- *storage aggregate scrub aggr0 start*

Step 4.

Review the possible keyword values for the –state parameter.

cluster1::storage aggregate> modify -state ?

TASK 3: Compare privilege levels

Step 1.

Look at the volume directory.

cluster1::> volume ?

The default privilege level is admin.

Step 2.

Note the commands that are available in this directory context at this privilege level.

Step 3.

Switch to the *advanced* privilege level.

cluster1::> set -privilege advanced

NOTE: Because *-privilege* is an optional positional parameter of the *set* command, you can also specify the desired privilege level as a positional parameter*: set advanced.*

Step 4.

While you are in the *advanced* privilege level, look again at the *volume* directory.

cluster1::> volume ?

Step 5.

Notice the additional commands that are available.

Each command and directory that is available for non-admin privilege levels has an asterisk (*) in front of its description.

Step 6.

Switch back to the *admin* privilege level.

cluster1::> set admin

TASK 4: Use partial commands and complete commands with the TAB key

Step 1.

Display the logical interfaces.

cluster1::> network interface show

Step 2.

Try the following command:

cluster1::> net int sho

The command fails because this form of the command is ambiguous. Multiple verbs in this command hierarchy begin with the letters "sho."

Step 3.

Retype the command using the full verb, *show*.

cluster1::> ne in show

Step 4.

Type the first two letters of the *network* command directory (**ne**) and press **Tab**.

If the substring that you type is unambiguous, when you press Tab, the command shell completes the substring.

Step 5.

Continue the command by typing *in* and **Tab** and then *sho* and **Tab**.

The network and interface substrings are completed, but because *sho* is ambiguous within the context, the command shell displays the options for *sho*.

Step 6.

This time, enter *ne*, **Tab**, *in*, **Tab**, *show*.

TASK 5: GUI management using web browser

You can directly use supported browsers to access GUI. Just enter cluster management interface IP address to your browser. Open browser and access GUI. Use https://10.34.x1.21

TASK 6: Explore package licensing

Open browser window and enter cluster management ip address (https://10.34.x1.21)

Step 1.

Within the NetApp OnCommand System Manager page, select **Configurations > Licenses**.

Step 2.

Verify that the Base package is licensed.

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B
IČO 357 034 66
IČ DPH SK2020341158

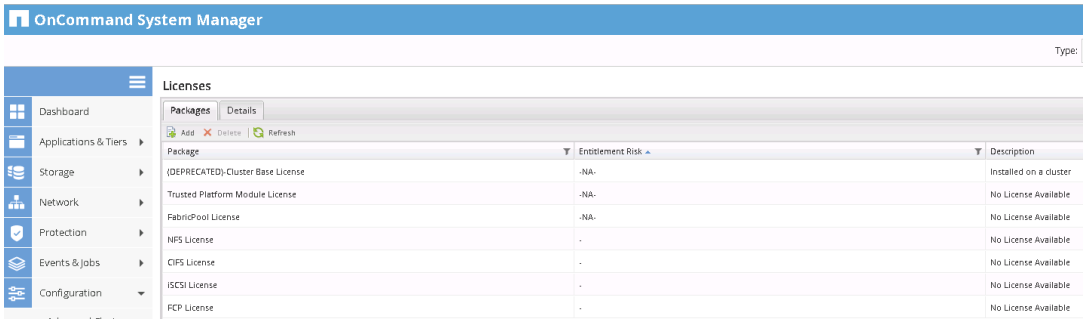

Step 3.

Click **Add** to begin adding new licenses for the SnapVault feature.

Step 4.

Verify that the Add License Packages dialog appears.

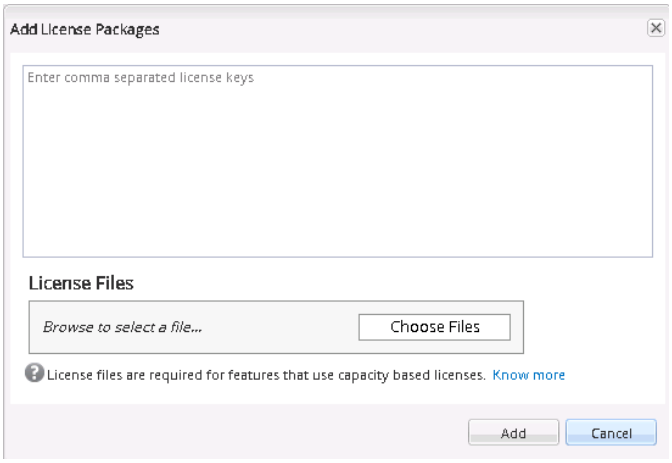

Step 5.

To follow best practices, you should add feature keys for each node in the cluster, so type the SnapVault key for each node.

```
INIIBQKBFDUFZGABGAAAAAAAAAAA
SUOYOUNFXMSMUCEZFAAAAAAAAAAA
```

Step 6.

Click **Add.**

Step 7.

Verify that the SnapVault package was correctly identified. These codes are node locked within the cluster.

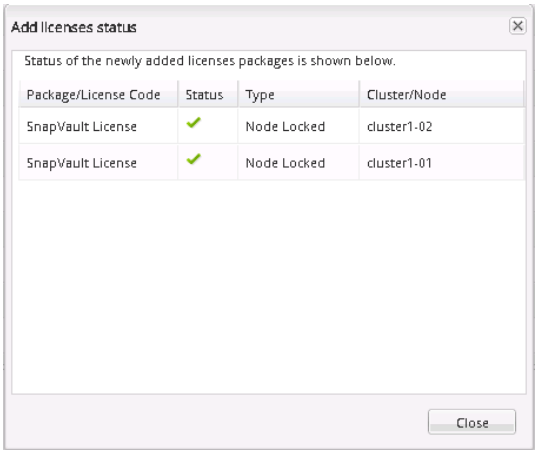

Step 8.

Click **Close**.

Step 9.

To verify that the new package is licensed in the list, click the package to see licensing details in the lower pane.

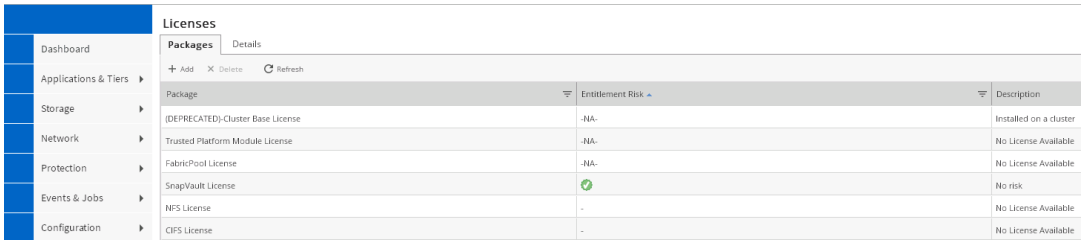

Step 10.

Switch back to the cluster management SSH PuTTY session and, if necessary, authenticate as *admin*.

Step 11.

Enter this command to view the cluster's serial id:

```
cluster1::> cluster identity show

          Cluster UUID: 4d0c7e7f-8b4a-11e8-a189-0050569fec97
          Cluster Name: cluster1
 Cluster Serial Number: 1-80-000011
      Cluster Location: Bratislava/Slovakia
       Cluster Contact:
```

For this sample output, the serial number is 1-80-000011.

Step 12.

Identify the serial numbers for each system within the cluster.

```
cluster1::> system node show -fields node, serialnumber
node         serialnumber
------------ ------------
cluster1-01 4082368511
cluster1-02 4034389062
2 entries were displayed.
```

Step 13.

Navigate to the license hierarchy.

```
cluster1::> license
cluster1::system license>
```

The prompt takes you to the *system license* command hierarchy.

Step 14.

List the available commands.

```
cluster1::system license> ?
  add                          Add one or more licenses
  clean-up                     Remove unnecessary licenses
  delete                       Delete a license
  entitlement-risk>            The entitlement-risk directory
  show                         Display licenses
  status>                      Display license status
```

Step 15.

View the currently licensed packages.

cluster1::system license> show

Step 16.

List details about the available license packages.

cluster1::system license> status show

Step 17.

Add the CIFS licenses for both nodes.

```
cluster1::system license> add -license-code
CAYHXPKBFDUFZGABGAAAAAAAAAAA,MHEYKUNFXMSMUCEZFAAAAAAAAAAA
```

```
License for package "CIFS" and serial number "1-81-
0000000000000004082368511" installed successfully.
License for package "CIFS" and serial number "1-81-
0000000000000004034389062" installed successfully.
(2 of 2 added successfully)
```

Step 18.

Verify the new CIFS licenses.

cluster1::system license> show

Step 19.

Add the FC license for a node that is not currently in the cluster.

NOTE: You can add license codes for nodes that do not exist in the cluster. This is often done to "preload" the license for nodes that are going to be added.

```
cluster1::system license> add -license-code
KOBYNDUCCLPKICAAAAAAAAAAAAAA
License for package "FCP" and serial number "1-81-
0000000000000000000333333" installed successfully.
(1 of 1 added successfully)
```

Step 20.

View the currently licensed packages.

cluster1::system license> show

Step 21.

Type the following command and then press the TAB key:

cluster1::system license> clean-up –<TAB>

-unused -expired -simulate

NOTE: You can select expired or unused licenses to be removed. An unused license is a license that is associated with nodes that don't belong to the cluster.

Step 22.

Verify what will be cleaned up with the *unused* command.

```
cluster1::system license> clean-up -unused -simulate
The following licenses can be cleaned up:
Serial number: 1-81-0000000000000000000333333
Owner: none
Package    Reason
---------- --------------------------------------------------
FCP        Serial number is not used by any node in the cluster
```

Step 23.

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B
IČO 357 034 66
IČ DPH SK2020341158

Clean up unused licenses

cluster1::system license> clean-up -unused

1 unused license deleted.

Step 24.

Verify that the FCP license was removed.

cluster1::system license> show

Step 25.

Use GUI or the command shell to enter the other licenses required on cluster1.

cluster1_01

```
FlexClone        WSKTAQKBFDUFZGABGAAAAAAAAAAA
iSCSI            OUVWXPKBFDUFZGABGAAAAAAAAAAA
NFS              QFATWPKBFDUFZGABGAAAAAAAAAAA
SnapMirror       KYMEAQKBFDUFZGABGAAAAAAAAAAA
SnapRestore      YDPPZPKBFDUFZGABGAAAAAAAAAAA
```

cluster1_02

```
FlexClone        GARJOUNFXMSMUCEZFAAAAAAAAAAA
iSCSI            YBCNLUNFXMSMUCEZFAAAAAAAAAAA
NFS              ANGJKUNFXMSMUCEZFAAAAAAAAAAA
SnapMirror       UFTUNUNFXMSMUCEZFAAAAAAAAAAA
SnapRestore      ILVFNUNFXMSMUCEZFAAAAAAAAAAA
```

END OF EXERCISE

LAB 4: Architecture

In this exercise, you identify kernel and user-space processes, answer questions about replicated database (RDB) and quorum concepts, and answer questions about vol0 volumes and limitations.

Objectives of this exercise

- By the end of this exercise, you should be able to:
- Identify kernel and user-space processes
- Explain RDB and quorum concepts
- Explain vol0 volumes uses and limitations

TASK 1: Identify kernel and user-space processes

Step 1.

Using the following list of cluster components, show the path of a network-attached storage (NAS) write request by labeling each component with a number, starting with 1 for the initiator of a write request and proceeding in order to the destination of the request.

___Cluster session manager (CSM)

___D-blade

___Data network port

___NFS or CIFS client

___N-blade

___Disks

___Nonvolatile RAM (NVRAM)

TASK 2: Explain RDB and quorum concepts

Step 1. List the names of the RDB units.

Step 2. Which RDB unit keeps track of the data that is used to operate the clustershell?

Step 3. Which RDB unit keeps track of the volumes and aggregates and which volumes are on which aggregates?

Step 4. How many healthy nodes are needed in a 16-node cluster to maintain a quorum?

TASK 3: vol0 and Vserver root volumes uses and limitations

Step 1. What is the minimum number of vol0 volumes in a 20-node cluster? What is the maximum?

Step 2. What is stored on a vol0 volume?

Step 3. For the following characteristics, write **vol0**, **vsroot**, or **both** to match the volumes with their characteristics.

a) Is a flexible volume _____

b) Can have Snapshot copies _____

c) Can have mirrors _____

d) Can be accessed by NFS or CIFS clients _____

e) Can be backed up to tape _____

f) Can be copied _____

g) Can be accessed by the systemshell _____

h) Can have junctions to other volumes _____

i) Cannot be moved to another aggregate _____

END OF EXERCISE.

LAB 5: Physical data storage

In this exercise, you use the CLI and NetApp OnCommand GUI to create and add disks to an aggregate.

Objectives of this exercise

- Create an aggregate
- Add disks to an aggregate
- Create a flash pool

TASK1: Create a new aggregate

Step 1.

From the command line, show the aggregates:

cluster1::> stor aggr show

Step 2.

Show the volumes

cluster1::> volume show

Step 3.

View the disks attached to each node, noticing the disks that belong to aggregates and the spare disks that are available to create additional aggregates.

cluster1::> storage disk show -owner cluster1-01

cluster1::> storage disk show -owner cluster1-02

Step 4.

Display disks by type.

cluster1::> storage disk show -type?

cluster1::> storage disk show -type ssd

cluster1::> storage disk show -type fcal -owner cluster1-01

cluster1::> storage disk show -type fcal -owner cluster1-02

cluster1::> storage disk show -type ssd -owner cluster1-01

cluster1::> storage disk show -type ssd -owner cluster1-02

Step 5.

Look at the usage information of the storage aggregate create command.

cluster1::> stor aggr create ?

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B
IČO 357 034 66
IČ DPH SK2020341158

NOTE: To help familiarize yourself with a command, you can type the command without parameters, followed by a question mark (?).

Step 6.

Create an aggregate with the unique name n01_aggr1 (as an abbreviation of aggregate 1 on the cluster1_01 node).

cluster1::> stor aggr create -aggr n01_aggr1 -node cluster1-01 -disktype fcal -diskcount 5

NOTE: The aggregate will have the storage capacity of three disks. (Two of the disks are used for parity for RAID-DP, which is the default RAID type.) This is only an example. In a production environment, a RAID-DP aggregate of this size is a very inefficient use of disks. Also, note that the –diskcount parameter cannot exceed the available number of spare disks.

Step 7.

Review the list of aggregates again.

cluster1::> stor aggr show

Step 8.

Review the details for the new aggregate, noticing that the new aggregate has an HA policy of sfo.

cluster1::> stor aggr show -aggr n01_aggr1

TASK 2: Add disks to aggregate

Step 1.

cluster1::> aggr add-disks -aggr n01_aggr1 -diskcount 2

The aggregate now has two parity disks and five data disks.

Step 2.

Verify the number of disks in the aggregate and the expanded disk capacity.

cluster1::> aggr show -aggr n01_aggr1

TASK 3: Use GUI to create an aggregate

Step 1.

In GUI select **Storage> Aggregates&Disks > Aggregates**.

You can see aggr0 for each of the nodes in your cluster, as well as the aggregate that you created by using the CLI:

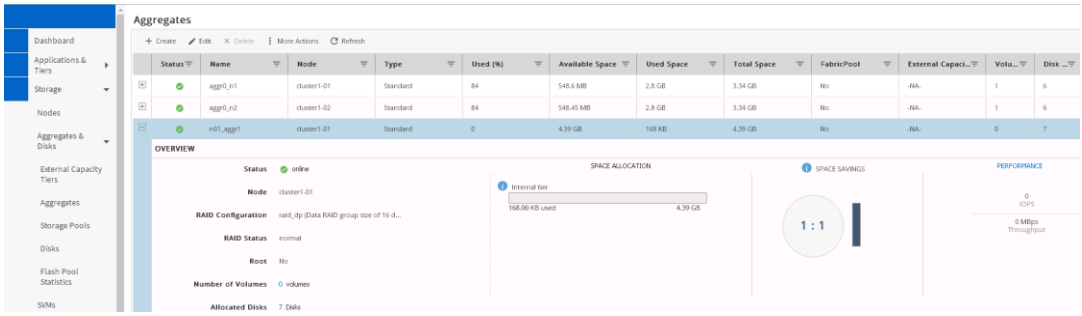

Step 2.

In the Aggregates toolbar, click Create. In the right upper corner click on “Manually Create Aggregate” button.

The Create Aggregate dialog window appears.

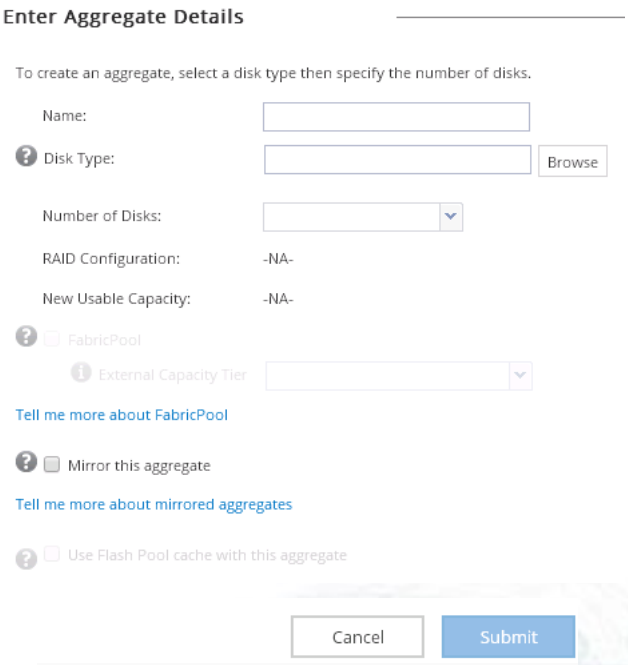

Step 3.

On the Create Aggregate dialog page, enter these aggregate properties:

a. Enter the Aggregate Name ***n02_aggr1***

b. Select disk type **FCAL** on node **cluster1-02**

c. Number of Disks: **10**

d. RAID Configuration: RAID-DP, RAID group size of 16 disks

e. Ensure that the Mirror this aggregate and Use Flash Pool cache checkbox is not selected.

Step 4.

Create the aggregate

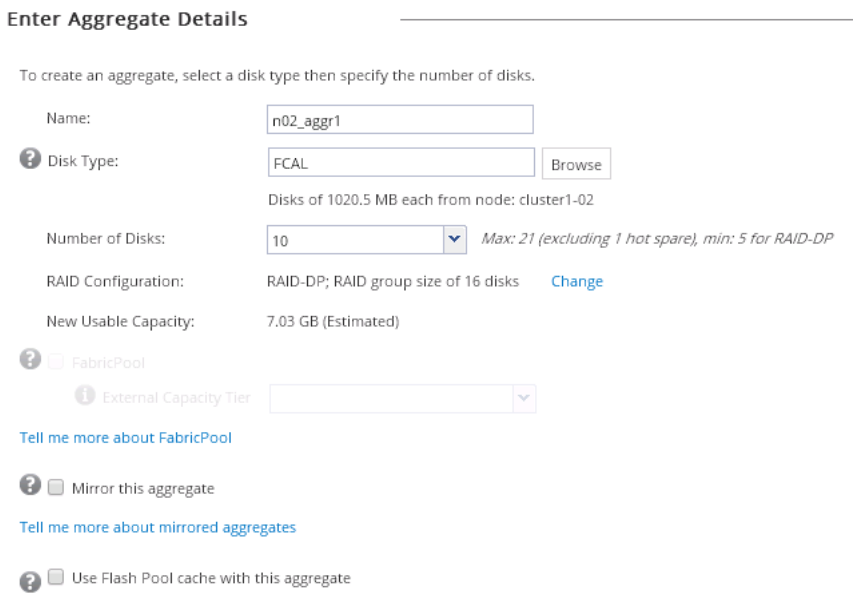

Click **Submit**.

Step 5.

Navigate to **Storage > Aggregates&Disks > Aggregates** and examine the new aggregate.

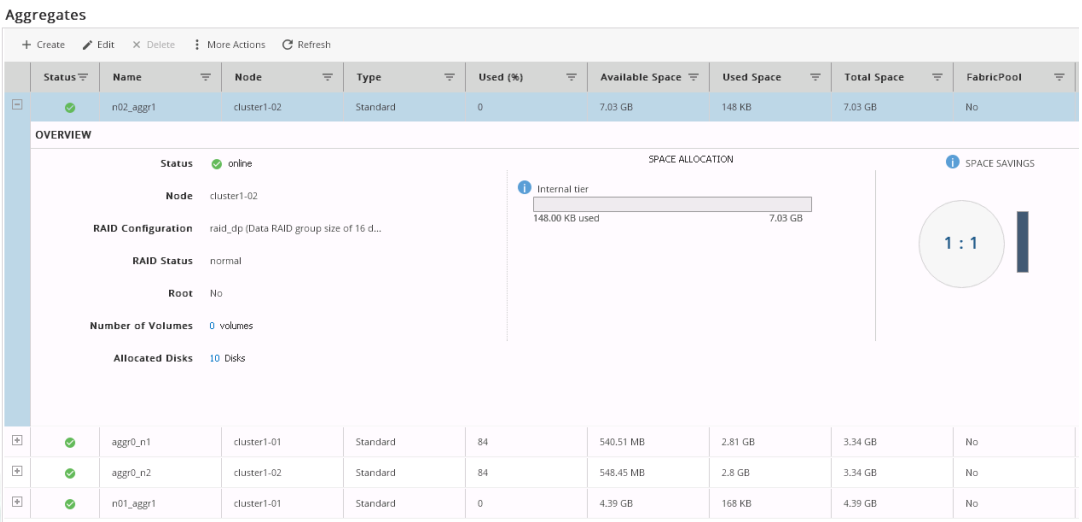

TASK 4: Create a flash pool

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B
IČO 357 034 66
IČ DPH SK2020341158

In this task, you will create a 10-disk aggregate of type “flex,” and then you will convert it to the HDD tier of a flash pool.

Step 1.

Type **stor aggr create** followed by a question mark (?).

cluster1::> stor aggr create ?

NOTE: There is no option to enable hybrids on the create command.

Step 2.

We create a standard aggregate of type flex first, then convert it to a Flash Pool. Create a new 10-disk aggregate on cluster1_01. This becomes the HDD tier of our Flash Pool:

cluster1::> stor aggr create -aggr n01_fp1 -node cluster1-01 -diskcount 10 -disktype fcal

Step 3.

Inspect the properties of the new aggregate.

cluster1::> stor aggr show -aggr n01_fp1

Step 4.

Enable flash pool by setting hybrid-enabled to true on n01_fp1.

cluster1::> stor aggr modify -aggr n01_fp1 -hybrid-enabled true

Step 5.

Check the status of your aggregate before adding the SSD tier, noticing the available size.

cluster1::> stor aggr show -aggr n01_fp1

Step 6.

Add SSDs to the aggregate and complete the flash pool.

cluster1::> stor aggr add-disk -aggr n01_fp1 -diskcount 5 -disktype SSD

NOTE: You can also specify a list of SSDs by using the -disklist option.

Step 7.

Explore the flash pool, noticing these properties:

- The number of disks
- The two RAID groups (in spite of the default RAID group size)
- The available size and the total hybrid cache size

cluster1::> stor aggr show -aggr n01_fp1

Step 8.

In GUI, perform these actions:

a. Select your new flash pool.

b. Below the lower pane, click the Disk Layout tab.

c. Inspect the RAID groups that were created.

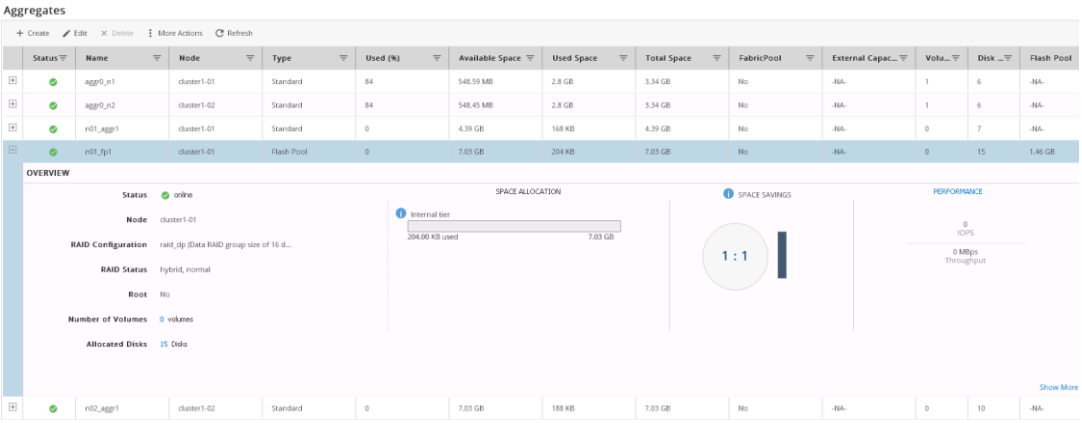

END OF EXERCISE

LAB 6: Logical data storage

In this exercise, you create a data storage virtual server (Vserver), create a flexible volume, and mount a volume in a SVM's namespace. You also create an infinite volume.

Objectives of this exercise:

- Create a data SVM
- Create a flexible volume
- Mount a volume in a SVM's namespace
- Create an flexgroup volume

TASK 1: Create a cluster SVM

Step 1.

View the SVMs and volumes in your cluster.

cluster1::> vserver show

cluster1::> volume show

Step 2.

View the "node" Vservers and "admin" Vserver.

The node SVMs are the scope in which entities like the vol0 volumes exist, along with the node management LIFs and cluster LIFs of each node. The administrative SVM is the scope in which the cluster management LIF exists.

Step 3.

Create a data SVM named vs1 and provide the name of a root volume (which will be created by this command).

cluster1::> vserver create -vserver vs1 -rootvolume vs1root -aggr n01_aggr1 -ns-switch file -rootvolume-security-style unix

NOTE: This command creates the root volume of the SVM, which means that it's the root of this namespace. You created the SVM root volume on node cluster1_01, but that does not mean that the SVM is assigned to node cluster1_01. SVMs are independent of physical nodes. The root of the SVM namespace happens to reside there.

Step 4.

View the SVM that you created and note that the new SVM type is "data."

cluster1::> vserver show

Step 5.

Get more details about the SVM and notice that all protocols are allowed by default.

cluster1::> vserver show -vserver vs1

Step 6.

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B

IČO 357 034 66
IČ DPH SK2020341158

View the list of volumes.

cluster1::> volume show

The SVM root volume that was just created is listed.

Step 7.

Notice the default volume size.

NOTE: You can't specify a volume size with the *vserver create* command, but you can use the *volume modify* command to change the size.

Step 8.

View all the attributes of this volume.

cluster1::> volume show -vserver vs1 -volume vs1root

Step 9.

Review the Junction Path value of this volume.

The slash (/) signifies that this is the root volume of this namespace. Giving this volume the name "vs1root" is merely a convention to make it obvious that this is a SVM root volume.

Step 10.

View the number of volumes on n01_aggr1.

cluster1::> stor aggr show

This reflects the new root volume and the size of it is reflected in (is subtracted from) the available field of the aggregate.

TASK 2: Create a flexible volume

Step 1.

In your SVM, on the aggregate n01_aggr1, create a volume.

cluster1::> volume create -vserver vs1 -volume volume1 -aggregate n01_aggr1 -junction-path /vol1

Warning: The export-policy "default" has no rules in it. The volume will therefore be inaccessible.

Do you want to continue? {y|n}: y

Each volume is associated with one SVM.

Step 2.

View the volumes.

cluster1::> vol show

Step 3.

View the details of the new volume.

cluster1::> vol show -vserver vs1 -volume volume1

Step 4.

View the assigned security style.

cluster1::> vol show -vserver vs1 -volume volume1 -fields security-style

vserver volume security-style

vs1 volume1 unix

TASK 3: Use GUI to create a flexible volume

Step 1.

In GUI, select **Storage > Volumes**.

Volume1 and the root volume for your SVM appear.

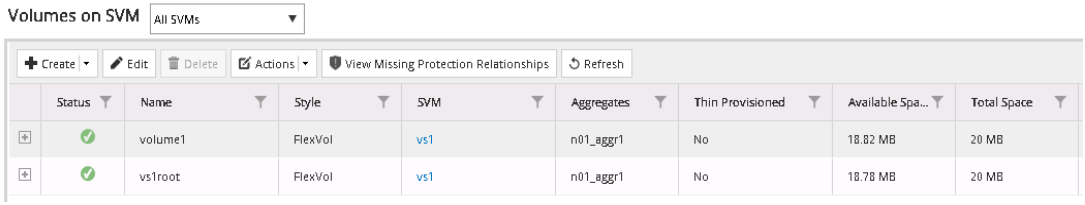

Step 2.

In the Volumes toolbar, click the **Create > Flexvol**.

Step 3.

Select SVM.

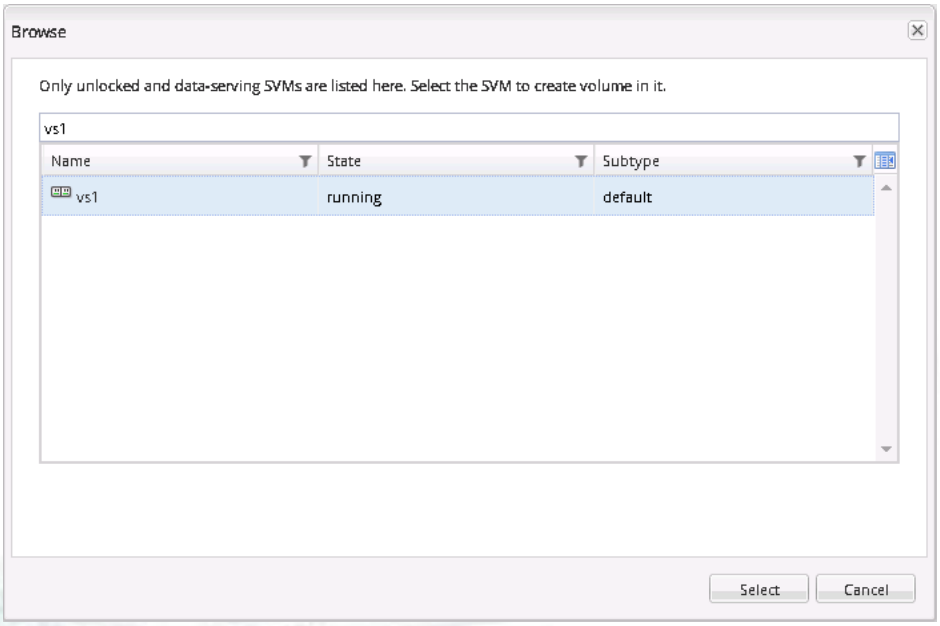

Step 4.

To select a home aggregate for the new volume, click the **Choose** button and select **n02_aggr1**, tiering policy **snapshot-only.**

NOTE: To host a user volume, never choose a node's aggr0 aggregate.

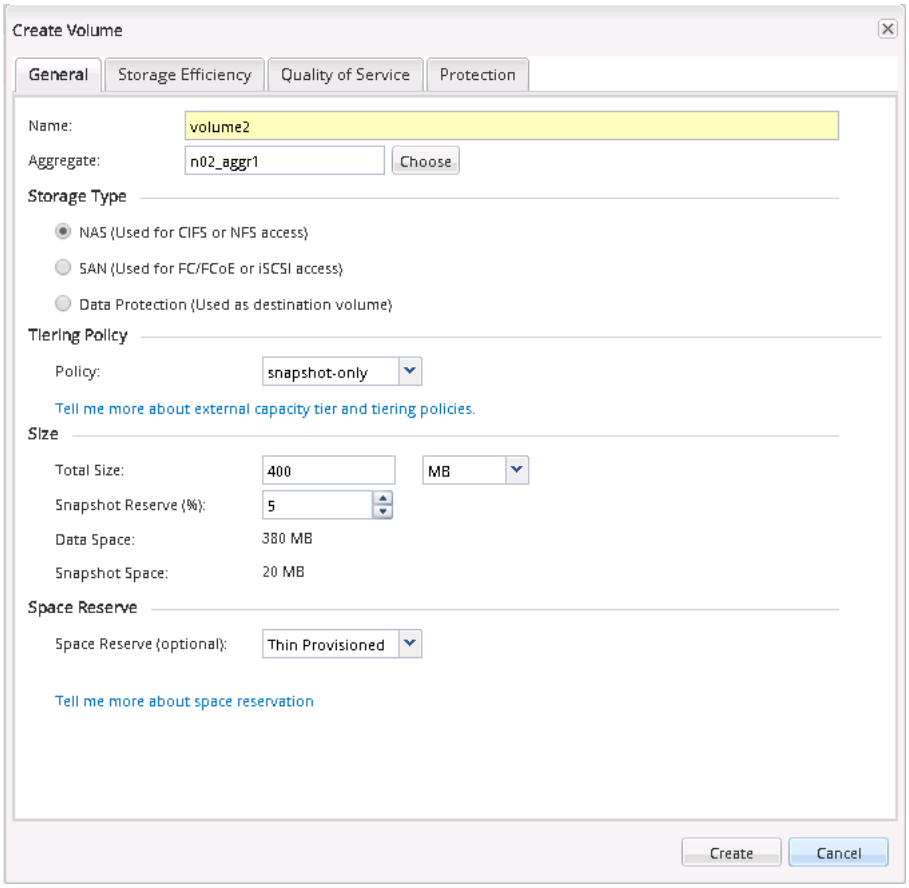

Step 5.

Under Storage Type, verify that NAS is selected.

Hint: What happens when protocols aren't licensed for your cluster?

Step 6.

Set the Total Size to **400 MB**.

Step 7.

Leave the Snapshot Reserve percentage at the default.

Step 8.

Select Space Reserve (optional): **Thin Provisioned**.

Step 9.

Click **Create**.

Step 10.

After the wizard is finished, verify the new volume in the volume list.

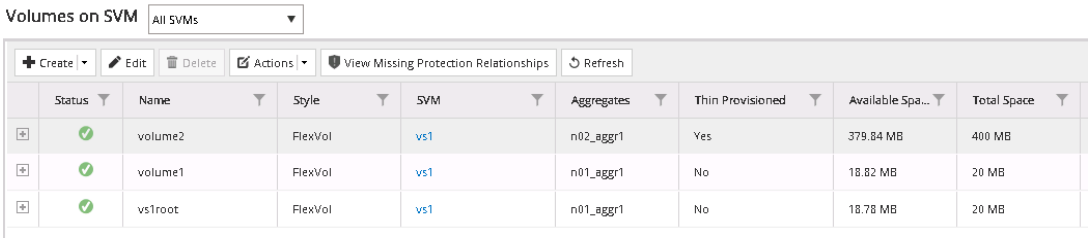

Step 11.

Select **Storage** > **Junction Path** and verify where the new volume is mounted in the namespace. The UI shows that the volume is accessed by clients as /volume2.

NOTE: The junction-path option was not specified when you created the volume with GUI. GUI automatically mounts new volumes at the SVM root by using the volume name as the junction name. The default Export Policy has been assigned to the volume.

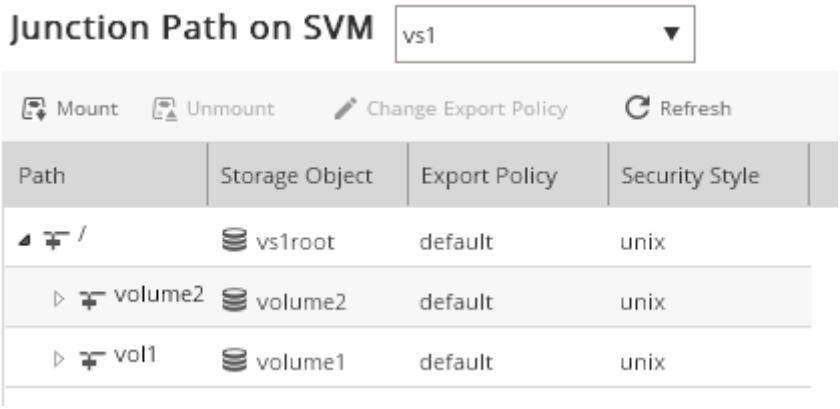

Step 12.

To unmount volume2, click **volume2** to select it and then, on the Junction Path toolbar, click **Unmount**.

NOTE: It is not necessary to force the unmount operation.

Step 13.

To begin a series of steps that will remount volume2 in the namespace as vol2 under volume1, on the Junction Path toolbar, click **Mount**.

The menu lists all unmounted volumes.

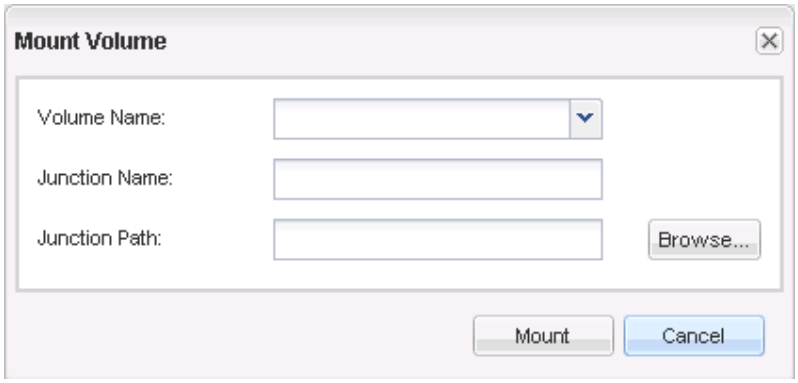

Step 14.

Select **volume2**.

Step 15.

Change the Junction Name to “vol2.”

Step 16.

To choose a point in the Junction Path to mount the volume, click **Browse**.

Step 17.

Select /**vol1** as the mountpoint.

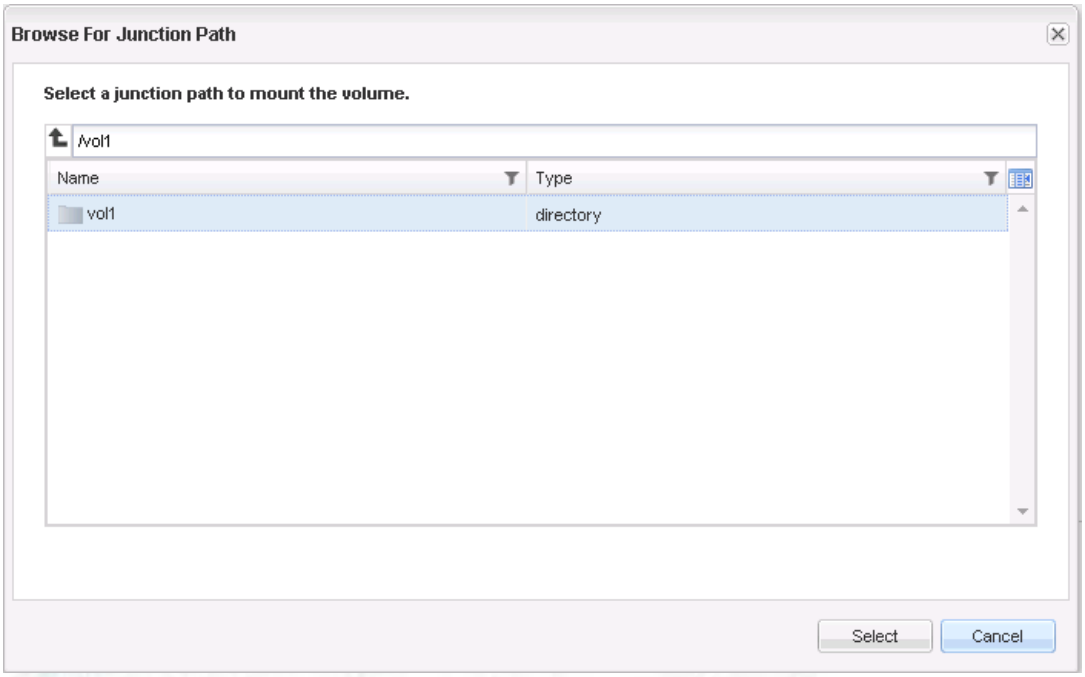

Step 18.

To complete the operation, click **OK** and then click **Mount**.

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B

IČO 357 034 66
IČ DPH SK2020341158

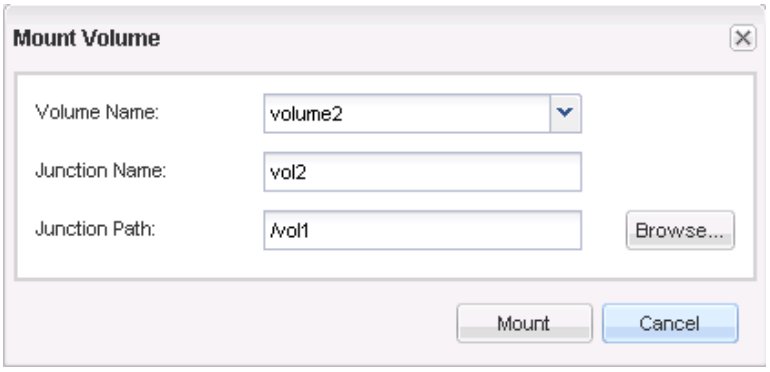

Step 19.

Verify that volume2 is mounted at /vol1/vol2.

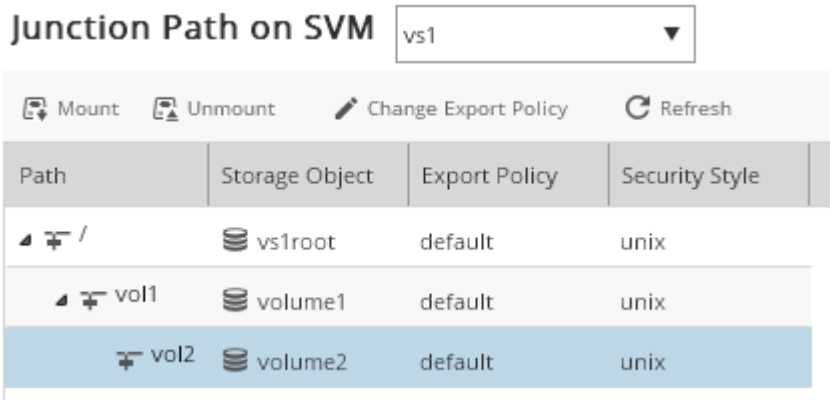

TASK 4: Configure a flexgroup volume

Step 1.

From CLI, create a SVM to host flexgroup volume.

Step 1.

From the CLI, create a SVM to host flexgroup volume

cluster1::> vserver create -vserver vsflxg -rootvolume vsflxgroot -aggregate n02_aggr1 -rootvolume-security-style unix -language C

Step 2.

Create 6 GB flexgroup volume

cluster1::> volume create -vserver vsflxg -volume flxg_vol -aggr-list n01_aggr1 ,n02_aggr1 -aggr-list-multiplier 3 -size 6gb

Warning: The FlexGroup "flxg_vol" will be created with the following number of constituents of size 1GB: 6.

Do you want to continue? {y|n}: y

Step 3.

Display the volumes.

cluster1::> vol show

cluster1::> vol show -vserver vsflxg

Step 4.

Display view a FlexGroup volume and its constituents, and monitor the space used by the FlexGroup volume.

*cluster1::> volume show -vserver vsflxg -is-constituent **

Step 5.

Set your session to diag privilege mode, and enter y in response to the prompt:

cluster1::> set diag

Warning: These diagnostic commands are for use by NetApp personnel only.

Do you want to continue? {y|n}: **y**

Step 6.

Display the list of volumes again.

cluster1::>vol show -vserver vsflxg -fields is-constituent, constituent-role*

Step 7.

Examine the aggregates hosting the constituents and the size of the constituents.

cluster1::> vol show -vserver vsflxg -fields aggregate,size*

Step 8.

Set your session back to administrative privilege mode.

cluster1::> set admin

Step 9.

Configure FlexGroup volumes to automatically grow and shrink their size. Starting with ONTAP 9.3, you can configure FlexGroup volumes to automatically grow and shrink according to how much space they currently require.

cluster1::> volume autosize -vserver vsflxg -volume flxg_vol -mode grow_shrink

You can also specify the maximum size, minimum size, and thresholds for growing or shrinking the volume.

END OF EXERCISE.

LAB 7: Physical networking

In this exercise, you identify port roles and create an interface group.

Objectives of this exercise

- Create ipspace
- Create broadcast domain
- Create subnet
- Create an interface group
- Create a virtual LAN (VLAN)

TASK 1: Create an ipspace

Step 1.

View existing ipspaces

cluster1::> network ipspace show

A list of the existing ipspaces appears.

Step 2.

You can see default created ipspaces: Cluster and Default

Step 3.

To which ipspaces are belonging existing vservers?

Step 4.

Create ipspace with name ipspace1.

cluster1::> network ipspace create ipspace1

Step 5.

What is broadcast domain for ipspace1? Try to create new broadcast domain. Go to **Network -> Broadcast Domains -> Create**

Step 6.

We cannot create new broadcast domain for the now because all existing ports are assigned to others broadcast domains. (From the CLI is possible to create broadcast domain without assigned ports, but not from GUI.)

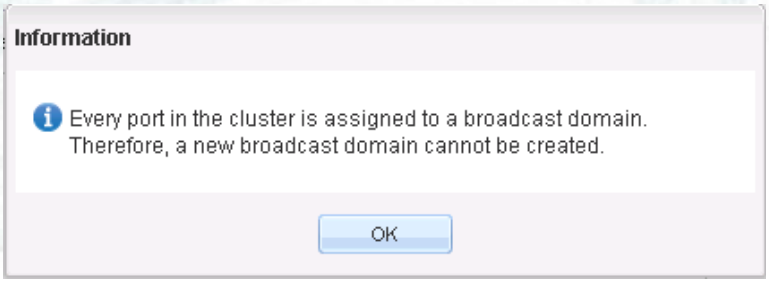

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B

IČO 357 034 66
IČ DPH SK2020341158

Step 7.

Remove ports cluster1-01:e0e, cluster1-01:e0f, cluster1-02:e0e, cluster1-02:e0f from Default broadcast domain. Go to **Network -> Broadcast Domains** and **Edit** Default. Uncheck mentioned ports.

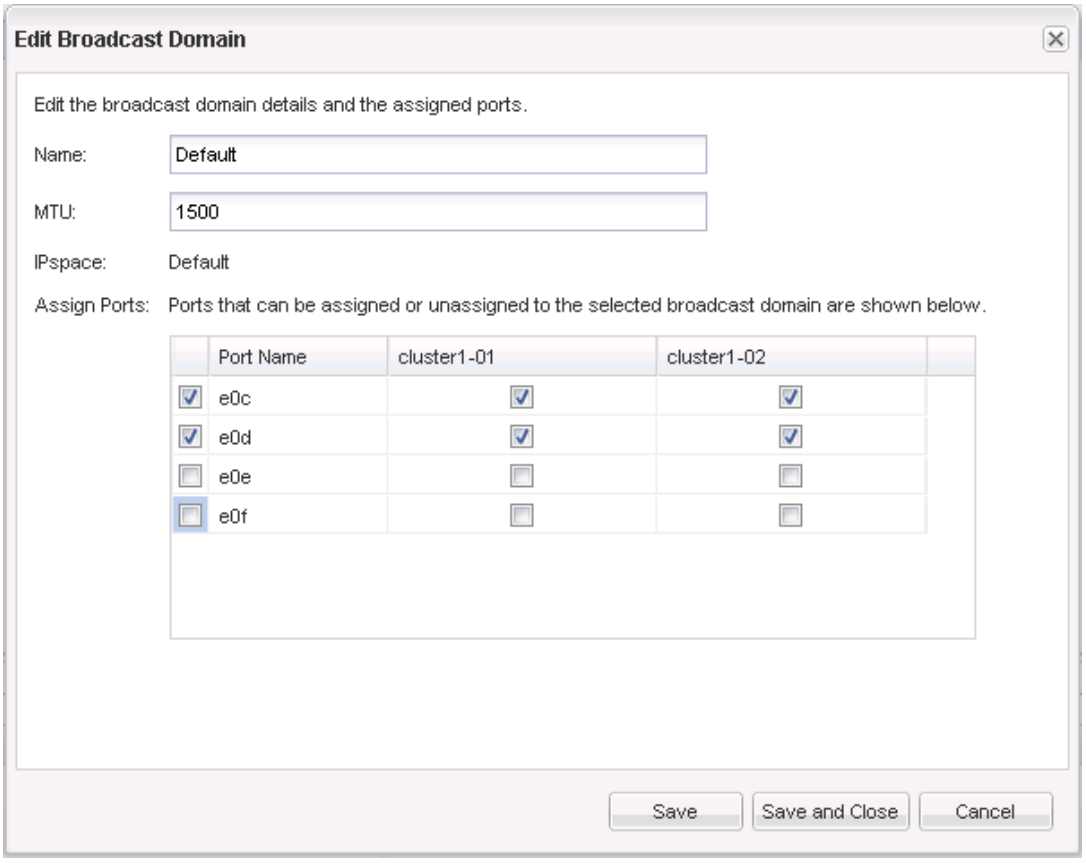

Click **Save and Close.**

TASK 2: Create an interface group

You can aggregate ports that aren't currently hosting logical interfaces and that aren't assigned to any broadcast domain.

Step 1.

In GUI, navigate to **Network -> Ethernet Ports**. You can see that ports e0e and e0f on both nodes haven't assigned Broadcast Domain.

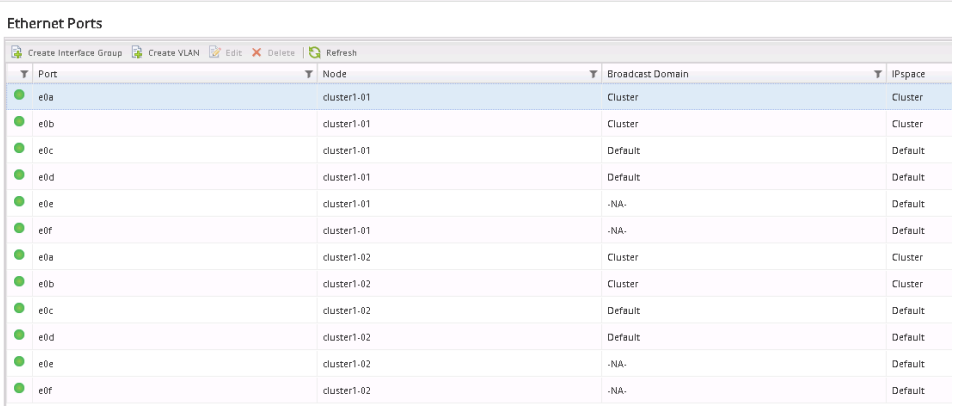

Step 2.

To begin creating an interface group, click **Create Interface Group**.

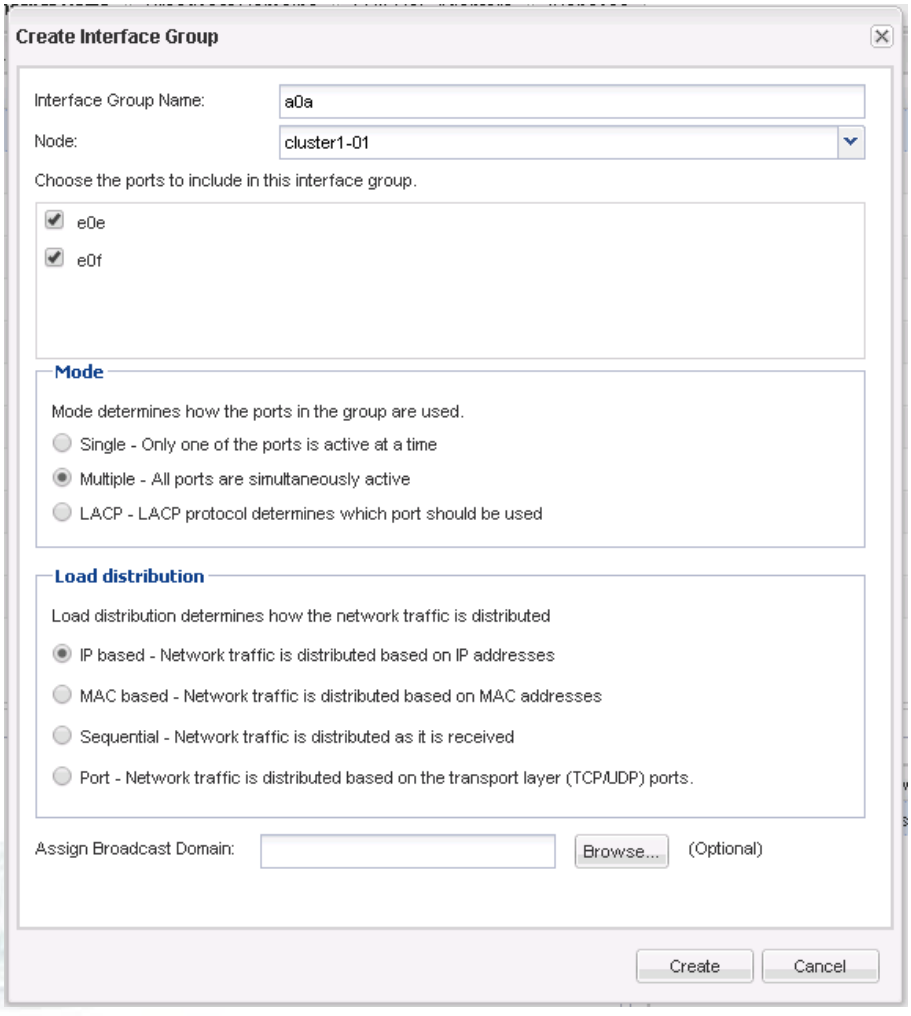

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B

IČO 357 034 66
IČ DPH SK2020341158

Step 3.

Select the **e0e** and **e0f** checkboxes.

Step 4.

Keep the default Interface group name.

Step 5.

Keep the node **cluster1-01.**

Step 6.

Choose **Multiple** for the Mode and **IP based** for the Load distribution.

Step 7.

Click **Create**.

Step 8.

Check the Ports list to verify the creation of the interface group.

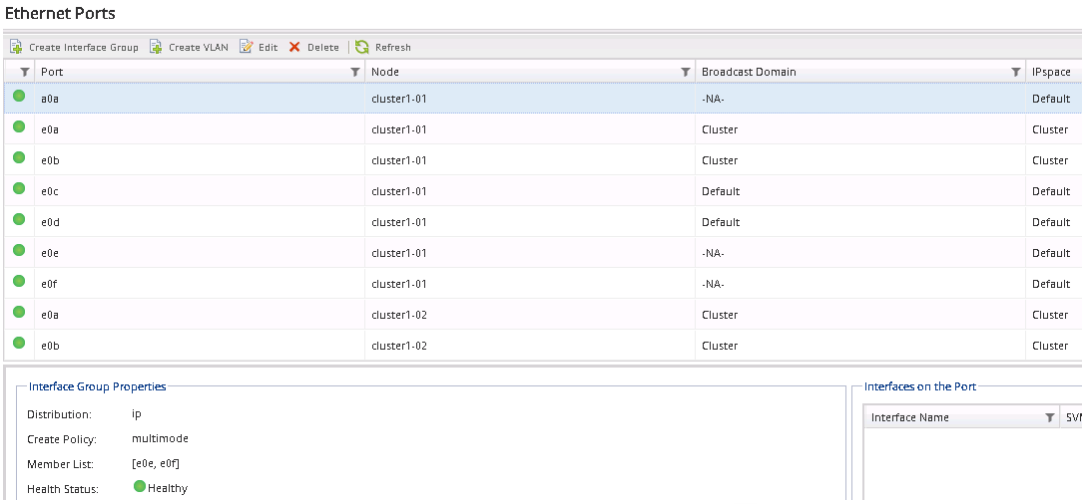

TASK 3: Create a VLAN

Step 1.

In GUI, go to **Network -> Ethernet Ports**.

Step 2.

To begin creating a VLAN, click **Create VLAN**:

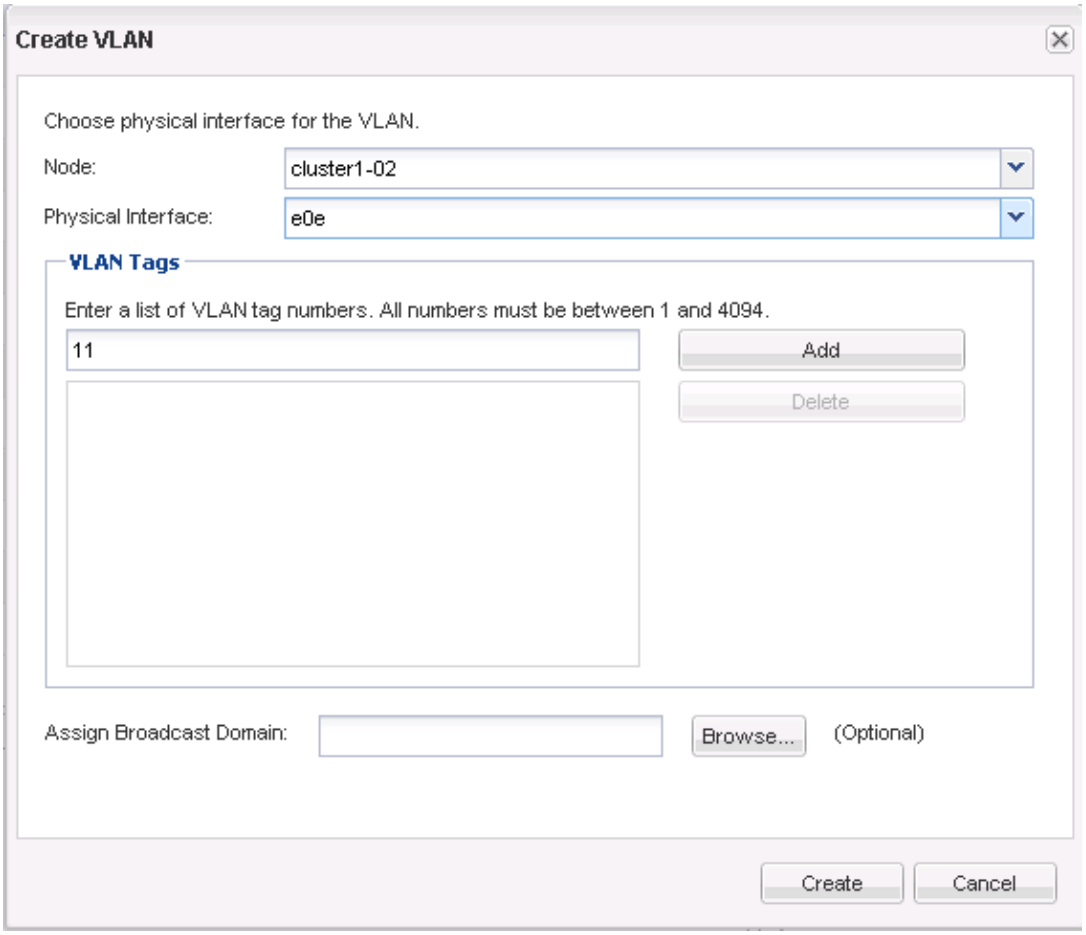

Step 3.

From the menu, choose port **e0e** on node **cluster1_02** to host the VLAN.

Step 4.

Choose a list of VLAN tags to identify each VLAN hosted on this port:

a. Type **11**, and then click **Add**.

b. Repeat with **22** and **33**.

Step 5.

Click **Create** to create the VLANs.

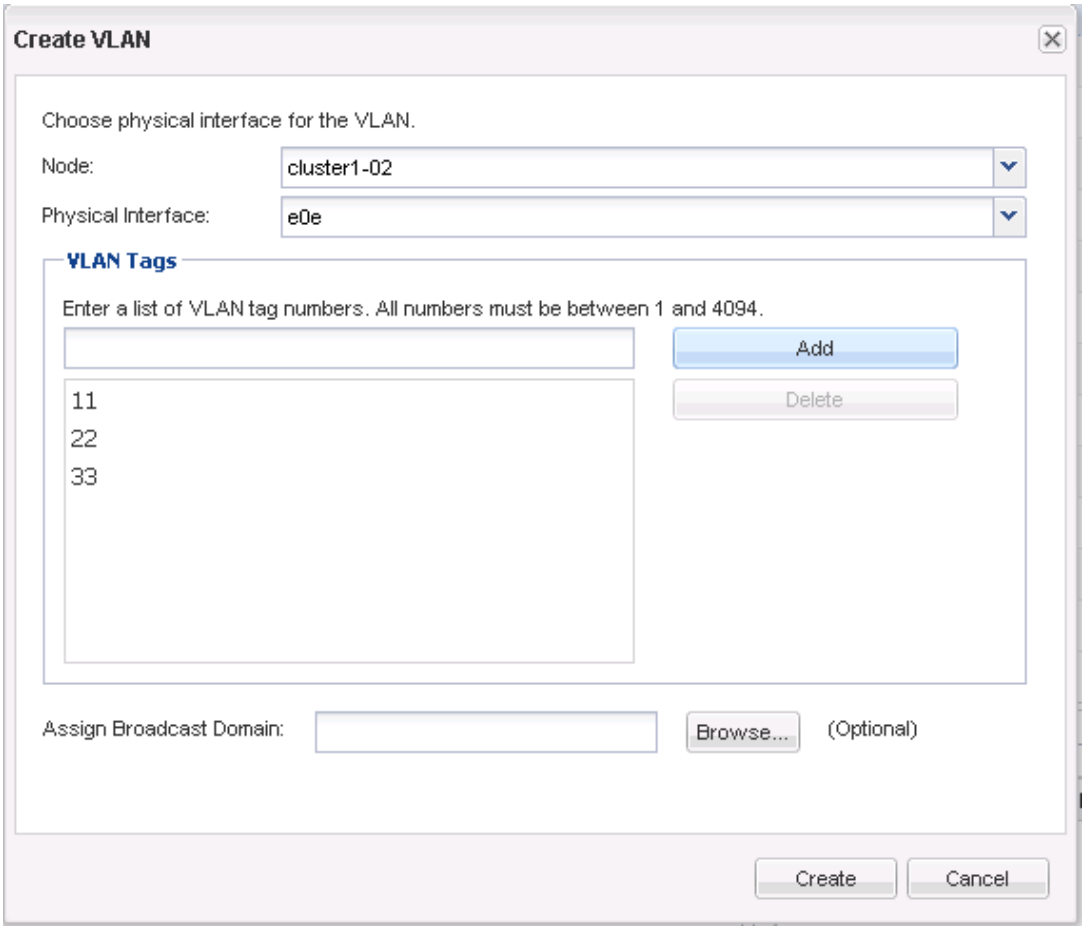

Step 6.

Check the Ports list to verify the creation of the VLANs. (If vlan ports are red, click refresh).

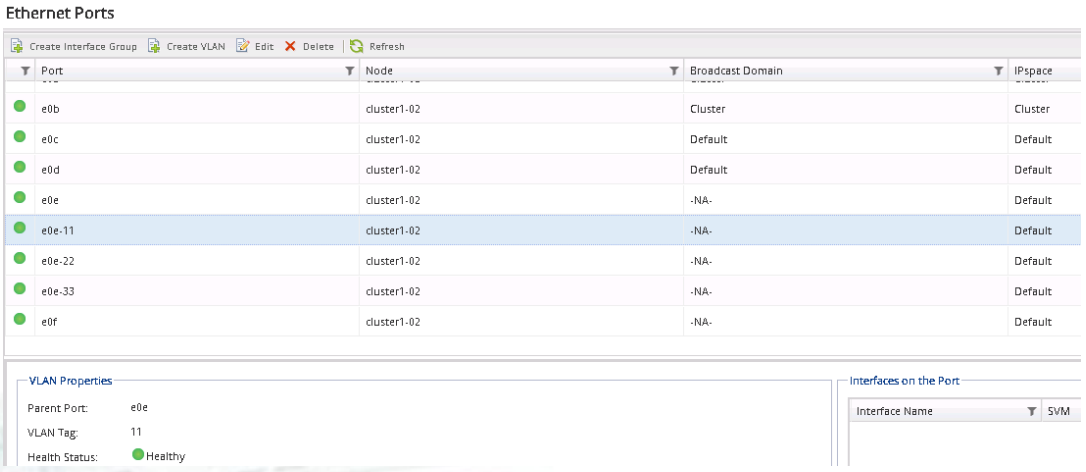

TASK 4. Create Broadcast Domain

Step 1.

In GUI, navigate to **Network -> Broadcast Domains**, and click **Create**.

Step 2.

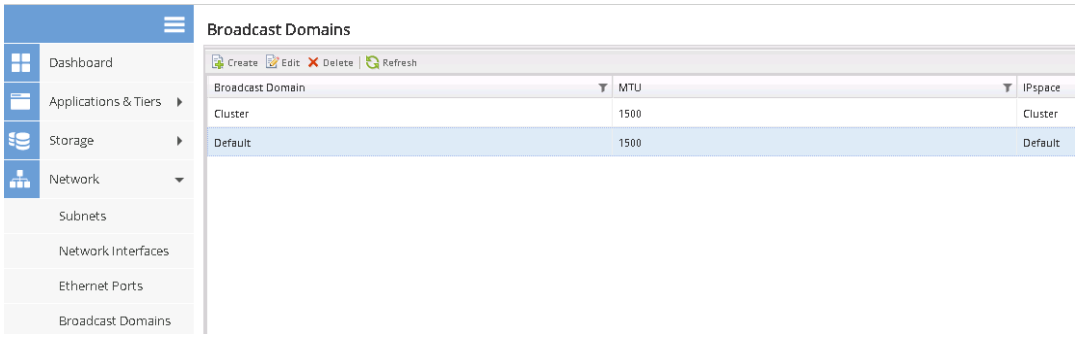

Step 3.

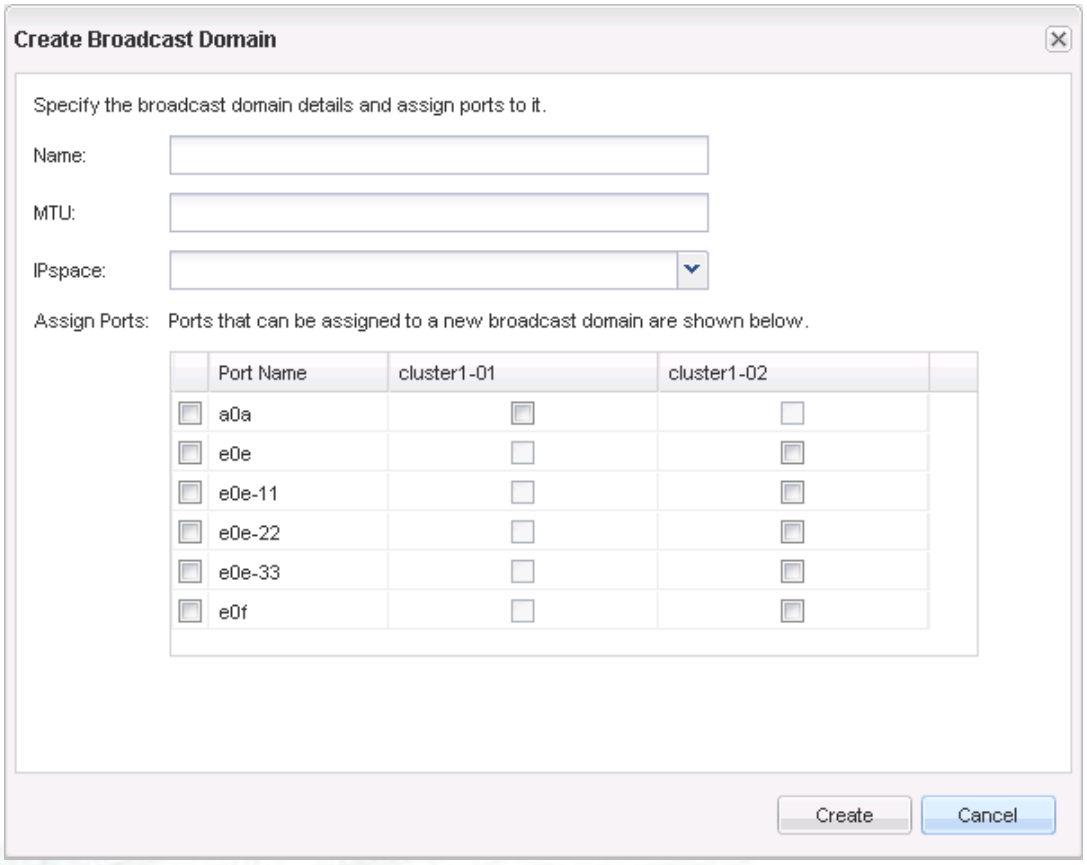

Step 4.

Enter the name **bd1**, MTU size **1500**, IPspace **ipspace1**, and assign all available ports.

Step 5.

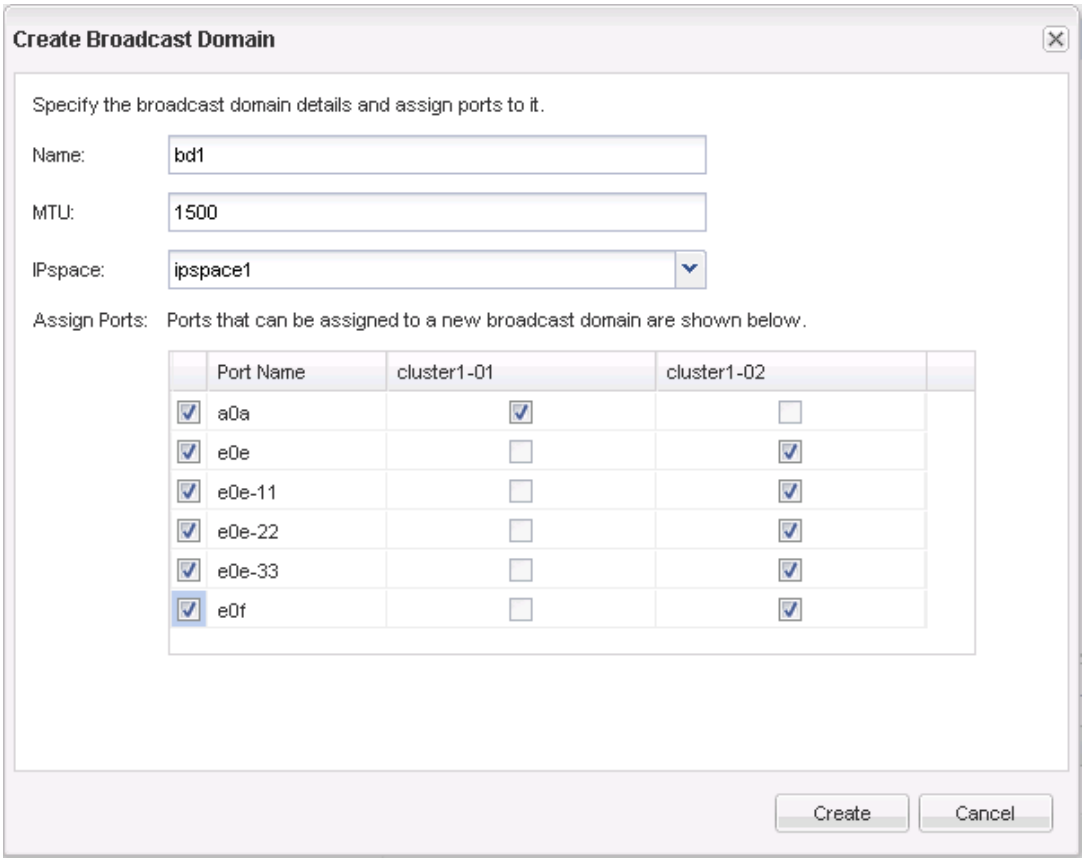

Click **Create.**

Step 6.

Why not all ports are available for assignment to broadcast domain?

TASK 5. Create subnet

Step 1.

In GUI, navigate to **Network -> Subnets**, and click **Create**.

Step 2.

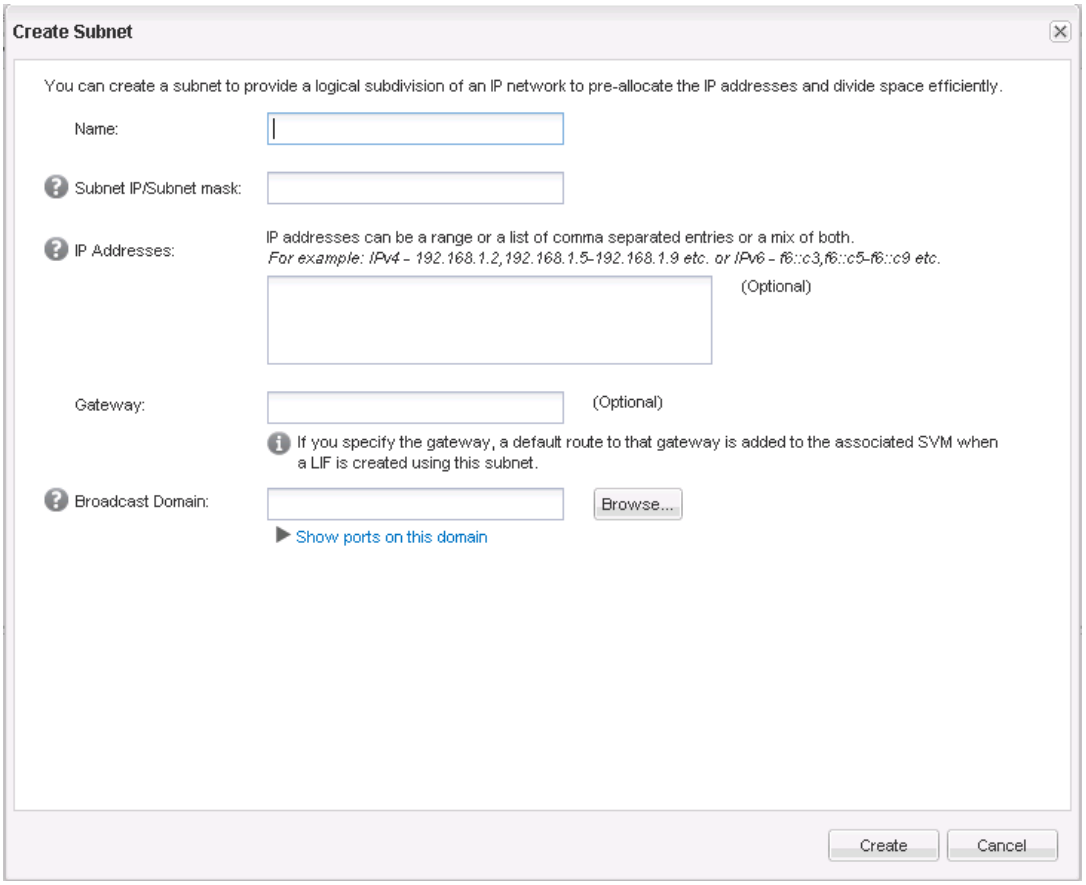

Enter the name **subnet1**, subnet IP/Subnet mask **10.34.x1.0/24**, IP address range will be **10.34.x1.40-10.34.x1.59**, and Broadcast Domain will be **Default**.

Step 3.

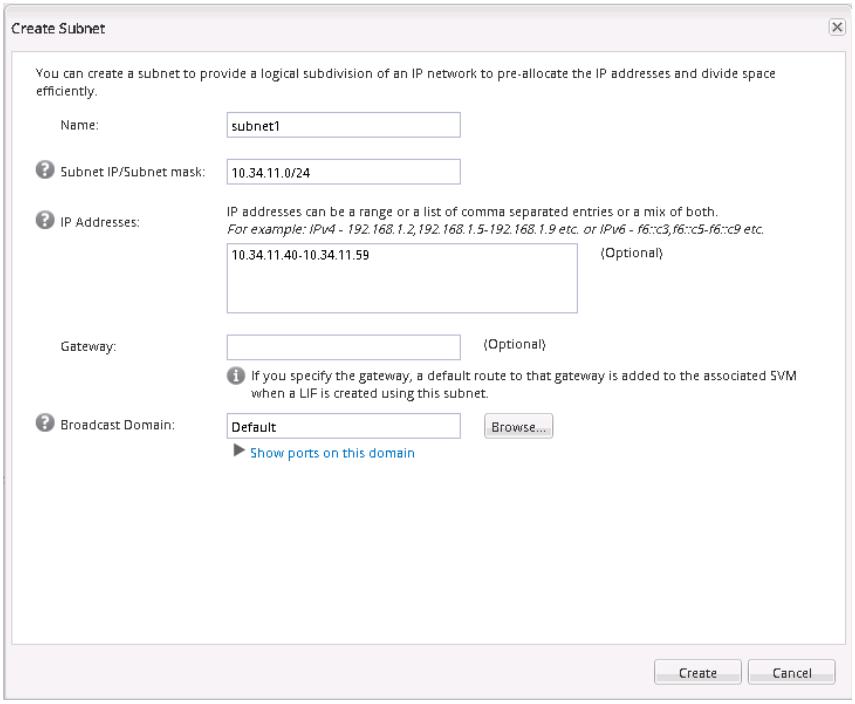

Then click **Create**.

Step 4.

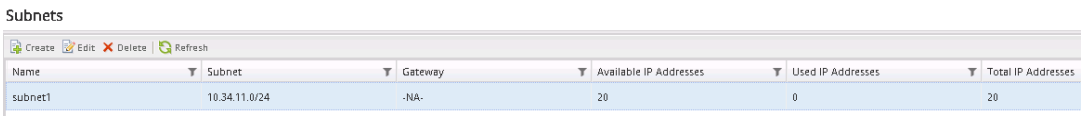

Subnet **subnet1** is ready.

END OF EXERCISE.

LAB 8: Logical networking

In this exercise, you create a NAS data LIF, explore routing groups, migrate a data LIF, send home a data LIF, and fail over a data LIF.

Objectives of this exercise

- Create a NAS data LIF
- Explore routing groups
- Migrate a data LIF
- Rehome a data LIF
- Fail over a data LIF
- Create a failover group

TASK 1: Delete the VLANs and the interface group

Because we need the ports for future exercises, in this task, you delete the VLANs and the interface group that you created in an earlier exercise.

Step 1.

Ensure that no LIFs are hosted by the either the interface group or the VLAN.

cluster1::> network interface show

cluster1::> network port show

cluster1::> net port vlan show

Step 2

Delete all VLANs.

cluster1::> net port vlan delete -node cluster1-02 -vlan-name e0e-11

cluster1::> net port vlan delete -node cluster1-02 -vlan-name e0e-22

cluster1::> net port vlan delete -node cluster1-02 -vlan-name e0e-33

Step 3.

Take interface group a0a offline.

cluster1::> set advanced

Do you want to continue? {y|n}: y

cluster1::> net port modify -node cluster1-01 -port a0a -up-admin false

cluster1::> set admin

Step 4.

Check status of interface a0a.

cluster1::> net port show -node cluster1-01 -port a0a

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B
IČO 357 034 66
IČ DPH SK2020341158

Step 5.

Delete the interface group.

cluster1::> net port ifgrp delete -node cluster1-01 -ifgrp a0a

Step 6.

Delete the broadcast domain **bd1**.

cluster1::> broadcast-domain delete -ipspace ipspace1 -broadcast-domain bd1

Step 6.

Delete the ipspace **ipspace1**.

cluster1::> ipspace delete -ipspace ipspace1

Step 7.

Assign free ports to Default Broadcast Domain.

cluster1::> broadcast-domain add-ports -ipspace Default -broadcast-domain Default -ports cluster1-01:e0e,cluster1-01:e0f,cluster1-02:e0e,cluster1-02:e0f

TASK 2: Create a NAS LIF

Step 1.

Create a data LIF called “data1” for one of your data IP addresses, setting e0e as the home port.

cluster1::> net int create -vserver vs1 -lif data1 -role data -home-node cluster1-01 -home-port e0e -address 10.34.x1.111 -netmask 255.255.255.0 –firewall-policy data

NOTE: The combination of home node and home port is what determines which physical port on which physical node will be the home of this LIF. The home concept exists because data LIFs can migrate to other ports on the same node or to any other node in the cluster.

Step 2.

Review the LIFs.

cluster1::> net int show

Now that a LIF is associated with your virtual server (SVM), the cluster SVM appears in the list. This view of the LIFs includes an Is Home field that indicates whether this LIF resides on its home node and home port, which were chosen when the LIF was created.

Step 3.

View the details of the data LIF.

cluster1::> net int show -vserver vs1 -lif data1

This view shows all the attributes of the new LIF. In addition to the Home Node and Home Port fields, Current Node and Current Port fields exist. If the LIF is migrated, one or both of these Current fields might differ from the Home fields.

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B
IČO 357 034 66
IČ DPH SK2020341158

TASK 3: Migrate a data LIF

Step 1.

Migrate your data LIF to the other node in your cluster.

cluster1::> net int migrate -vserver vs1 -lif data1 -dest-port e0c -dest-node cluster1-02

NOTE: The node and port parameters of the *net int migrate* command are for the destination node and port. You could have migrated to any other data port (on either node) instead. Remember that data LIFs aren't owned by nodes, but they do reside on a node at any given time, transparent to any NFS and CIFS clients.

Step 2.

Review the LIFs.

cluster1::> net int show

In this summary view, your data LIF's Current Node field is the other node, and its Is Home field is false.

Step 3.

View the details.

cluster1::> net int show -vserver vs1 -lif data1

NOTE: Unlike storage failover (SFO), this LIF migration does not cause a reboot of the node from which the LIF is moving. Also unlike SFO, LIFs can migrate to any node in the cluster, not just the high-availability (HA) partner.

Step 4.

Revert the data LIF (send it "home").

cluster1::> net int revert -vserver vs1 -lif data1

NOTE: You don't have to specify the destination node and port. A LIF knows where its home is.

Step 5.

View the LIFs to verify that this one went home.

TASK 4: Rehome a data LIF

Step 1.

Assign the LIF a new home port on port e0d on the opposite node.

cluster1::> net int modify -vserver vs1 -lif data1 -home-node cluster1-02 -home-port e0d

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B
IČO 357 034 66
IČ DPH SK2020341158

Step 2.

Check the status of your data LIF again.

cluster1::> net int show

It has not yet moved to its home port, but because it is in a migrated state (because the current port is no longer the home port), the Is Home status is false.

Step 3.

Revert the data LIF to send it to its new home.

cluster1::> net int revert -vserver vs1 -lif data1

Step 4.

Verify that your data LIF has moved to its new home port.

cluster1::> net int show

TASK 5: Fail over a data LIF

Step 1.

View the failover group that was chosen by default.

cluster1::> net int show -vserver vs1 -lif data1

Step 2.

View the Failover Group values for the LIF.

It is set to *system-defined*.

Step 3.

Show the current LIF failover groups and view the targets defined for the data and management LIFs.

cluster1::> net int show -failover

Groups were automatically created for the data LIF.

Step 4.

To see how it fails over, reboot the node that hosts LIF data1.

cluster1::> system node reboot -node cluster1-02

Step 5.

Type **Y** to confirm.

Step 6.

If you are accessing the cluster via the node management interface on cluster1-02, sign in to the cluster management interface or the node management interface on cluster1-01.

From there you can view all the LIFs by using *net int show* to see that the data LIF fails over to the other node.

NOTE: Because the Auto Revert value of the LIF is set to false, the LIF does not revert to its home port automatically. If the cluster-management LIF is on the node that booted, it also fails over to the opposite node.

Step 7.

When the rebooted node is up and running again, revert the data LIF by issuing this command from either node:

*cluster1::> net int revert -vserver vs1 -lif data**

Step 8.

Look at the data LIF to see that it went home.

If the cluster management LIF did indeed fail over, there is no compelling reason to revert it.

TASK 6: Create a failover group

Step 1.

Create a custom failover group that excludes data ports e0e and e0f on both nodes.

cluster1::> net int failover-groups create -failover-group newfogroup -vserver vs1 -targets cluster1-01:e0c

cluster1::> net int failover-groups show

cluster1::> net int failover-groups add-targets -failover-group newfogroup -vserver vs1 -targets cluster1-01:e0d

cluster1::> net int failover-groups add-targets -failover-group newfogroup -vserver vs1 -targets cluster1-02:e0d

cluster1::> net int failover-groups add-targets -failover-group newfogroup -vserver vs1 -targets cluster1-02:e0c

Step 2.

Assign the cluster failover group to the LIF data1.

cluster1::> net int modify -vserver vs1 -lif data1 -failover-group newfogroup

Step 3.

Check the failover groups assigned to each LIF.

cluster1::> net int show -fields failover-group

Step 4.

Verify the failover targets for data1.

cluster1::> net int show -failover -vserver vs1 -lif data1

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B
IČO 357 034 66
IČ DPH SK2020341158

Step 5.

Set your session to diagnostic privilege mode.

cluster1::> set diag

Warning: These diagnostic commands are for use by NetApp personnel only.

Do you want to continue? {y|n}: y

Step 6.

Deactivate the port that hosts LIF data1.

cluster1::> net port modify -node cluster1-02 -port e0d -up-admin false*

Step 7.

Verify that the port is down.

cluster1::> net port show*

Step 8.

To which port did data1 migrate?

cluster1::> net int show*

Step 9.

Reactivate the inactive port.

cluster1::> net port modify -node cluster1-02 -port e0d -up-admin true*

cluster1::> net port show*

Step 10.

Revert LIF data1 back to its home port.

cluster1::> net int revert -vserver vs1 -lif data1*

cluster1::> net int show*

Step 11.

Set your session back to administrative privilege mode.

cluster1::> set admin*

END OF EXERCISE.

LAB 9: NAS protocols

In this exercise, you create a SVM, configure CIFS and NFS, and access the SVM namespace with CIFS and NFS clients.

Objectives of this exercise

- Create a SVM
- Use OnCommand GUI to configure a SVM and NAS protocols
- Create an export policy
- Create a CIFS share
- Access a CIFS share from a Windows client
- Access the namespace from an NFS client

TASK 0: Check and/or set up system time

Step 1. Set NTP server.

cluster1::> cluster time-service ntp server show

This table is currently empty.

cluster1::> cluster time-service ntp server create -server 10.34.x1.10

cluster1::> cluster time-service ntp server show

Server Version

------------------------------ -------

10.34.x1.10 auto

Step 2.

Set date and time (to be synchronized with your AD server). Time difference between AD controller and clusters have to be less than 5 minutes, otherwise you'll be not able register SVMs to domain. Even you set up NTP server, the system time is not moved immediately but in some steps. So better is set up the time manually.

cluster1::> cluster date show
Node Date Time zone
--------- ------------------------- -------------------------
cluster1-01
1/12/2016 13:25:10 +00:00 Etc/UTC
cluster1-02
1/12/2016 13:25:10 +00:00 Etc/UTC
2 entries were displayed.

Step 3.

cluster1::> cluster date modify -timezone Europe/Bratislava -date "1/12/2016 12:17:25"

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B

IČO 357 034 66
IČ DPH SK2020341158

```
cluster1::> cluster date show                                   Node      Date
Time zone
--------- ------------------------ ------------------------
cluster1-01
        1/12/2016 12:17:29 +01:00 Europe/Bratislava
cluster1-02
        1/12/2016 12:17:27 +01:00 Europe/Bratislava
2 entries were displayed.
```

TASK 1: Configure a SVM to serve CIFS and NFS

Step 1.

In the GUI navigation frame, click **Storage > SVMs (Storage Virtual Machines)**.

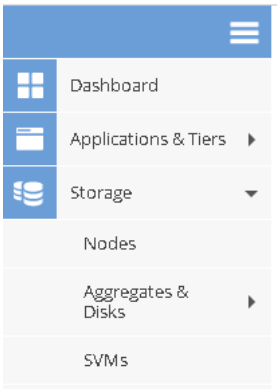

Step 2.

In the SVMs toolbar, click **Create** to create a new SVM.

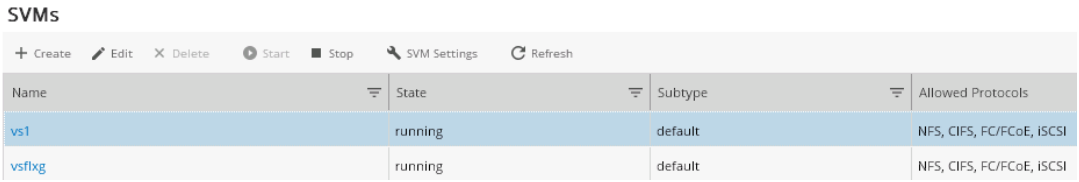

Step 3.

Enter **vs2** as the name for your new SVM.

Step 4.

Select the aggregate **n02_aggr1**.

NOTE: Keep in mind that SVMs are independent of aggregates and nodes. Here you are simply choosing an aggregate to host the new SVM's root volume. Best practices suggest that you do not use a node root aggregate to host a SVM root. If you do so, you receive a warning.

Step 5.

To complete the SVM Details page of the wizard, take these actions:

a. Data protocols: **CIFS** and **NFS**
b. Language: **C.UTF-8 [c.utf_8]**
c. Security style: **UNIX**
d. Root aggregate**: n02_aggr1**
e. DNS configuration: Leave defaults
f. Click **Submit & Continue**.

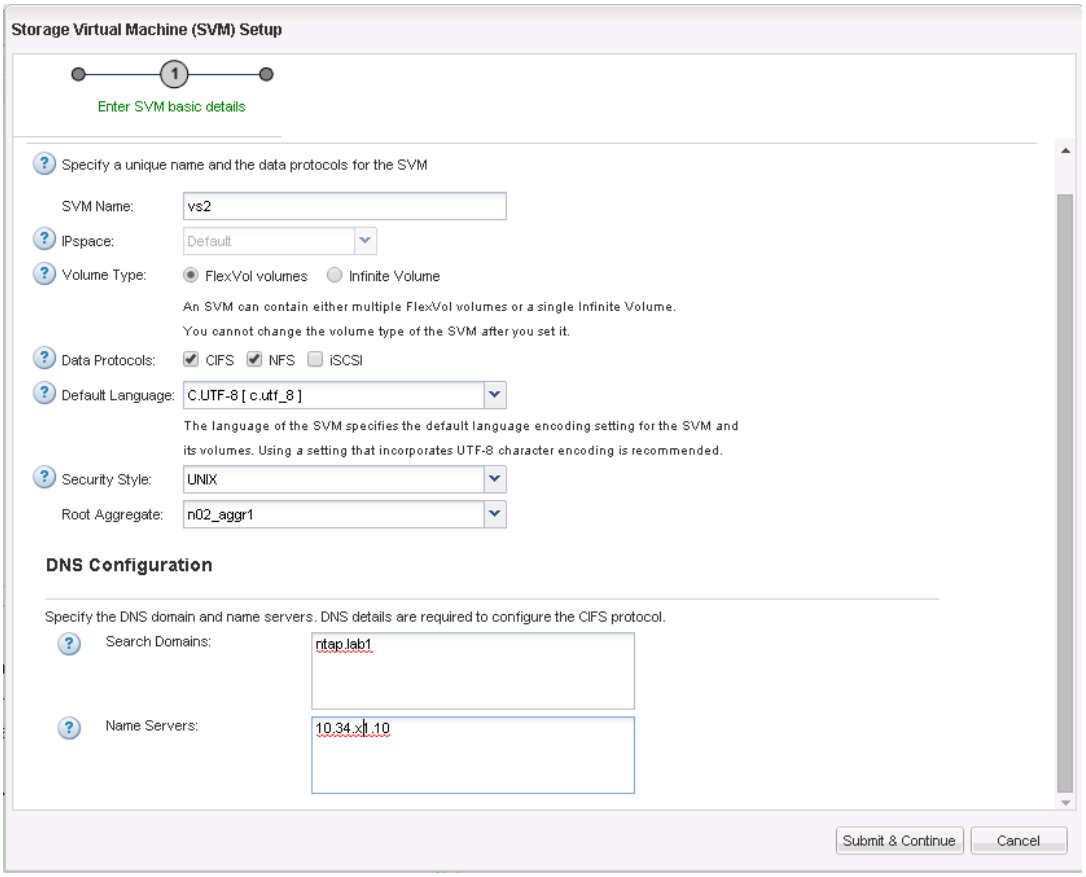

Step 6.

Select the **Retain the CIFS data LIFs** checkbox.

This option specifies that the data LIF supports both CIFS and NFS sessions. Assign IP address select “Using a subnet”. Subnet will be **subnet1** and let radio-button “Auto-select the IP address from this subnet” selected.

As a home port select **cluster1-02:e0d**.

Step 7.

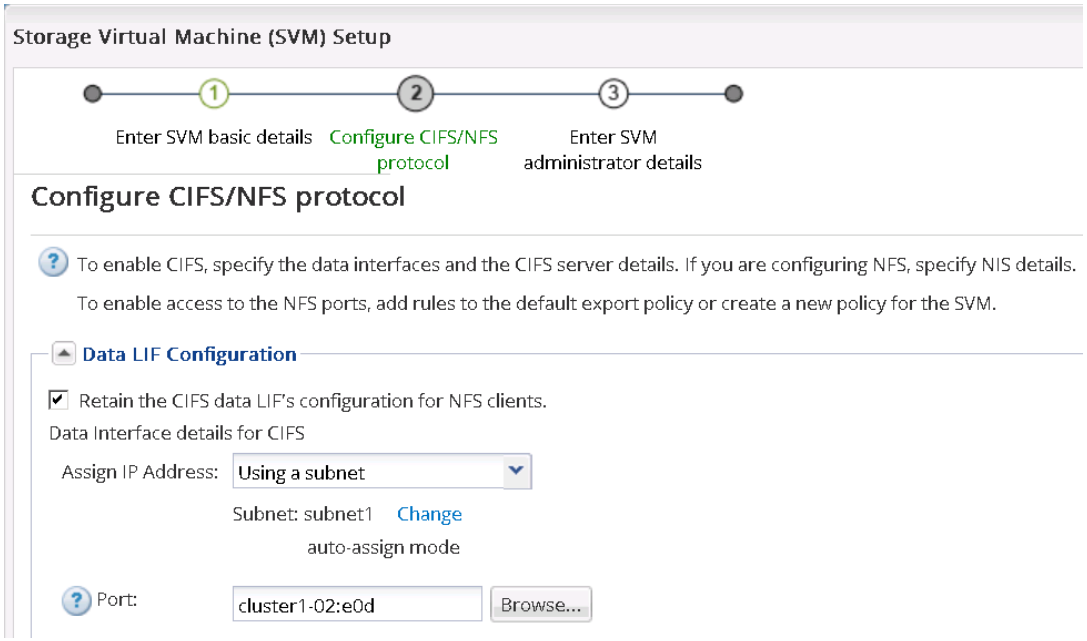

Step 8.

In the CIFS Server Configuration type **ntapvs2** as the name for your CIFS server, and enter your fully qualified domain name **ntap.lab1**.

Step 9.

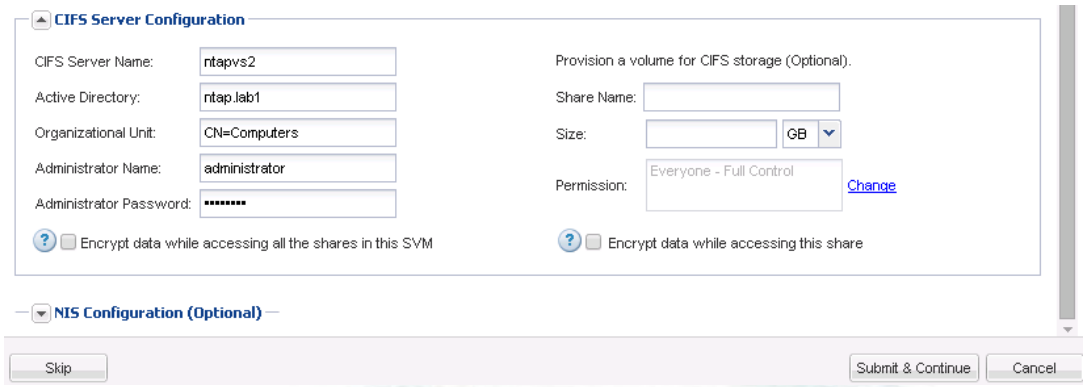

Enter the domain administrator name **administrator** and the password supplied by your instructor, and then click **Submit & Continue**.

Step 10.

On the SVM administration page, click **Skip**.

Step 11.

Review the configuration summary and click **OK**.

Step 12.

In GUI, select **SVMs (Storage Virtual Machines).**

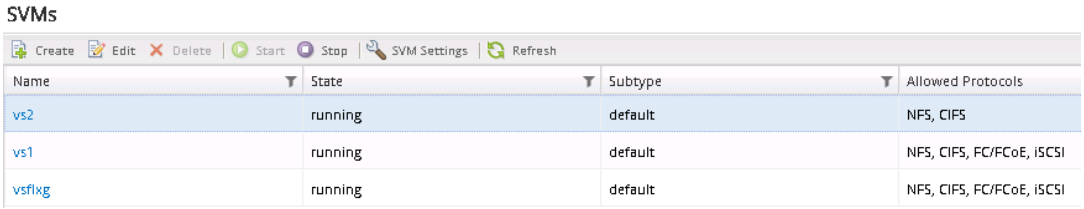

Step 13.

In the list of **SVMs**, select **vs2** and, on the toolbar, click the **Edit** and then **Services**.

Step 14.

Be sure that only the **files** in first column are selected

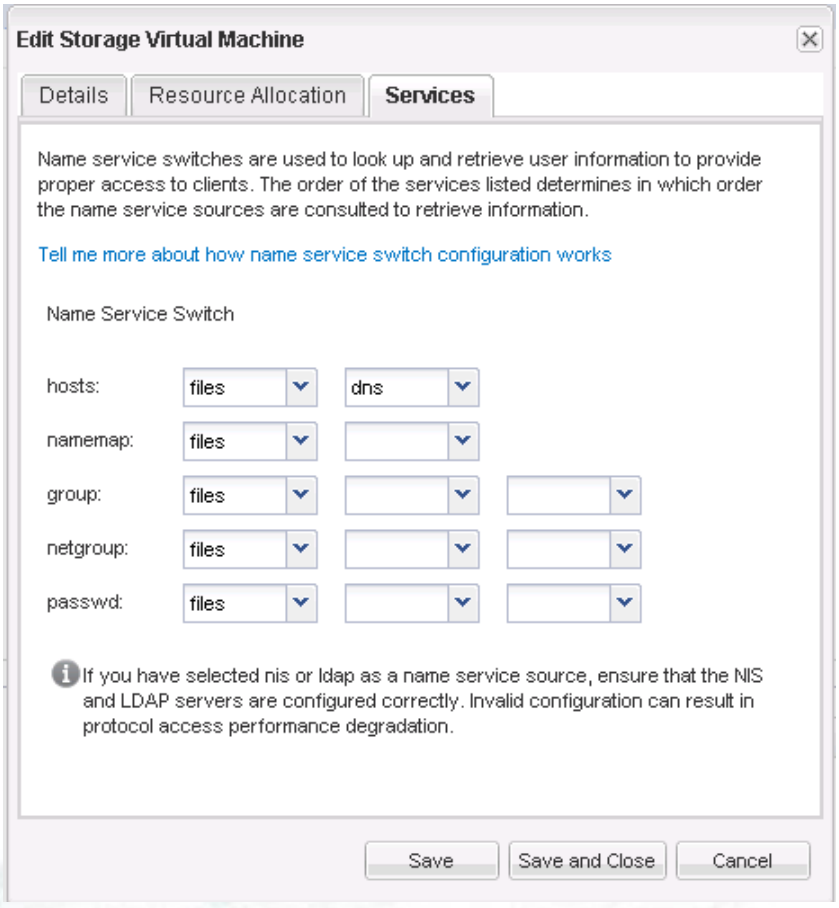

Step 15.

Click the **Resource Allocation** tab.

Step 16.

Select **Delegate volume creation**, and then choose aggregates that are available to host volumes for this SVM by clearing the checkboxes for the aggr0 aggregates and selecting the rest.

Best practices suggest that data volumes should never be stored on a node's aggr0 aggregate.

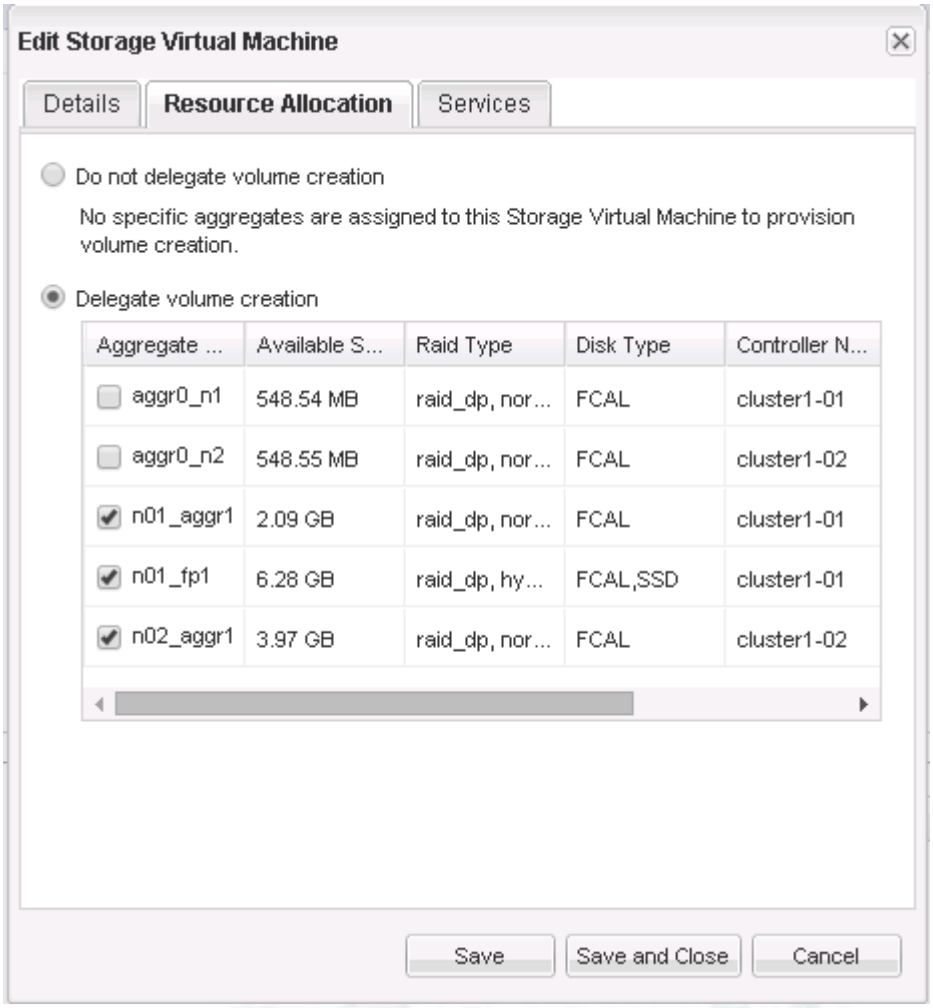

Step 17.

Click **Save and Close**.

Step 18.

If the Configure Services warning appears, click **OK**.

TASK 2: Create a NAS data LIF

Step 1.

Select **Network** > **Network Interfaces**.

Step 2.

Click the **Create** button to begin creating a data LIF.

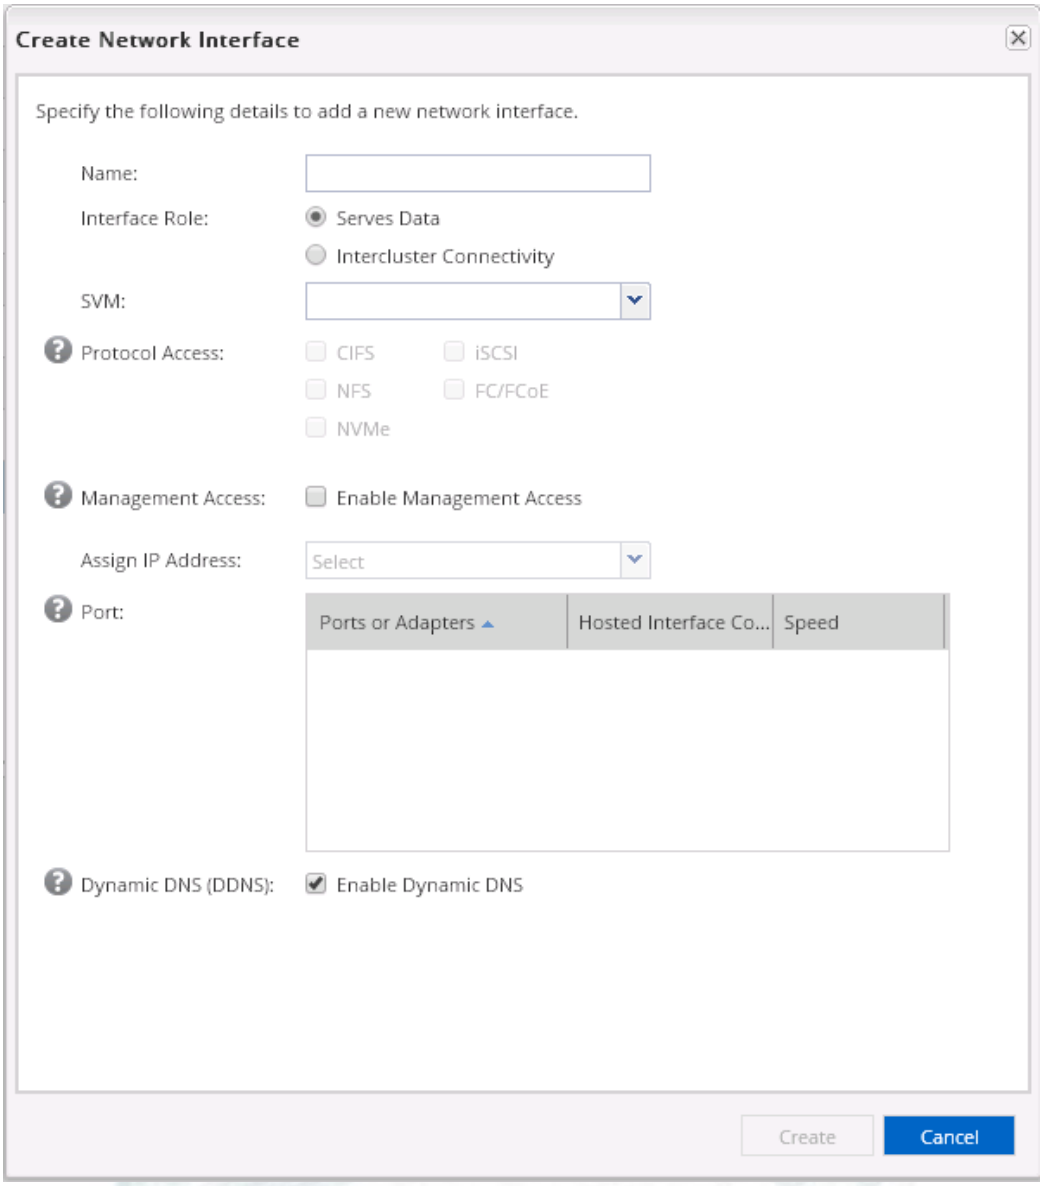

Step 3.

Name the LIF **vs2_cifs_nfs_lif2** and select the **Serves Data** role. Select SVM **vs2** and choose NAS protocols **CIFS** and **NFS**. Assign IP address **Using a Subnet**. Subnet will be **subnet1** and let radio-button **"Auto-select the IP address from this subnet**" selected.

As a home port select **cluster1-01:e0e**.

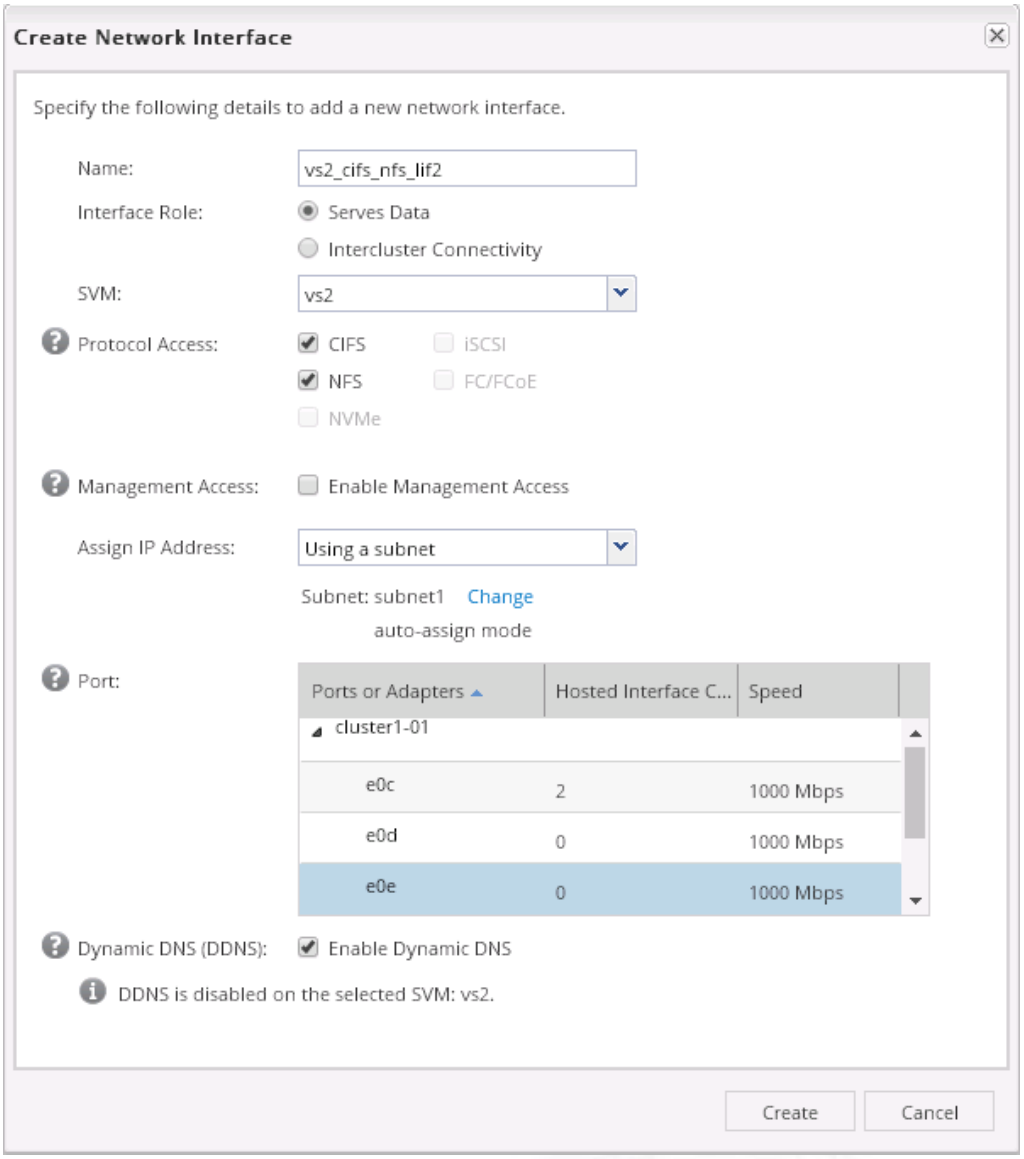

Click **Create**.

Step 4.

Verify creation of the data LIF.

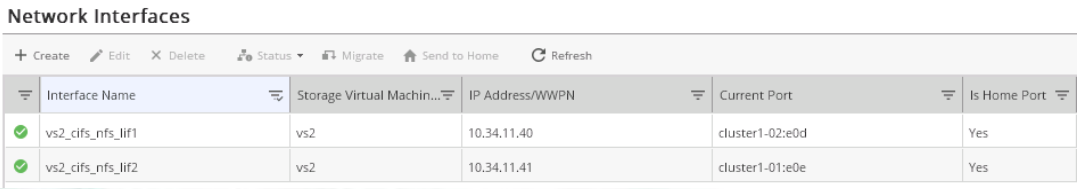

TASK 3. Create an export policy

Step 1.

Select **Storage > SVMs** > **vs2** > **SVM Settings** > **Export Policie**s.

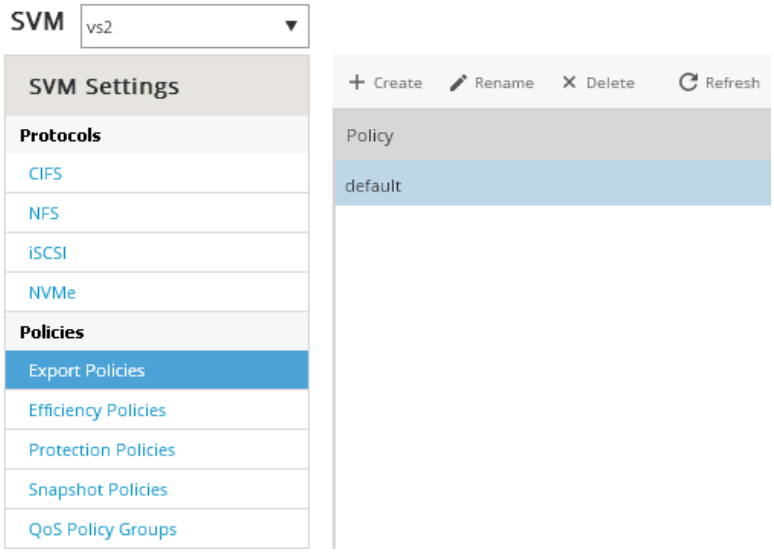

Step 2.

Click the **Create Policy** button.

Step 3.

Name the policy **vs2def** and click **Add** to add a rule.

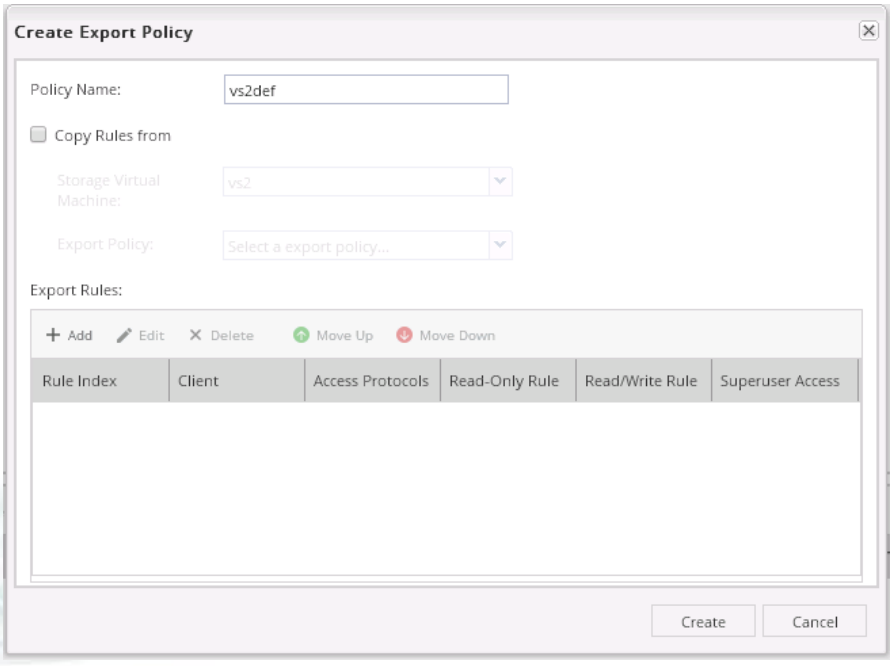

Step 4.

Specify **0.0.0.0/0** for the client specification.

Select both **CIFS** and **NFS** for protocols.

Select all checkboxes for each access type and click **OK**

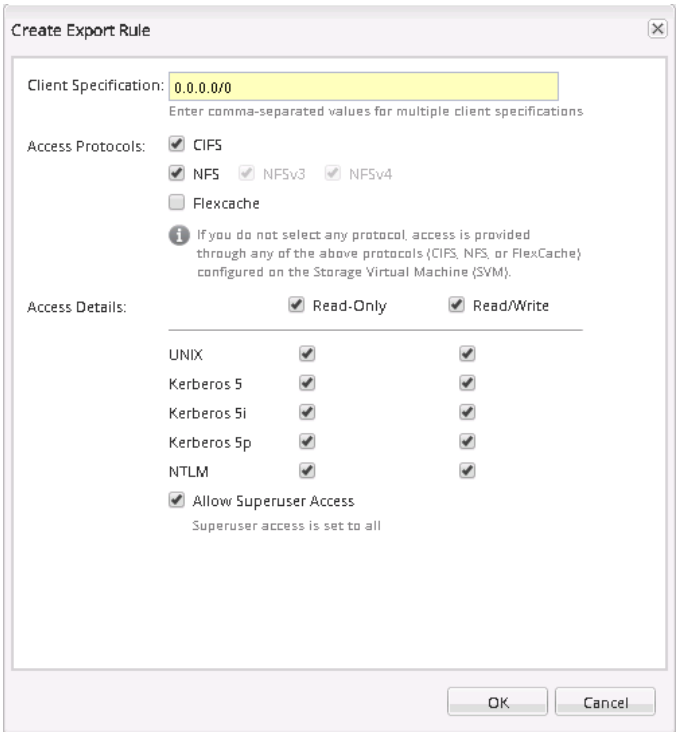

Step 7.

Click **Create**.

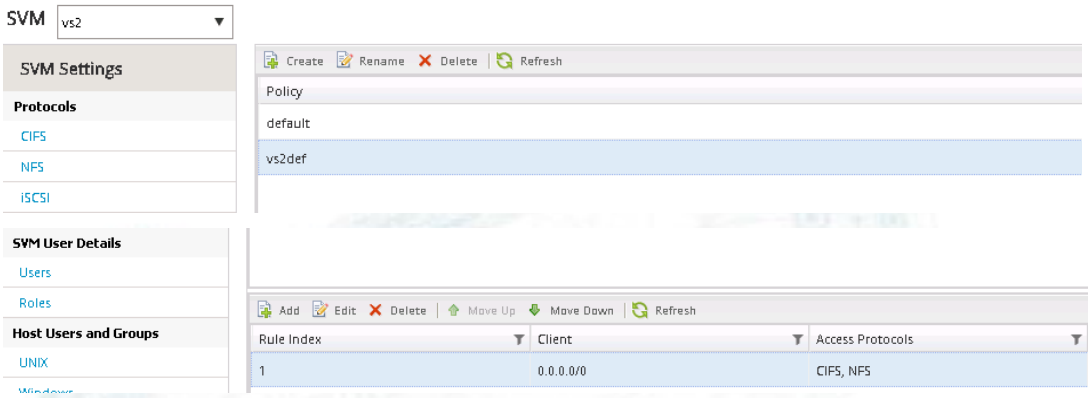

TASK 4: Create and export of a volume

Step 1.

Select **Storage** > **Volumes**

Step 2

In the **Volumes** toolbar, click the **Create FlexVol** button to create a new volume.

Step 3.

Name the volume **vs2_vol01**. Click the **Choose** button to select a home aggregate for the new volume. Select **n01_aggr1** and click **OK**.

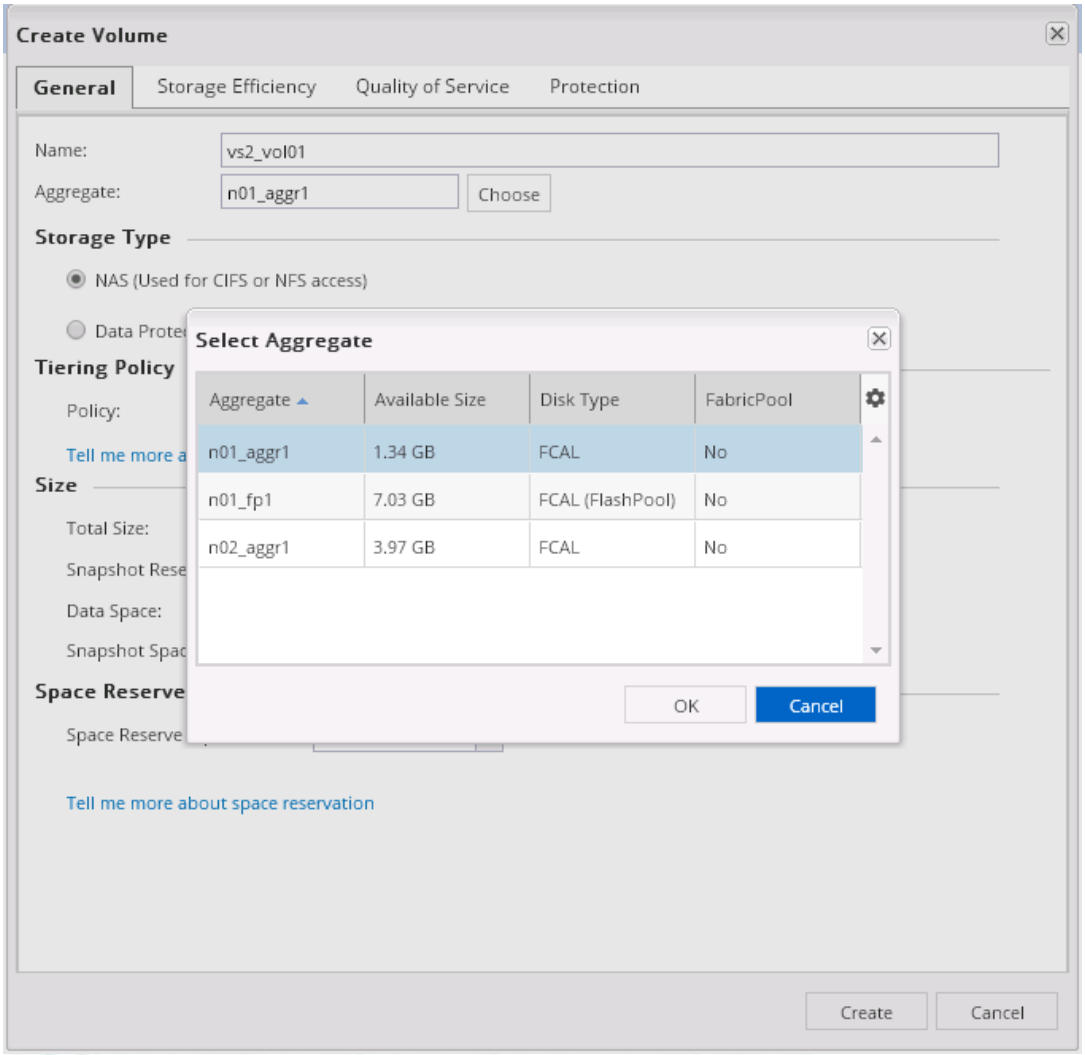

Step 4.

In the Storage Type area, ensure that **NAS** is selected and tiering policy let default **snapshot-only.**

Step 5.

To complete the volume specifications, enter these values:

a) Set the Total Size to **400 MB**.

b) Leave the Snapshot Reserve at the default.

c) Choose Space reserve as **Thin Provisioned**.

d) Click **Create**.

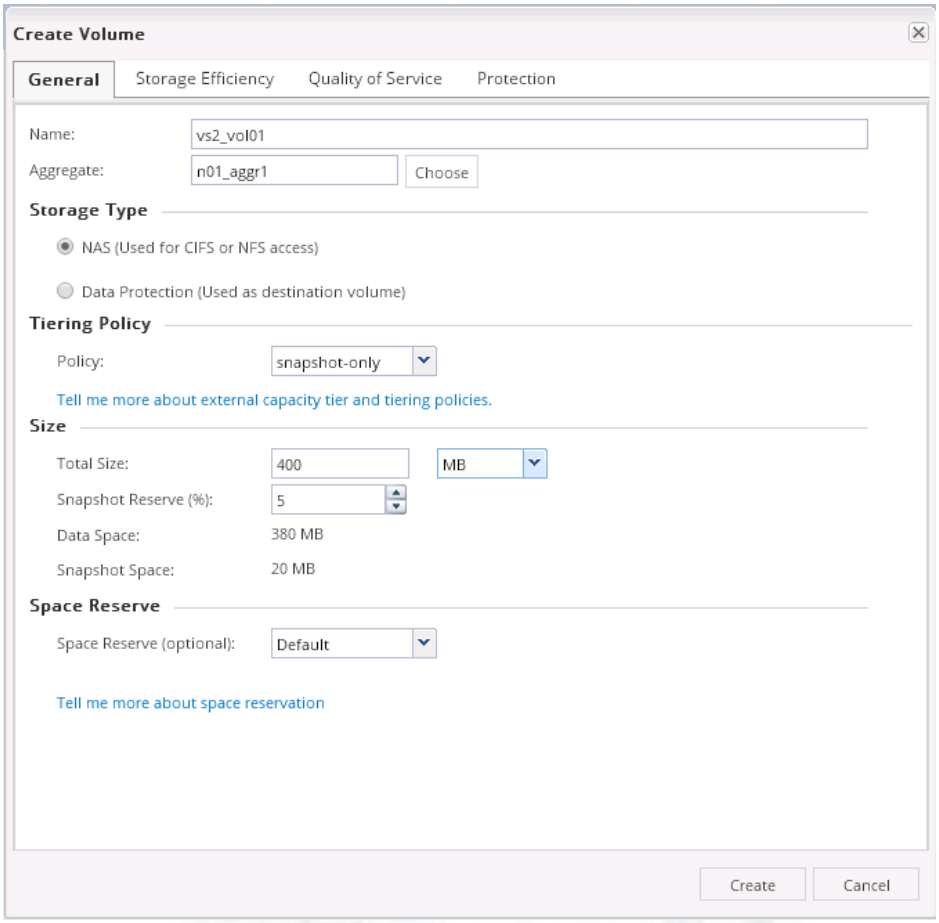

Step 5.

After the wizard completes the volume creation, verify the new volume in the volume list.

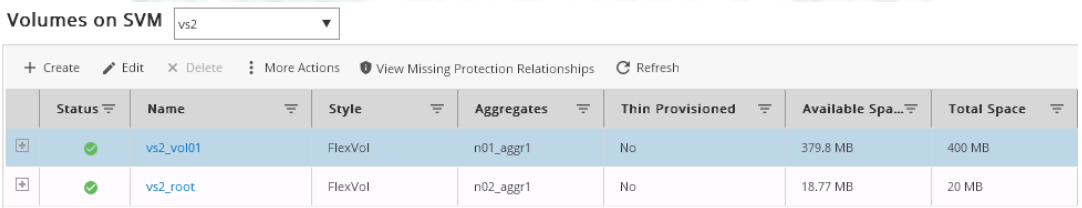

Step 6.

Perform these actions to modify permissions on the volume that was created at the beginning of this task:

a. With **vs2_vol01** selected, click **Edit** and check Configure UNIX permissions (Optional)

b. Select all of the **Read**, **Write**, and **Execute** checkboxes for **Owner**, **Group**, and **Others**.

c. Clear the **Thin Provisioned** checkbox.

d. Click **Save and Close**.

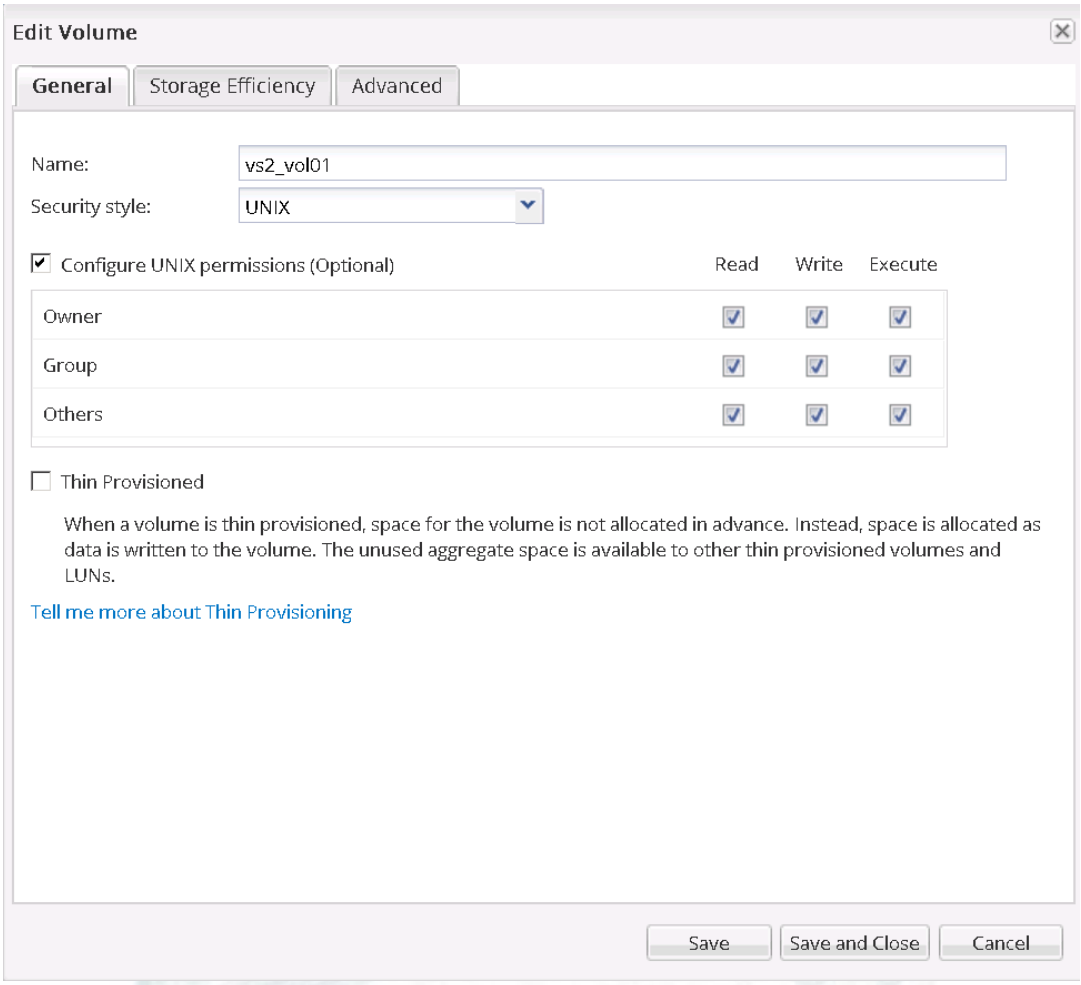

Step 7.

Repeat the previous step for the **vs2** SVM root volume.

Step 8.

Select **Storage** > **vs2** > **Junction Path** and verify where the new volume has been mounted in the namespace.

This page indicates that the volume is accessed by clients as /vs2_vol01 and that the default export policy has been assigned to the volume.

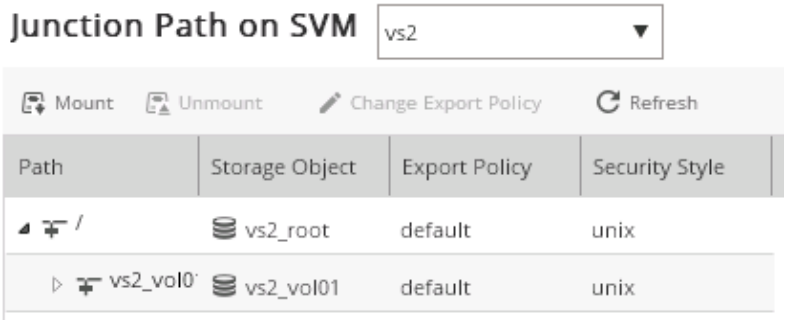

Step 12.

To unmount the volume and remount it with a different junction name, select the volume on the **Junction Path** page, click **Unmount** and, leaving the "Force volume unmount operation" checkbox unselected, click **Unmount**.

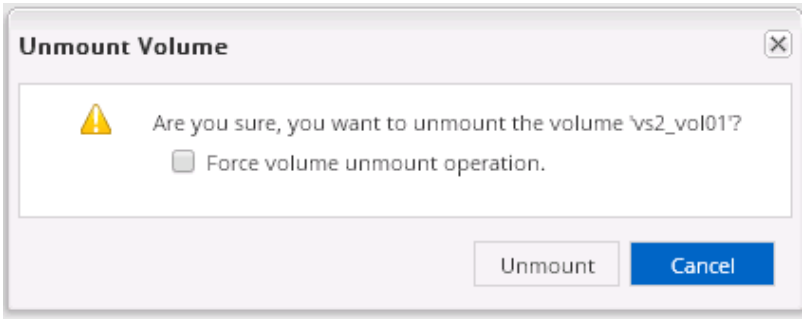

Step 13.

On the Junction Path toolbar, click **Mount**.

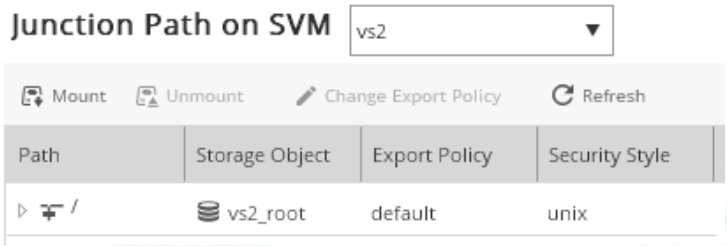

Step 14.

Perform the following actions:

a. Select volume **vs2_vol01**.

b. Type junction name **vs2vol01**.

c. Click **Browse** and select the root directory as the junction path.

d. Click **Mount**.

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B

IČO 357 034 66
IČ DPH SK2020341158

NOTE: In this exercise, the junction name is slightly different from the volume name. It is not necessary for the names to be the same. The volume name is used to reference the volume within the cluster. The junction name is used to reference the root of the volume in the namespace.

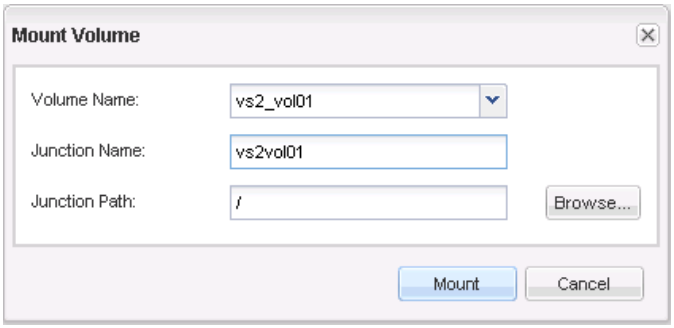

Step 15.

Verify that the junction path in the **vs2** namespace is correct.

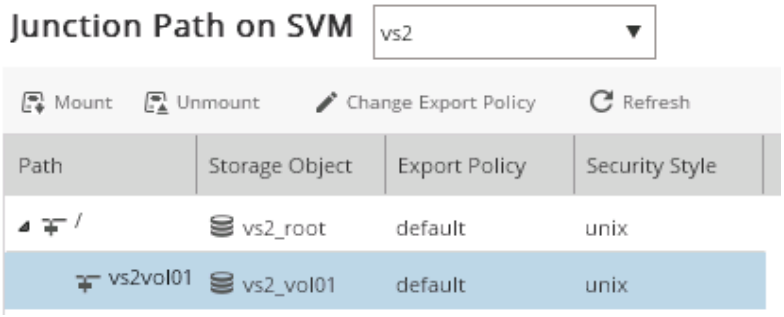

Step 16.

Select the volume **vs2_root** and click **Change Export Policy**.

Step 17.

Select export policy **vs2def** and click **Change**.

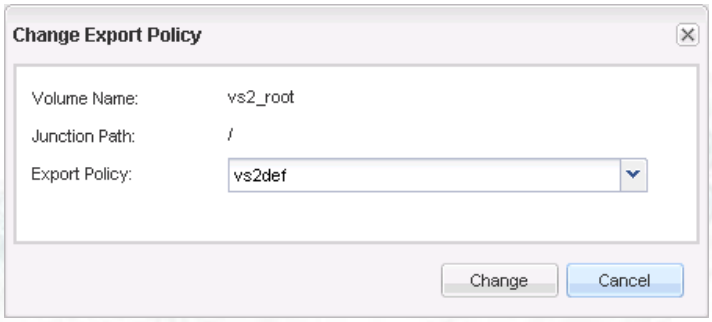

Step 18.

Repeat the process for **vs2_vol01**.

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B

IČO 357 034 66
IČ DPH SK2020341158

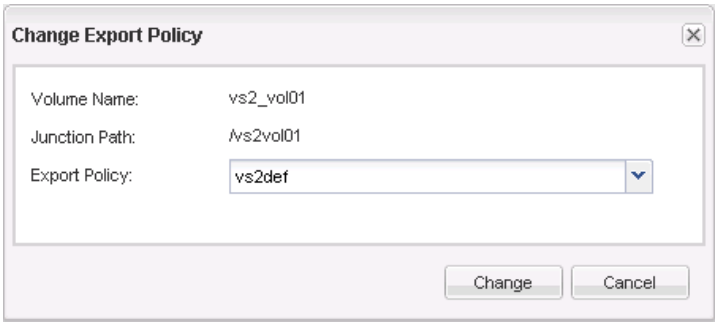

TASK 5: Create CIFS share

Step 1.

To create a CIFS share for the new volume, select **Storage** > **Shares** -> **vs2**, and click the **Create Share** button.

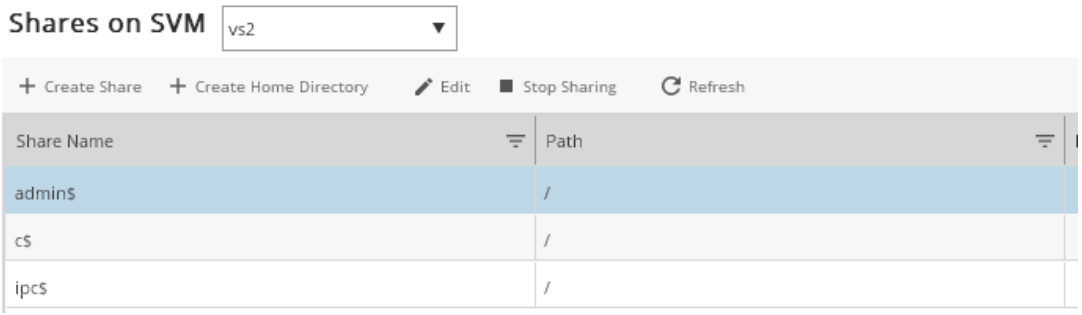

Step 2.

When the Create Share window appears, click the **Browse** button to select a folder to share.

NOTE: It is helpful to remember that you share folders, not volumes. You can share a subdirectory within a volume or any junction path.

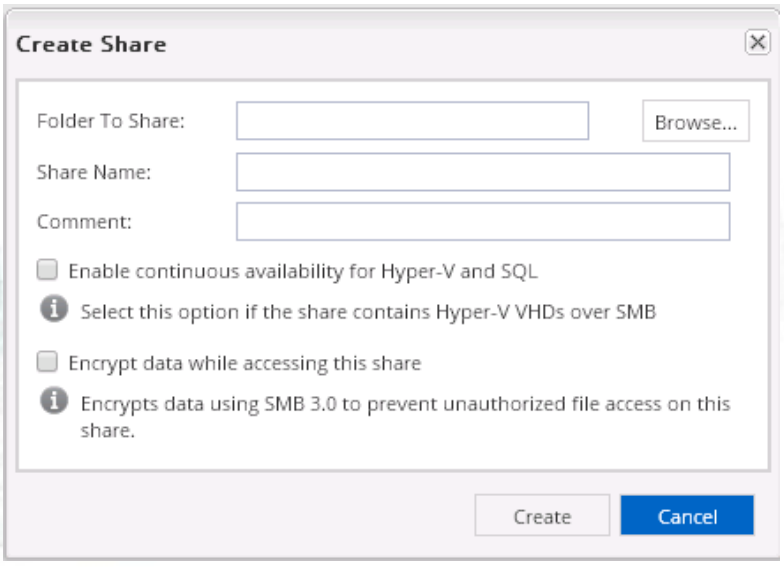

Step 3.

Expand the root directory to view the volume that was recently mounted in the namespace, select the folder that represents the junction for the volume, and click **OK**.

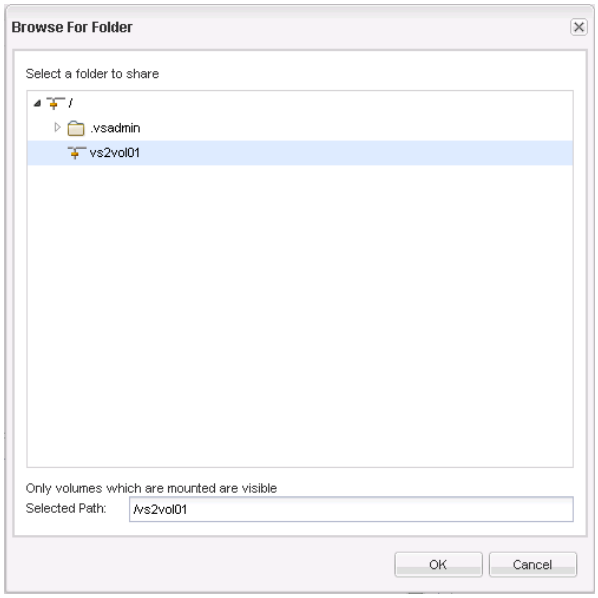

Step 4.

Change the Share Name to **vol1** and click **Create**.

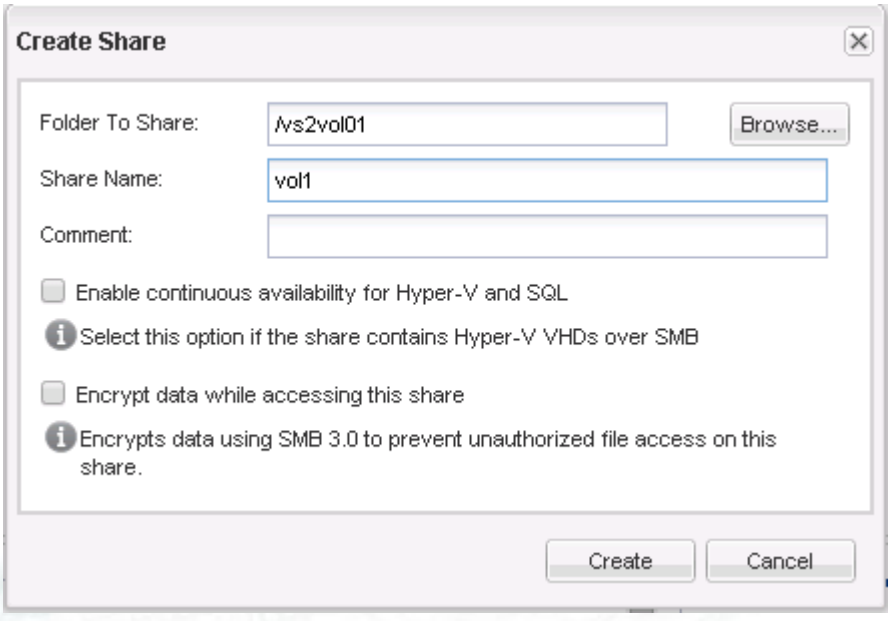

Step 5.

Verify that the new CIFS share appears in the **Shares** list.

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B

IČO 357 034 66
IČ DPH SK2020341158

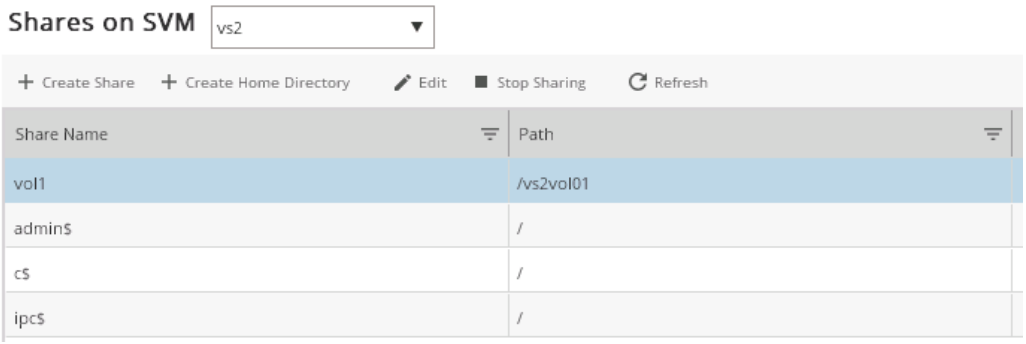

Step 6.

Create a second CIFS share, sharing the root directory as **rootdir**.

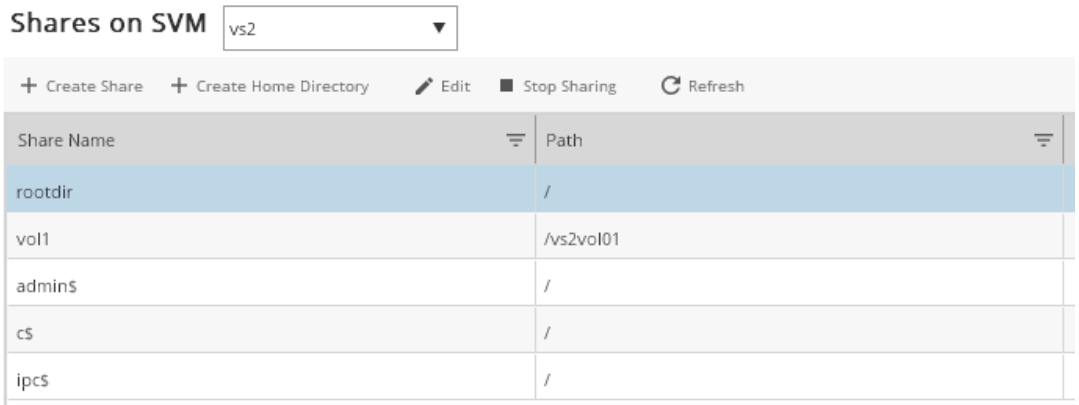

TASK 6: Access your CIFS share from a Windows client

Step 1.

From the Windows command line, access the CIFS share by entering the following command:

PS C:\> net view ntapvs2

You access the CIFS server by the name that you assigned when you used the *cifs create* command to create the server.

Step 2.

If you encounter "Error 53. The Network Path was not found. " attempt to identify the problem by performing one or more of the following actions:

- Verify that the export policy allows CIFS access.
- Verify that CIFS access is enabled for the SVM.
- Review the data LIF setup to ensure that the LIF has the proper routing group and that you can ping the IP address from the Windows client.

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B
IČO 357 034 66
IČ DPH SK2020341158

- Verify that can you ping the CIFS server by name.
- If you cannot ping the CIFS server by name (the DNS is not set up to resolve the CIFS server**), attempt to access the CIFS server with the IP address of a data LIF**.

Step 3.

Map a drive letter to the CIFS shares.

*PS C:\> net use * \\ntapvs2\vol1*

*PS C:\> net use * \\ntapvs2\rootdir*

Step 4.

Verify successful access.

PS C:\> dir z:

PS C:\> dir y:

TASK 7: Configure CIFS home directories

Step 1.

In GUI, select **Storage > Shares on SVM -> vs2**, and click the **Create Home Directory** button.

Step 2.

Enter the share name and the relative path name **%w**, and click **Manage**.

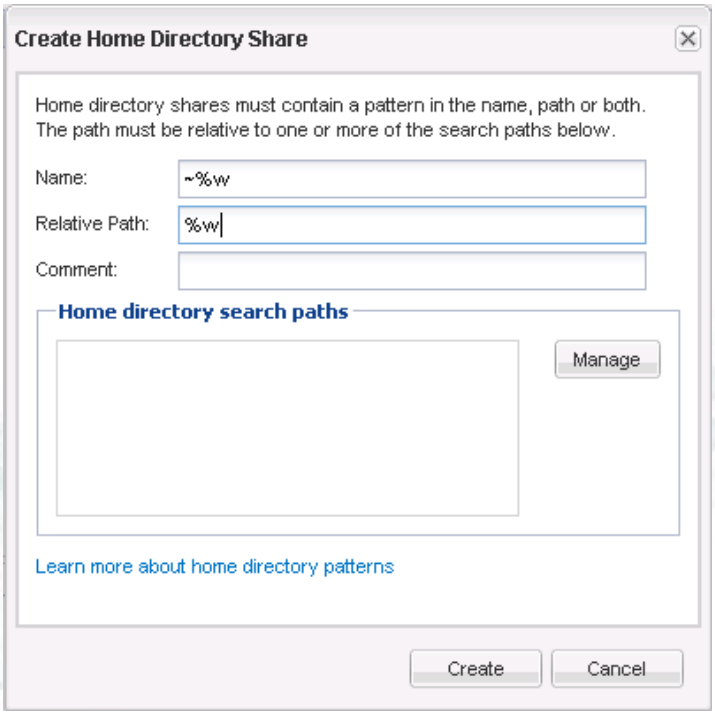

Step 3.

Click **Browse**.

NOTE: You can add one or more search paths for a directory name match.

Step 4.

Select **/vs2vol01** and click **OK**.

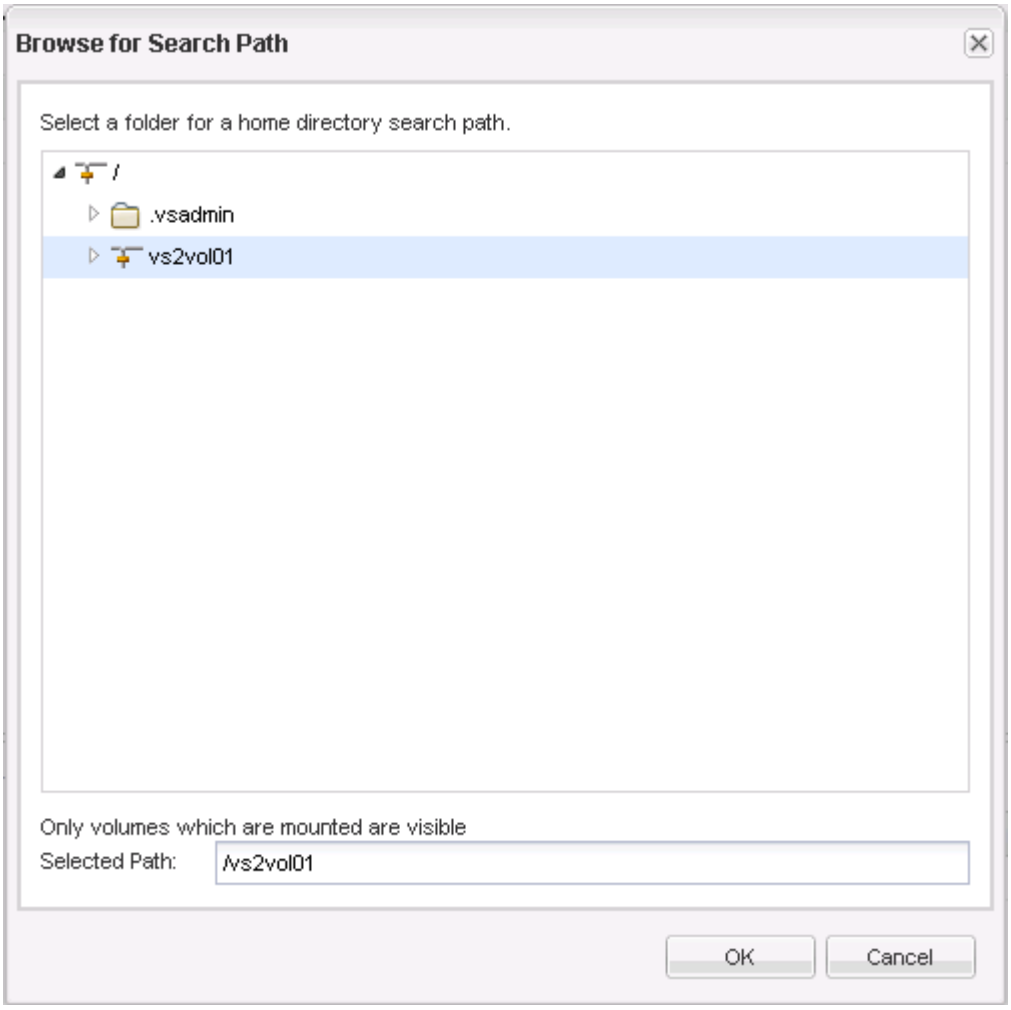

Step 5.

Click **Add**.

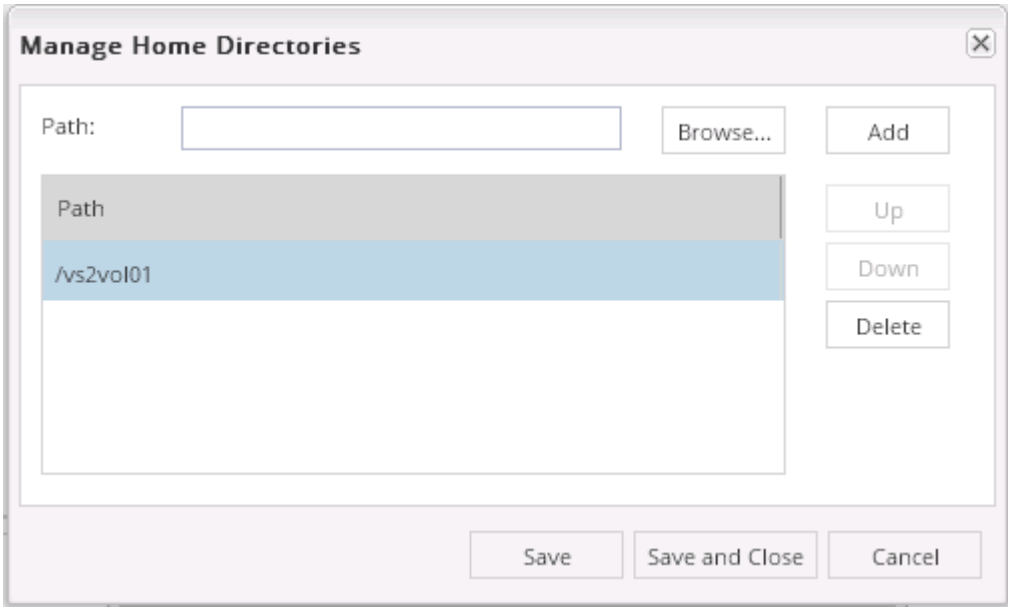

Step 6.

Click **Save & Close**.

Step 7.

Click **Create**.

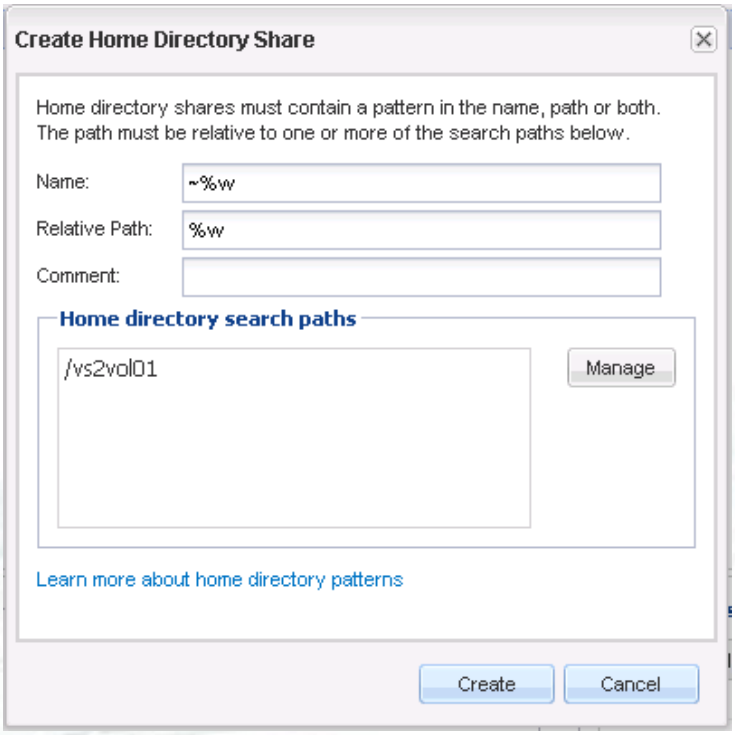

Step 8.

Notice that the home directory share is among your list of CIFS shares.

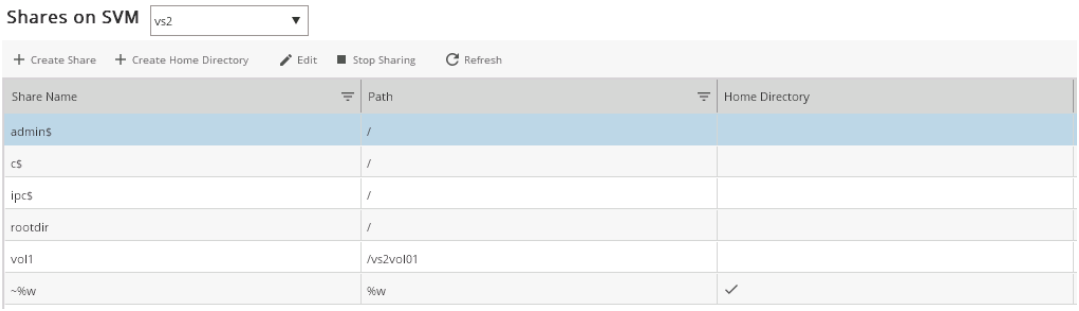

Step 9.

From your Windows desktop, open Windows Explorer.

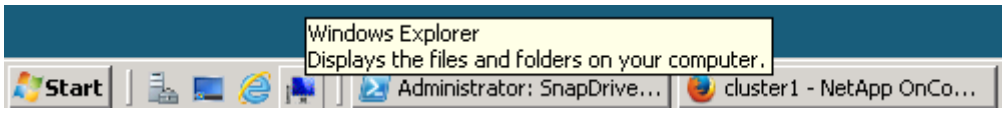

Step 10.

Navigate to the **vol1** share on your z: drive.

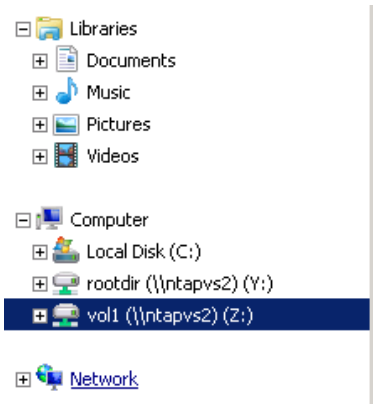

Step 11.

On the vol1 share, create a directory called **administrator** and, in that directory, create an empty text file called **adminfile**.

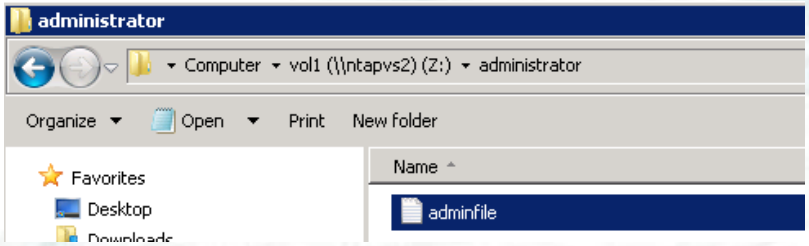

Step 12.

Map a drive to the home-directory share and verify the contents.

PS C:\> net view \\ntapvs2

*PS C:\> net use * \\ntapvs2\~administrator*

```
PS C:\Users\Administrator> net view \\ntapvs2
Shared resources at \\ntapvs2

(null)

Share name  Type  Used as  Comment

-------------------------------------------------------------------------------
rootdir     Disk
vol1        Disk
The command completed successfully.

PS C:\Users\Administrator> net use * \\ntapvs2\~administrator
Drive X: is now connected to \\ntapvs2\~administrator.

The command completed successfully.

PS C:\Users\Administrator> x:
PS X:\> dir


    Directory: X:\


Mode                LastWriteTime     Length Name
----                -------------     ------ ----
-a---        9. 12. 2014      7:36          0 adminfile.txt
```

TASK 8: Access your data from a NFS client

Step 1.

Verify the export policy that was created for your volume.

cluster1::> vserver export-policy rule show -vserver vs2

Step 2.

Using Putty log in to the Linux (CentOS) machine.

[root@linsrv ~]# mkdir /mnt/vs2

[root@linsrv ~]# mkdir /mnt/path01

Step 3.

Using the IP address of either data LIF within vs2, access vs2 exports through NFS.

[root@linsrv ~]# mount -t nfs 10.34.x1.40:/ /mnt/vs2

[root@linsrv ~]# mount -t nfs 10.34.x1.40:/vs2vol01 /mnt/path01

Step 4.

Explore both NFS mounts, which are mounted at different points in the vs2 namespace, and locate the home directory that you created earlier in the exercise.

[root@linsrv ~]# cd /mnt/vs2/vs2vol01/administrator/

[root@linsrv administrator]# ls

adminfile.txt

Step 5.

Copy some data into the vs2_vol01 volume, and compare the amount of space that was used before the copy with the amount of space that is used after the copy:

a. From the cluster shell:

   cluster1::> vol show -volume vs2_vol01 -fields used

b. Then from your Linux client:

   [root@linsrv /]# cd /mnt/path01

   [root@linsrv /]# cp /usr/include/ .*

   [root@linsrv /]# ls

c. Then from the cluster shell:

   cluster1::> vol show -volume vs2_vol01 -fields used

The space consumed before the copy represents the Snapshot copy reserve space.

TASK 9: Create fpolicy to block mp3 files on file system

We will use the native engine, so we don't need to create external engine.

Step 1.

We need to create events for all protocols

cluster1::> vserver fpolicy policy event create -vserver vs2 -event-name mp3_cifs -protocol cifs -file-operations create,open,rename -volume-operation false

cluster1::> vserver fpolicy policy event create -vserver vs2 -event-name mp3_nfsv3 -protocol nfsv3 -file-operations create,write,rename,symlink -volume-operation false

cluster1::> vserver fpolicy policy event create -vserver vs2 -event-name mp3_nfsv4 -protocol nfsv4 -file-operations create,open,rename,symlink -volume-operation false

cluster1::> vserver fpolicy policy event show -vserver vs2

Step 2.

In second step we will create fpolicy and assign event to fpolicy container

cluster1::> vserver fpolicy policy create -vserver vs2 -policy-name mp3blocker -events mp3_cifs, mp3_nfsv3, mp3_nfsv4 -engine native -is-mandatory true -allow-privileged-access no

cluster1::> vserver fpolicy policy show -vserver vs2

Step 3.

Scope defines on which shares and volumes is fpolicy active using include and exclude sections.

*cluster1::> fpolicy policy scope create -vserver vs2 -policy-name mp3blocker -shares-to-include * -file-extensions-to-include mp3 -volumes-to-include **

cluster1::> vserver fpolicy policy scope show -vserver vs2

Step 4.

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B
IČO 357 034 66
IČ DPH SK2020341158

When fpolicy definition is ready we need to enable it.

cluster1::> vserver fpolicy enable -vserver vs2 -policy-name mp3blocker -sequence-number 1

cluster1::> vserver fpolicy show -vserver vs2

Step 5.

Try to create a file with extension mp3 in the NFS mount on the Linux server or on the Windows share.

[root@linsrv path01] cd /mnt/vs2

[root@linsrv vs2]# dd if=/dev/urandom of=file.mp3 bs=1k count=30

dd: opening `file.mp3': Permission denied

END OF EXERCISE

LAB 10: SAN protocols

In this exercise, you experiment with scalable SAN by configuring a SVM as an iSCSI target and connecting a LUN to a Windows host.

Objectives of this exercise:

- Use NetApp GUI to create a SVM for iSCSI
- Enable an aggregate as a resource for a SVM
- Configure Windows for multipath I/O (MPIO)
- Install the Windows Host Utilities Kit (HUK)
- Configure the iSCSI software initiator in Windows
- Use NetApp GUI to create an iSCSI-attached LUN
- Access the iSCSI-attached LUN on the initiator

TASK 1: Use NetApp GUI to create a SVM for iSCSI

Step 1.

In the GUI navigation frame, click **Storage > SVMs**.

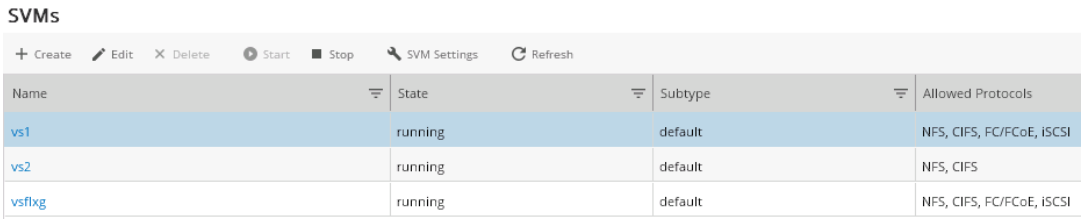

Step 2.

Click **Create** to create a SVM for iSCSI.

Step 3.

On the SVM Details page, enter and select the following, and then click **Submit & Continue.**

- SVM Name: **vsISCSI1**
- IPspace: **Default**
- Volume Type: **FlexVol volumes**
- Data Protocols: **iSCSI**
- Language: **C.UTF-8 [c.utf_8]**
- Security Style: **UNIX**
- Aggregate: **n01_aggr1**

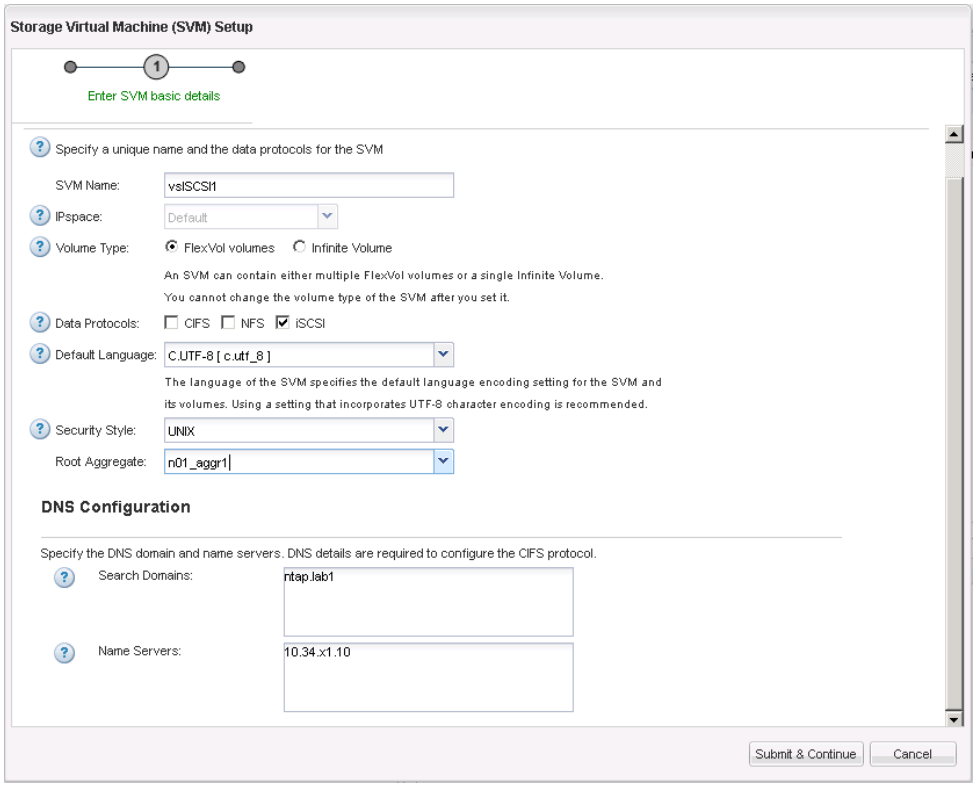

Step 4.

On the iSCSI configuration page, enter and select the following, and then click **Submit & Continue.**

- Target Alias: **vsISCSI1_target**
- LIFs Per Node: **2**
- Assign IP address: **Using a subnet**
    - Subnet: **subnet1**
    - Radio button: **Auto-select the IP address from this subnet**
- Review or Modify LIFs configuration checkbox: **Select**
- Number of portsets: **1**
- Click **Submit and Continue**

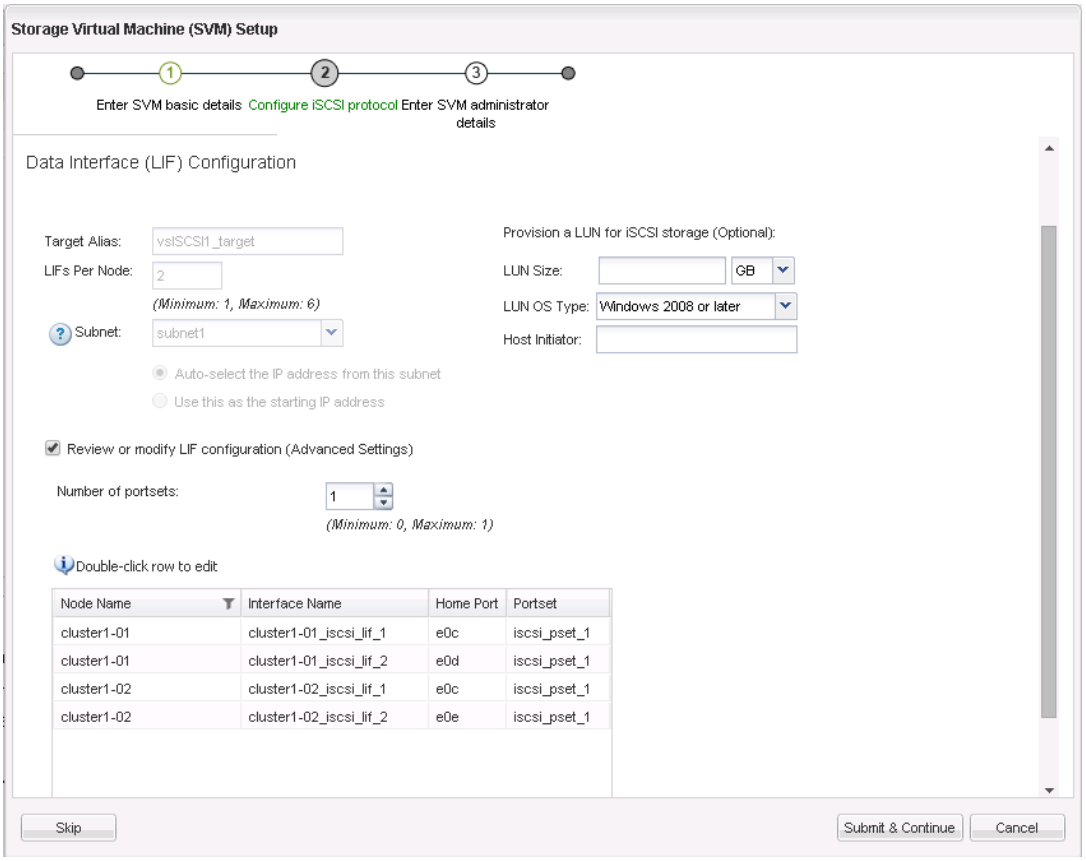

Step 5.

On the SVM administration page, click **Skip**.

Step 6.

Review the summary and, click **OK**.

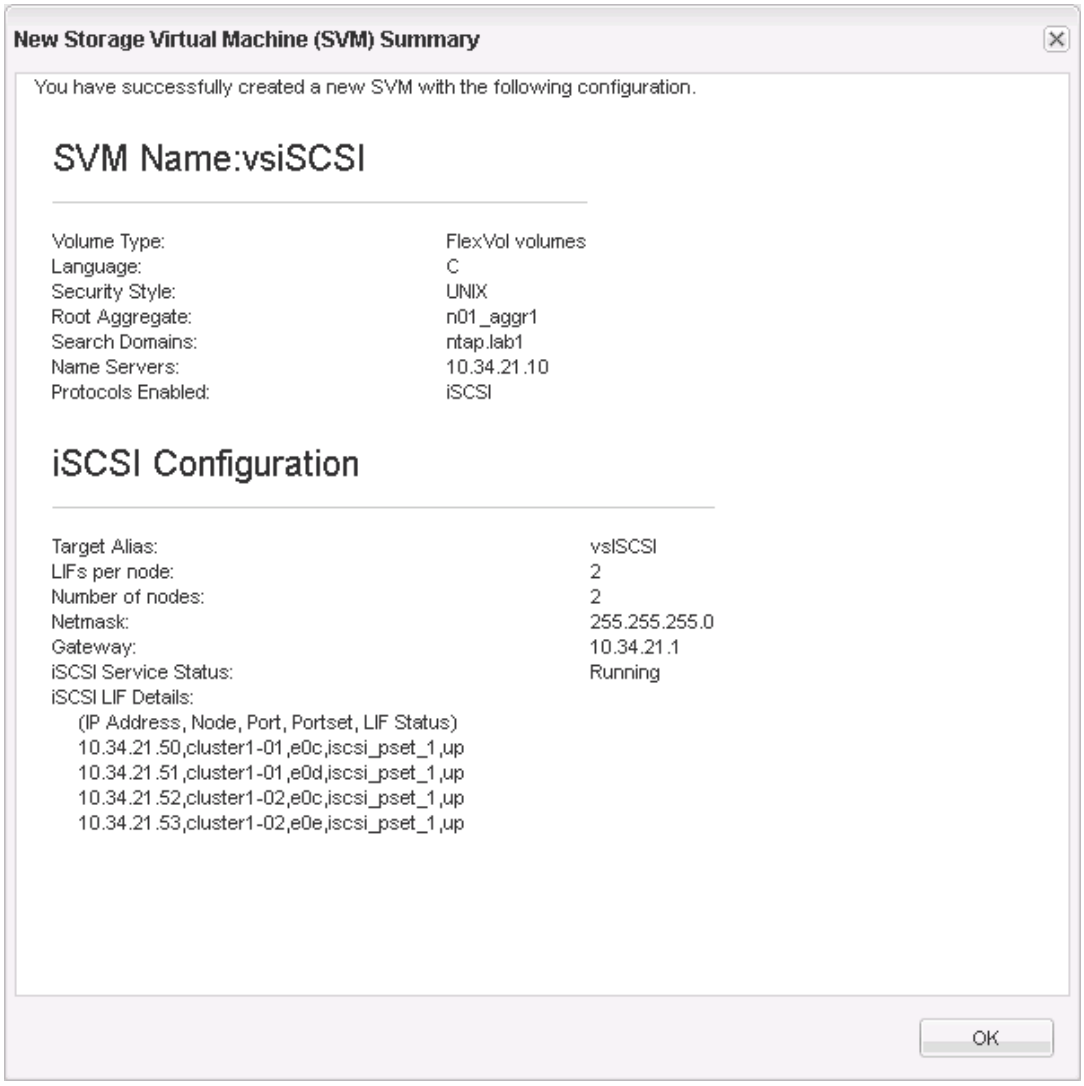
New Storage Virtual Machine (SVM) Summary
You have successfully created a new SVM with the following configuration.
SVM Name:vsiSCSI
Volume Type: FlexVol volumes
Language: C
Security Style: UNIX
Root Aggregate: n01_aggr1
Search Domains: ntap.lab1
Name Servers: 10.34.21.10
Protocols Enabled: iSCSI
iSCSI Configuration
Target Alias: vsISCSI
LIFs per node: 2
Number of nodes: 2
Netmask: 255.255.255.0
Gateway: 10.34.21.1
iSCSI Service Status: Running
iSCSI LIF Details:
(IP Address, Node, Port, Portset, LIF Status)
10.34.21.50,cluster1-01,e0c,iscsi_pset_1,up
10.34.21.51,cluster1-01,e0d,iscsi_pset_1,up
10.34.21.52,cluster1-02,e0c,iscsi_pset_1,up
10.34.21.53,cluster1-02,e0e,iscsi_pset_1,up
OK

Step 7.

Select **Storage** > **SVMs** > **vsISCSI1** > **SVM Settings** > **iSCSI** to review your iSCSI configuration.

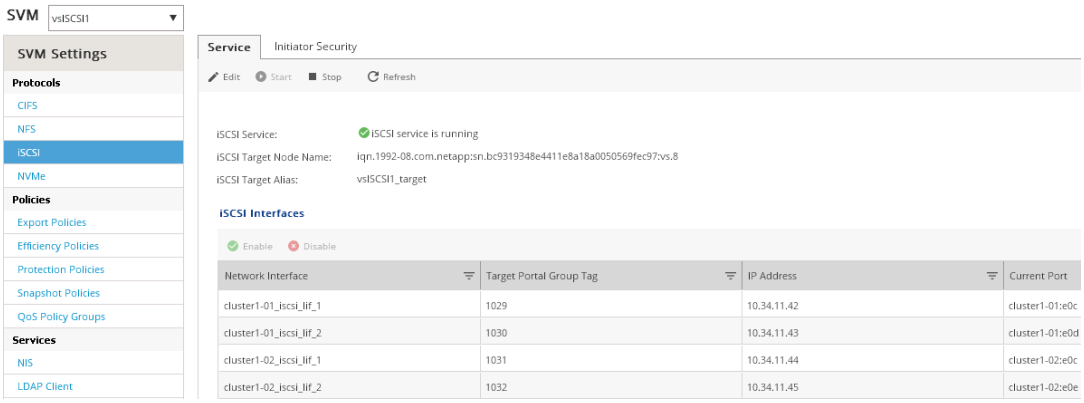

TASK 2: Enable an aggregate as a resource for a SVM

Step 1.

Select **SVMs**, and then select **vsISCSI1**.

Step 2.

In the menu bar at the top of the page, click **Edit**.

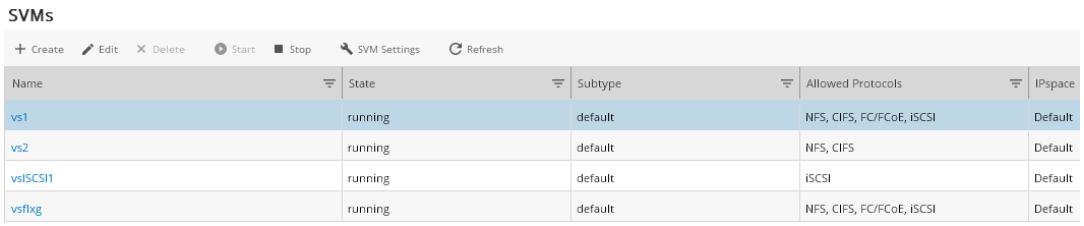

Step 3.

When the Edit Storage Virtual Machine dialog box appears, click the **Resource Allocation** tab.

Step 4.

Select **Delegate volume creation, n01_aggr1**, and **n02_aggr1**.

Selecting these values enable this Storage Virtual machine to provision volumes in these aggregates.

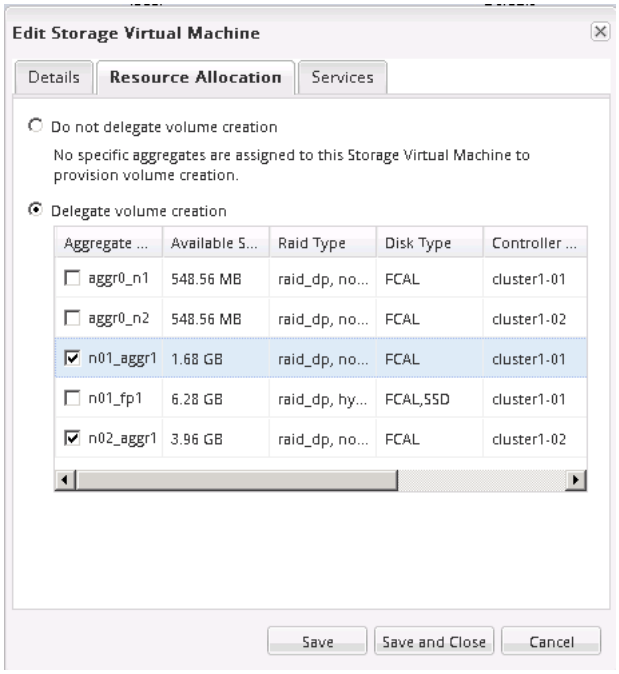

Step 5.

Click **Save and Close** to complete the process.

TASK 3: Configure Windows for MPIO (MPIO is already installed)

NOTE: This exercise uses the Microsoft device-specific module (DSM) instead of the NetApp DSM.

Step 1.

On the Windows desktop, click the Server Manager icon.

Server Manager opens.

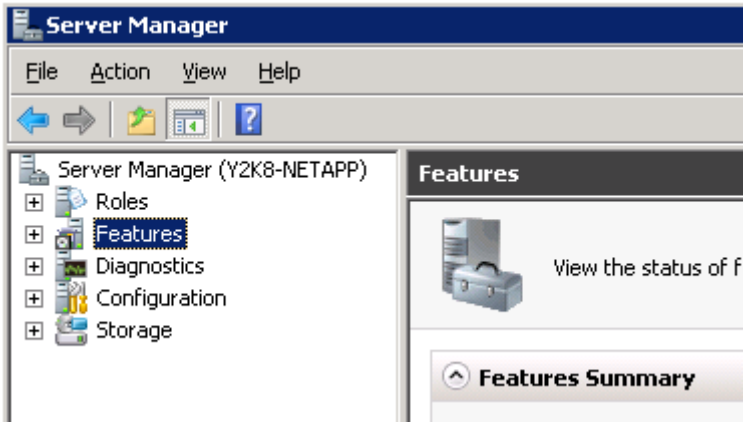

Step 2.

On the toolbar at the top right of the page, right click on **Features** and then **Add Features**.

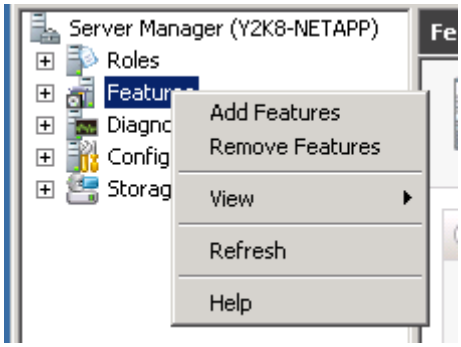

Step 3.

On the Select features page, select **Multipath I/O** and click **Next**.

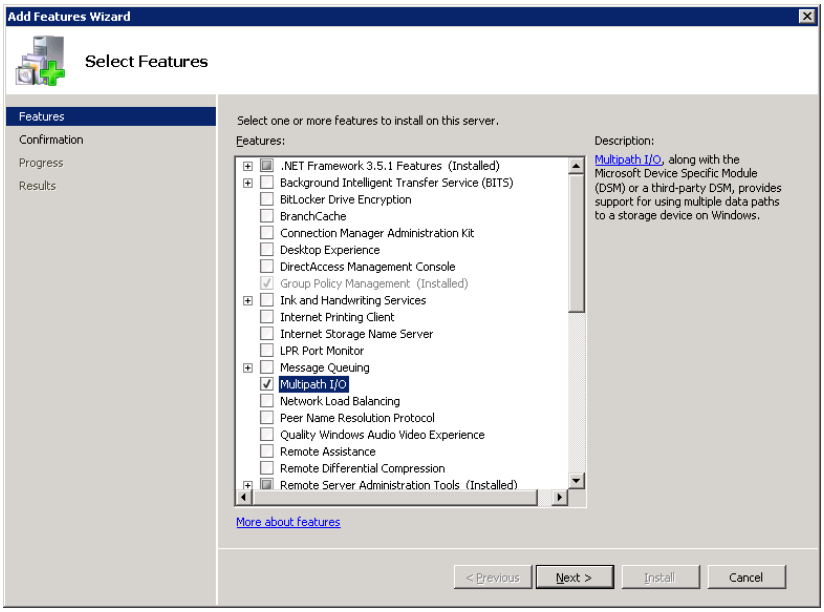

Step 4.

On the Confirm installation selections page, select the "Restart the destination server" checkbox, reply **Yes** to the warning, and then click **Install**.

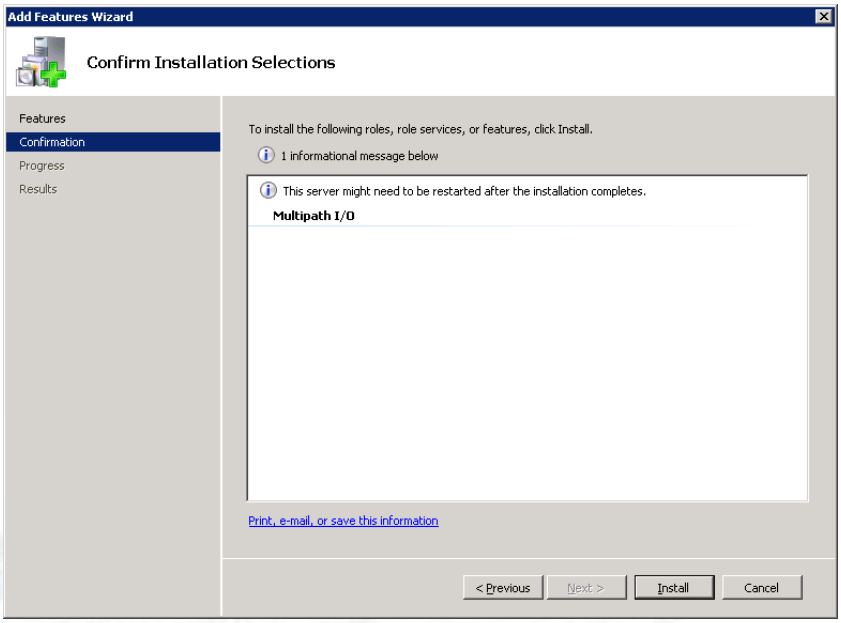

Step 5.

After the feature is installed and the Results window appears, confirm that the installation was successful and click **Close** and close Server Manager.

TASK 4: Install the Windows host utilities (Optional)

Step 1.

On the desktop of your Windows system, open the **Downloads** folder.

Step 2.

Double-click the NetApp Windows Host Utilities installation file.

Step 3.

If are prompted with a security warning, confirm that you want to continue.

Step 4.

On the title page of the installation wizard, click **Next**.

Step 5.

Confirm the license agreement and click **Next**.

Step 6.

On the Support for Multipathing page, select **Yes, install support for Multipath I/O** and click **Next**.

Step 7.

Confirm the destination folder and click **Next**.

Step 8.

Click **Install**.

Step 9.

If an FC Configuration Tool warning appears, click **OK**.

Step 10.

After the installation is complete, click **Yes** to restart the Windows system.

Step 11.

Wait about five minutes before you attempt to reconnect to your Windows Server system.

TASK 5: Configure the iSCSI software initiator (in Windows)

In this task, you use the MPIO instead of Multiple Connections per Session (MCS) technique for multipathing.

Step 1.

On your Windows desktop, open the Control Panel.

Step 2.

Select **View by small icons** and double-click **iSCSI Initiator**.

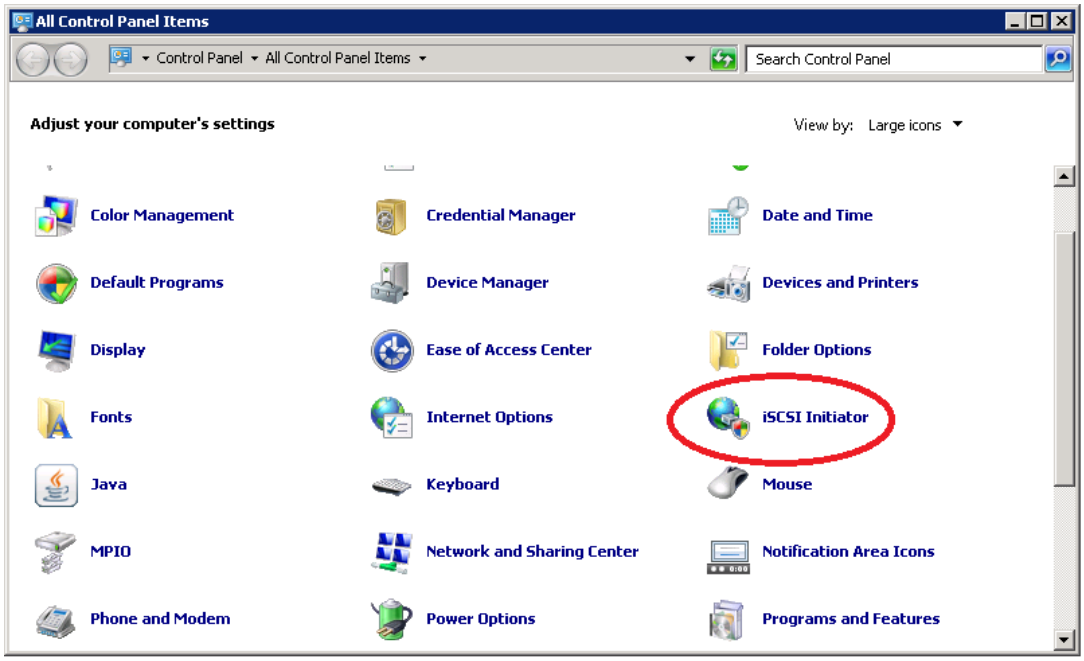

Step 3.

If an error message appears to indicate that the Microsoft iSCSI service is not running, click **Yes** to start the service.

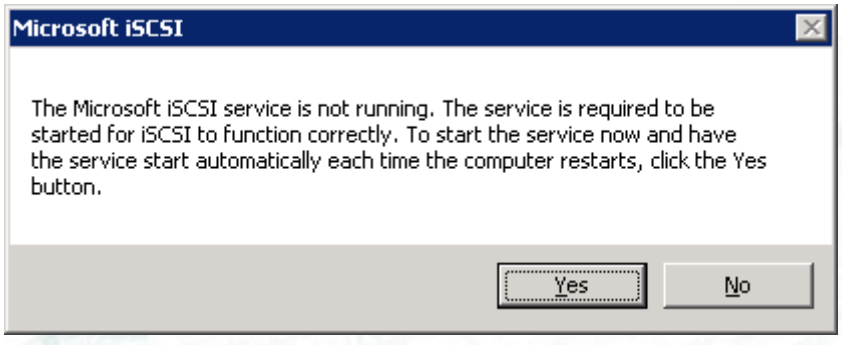

Step 4.

If a message asks if you want to unblock the Microsoft iSCSI service through the Windows Firewall, click **Yes**.

Step 5.

When the iSCSI Initiator Properties dialog box appears, click the **General** tab.

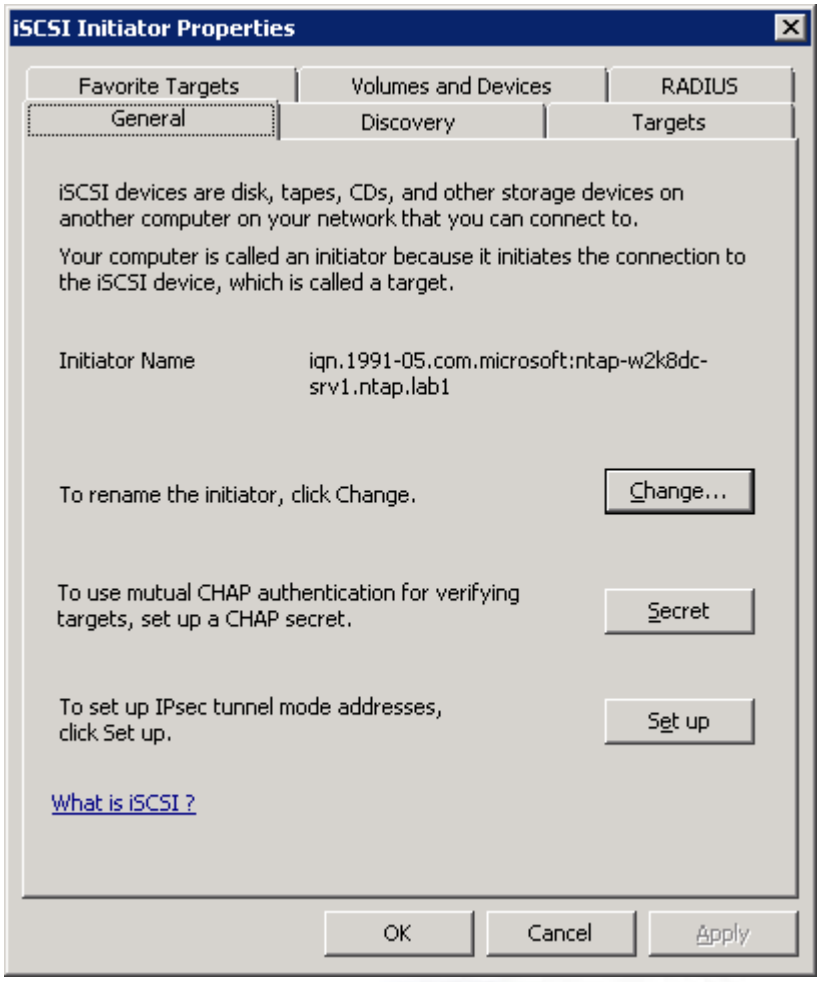

Step 6.

Record the Initiator Name (IQN): _____

Step 7.

Click the **Discovery** tab.

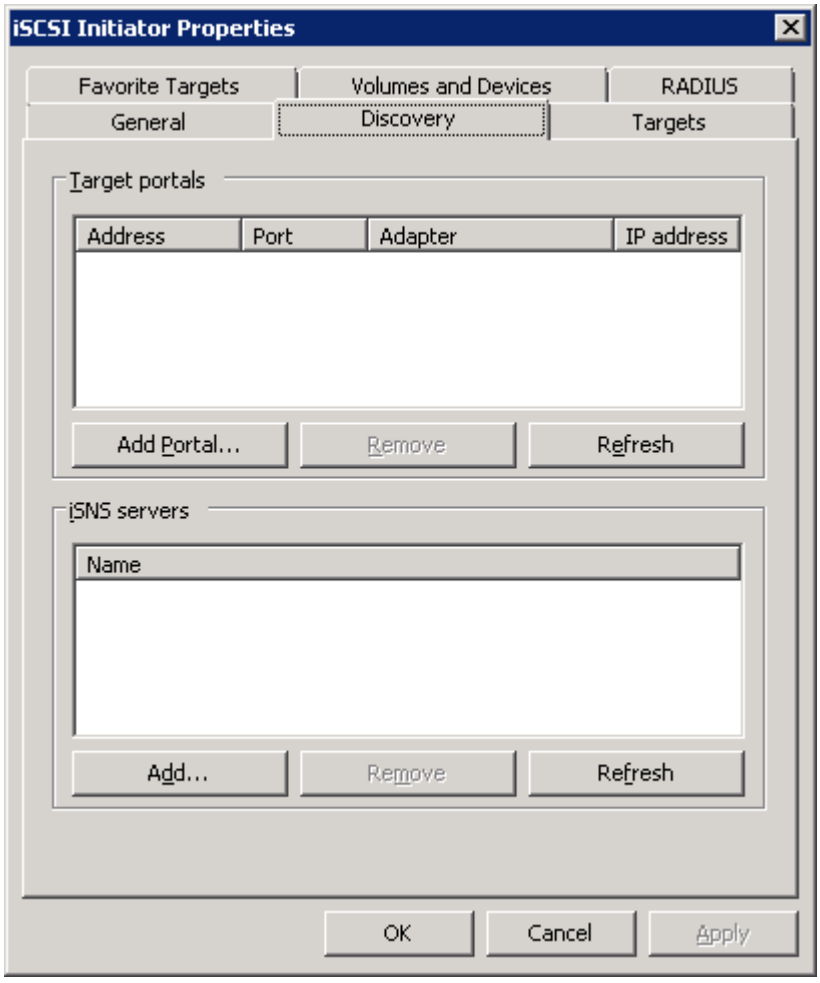

Step 8.

Click **Add Target Portal**, enter the IP address of one of ports within the vsISCSI1 port set, and click **OK**.

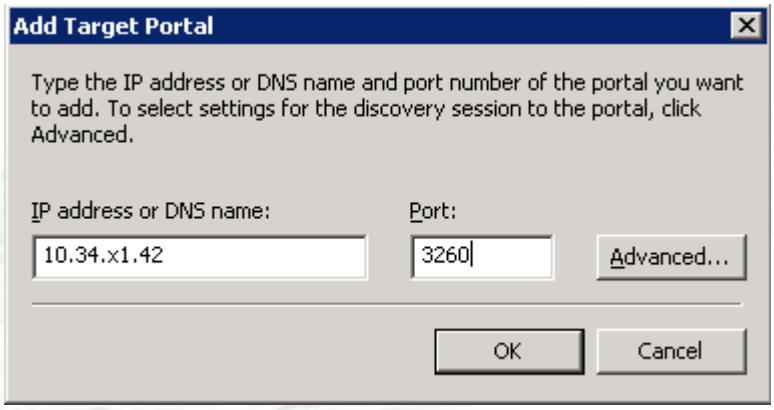

Step 9.

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B

IČO 357 034 66
IČ DPH SK2020341158

Click the **Targets** tab.

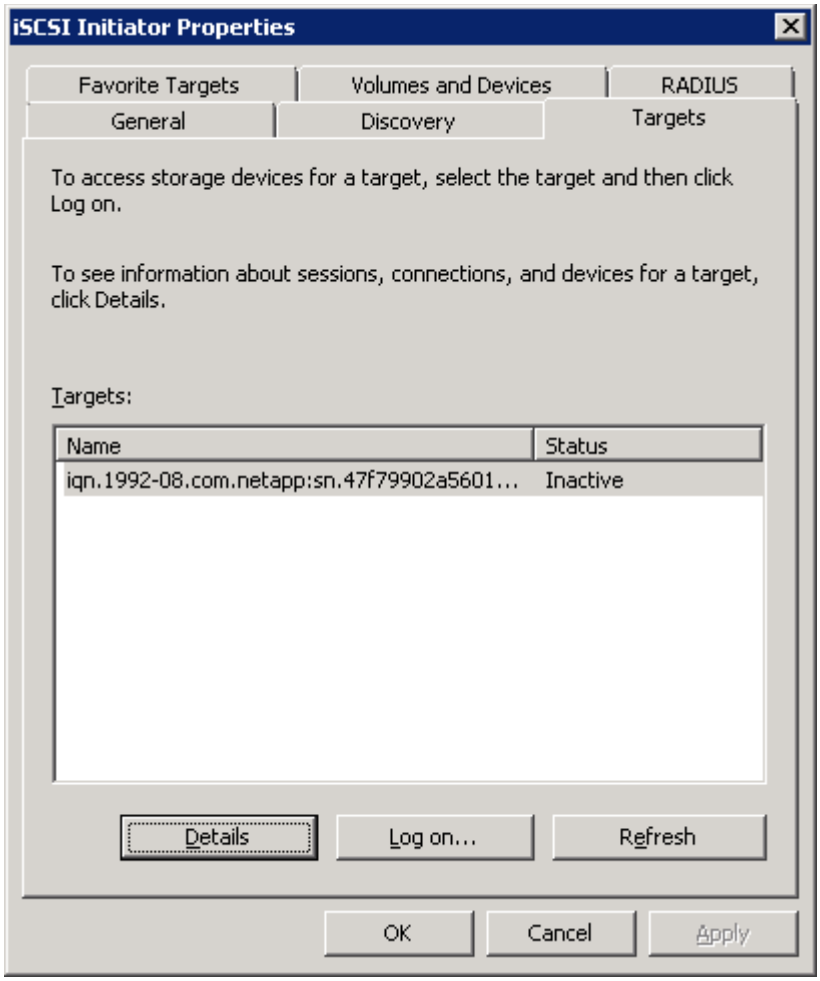

Step 10.

Verify that the discovered target appears in the list and click **Log On to Target**.

Step 11.

In the Connect To Target dialog box, select **Enable multi-path** and click **Advanced**.

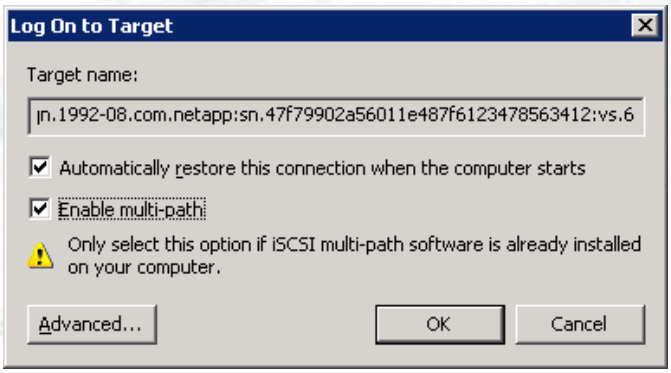

Step 12.

In the Advanced Settings dialog box, from the **Target portal IP** list, select the lowest target portal IP address, and click **OK**.

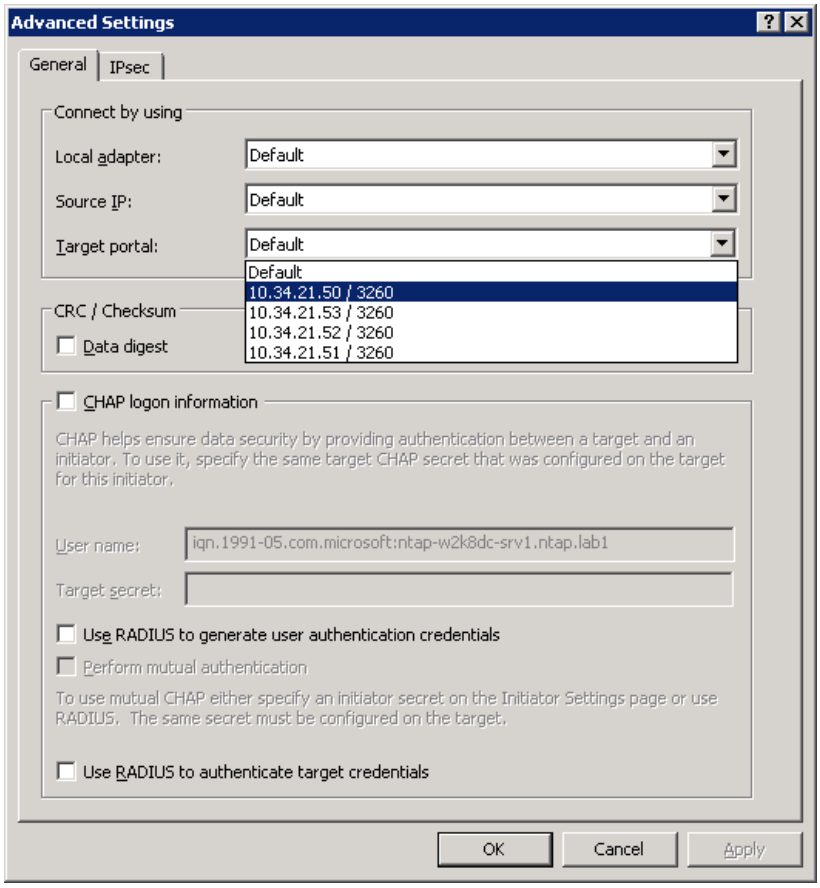

Step 13.

Click **OK** to close the Connect to Target dialog box and start a new iSCSI session between the initiator and target.

Step 14.

In the **Target->Details-> Properties** dialog box, clicking on the Properties tab you can create additional sessions with all of the iSCSI LIFs within the port set.

Step 15.

Click **OK** to close the Properties window.

Step 16.

Click **OK** to close the iSCSI Initiator Properties window.

Step 17.

Close Control panel.

TASK 6: Use NetApp GUI to create an iSCSI attached LUN

Step 1.

In GUI, select **Storage > (vsISCSI) > LUNs**.

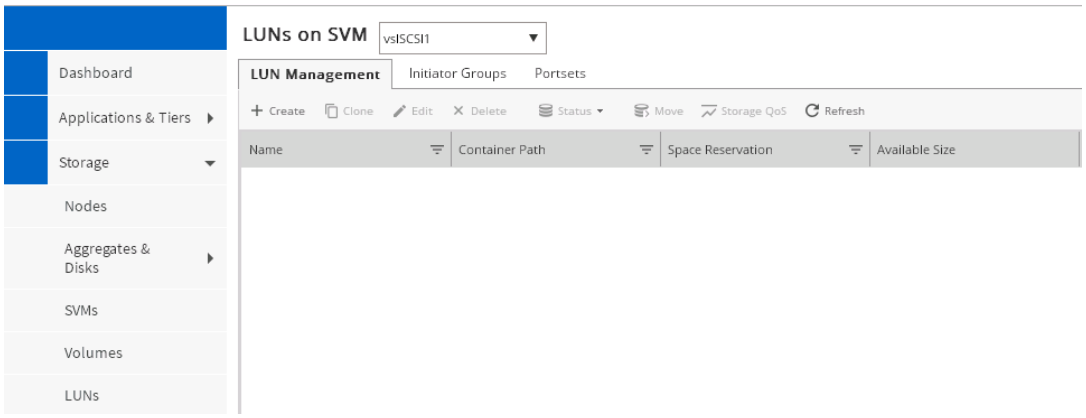

Step 2.

Click the **Create** button. The Create LUN Wizard appears.

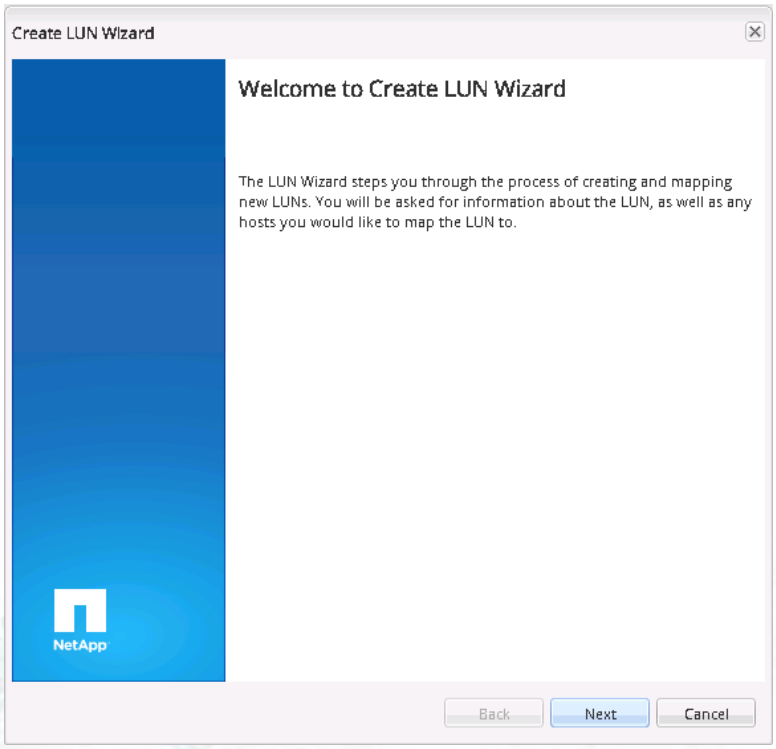

Step 3.

Click **Next**.

Step 4.

On the next page, enter and select the following:

- Name: **lun_vsISCSI1_1**
- LUN Type: **Windows 2008 or later**
- LUN size: **100 MB**
- Thin Provisioned checkbox: **clear**

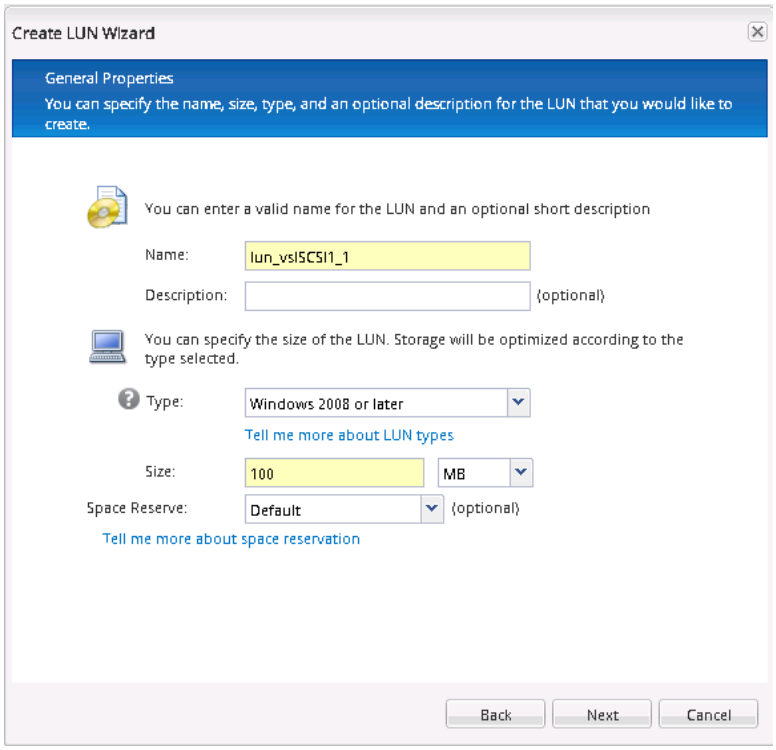

Step 5.

Click **Next**.

Step 6.

On the LUN Container page, first select **Create a new flexible volume in,** then click **Choose**, select the **n01_aggr1** aggregate, and click **OK**.

Step 7.

In the Volume Name text box, accept the default name and click **Next**.

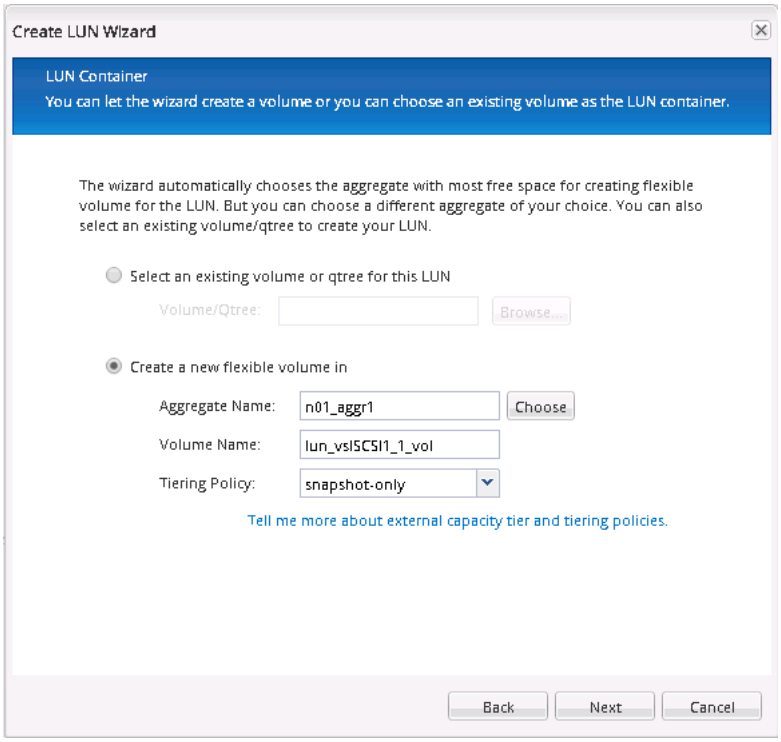

Step 8.

On the Initiators Mapping page, click **Add Initiator Group** to begin creating an initiator group (igroup).

Step 9.

In the Create Initiator Group dialog box, on the General tab, enter or select the following:

- Name: **ig_myWin**
- Operating System: **Windows**
- Type: **iSCSI**

Step 10.

In the Portset area, click **Choose**.

Step 11.

Select the port set that you created with the Vserver in task 1 and click **OK**.

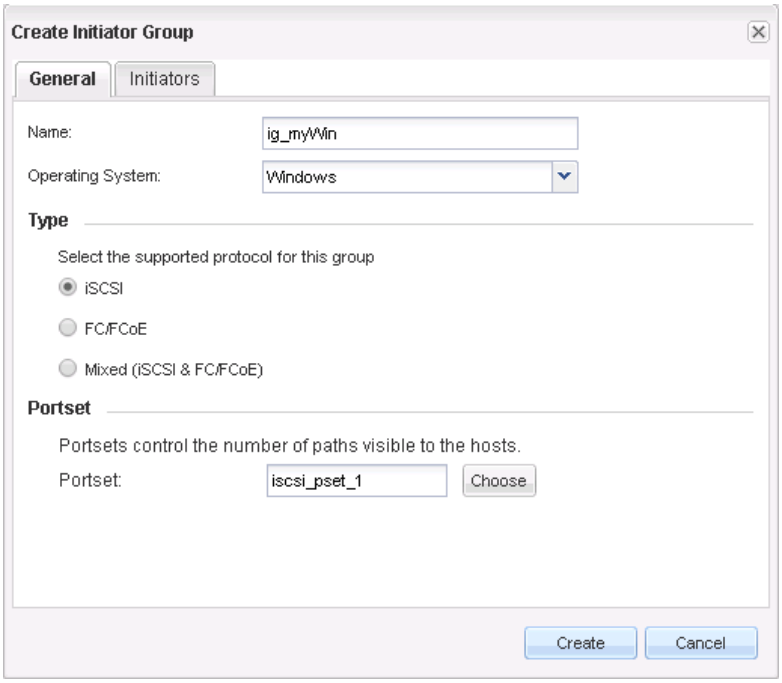

Step 12

Click the **Initiators** tab, and click **Add**.

Step 13.

Enter the initiator **IQN** that you recorded earlier (Task 5, Step 5) (or select it from pull down menu) and click **OK**.(Or you can use command *vserver iscsi initiator show* from CLI.)

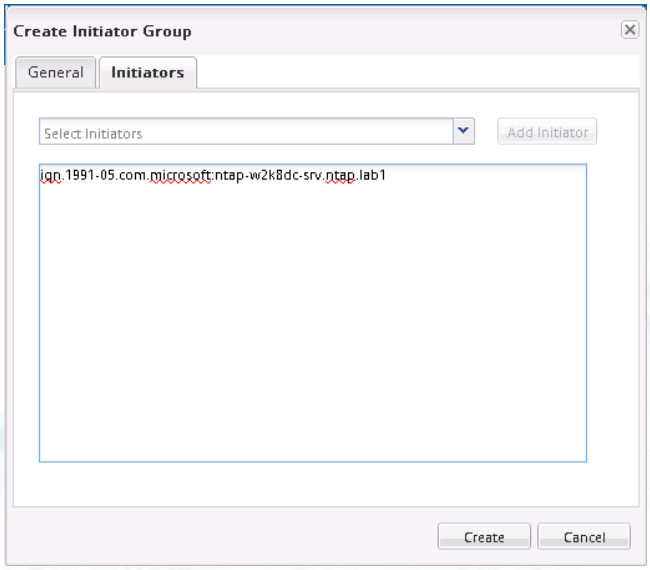

Step 14.

Click **Create**.

Step 15.

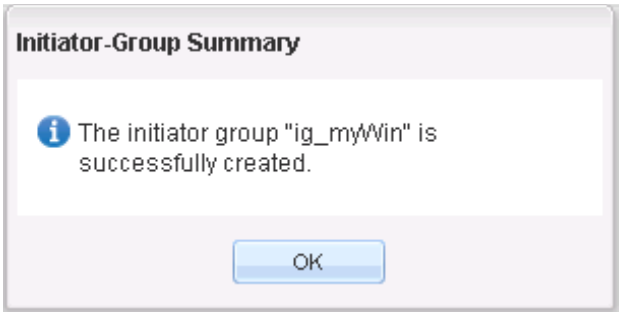

Step 16.

Back on the Initiators Mapping page, verify that the new igroup has been added to the list; then select the **Map** checkbox to the left of the igroup and click **Next**.

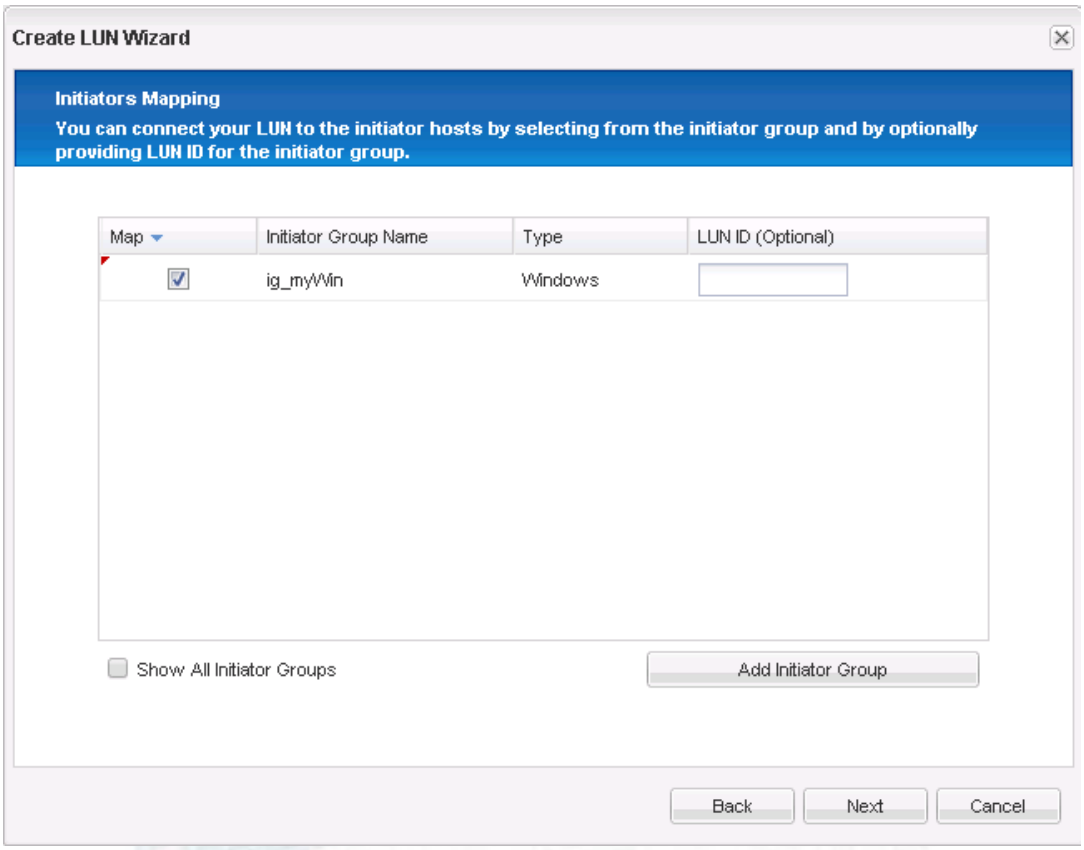

Step 17.

On the page Storage Quality of Service Properties click **Next**. Then review the LUN Summary page and click **Next**.

Step 18.

Review the Create LUN Wizard completion page and click **Finish**. Creation of the iSCSI-attached LUN is now complete.

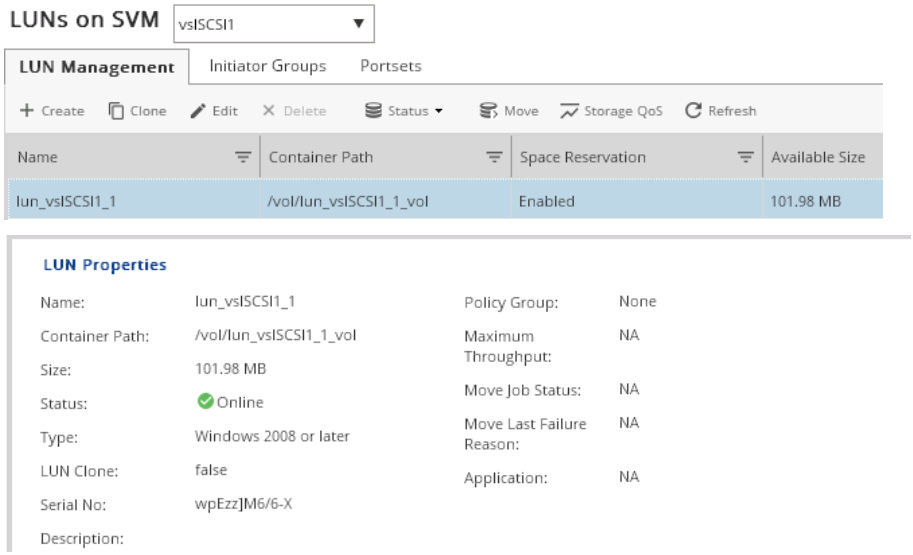

TASK 7: Access the iSCSI attached LUN on the Windows host

Step 1.

On the desktop right click on Computer icon, and click on Manage.

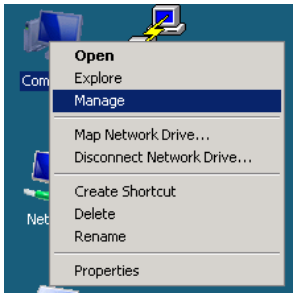

Step 2

In **Server Manager** window click on **Storage > Disk Management.**

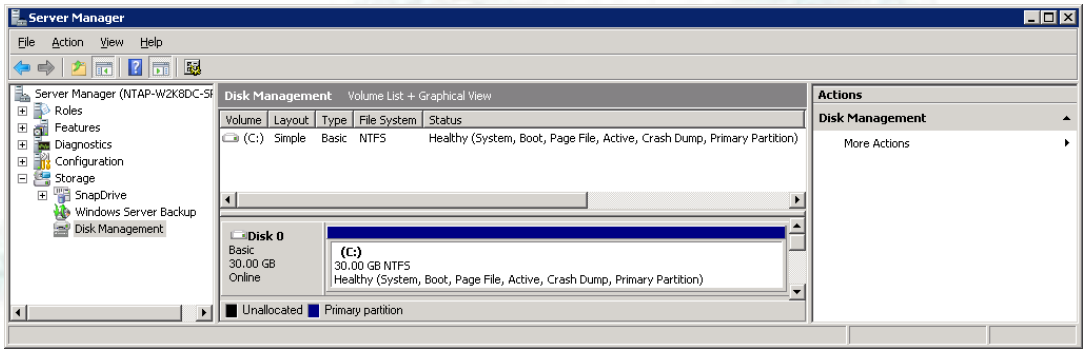

Step 3.

If you do not see the LUN disk in the bottom section of the center pane, right-click the Disk Management node in the left pane and select **Rescan Disks**. New disk will appear.

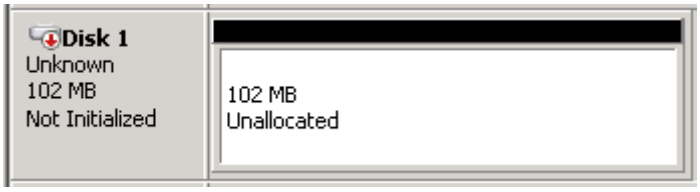

Step 4.

Right-click the disk header and, if the disk is offline, select **Online**.

Step 5.

Right-click the disk header again and select **Initialize Disk**.

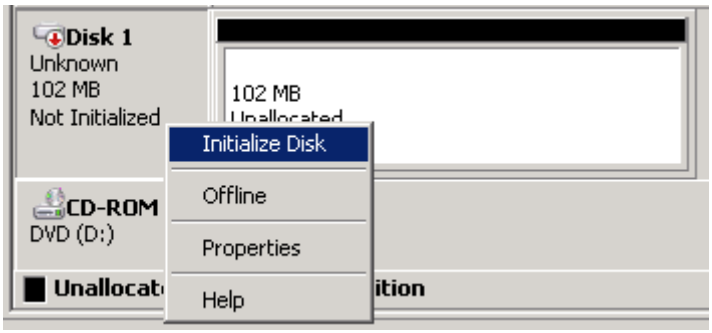

Step 6.

Review the Initialize Disk dialog box and click **OK**.

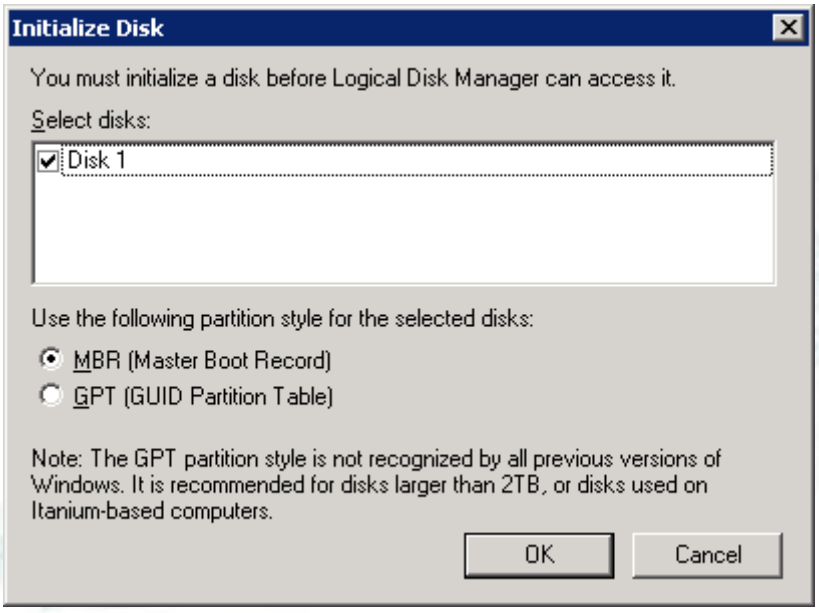

Step 7.

In the Disk Management pane, right-click the **Unallocated** partition and select **New Simple Volume.**

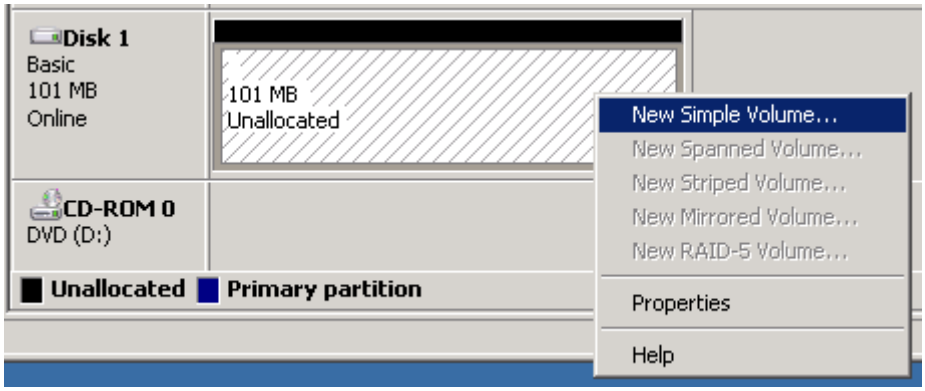

Step 8.

On the introduction page of the New Simple Volume wizard, click **Next**.

Step 9.

On the Specify Volume Size page, click **Next**.

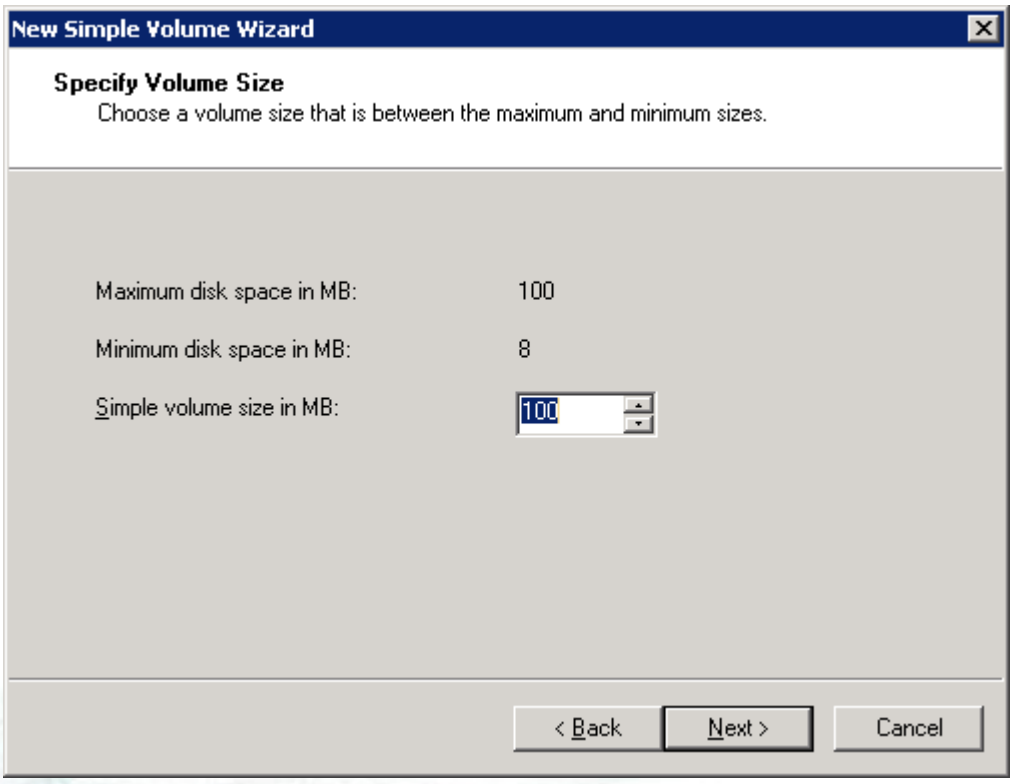

Step 10.

On the Assign Drive Letter or Path page, click **Next**.

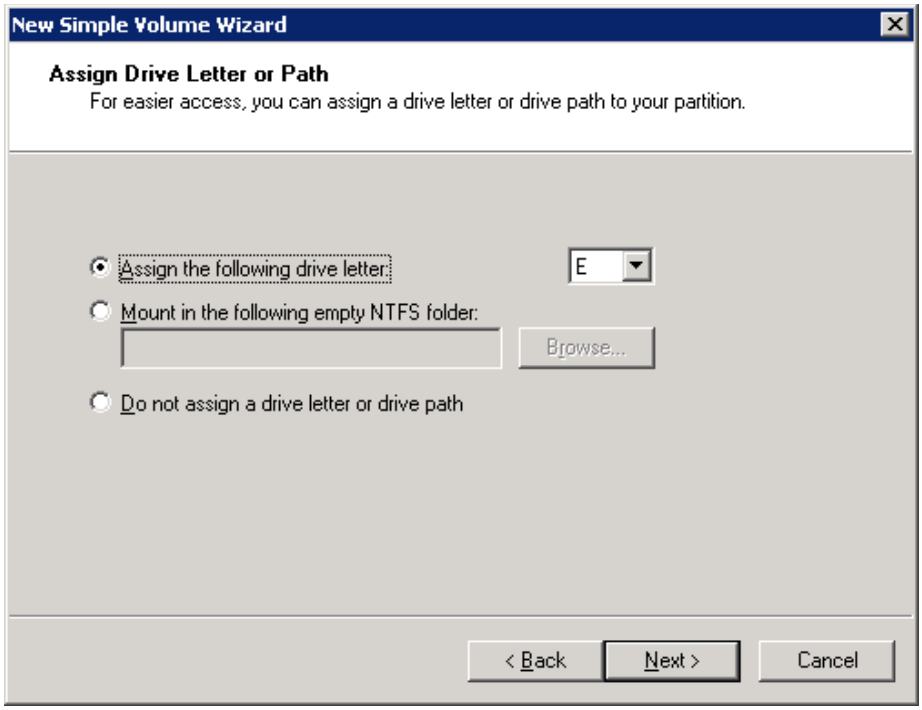

Step 11.

On the Format Partition page, perform one of these two actions:

- Click **Next** to accept the default values.
- In the “Volume label” text box, enter a label (such as the one shown here), and then click **Next**.

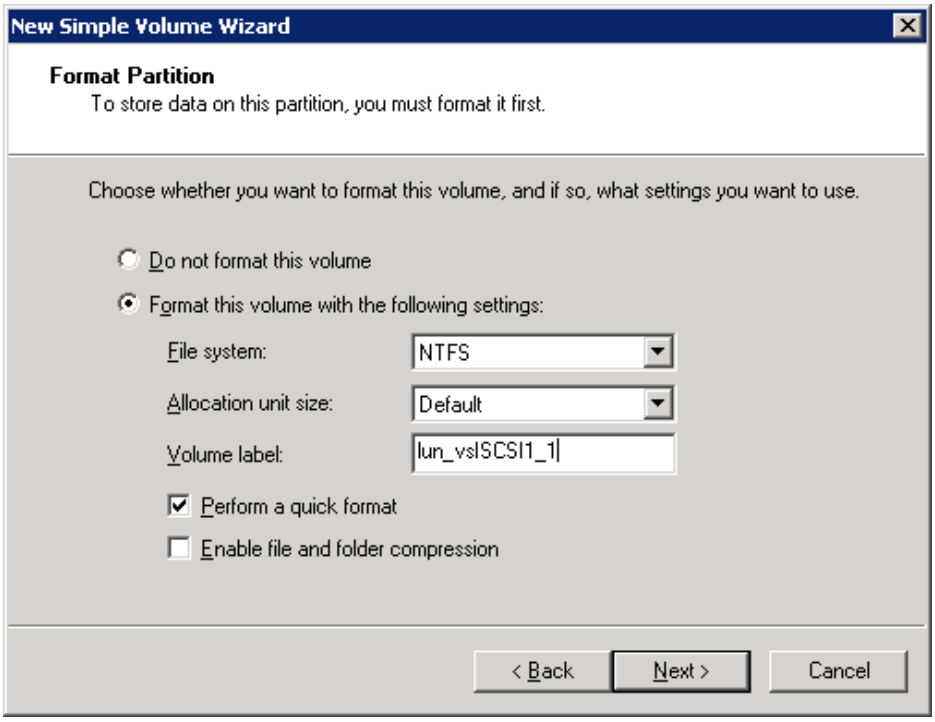

Step 12.

Review the Completing page and click **Finish**.

Step 13.

Verify that the new LUN is now provisioned and, when you are finished, close the Computer Management window.

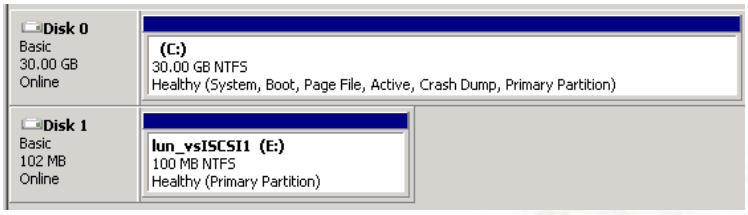

Step 14.

Navigate to the mount location of the LUN and verify that you can create a file in the LUN.

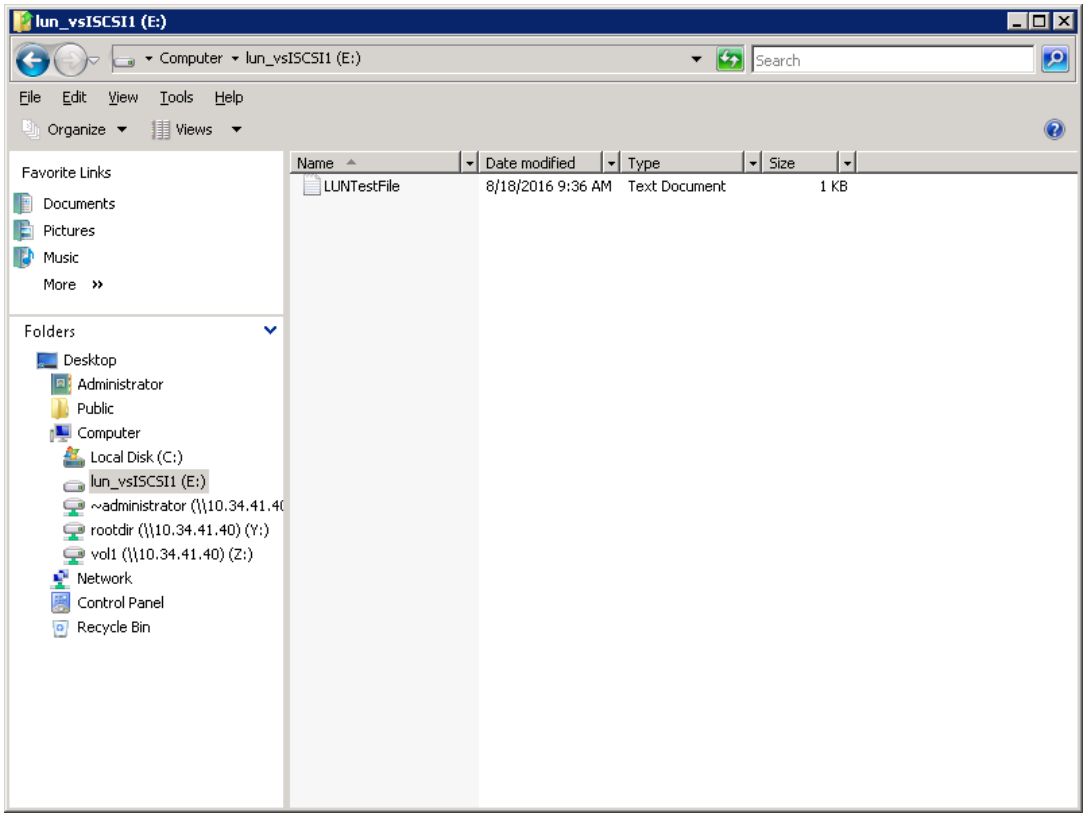

END OF EXERCISE

LAB 11: Storage efficiency

In this exercise, you create a volume so that you can learn about resizing, working with qtrees and quotas, creating FlexClone volumes, and configuring storage efficiency.

Objectives of this exercise

By the end of this exercise, you should be able to:

- Resize a volume
- Create a qtree and set user quotas
- Work with FlexClone volumes
- Enable deduplication and data compression

Step 1.

In GUI, select **Storage > (Volumes on SVM: vs2) > Volumes**, and then click **Create FlexVol**.

Step 2.

In the Create Volume window, enter these volume attributes, and then click the **Create** button:

- Name: **dataefficiency**
- Aggregate: **n02_aggr1**
- **20 MB** size, with **0%** Snapshot copy reserve, and Space reserve: **Thin-provisioned**

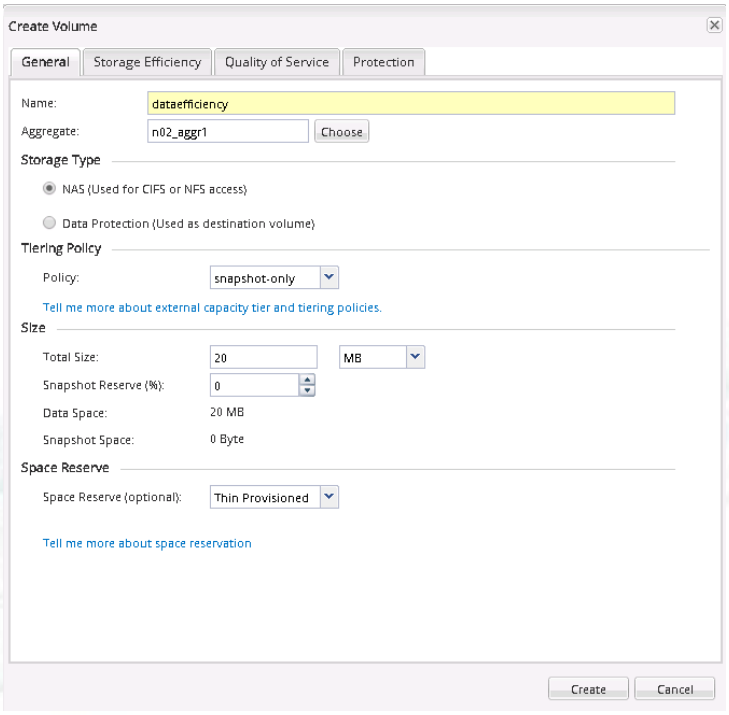

Step 3.

Check the **Namespace** page to see where the volume was mounted. Remember that GUI automatically mounts new volumes at /<volname>.

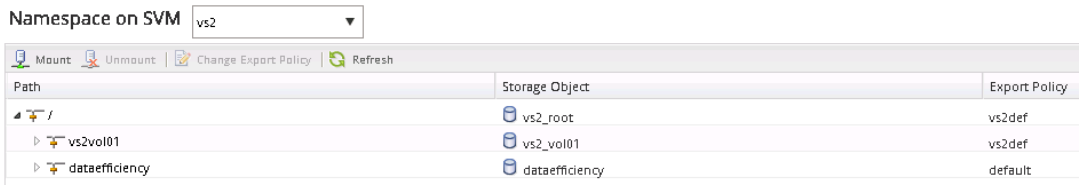

Step 4.

Select the dataefficiency volume from the list, click **Change Export Policy**, and change the export policy to **vs2def**.

Step 5.

Navigate to the **Volumes** page for vserver vs2, select the volume **dataefficiency**, and click **Edit**.

Step 6.

In the Edit Volume window, on the General tab, click Configure UNIX permissions and give the Group and Others read, write, and execute permissions, and then click **Save and Close**.

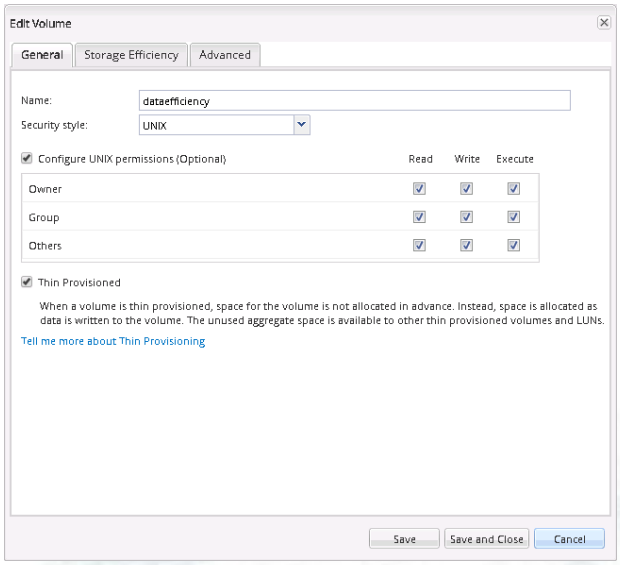

Step 7.

From your NFS client, create a file in the dataefficiency volume.

cd /mnt/vs2/dataefficiency/

echo "This is the content of datafile in volume dataefficiency" >datafile

ls

cat datafile

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B

IČO 357 034 66
IČ DPH SK2020341158

Step 8.

In GUI, on the Shares page, click **Create Share** to create a **CIFS** share for the new volume.

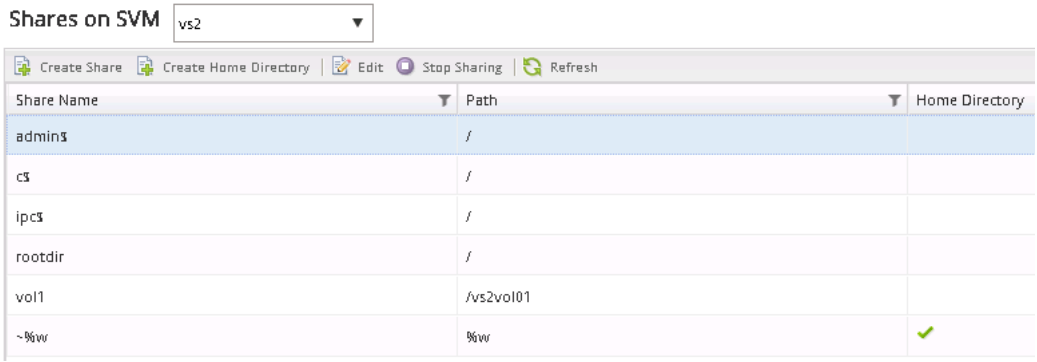

Step 9.

Click **Browse** to select a folder to share, select the dataefficiency junction, and click **OK**.

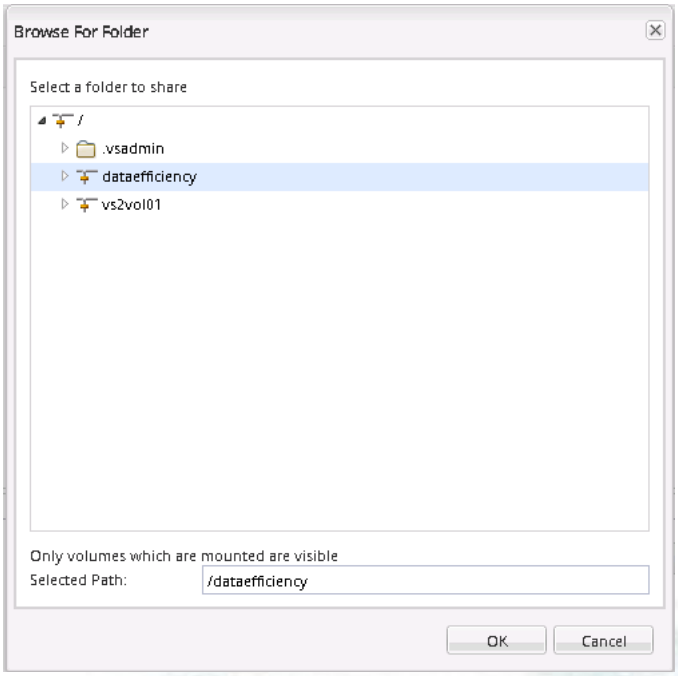

Step 10.

Name the share **de** and click **Create**.

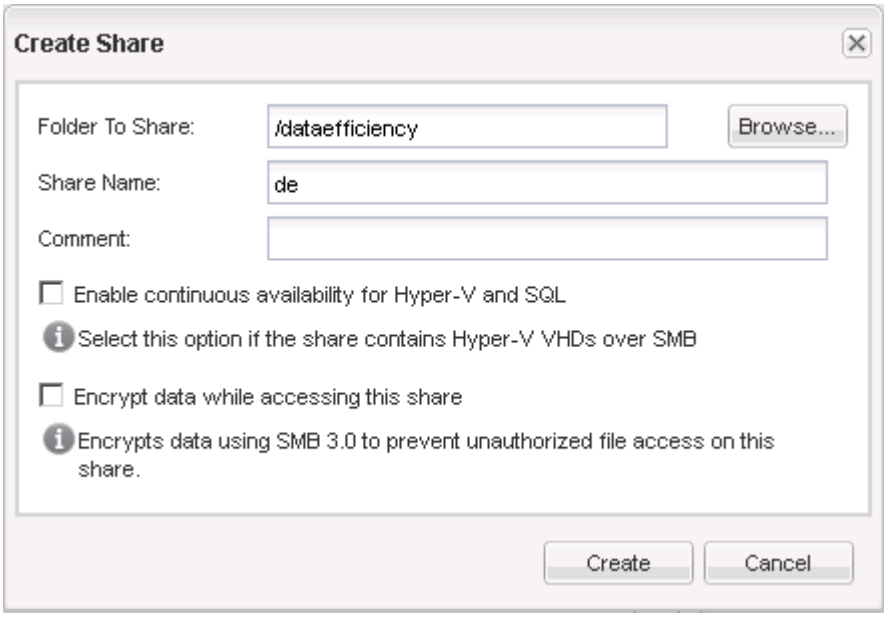

Step 11.

In Windows, map the new share as the T: drive.

```
PS C:\> net use t: \\ntapvs2\de
The command completed successfully.
PS C:\> t:
PS T:\> dir
    Directory: T:\
Mode                 LastWriteTime     Length Name
----                 -------------     ------ ----
-a---       10. 12. 2014     11:25         63 datafile
```

TASK 2: Resize a volume

Step 1.

Using Putty log in to the Linux (CentOS) machine and issue command

cd /mnt/vs2/dataefficiency/

Step 2.

Then we will try to create 30MB file.

dd if=/dev/urandom of=30mfile bs=1k count=30000

Step 3.

Action ended with message “No space left on device”. Remove partial file.

ls -la

rm 30mfile

rm: remove regular file `30mfile'? y

ls -la

Step 4.

In GUI, on the Volumes page, select the **dataefficiency** volume and click the **Actions** > **Resize** option to start the Volume Resize wizard.

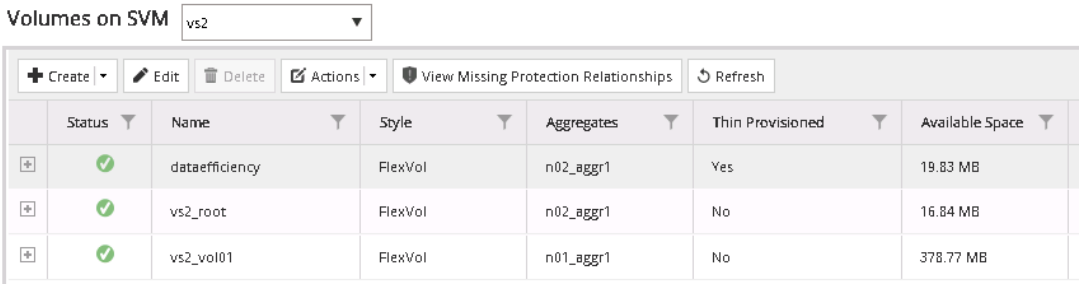

Step 5.

When the wizard starts, click **Next**.

Step 6.

In the Resize Volume Settings window, change the Total capacity value to **100 MB** without changing the value in the Snapshot reserve field, and click **Next**.

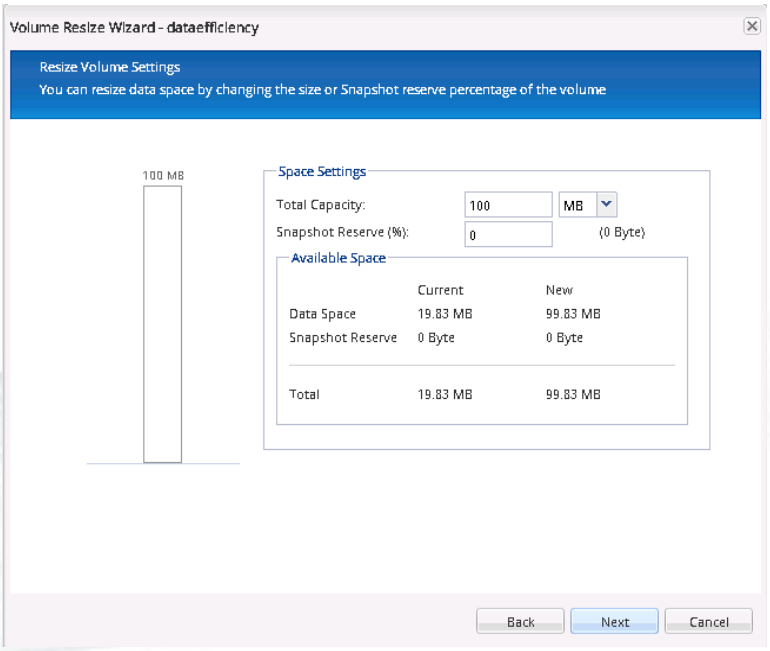

Step 7.

On the Delete Snapshot Copies page, click **Next**.

Step 8.

On the Summary page, confirm the new capacity, and then click **Next** and **Finish**.

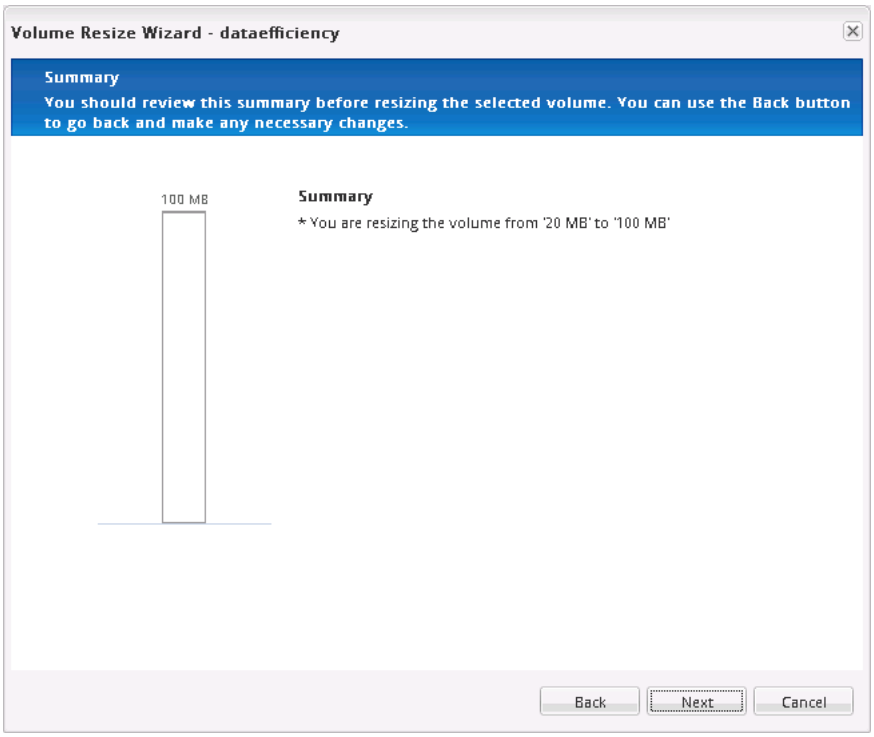

Step 9.

On the Linux (CentOS) machine create 30MB file.

```
dd if=/dev/urandom of=30mfile bs=1k count=30000
ls -la
total 30132
drwxrwxrwx. 2 root root     4096 Dec 10  2014 .
drwxrwxrwx. 4 root root     4096 Dec 10  2014 ..
-rw-r--r--. 1 root root 30720000 Dec 10  2014 30mfile
-rw-r--r--. 1 root root       63 Dec 10  2014 datafile
```

TASK 3: Create a qtree and set quotas

Step 1.

In GUI, select **Storage > Qtrees > (Qtrees on SVM: vs2)**.

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B
IČO 357 034 66
IČ DPH SK2020341158

Step 2.

On the Qtrees page, click **Create**.

Step 3.

In the Create Qtree window, on the Details tab, name the qtree **qtdemo** and ensure that it is in the dataefficiency volume.

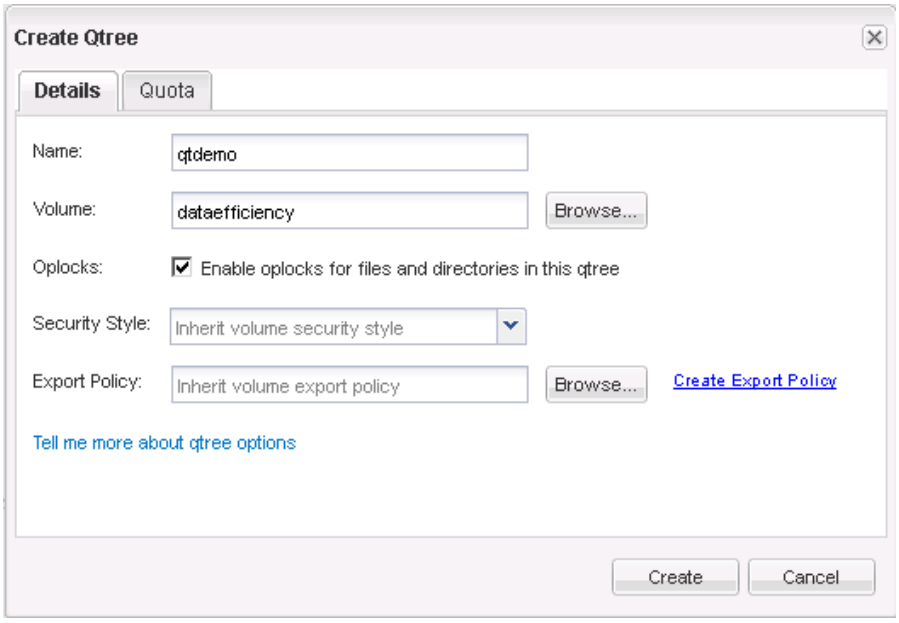

Step 4.

On the Quota tab, add a user quota of **5 MB**, and then click **Create**.

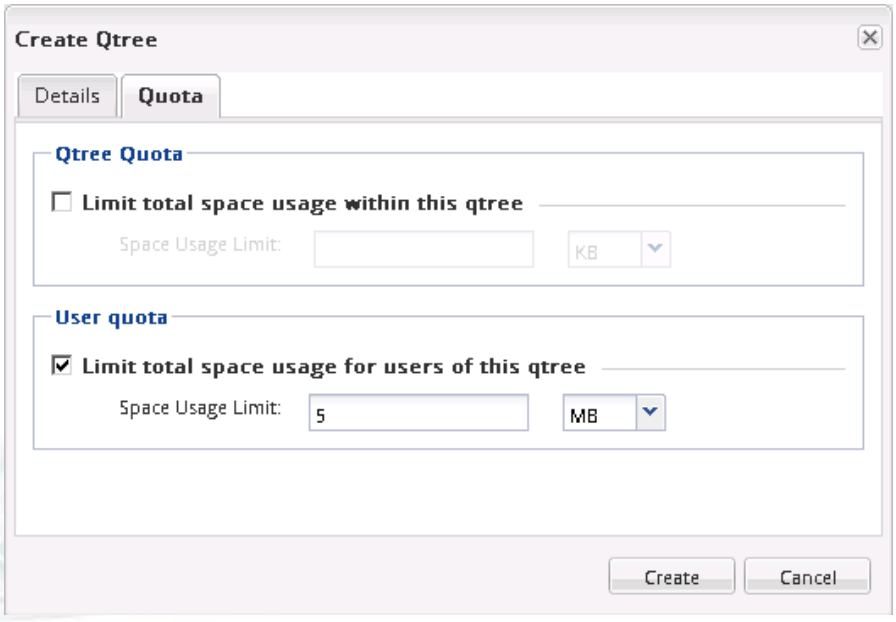

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B

IČO 357 034 66
IČ DPH SK2020341158

Step 5.

In quota Activate window check "**OK to activate the quota**" and click **Activate**.

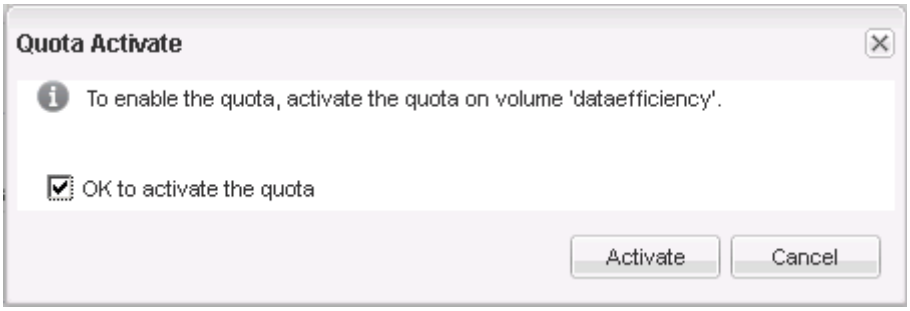

Step 6.

On the Quotas page, check the **dataefficiency** volume. Initialization can take several minutes.

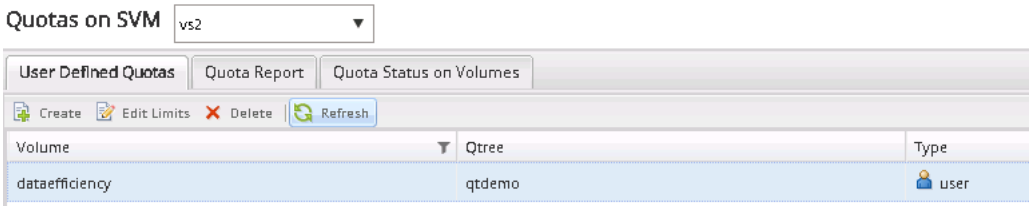

Step 7.

Verify that the activation was successful.

You might need to click the Refresh button.

Step 8.

On your Windows desktop, from the **Desktop**, copy the file **10mfile** into **T:\qtdemo**.

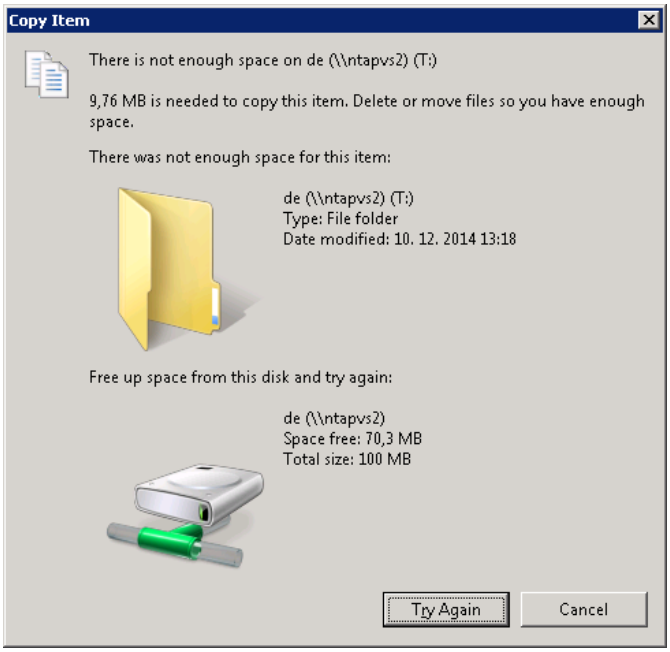

Step 9.

Notice that the copy operation puts the qtree over quota and click **Cancel**.

Step 10.

In GUI, click **the User Defined Quotas**, and then click the **Edit Limits** button.

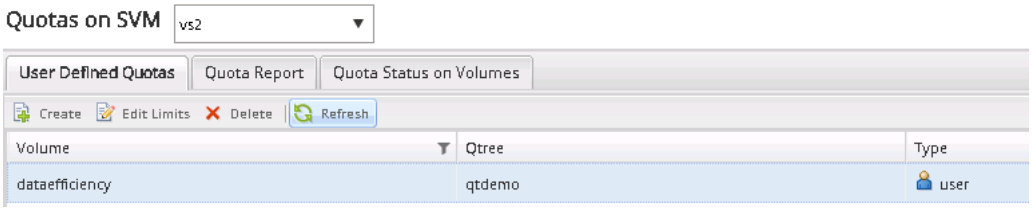

Step 11.

In the **Edit Limits** window, click the **Space Limit** tab, change the Hard Limit to **15 MB**, and then click **Save and Close**.

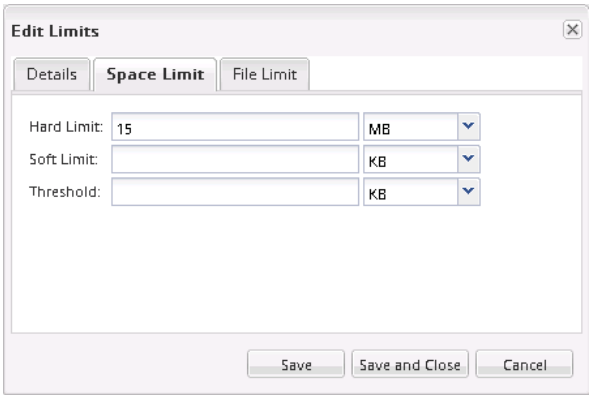

Step 12.

Go back to the **Quota Status on Volumes** tab, select the **dataeffeciency** volume, and click **Refresh**.

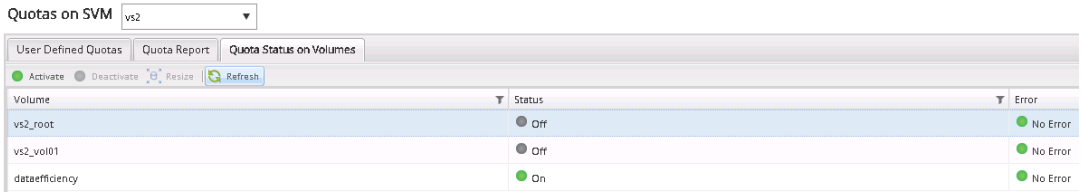

Step 13.

When resizing is over, copy the 10mfile to T:\qtdemo again. The operation should succeed this time.

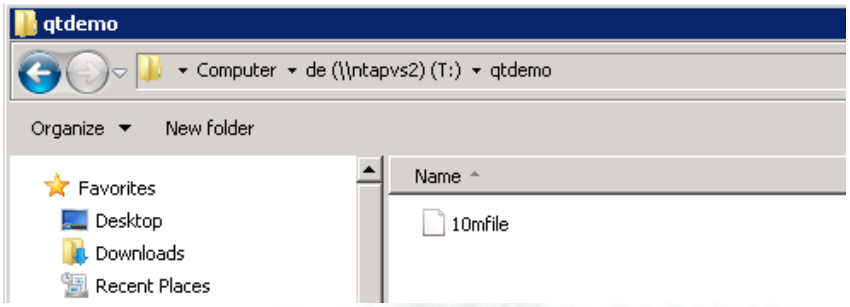

TASK 4: Use FlexClone volumes

Step 1.

On the Volumes page, select the **dataefficiency** volume, click **Actions > Clone**, and then click **Create** and **Volume** to create a FlexClone copy of the volume.

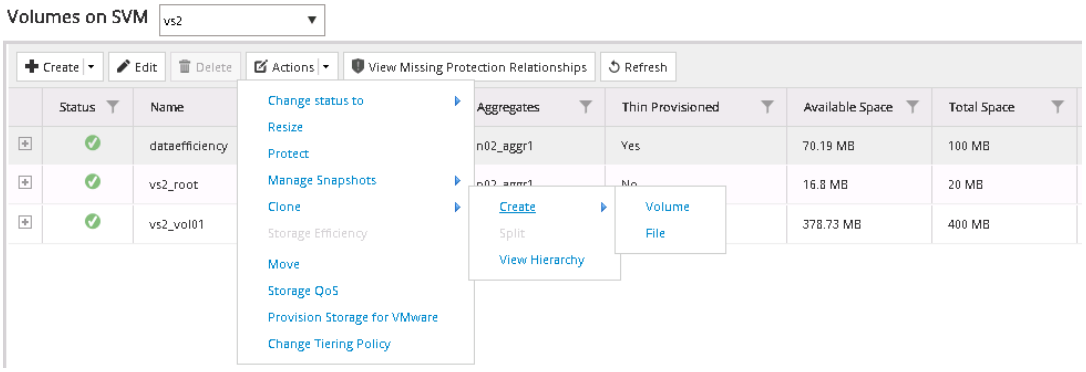

Step 2.

In the Create FlexClone Volume window, accept the default name, select **Create new Snapshot copy now**, and then click **Clone**.

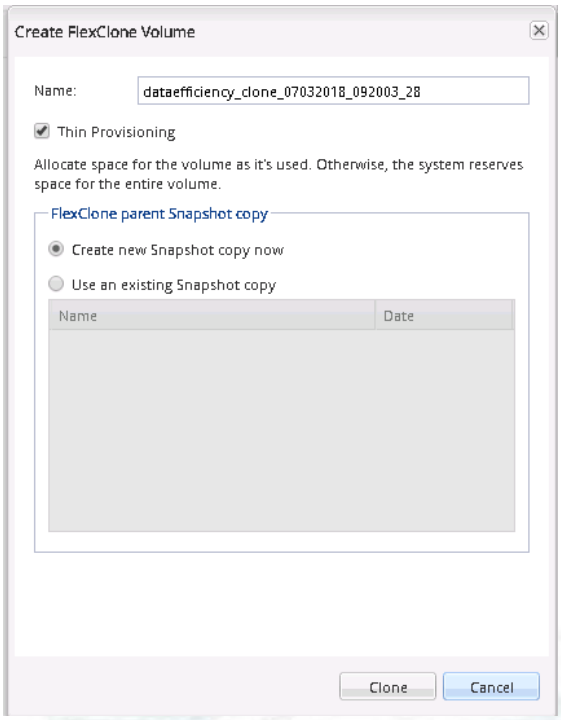

Step 3.

In the top-right corner of the Volume page, click the small icon (cogwheel) at the edge to display a selection list that enables you to add additional fields.

Step 4.

Select the **Type** (if not selected) and **Clone** checkboxes.

Step 5.

In GUI, navigate to the **Storage** > **Namespace** page and notice that clone1 does not appear in the Namespace list.

Step 6.

Mount your clone volume under the path */vs2vol1* as *clone1*.

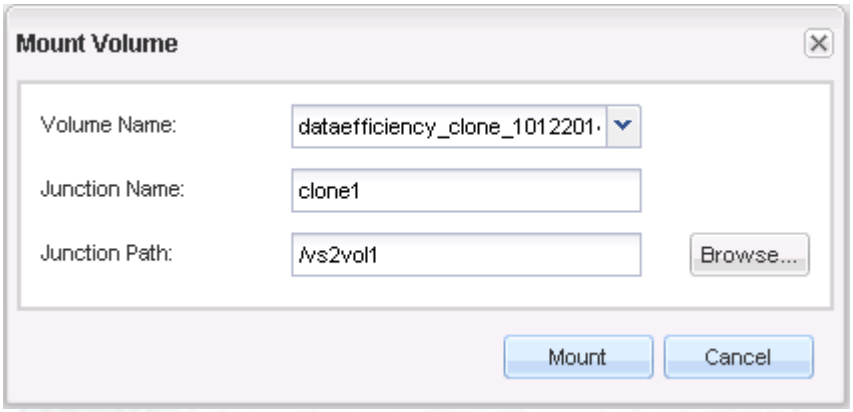

Step 7.

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B

IČO 357 034 66
IČ DPH SK2020341158

Determine where the clone is mounted in the namespace.

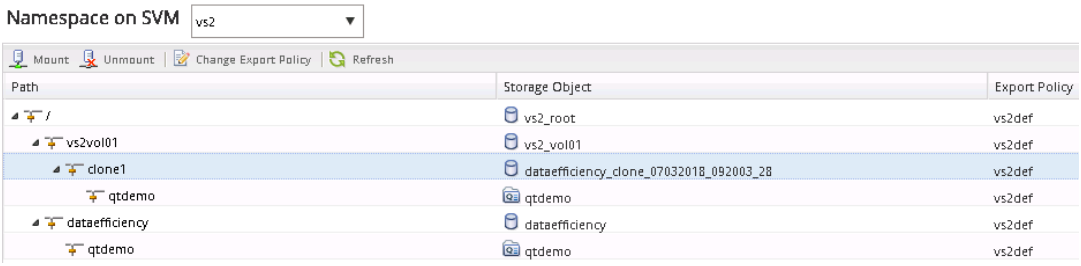

Step 8.

Check the export policy on the clone and set to *vs2def*.

Step 9.

In Windows File Explorer, from your vs2 root volume drive, navigate to */vs2vol01/clone1*.

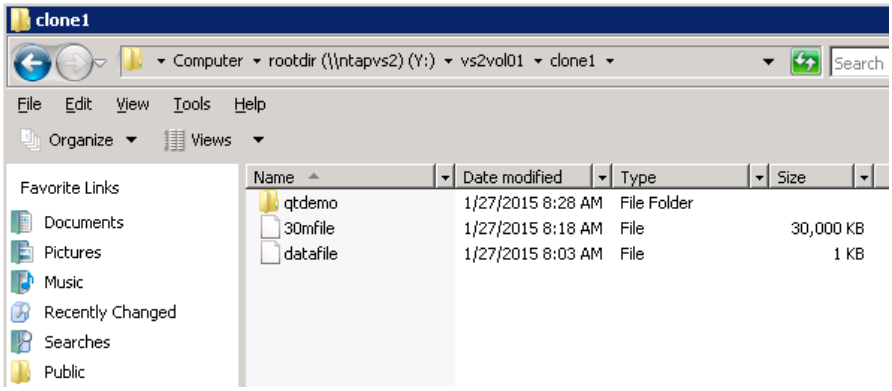

Step 10.

From your NFS client, append some text to datafile on clone1.

cd /mnt/vs2/vs2vol01/clone1

ls

30mfile datafile qtdemo

echo "Edited from a FlexClone..." >> datafile

cat datafile

"This is the content of datafile in volume dataefficiency"

Edited from a FlexClone...

Step 11.

Check the contents of the file on the parent volume, and notice that, although the clone and parent share common data blocks, they function as separate volumes.

cd /mnt/vs2/dataefficiency/

cat datafile

"This is the content of datafile in volume dataefficiency"

Step 12.

In GUI, on the Volumes page, select the clone, and then select **Actions > Clone > Split**.

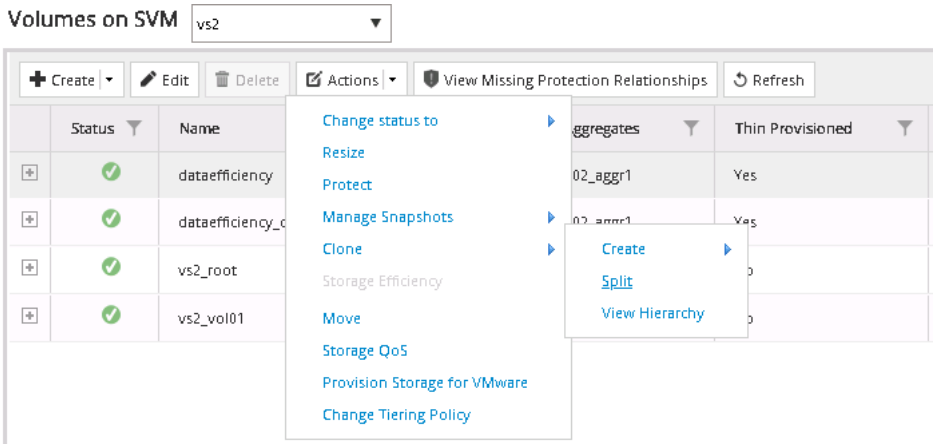

Step 13.

In the Clone Split window, select **the OK to split the clone and delete all its Snapshot copies** checkbox, and then click **Start Split**.

The split takes some time to complete.

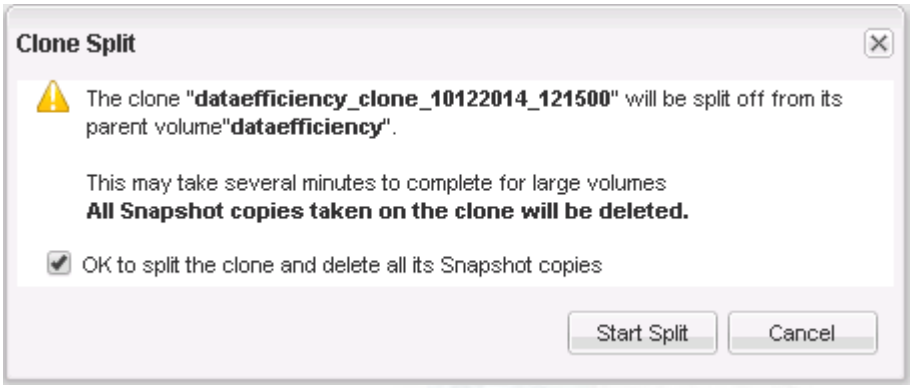

Step 14.

On the Volumes page, notice the status in the **Clone** column. The clone no longer shares blocks with the parent and is now a separate volume.

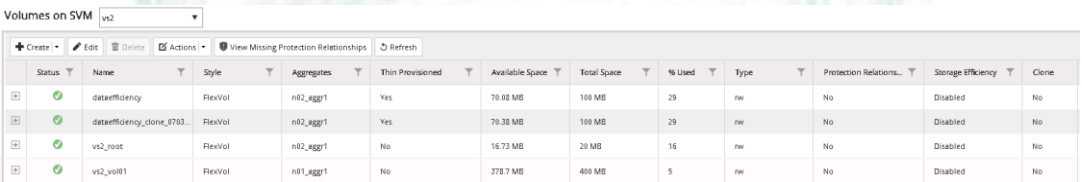

Step 15.

(Optional) If you wish, you can rename the volume to reflect its new status.

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B
IČO 357 034 66
IČ DPH SK2020341158

TASK 5: Enable deduplication and data compression

Step 1.

In GUI, on the **Storage > Volumes** page, select the **dataefficiency** volume and click the **Edit** and **Storage Efficiency** tab.

Step 2.

Select the **Enable Background Deduplication, Policy based** checkboxes, select **Storage Efficiency Policy: Default, Background Compression** and click **Save and Close**.

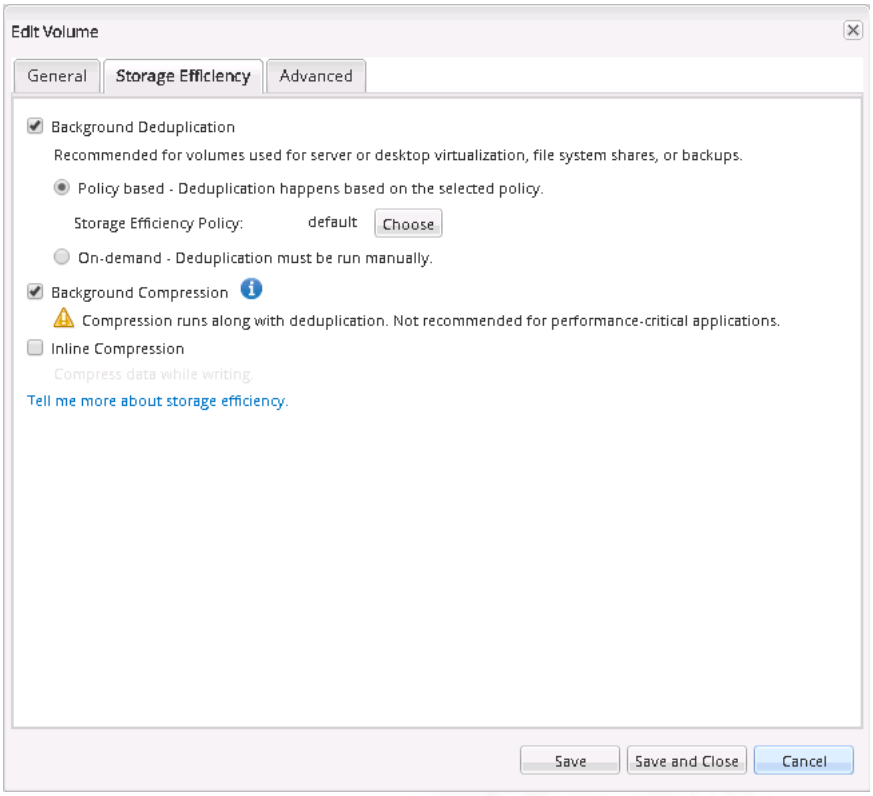

Step 3

On the Volumes page, notice the dataefficiency volume's status in the **Storage Efficiency** column.

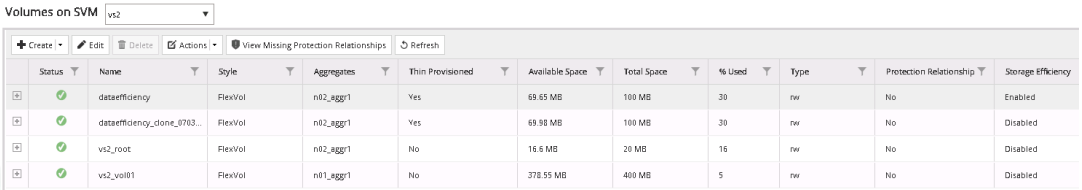

Step 4.

With the dataefficiency volume selected, notice % Used and Available space.

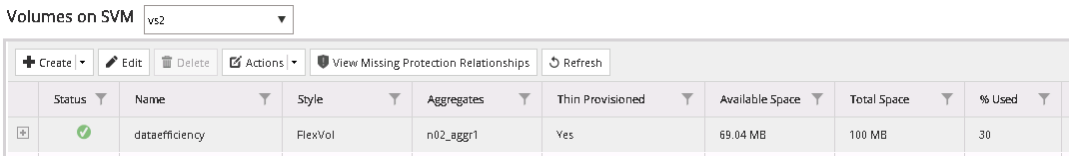

Step 5.

From the Y:\vs2vol01\clone1 directory, copy the 30mfile to the T: drive, which is mounted to the dataefficiency volume.

Step 6.

When prompted, select **Copy, but keep both files**

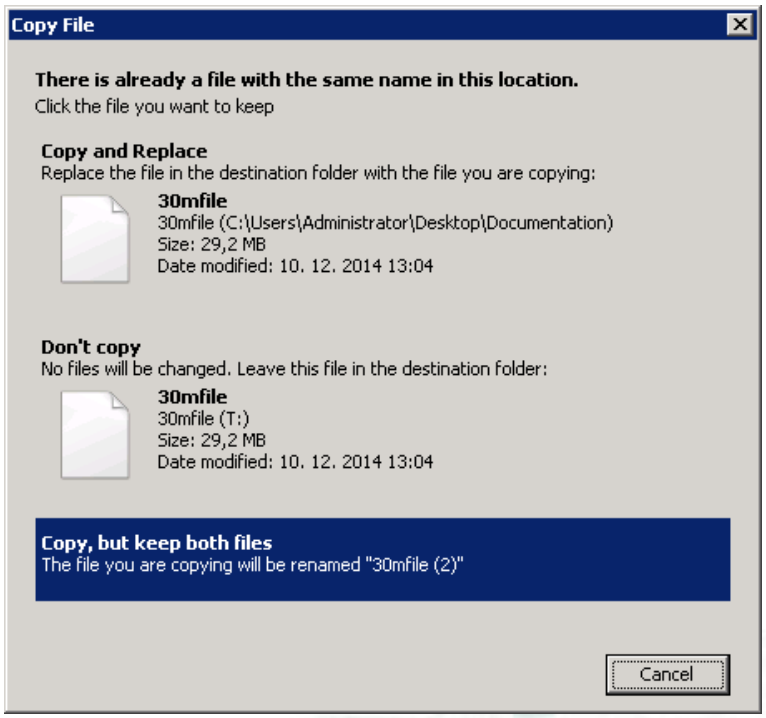

Step 7.

Again with the dataefficiency volume selected, notice % Used and Available space.

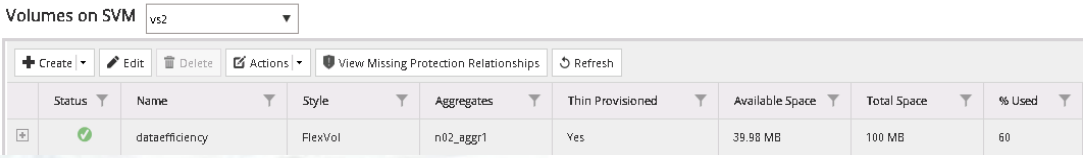

Step 8.

With the dataefficiency volume selected, click the **Actions > Storage Efficiency** and check box **Scan Entire Volume** and than click **Start**.

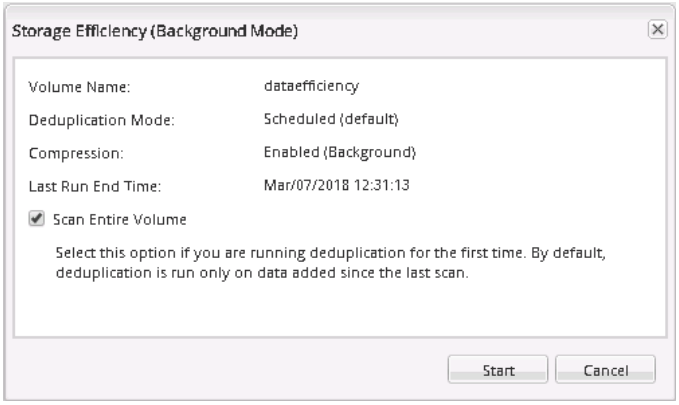

Step 9.

Refresh volume page and notice % Used and Available space.

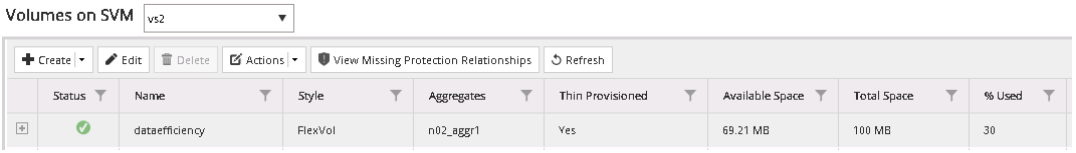

On production systems, it can take longer time.

END OF EXERCISE.

LAB 12: Data Protection: SnapShot and SnapMirror copies

Objectives of this exercise:

- Create and initialize LS and DP SnapMirror replications
- Compare DP mirror replication times
- Add volumes and files to a replicated namespace
- Schedule periodic SnapMirror replications
- Promote an LS mirror
- Use System Manager to configure a SnapMirror relationship
- Use the CLI to configure a SnapMirror relationship
- Update SnapMirror relationships

TASK 1: Create and initialize LS and DP snapmirror replications

Step 1.

Create a mirror for the root volume of your virtual server on the same aggregate as the root volume.

cluster1::> vol create -vserver vs2 -volume vs2_root_ls1 -aggregate n01_aggr1 -type dp

NOTE: Although this is an LS mirror, it must be created as a DP mirror volume and then changed.

Step 2.

Establish the mirror relationship between the source volume and the destination volume and change the destination to an LS mirror.

cluster1::> snapmirror create -source-cluster cluster1 -source-vserver vs2 -source-volume vs2_root -destination-cluster cluster1 -destination-vserver vs2 -destination-volume vs2_root_ls1 -type ls

Step 3.

Create another LS mirror of the same volume on the other node and establish the mirror relationship with the same source.

cluster1::> vol create -vserver vs2 -vol vs2_root_ls2 -aggregate n02_aggr1 -type dp
cluster1::> snapmirror create -source-path cluster1://vs2/vs2_root -destination-path cluster1://vs2/vs2_root_ls2 -type ls

Step 4.

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B
IČO 357 034 66
IČ DPH SK2020341158

Use a summary view and then an instance (detailed) view to view the mirrors.

cluster1::> snapmirror show

cluster1::> snapmirror show -instance

The state of each mirror is “Uninitialized.” No data has been transferred to the mirrors.

Step 5.

Perform the initial (baseline) replication to the set of LS mirrors of this source volume.

cluster1::> snapmirror initialize-ls-set -source-path cluster1://vs2/vs2_root

cluster1::> snapmirror show

NOTE: Now that the first replication of the LS mirrors has occurred, all requests that are sent through existing NFS mounts or CIFS shares to the Vserver root volume are transparently routed to one of the LS mirrors of this Vserver root volume rather than to the read and write volume. This includes write requests, which will fail, because mirrors are read-only.

Step 6.

Create two DP mirrors.

cluster1::> volume create -vserver vs2 -volume vs2_root_dp1 -aggregate n01_aggr1 -type dp

cluster1::> volume create -vserver vs2 -volume vs2_root_dp2 -aggregate n02_aggr1 -type dp

Step 7.

Establish the DP mirror relationships.

cluster1::> snapmirror create -source-path cluster1://vs2/vs2_root -destination-path cluster1://vs2/vs2_root_dp1 -type dp

cluster1::> snapmirror create -source-path cluster1://vs2/vs2_root -destination-path cluster1://vs2/vs2_root_dp2 -type dp

cluster1::> snapmirror show

Step 8.

Perform the initial (baseline) replication to one of the DP mirrors.

cluster1::> snapmirror initialize -source-path cluster1://vs2/vs2_root -destination-path cluster1://vs2/vs2_root_dp1

Step 9.

View the volume Snapshot copies of the source volume.

cluster1::> volume snapshot show -vserver vs2 -volume vs2_root

NOTE: The Snapshot copies are named with a prefix of "snapmirror." These Snapshot copies are specifically for the mirrors and are kept so that future replications to the mirrors can determine which data is new since the last replication. Two SnapMirror Snapshot copies of this volume exist. One is for the set of two LS mirrors. The other is for the DP mirror.

Step 10.

Review the mirrors.

cluster1::> snapmirror show -inst

NOTE: The mirror root_dp1 has the state "Snapmirrored," and the Mirror Timestamp field shows the date and time of that replication. Because this is asynchronous mirroring, mirrors are only as up to date as the manual replications or the scheduled replications keep them.

TASK 2: Compare DP mirror replication times

Step 1.

Replicate the DP mirror that has not been replicated.

cluster1::> snapmirror initialize -source-path cluster1://vs2/vs2_root -destination-path cluster1://vs2/vs2_root_dp2

Step 2.

Review the mirrors.

cluster1::> snapmirror show -inst

The two DP mirrors have different replication times.

Step 3.

View the volume Snapshot copies of the source volume.

cluster1::> volume snapshot show -vserver vs2 -volume vs2_root

NOTE: Three SnapMirror Snapshot copies of this volume exist. They enable SnapMirror software to perform incremental transfers to the mirrors of this source volume. One of these SnapMirror Snapshot copies is for the set of two LS mirrors. The other two are for the two independent DP mirrors.

Step 4.

Replicate to both DP mirrors.

cluster1::> snapmirror update -source-path cluster1://vs2/vs2_root -destination-path cluster1://vs2/vs2_root_dp1

cluster1::> snapmirror update -source-path cluster1://vs2/vs2_root -destination-path cluster1://vs2/vs2_root_dp2

Step 5.

Review the mirrors.

cluster1::> snapmirror show -inst

Step 6.

Review the volume Snapshot copies of the source volume.

cluster1::> volume snapshot show -vserver vs2 -volume vs2_root

Three SnapMirror Snapshot copies of this volume still exist.

TASK 3: Add volumes and files to a replicated namespace

Step 1.

In your Vserver on the aggregate vs2_aggr1, create a volume.

cluster1::> volume create -vserver vs2 -volume vs2_vol03 -aggregate n01_aggr1 -junction-path /vs2vol3 -policy vs2def

Step 2.

Access your Linux client machine.

Step 3.

Assuming that your data LIF is still mounted from that NFS client, view the root of your namespace.

cd /mnt/vs2

ls

Step 4.

From your cluster shell, perform an incremental replication to the set of LS mirrors of the Vserver root volume.

cluster1::> snapmirror update-ls-set -source-path cluster1://vs2/vs2_root

Step 5.

After the mirror jobs are completed, from your NFS client, review the mirror of the root of your namespace.

ls /mnt/vs2

The junction for vs2_vol03 appears.

Step 6.

In the root of the namespace, begin to create a file called “myfile.”

touch /mnt/vs2/myfile

Because the NFS client references one of the LS mirrors, which is read-only, the file creation fails. No writes can be done to an LS mirror.

Step 7.

Using the special .admin path, mount (as “root”) your namespace.

mkdir /mnt/vs2rw

mount 10.34.x1.40:/.admin /mnt/vs2rw

The .admin path forces your mount to use the source volume, which is read and write, rather than allowing automatic routing to LS mirrors.

Step 8.

In the read and write root of the namespace, create a file called “myfile.”

touch /mnt/vs2rw/myfile

ls /mnt/vs2rw/myfile

The new file should appear.

Step 9.

Using the “normal” path that is routed to the LS mirrors, view the Vserver root.

ls /mnt/vs2/myfile

Because the new file is on the read/write volume and the LS mirrors have not been re-replicated, the “myfile” file is not visible.

Step 10.

From your cluster shell, perform an incremental replication to the set of LS mirrors of the Vserver root volume.

cluster1::> snapmirror update-ls-set -source-path cluster1://vs2/vs2_root

Step 11.

From your NFS client, review the mirror of the root of your namespace.

ls /mnt/vs2/myfile

The file should appear.

TASK 4: Schedule periodic snapmirror replications

For this exercise, you use the 5min schedule for LS mirrors and the hourly schedule for a DP mirror.

NOTE: In a real-world situation, it might be sufficient to replicate to LS mirrors hourly and DP mirrors daily.

Step 1.

View the schedules that were created by default.

cluster1::> job schedule show

NOTE: The schedule that you use for replication depends on the data that is contained in each volume and the requirements of the particular mirror. Some volumes might not need to be replicated, while other volumes do; for example, the DP mirrors don't have to be synchronized as often as the LS mirrors. For this exercise, we'll use the 5min schedule for LS mirrors and the hourly schedule for a DP mirror. In the real world, it may be sufficient to replicate to LS mirrors hourly and DP mirrors daily.

Step 2.

Modify one of the LS mirrors to use the 5min schedule.

cluster1::> snapmirror modify -destination-path cluster1://vs2/vs2_root_ls1 -schedule 5min

Step 3.

View the details of the LS mirrors.

cluster1::> snapmirror show -destination-path cluster1://vs2/vs2_root_ls -instance*

cluster1::> snapmirror show -destination-path cluster1://vs2/vs2_root_ls -fields schedule*

The SnapMirror schedule of each LS mirror is now set to 5min.

Step 4.

Modify one of the DP mirrors to use the hourly schedule.

cluster1::> snapmirror modify -destination-path cluster1://vs2/vs2_root_dp1 -schedule hourly

Step 5.

View the details of the DP mirrors.

cluster1::> snapmirror show -fields schedule

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B
IČO 357 034 66
IČ DPH SK2020341158

NOTE: Only the schedule for the DP mirror that you explicitly modified is set to use this schedule. Each DP mirror of a read and write volume is separate from the other DP mirrors, except that they're associated with the same read and write. They can each have separate replication schedules, or one DP can have a schedule while another does not.

Step 6.

Check the time on a node.

cluster1::> system date show

Step 7.

While waiting for the scheduled mirror update to run, continue to run the *system date show* command and watch for the system clock to reach a multiple of five, which triggers the replication of all the LS mirrors of this volume.

NOTE: After the system clock reaches the five-minute mark, the mirror update occurs but is not visible, which means that no job is created for it.

Step 8.

To determine which ones were replicated, view the summary of the mirrors.

cluster1::> snapmirror show -instance

cluster1::> snapmirror show -fields newest-snapshot-timestamp

TASK 5: Promote an LS mirror

Step 1.

Examine the read and write copy of the Vserver root volume and all its mirrors.

*cluster1::> volume show -volume vs2_root**

The Type field shows the values RW, LS, and DP.

Step 2.

Promote one of the LS mirrors.

cluster1::> snapmirror promote -source-path cluster1://vs2/vs2_root -destination-path cluster1://vs2/vs2_root_ls2

Step 3.

At the prompt to proceed, type **y**.

Step 4.

While the command runs in the foreground, wait for the command to complete.

Step 5.

Review this volume family.

*cluster1::> volume show -volume vs2_root**

The old read and write volume (vs2_root) is gone and the volume vs2_root_ls2 has the type RW. The name of the volume didn't change, but it is the read and write volume now, and all the mirrors replicate from that volume.

Step 6.

Review the SnapMirror relationships.

cluster1::> snapmirror show

Only one LS mirror (vs2_root_ls2) exists.

Step 7.

From the newly appointed read/write volume, replicate everything again.

cluster1::> snapmirror update-ls-set -source-path cluster1://vs2/vs2_root_ls2

TASK 6. Set up an intercluster peer relationship

Step 1.

Start a PuTTY session Console node3.

Step 2.

Create a single node cluster2. Accept the defaults for all values except for the following, and use specific IP addresses and netmasks.

Cluster name:	**cluster2**
CIFS	**COXCBSUQACAAAAABGAAAAAAAAAAA**
FlexClone	**WGKOESUQACAAAAABGAAAAAAAAAAA**
iSCSI	**OIVRBSUQACAAAAABGAAAAAAAAAAA**
NFS	**QTZNASUQACAAAAABGAAAAAAAAAAA**
SnapMirror	**KMMZDSUQACAAAAABGAAAAAAAAAAA**
SnapProtect	**SKBWGSUQACAAAAABGAAAAAAAAAAA**
SnapRestore	**YROKDSUQACAAAAABGAAAAAAAAAAA**
SnapVault	**IBIDFSUQACAAAAABGAAAAAAAAAAA**

Cluster management port:	**e0c**
Cluster management IP:	**10.34.x1.22**
Cluster management subnet mask:	**255.255.255.0**
Cluster management gateway:	**10.34.x1.1**
DNS domain:	**ntap.lab1**
Name server:	**10.34.x1.10**
Node management port:	**e0c**
Node management IP:	**10.34.x1.33**
Node management netmask:	**255.255.255.0**
Node management gateway:	**10.34.x1.1**

Step 3.

To check the status of your new cluster, enter the **cluster show** command at the console.

cluster2::> cluster show

Node	*Health*	*Eligibility*
---	---	---
cluster2-01	*true*	*true*

Step 4.

Disable autosupport

cluster2::> autosupport modify -support disable

Step 5.

Log off from terminal (telnet) session and log on trough SSH session to cluster management.

Step 6.

Log on cluster2 using ssh and cluster management IP address.

Step7.

This step is necessary only when running a virtualized cluster. The virtual disks attached to each node are small compared to real hard disks, resulting in unrealistic sizes for aggr0 and vol0.

Turn off snapshots on vol0 on both nodes.

cluster2::> system node run -node cluster2-01 vol options vol0 nosnap on

Step 8.

Set snap reserve on vol0 to 0%.

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B
IČO 357 034 66
IČ DPH SK2020341158

cluster2::> system node run -node cluster2-01 snap reserve vol0 0

Step 9.

From the command line, show the aggregates:

cluster2::> stor aggr show

Step 10.

Rename the aggr0 aggregates so that they are more easily identifiable.

cluster2::> aggr rename -aggregate aggr0 -newname aggr0_n3

Step 11.

Verify the new names.

cluster2::> stor aggr show

The aggr0 aggregates have only one data disk each. The size of the data disk is 1GB.

Step 12.

Next steps till end of the exercise are specific only for the simulator environment. Add 3 data disks to each aggr0 aggregate.

cluster2::> aggr add-disk -aggregate aggr0_n3 -diskcount 3

Step 13.

Increase the size of vol0 to 3GB.

cluster2::> system node run -node cluster2-01 vol size vol0 +2g

Step 14.

Verify the vol0 settings.

cluster2::> volume show -vserver cluster2-01 -volume vol0

Step 15.

Assign all unassigned disk drives

cluster2::> storage disk show -container-type unassigned

…

cluster2::> storage disk assign -all -node cluster2-01

cluster2::> storage disk show -container-type unassigned

There are no entries matching your query.

Step 16. Set NTP server.

cluster2::> cluster time-service ntp server show

This table is currently empty.

cluster2::> cluster time-service ntp server create -server 10.34.x1.10

cluster1::> cluster time-service ntp server show

Server Version

------------------------------ -------

10.34.x1.10 auto

Step 17.

Set date and time (to be synchronized with your AD server). Time difference between AD controller and clusters have to be less than 5 minutes, otherwise you'll be not able register SVMs to domain. Even you set up NTP server, the system time is not moved immediately but in some steps. So better is set up the time manually.

cluster2::> cluster date show
Node Date Time zone
--------- ------------------------- -------------------------
cluster2-01
3/9/2018 10:33:46 +00:00 Etc/UTC

Step 18.

cluster1::> cluster date modify -timezone Europe/Bratislava -date "1/12/2016 12:17:25"

Step 19.

From the cluster shell, create an aggregate.

cluster2::> aggr create -aggr aggr1 -diskcount 16 -node cluster2-01

Step 20.

From the cluster2 tab in GUI, access Storage > SVMs window and click Create to start SVM creation wizard.

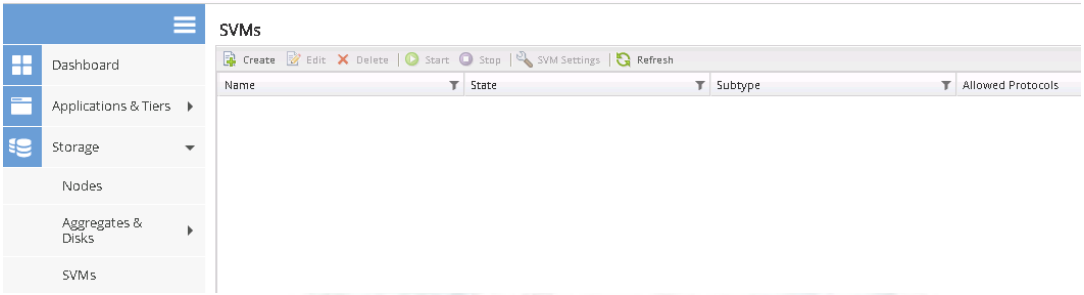

Step 21.

Accept all defaults except for those that follow.

SVM name: **vs9**

Language: **C**

Protocols: **NFS**

Security style: **UNIX**

Aggregate: **aggr1**

Primary domain: **ntap.lab1**

Name server: **10.34.x1.10**

Data LIF home ports: **cluster2_01:e0d**

Data LIF IP addresses: **10.34.x1.61**

Data LIF netmask: **255.255.255.0**

Data LIF gateway: **10.34.x1.1**

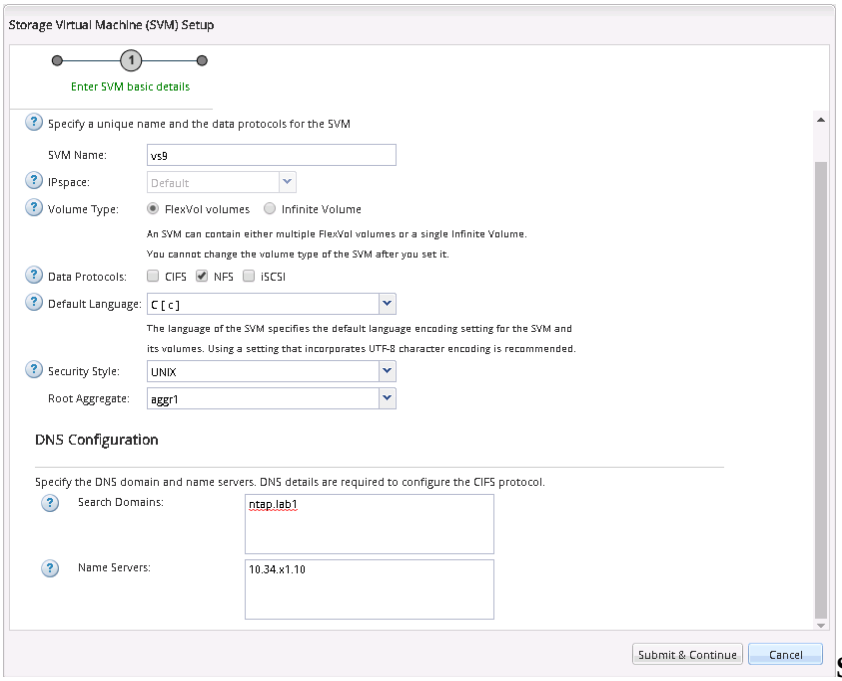

Step 22.

On the SVM administration page, click **Skip**.

NOTE: Unless the destination SVM has the same language type as the source SVM, SnapMirror software does not let you create an intercluster SnapMirror relationship.

Step 23.

Select **Network > Network Interfaces**, click **Create**, and create a second data LIF, using following values:

Name: **vs9_nfs_lif2**

Role: **Data**

Home port: **cluster2_01:e0c**

IP address: **10.34.x1.62**

Netmask: **255.255.255.0**

Gateway: **10.34.x1.1**

Step 24.

Create intercluster LIF on cluster2.

cluster2::> net int create -vserver cluster2 -lif ic1 -role intercluster -home-node cluster2-01 -home-port e0e -address 10.34.x1.71 -netmask 255.255.255.0

Step 25.

On both nodes in cluster1, create an intercluster LIF.

cluster1::> net int create -vserver cluster1 -lif ic1 -role intercluster -home-node cluster1-01 -home-port e0e -address 10.34.x1.81 -netmask 255.255.255.0

cluster1::> net int create -vserver cluster1 -lif ic2 -role intercluster -home-node cluster1-02 -home-port e0e -address 10.34.x1.82 -netmask 255.255.255.0

Step 26.

To create peer relationship between clusters, on the **cluster1** go to **Configuration > Cluster Peers.**

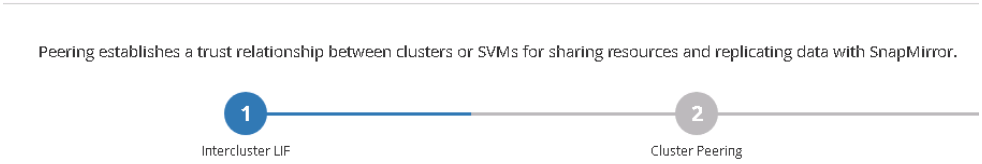

IP space will be **Default**, next click **Submit and Continue**.

Step 27.

Insert Intercluster LIFs address from cluster2: **10.34.x1.71**

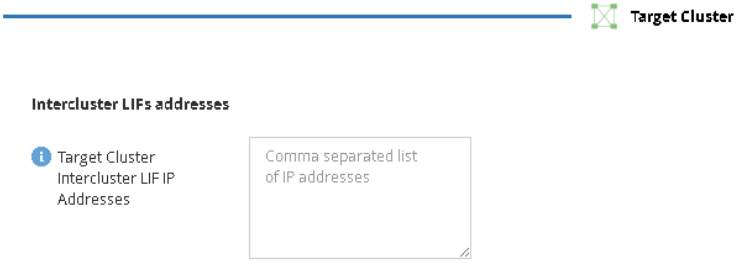

Step 28.

To get a passphrase, insert cluster2 management IP address: **10.34.x1.22,** and then click on the shortcut, which will take you to peering management on the cluster2.

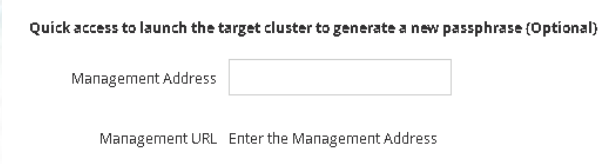

Step 29.

Click on the button Generate Peering Passphrase and generate passphrase.

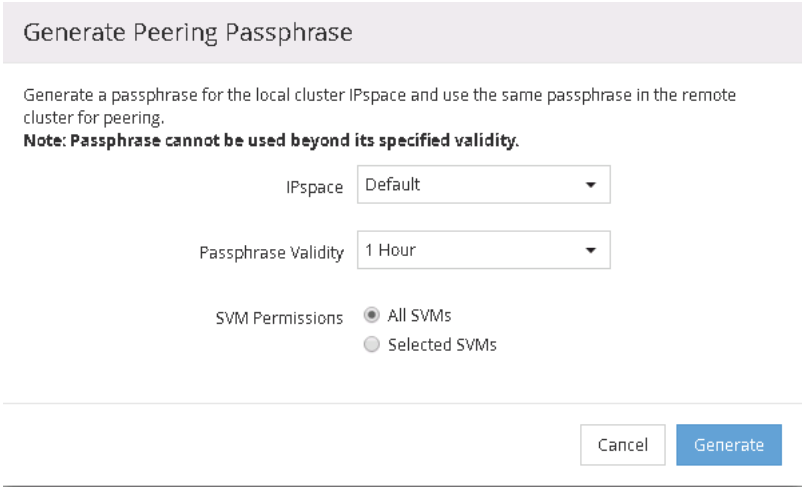

Step 30.

Copy passphrase.

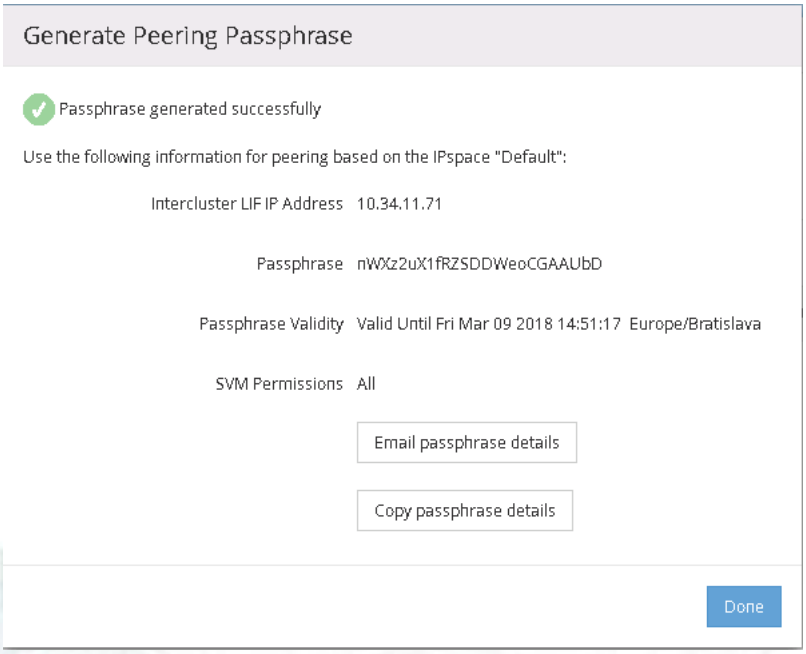

and enter it to field on page for Peering established relationship.

Step 31.

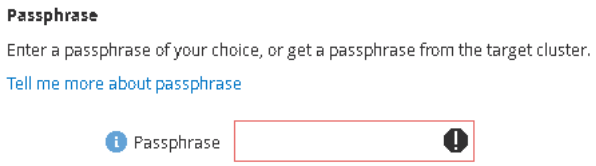

and then click **Initiate Cluster Peering.**

Step 32.

From the cluster shell of cluster1 and cluster2, verify the relationship.

cluster peer show

cluster peer health show

Step 33.

Enter the following command to troubleshoot connectivity issues and long response times between specific intercluster LIFs.

cluster peer ping

TASK 7: Configure a vserver peer relationship

Step 1.

From your putty session on cluster1, verify that the cluster peer relationship is currently configured.

cluster1::> cluster peer show

Step 2.

Verify that the source Vserver has the language set to C.

cluster1::> vserver show -vserver vs2 -fields language

Step 3.

From your PuTTY session on cluster2, verify that the cluster peer relationship is currently configured.

cluster2::> cluster peer show

Step 4.

cluster2::> vserver show -vserver vs9 -fields language

Step 5.

From cluster1, create a new Vserver peer relationship between vs2 in cluster1 and vs9 in cluster2.

cluster1::> vserver peer create -vserver vs2 -peer-vserver vs9 -applications snapmirror -peer-cluster cluster2

Step 6.

From cluster2, accept the peer request.

cluster2::> vserver peer accept -vserver vs9 -peer-vserver vs2

Step 7.

From both clusters, verify that the new Vserver peer relationship is successfully configured.

vserver peer show-all

TASK 8.

Step 1.

In System Manager, click the Protection > Relationship tab.

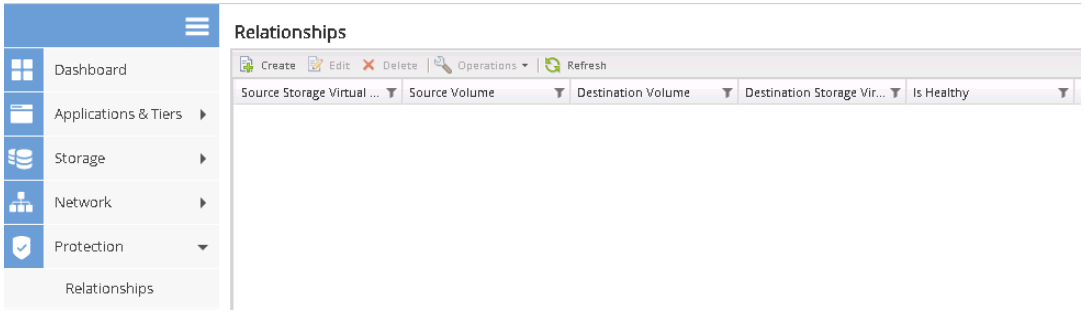

Step 2.

Open the Protection window, click **Create** and select **vs9**.

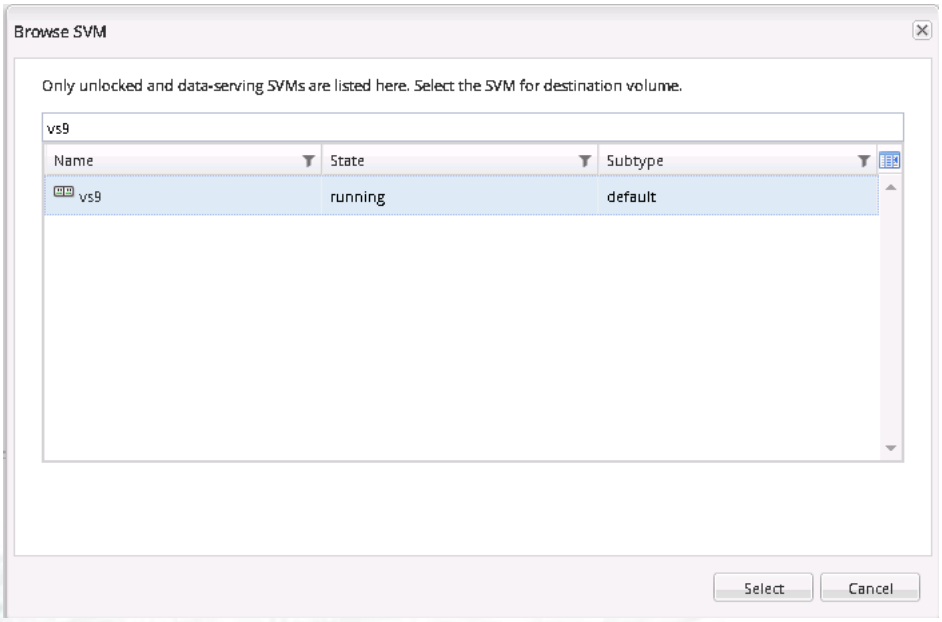

Step 3.

Configure the mirror relationship as follows:

Relationship type: **Mirror**

Source Cluster: **cluster1**

Source SVM: **vs2**

Source Volume: **vs2_vol01**

Destination SVM: **vs9**

Destination Volume Name Suffix: **_vs9_vol01**

Mirror Policy: **MirrorAllSnapshots**

Schedule: **hourly**

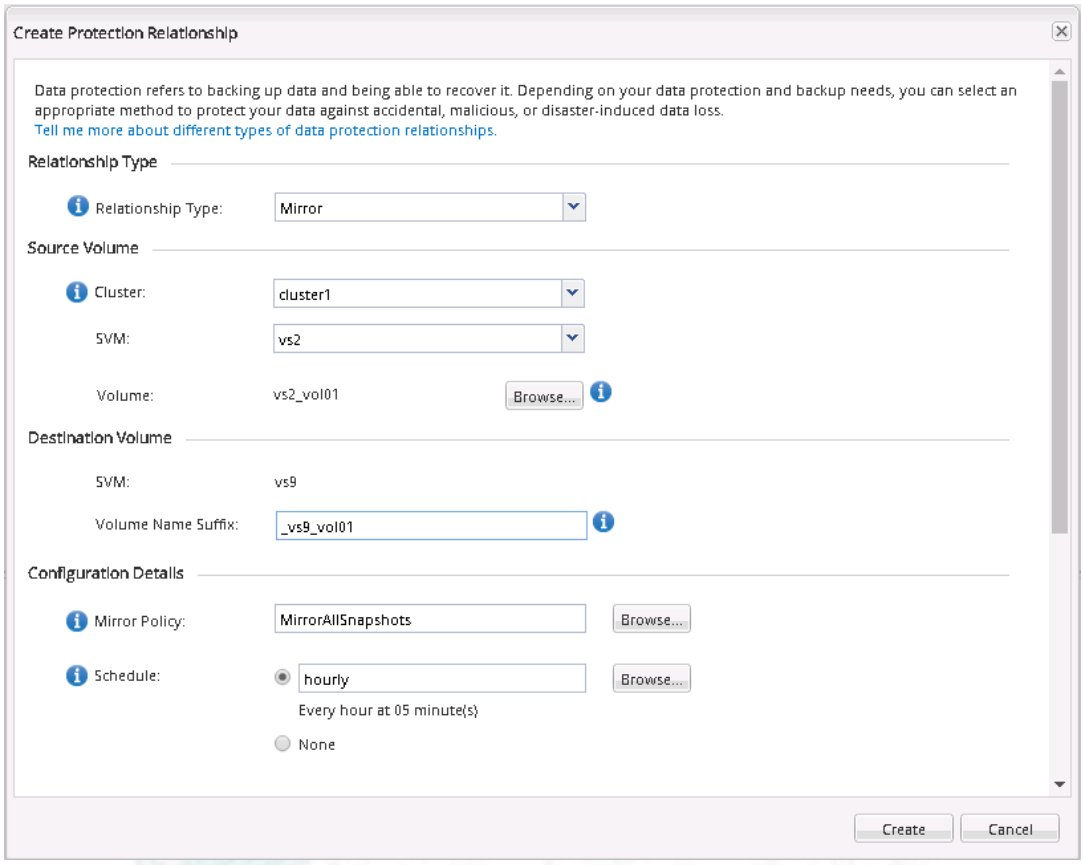

Click **Create.**

Step 4.

Go to tab **Storage** > **Volumes** and change the name to **vs9_vol01**.

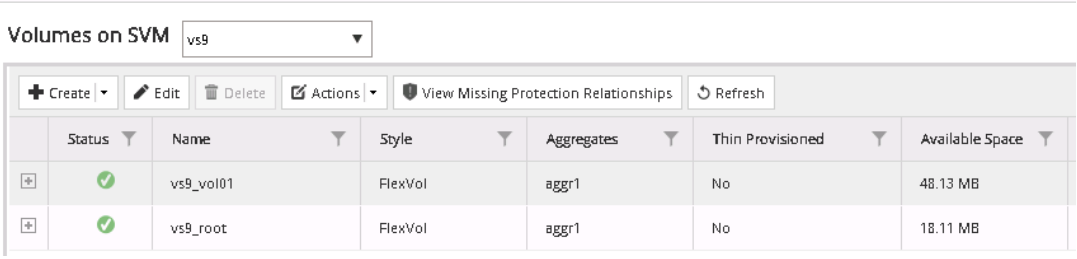

Step 5.

Return to **Storage > Relationships** tab and verify successful initialization of the SnapMirror relationship. The intercluster initialization might take a few minutes.

Step 6.

To view the relationship status from cluster1, select **Storage > Volumes > vs2 > vs2_vol01**

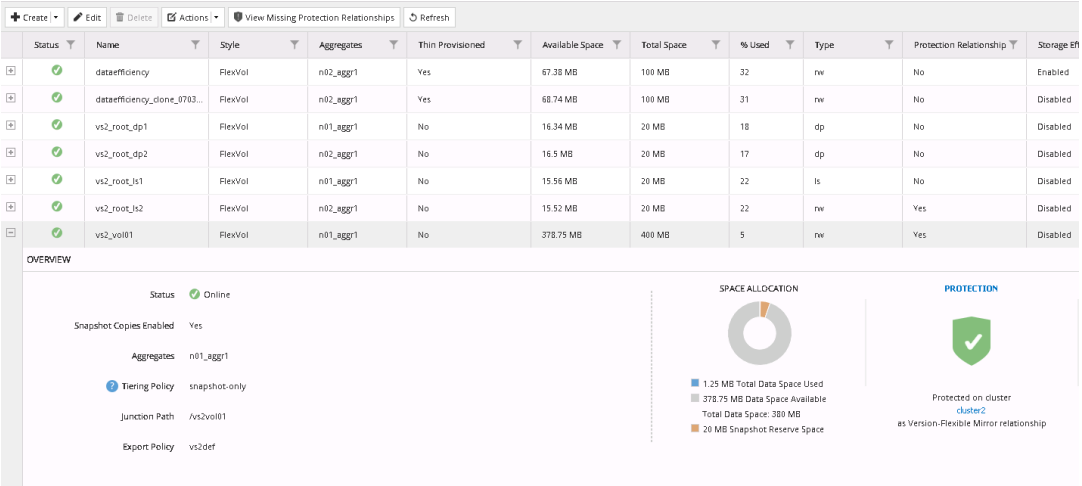

TASK 9.

Use the CLI to configure a snapmirror relationship

Step 1.

Check the language on the source volume.

cluster1::> volume language -vserver vs2 -volume vs2_vol03 -fields language

Step 2.

On cluster2, create a destination volume of type DP with the same language type.

cluster2::> vol create -vserver vs9 -volume vs2_vol03 -aggregate aggr1 -size 25mb -type dp -language C.UTF-8

Step 3.

Create the SnapMirror relationship for vs2_vol03.

cluster2::> snapmirror create -source-path cluster1://vs2/vs2_vol03 -destination-path cluster2://vs9/vs2_vol03 -type DP -schedule daily

NOTE: The source and destination volume names are not required to be the same. In this exercise, you chose the names for clarity. This command can be issued from either cluster without modification.

Step 4.

Initialize the SnapMirror relationship.

cluster2::> snapmirror initialize -destination-path cluster2://vs9/vs2_vol03

NOTE: When you create Storage Virtual Machine (SVM), the language is set for the SVM. The language of the SVM determines the default language setting for volumes in that SVM. You can modify the language of an SVM. You can specify the language for a volume when creating a volume and it can be different from the language of an SVM. If you do not specify the language for a volume then it inherits the language setting of its SVM. After the volume is created, you cannot modify the language of a volume. Therefore, you must be aware of the available language options.

TASK 10. Update snapmirror relationship

Step 1.

From your CIFS client, copy 30mfile into the vol1 share.

net use
New connections will be remembered.
Status Local Remote

OK T: \\ntapvs2\de
OK X: \\ntapvs2\~administrator
OK Y: \\ntapvs2\rootdir
OK Z: \\ntapvs2\vol1

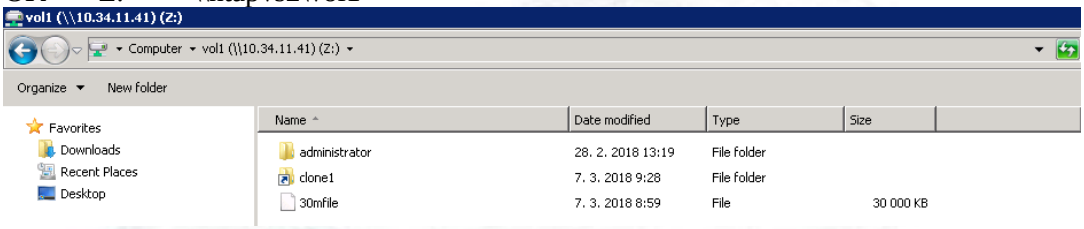

The scheduled update occurs daily, but you can perform an on-demand transfer from the CLI or from System Manager.

Step 2.

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B
IČO 357 034 66
IČ DPH SK2020341158

To perform an on-demand transfer in System Manager, from cluster 2, in the **Protection** > **Relationship** pane, select the relationship for **vs2_vol01**, click **Operations** to see the menu, and then click **Update**.

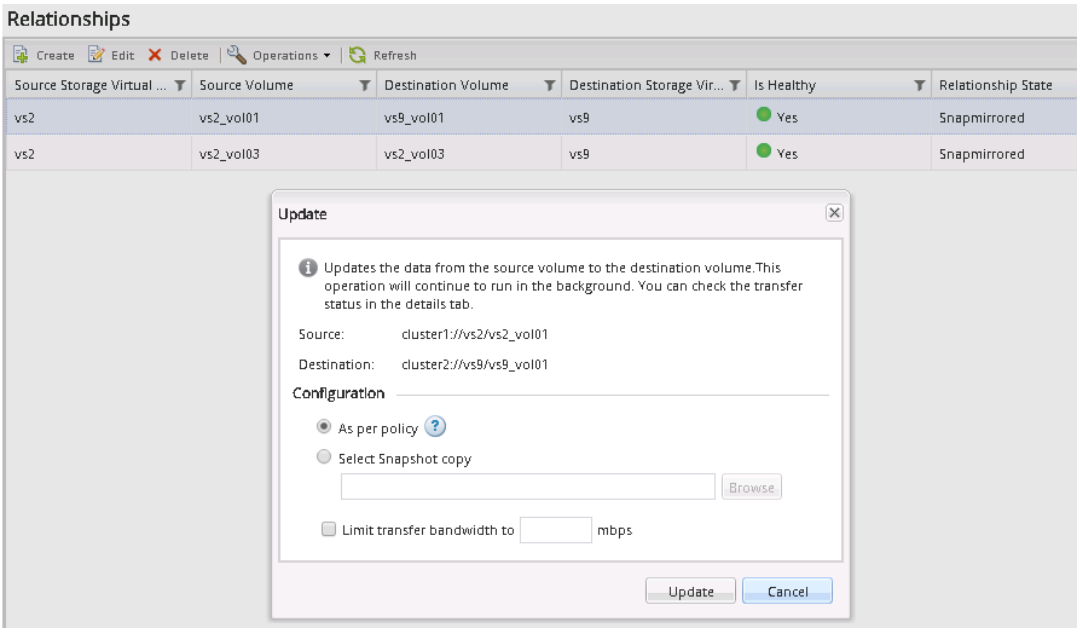

Step 3.

In the Details frame at the bottom of the SnapMirror window, view the status of each SnapMirror relationship.

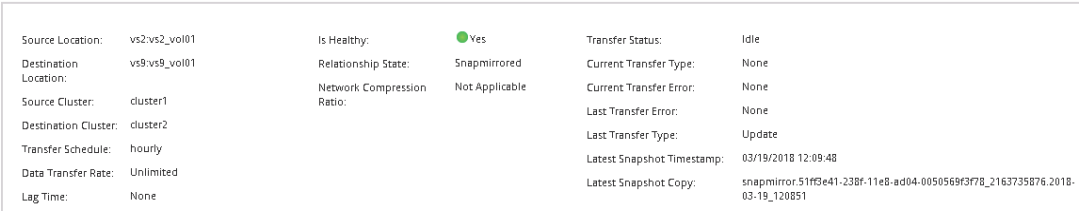

END OF EXERCISE.

LAB 13: Data Protection: Backups and disaster recovery

In this exercise, perform a SnapVault backup restore data from a SnapVault backup.

OBJECTIVES

By the end of this exercise, you should be able to:

- Configure SnapMirror replication
- Perform a SnapVault backup
- Restore data from a SnapVault backup

TASK 1. Configure a snapvault relationship

Step 1.

In System Manager, click the cluster2 tab and select Protection > Protection Policies.

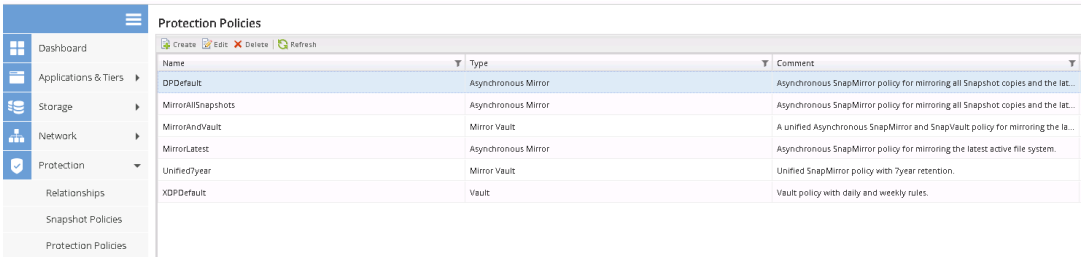

Step 2.

Click **Create**.

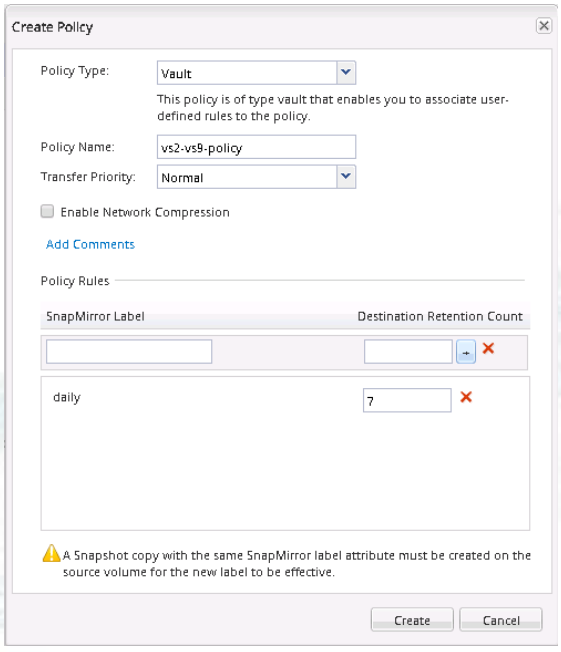

Policy type: **Vault**

Policy name: **vs2-vs9-policy**

Policy rules:

SnapMirror label: **daily**

Destination Retention Count: **7**

Click + to add rule.

Click **Create**.

Step 3.

Verify creation of the policy.

Step 4.

Select **Protection** > **Schedule** and review the current schedules.

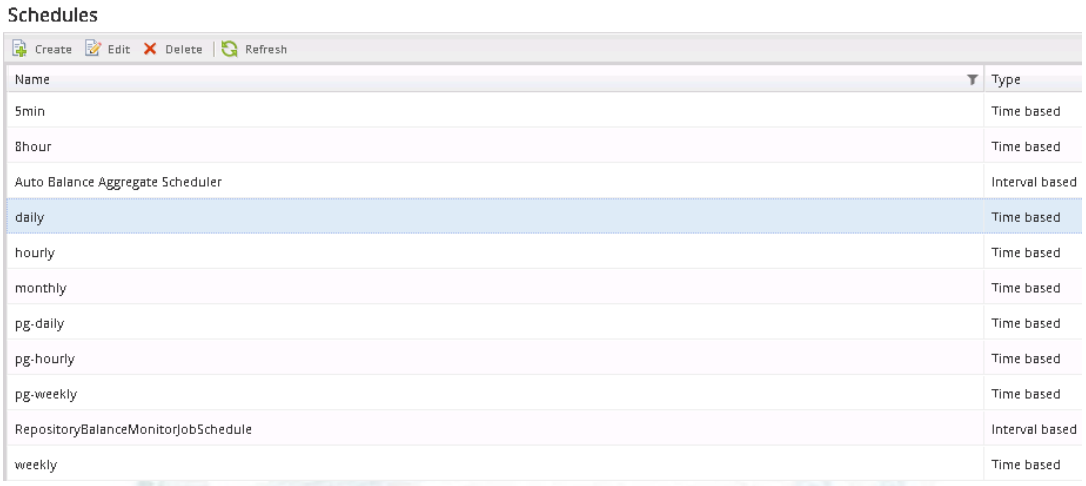

Step 5.

Select **Protection** > **Relationship** then **Create**, and select SVM **vs9**.

Step 6.

On the Create Protection Relationship page, enter and select following:

Relationship Type: **Vault**

Source Cluster: **cluster1**

Source SVM: **vs2**

Source volume: **vs2_vol01**

Destination SVM: **vs9**

Volume Name Suffix: **_dst**

Vault Policy: **vs2-vs9-policy**

Schedule: **hourly**

Initialize Relationship: **selected**

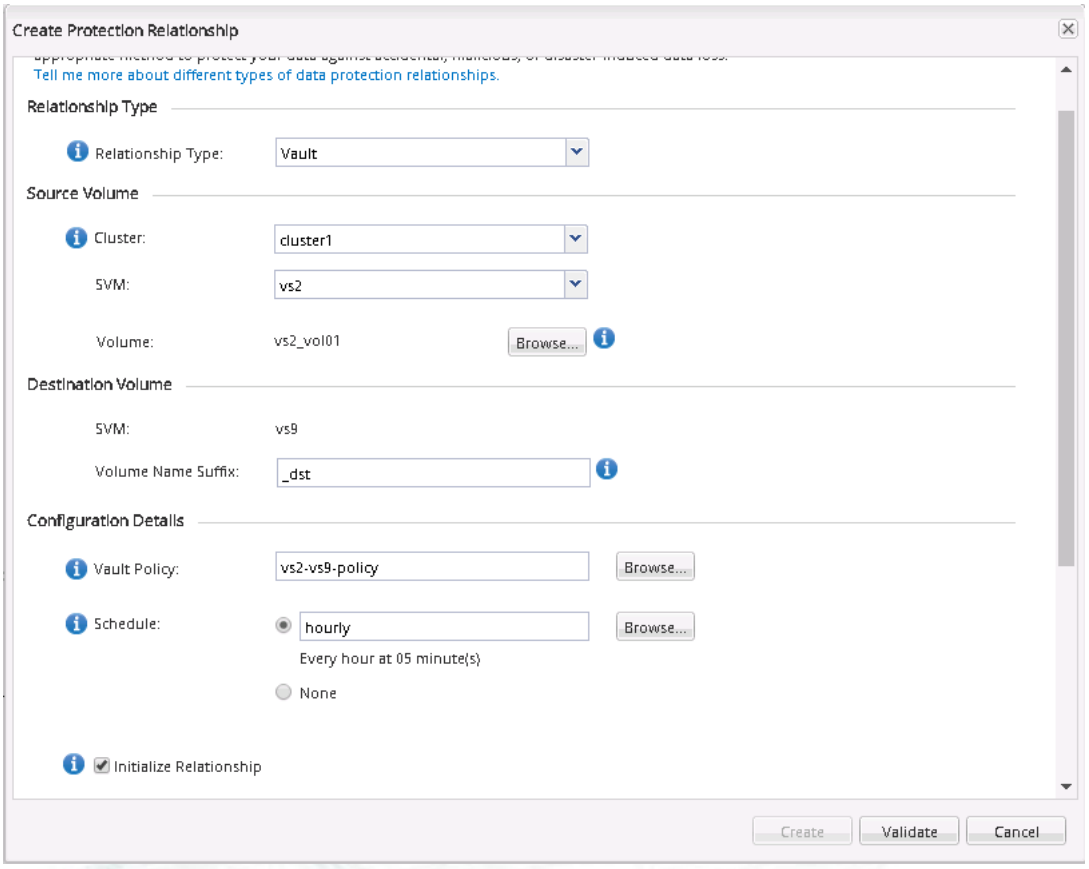

Click **Validate.**

Step 7.

After successful validation, click **Create**.

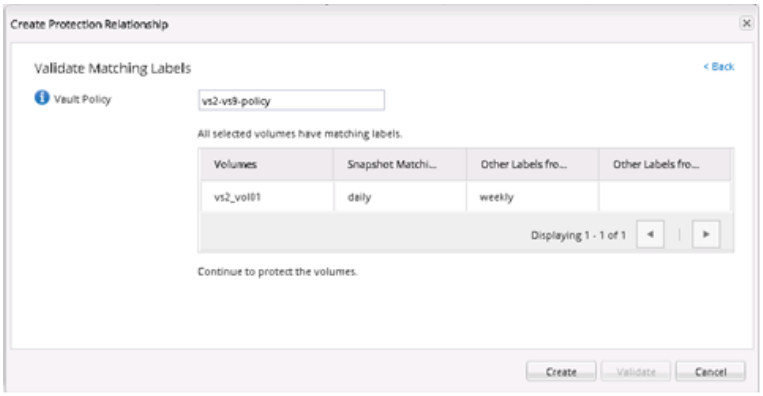

Step 8.

Verify initialization from **cluster2**.

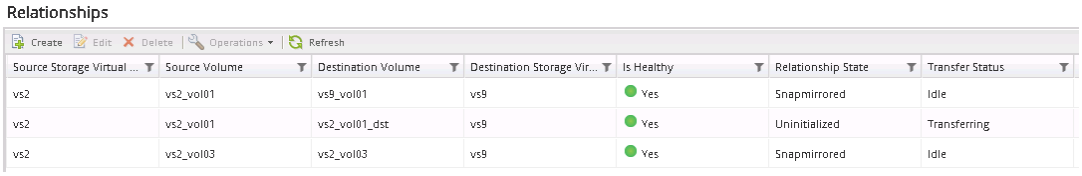

Step 9.

To view the SnapVault relationship status from cluster1, select Storage > Voumes > select vs2_vol01 from the list, and click + to see details.

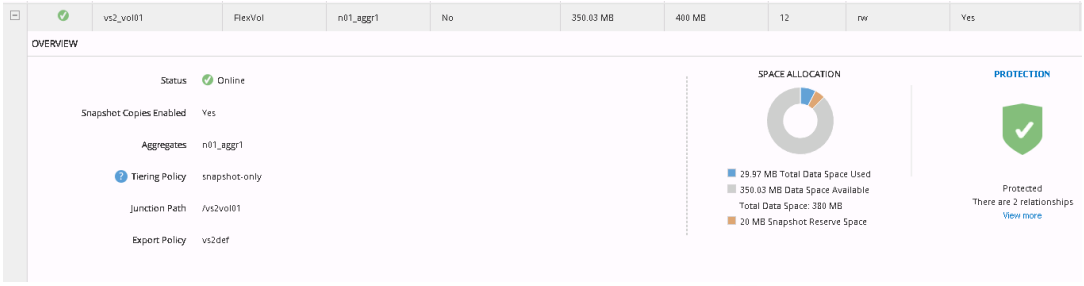

Step 10.

Click **View more**.

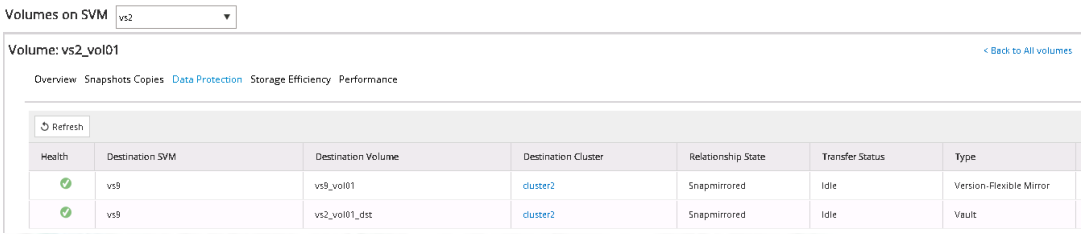

TASK 2. Restore from a SnapVault backup

Step 1.

On cluster2, navigate to the **Protection > Relationship** page, select **vs2_vol01**, with Relationship Type **Vault** and click the **Operations button**, and select **Restore**.

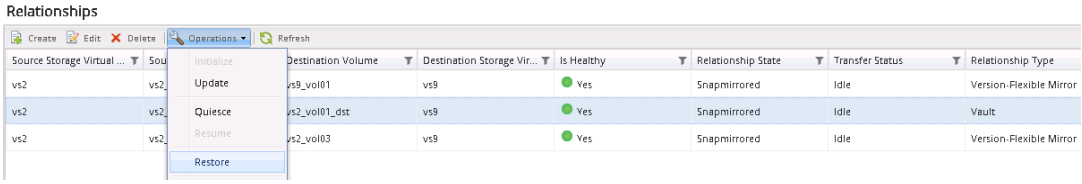

Step 2.

On the restore page, enter and select following.

- Restore to: **Other volume**
- Cluster: **cluster1**
- SVM: **vs2**
- Volume (new): **vs2_vol01_vault_restore**
- Aggregate: **n01_fp1**
- Configuration: **Latest Snapshot Copy**
- OK to restore the volume from the SnapShot copy: **select**

NOTE: The Restore button is activated when you select the “OK to restore the volume from the snapshot copy” checkbox.

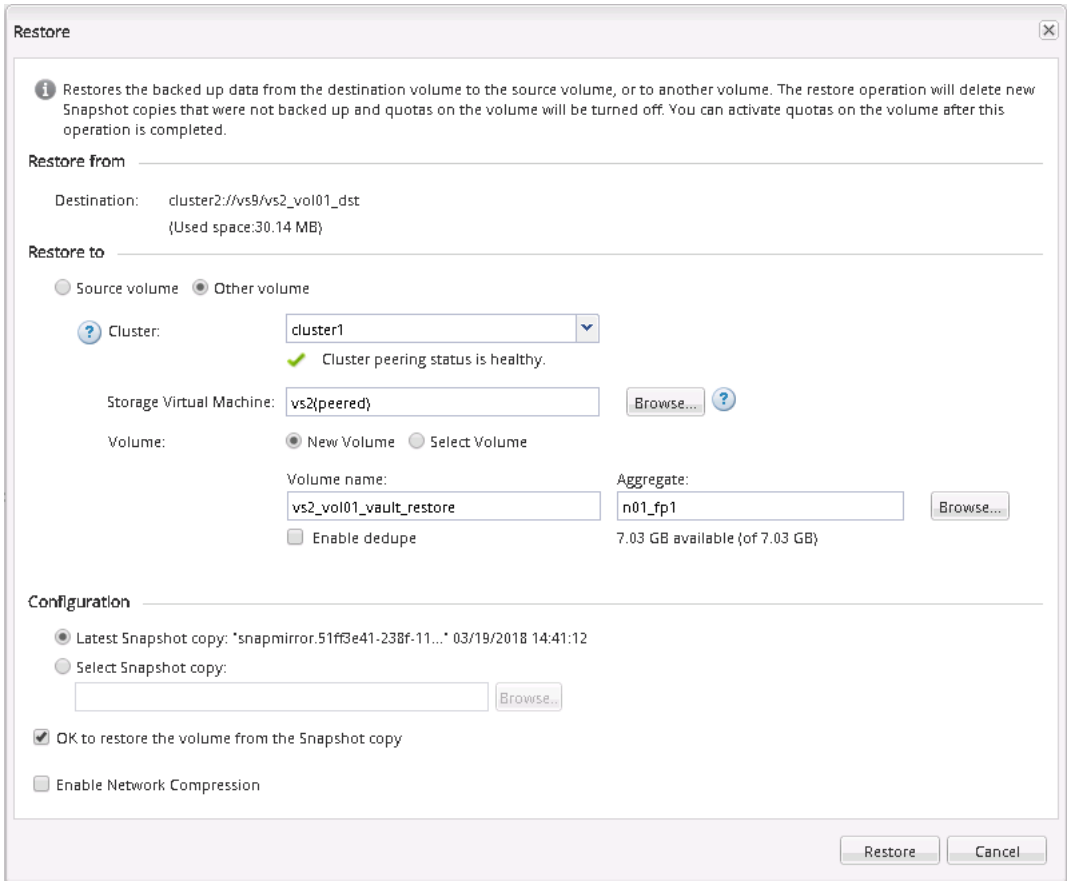

Step 3.

Click **Restore**.

Step 4.

On the cluster1 tab, select **Storage > Volumes > volumes on SVM: vs2**

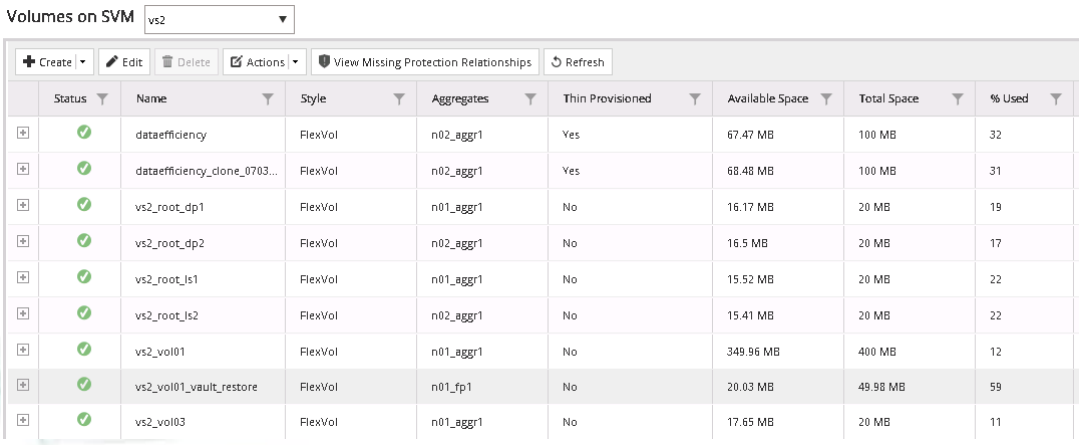

Status	Name	Style	Aggregates	Thin Provisioned	Available Space	Total Space	% Used
	dataefficiency	FlexVol	n02_aggr1	Yes	67.47 MB	100 MB	32
	dataefficiency_clone_0703...	FlexVol	n02_aggr1	Yes	68.48 MB	100 MB	31
	vs2_root_dp1	FlexVol	n01_aggr1	No	16.17 MB	20 MB	19
	vs2_root_dp2	FlexVol	n02_aggr1	No	16.5 MB	20 MB	17
	vs2_root_ls1	FlexVol	n01_aggr1	No	15.52 MB	20 MB	22
	vs2_root_ls2	FlexVol	n02_aggr1	No	15.41 MB	20 MB	22
	vs2_vol01	FlexVol	n01_aggr1	No	349.96 MB	400 MB	12
	vs2_vol01_vault_restore	FlexVol	n01_fp1	No	20.03 MB	49.98 MB	59
	vs2_vol03	FlexVol	n01_aggr1	No	17.65 MB	20 MB	11

Step 5.

Mount the volume at /vault_restore. Go to **Storage > Namespace** and click Mount.

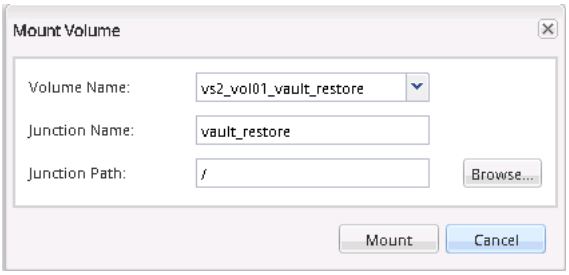

Step 6.

Click **Mount**.

Step 7.

Click **Change Export Policy** button and change the export policy to **vs2def**.

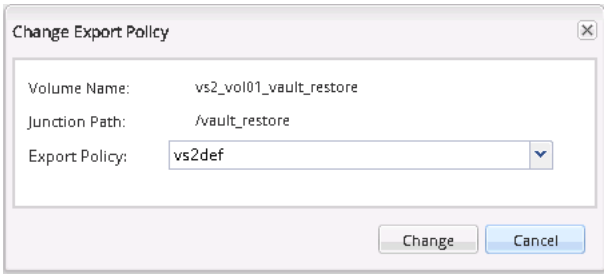

Step 8.

Click **Change**.

Step 9.

From your Linux client, inspect the contents of the restored volume.

ls /mnt/vs2/vault_restore/

ls /mnt/vs2/vs2vol01

NOTE: The restored volume will not appear in the namespace until the Vserver root volume replicates its LS mirrors. This occurs at five-minute intervals. You can manually update the LS mirrors by using the snapmirror update-ls-set command if you prefer.

END OF EXERCISE.

LAB 14: Cluster Management

Objectives of this exercise

- Move a volume
- Use roles to delegate administrative tasks using

TASK 1: Move a volume

Step 1.

From your cluster1 PuTTY session, look at your volumes and note the aggregate on which vs2_vol01 is found.

cluster1::>volume show

Step 2.

Check for Snapshot copies of this volume and notice that there are some Snapshot copies.

cluster1::> volume snapshot show -volume vs2_vol01

Step 3.

Move the volume from its aggregate to an aggregate on the opposite node, and note the job ID that is displayed.

cluster1::> vol move start -vserver vs2 -volume vs2_vol01 -destination-aggregate n02_aggr1

Step 4.

While the move job is running, continually show the volumes (vol show) to watch the temporary (type TMP) volume that is created (on the destination aggregate) as part of this process.

cluster1::> vol show -vserver vs2

After the data movement is complete, an automatic action is taken to change the identity of this temporary volume to become the flexible volume that was moved, and then the old volume on the source aggregate is deleted.

Step 5.

Look at the jobs, then look at the move job by using its job ID while it is running and after it finishes.

cluster1::> job show

cluster1::> job show -id <jobid>

Jobs only remain in the job show output for a few minutes after they are complete.

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B
IČO 357 034 66
IČ DPH SK2020341158

Step 6.

When the move operation is finished, look at the details of the volume and assure yourself that the only thing that changed was the aggregate (and node).

cluster1::> vol show -vserver vs2 -volume vs2_vol01

The junction path is the same as it was before, which means that the volume retained its place in the namespace hierarchy, and because of that, the path known to all NFS and CIFS clients is the same. The change of location is transparent to clients.

Step 7.

Look at the Snapshot copies of the volume that was just moved and notice that the volume move kept all the Snapshot copies as they were.

cluster1::> vol snapshot show -volume vs2_vol01

Step 8.

In GUI, select **vs2_vol01** and click the **Move** button to move the same volume from that node and aggregate back to its original node and aggregate.

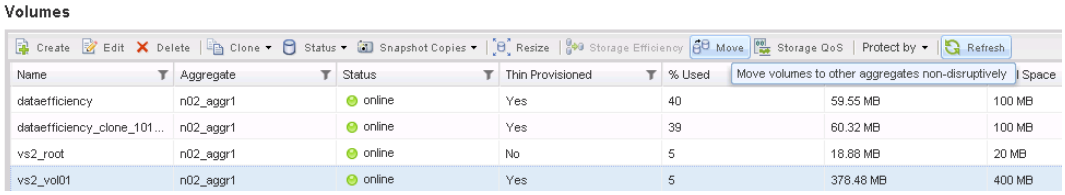

Step 9.

Select **n01_aggr1** and click **Move**.

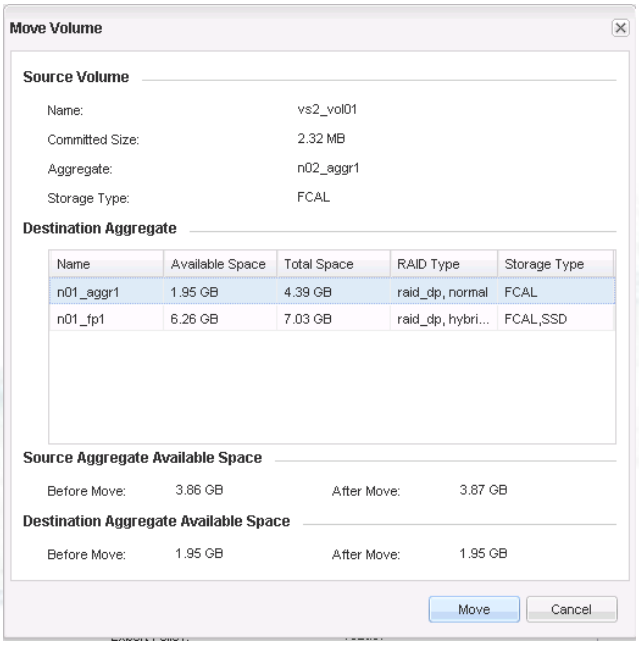

Step 10.

Click **Move** again and then **OK**.

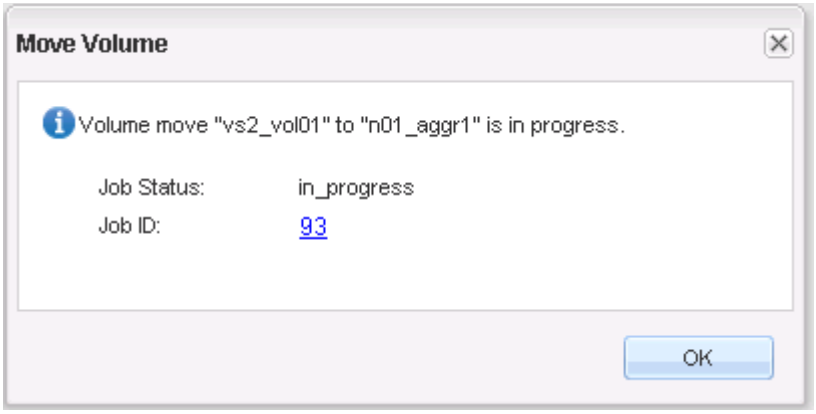

TASK 2: Use roles to delegate administrative tasks

Step 1.

Enable a data LIF owned by Vserver vs2 to support Vserver-level management connections.

cluster1::> net int modify -vserver vs2 -lif vs2_cifs_nfs_lif2 -firewall-policy mgmt

Step 2.

In GUI, select **Configuration > Roles** to display the predefined cluster-scoped roles.

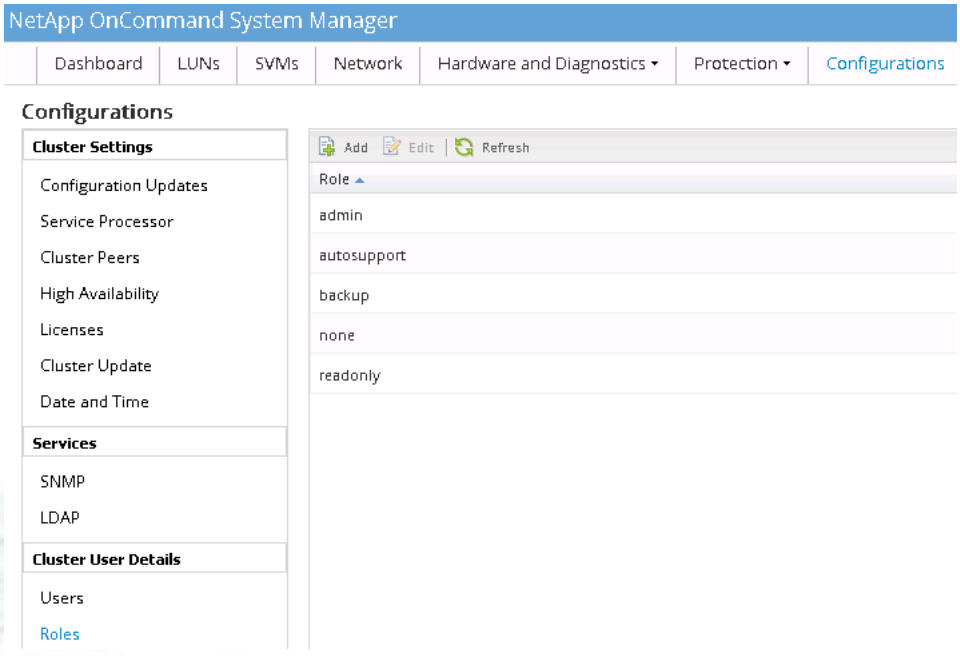

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B

IČO 357 034 66
IČ DPH SK2020341158

Step 3.

Select **SVMs > vs2 > SVM Settings > Roles** to display the Vserver-scoped, predefined roles.

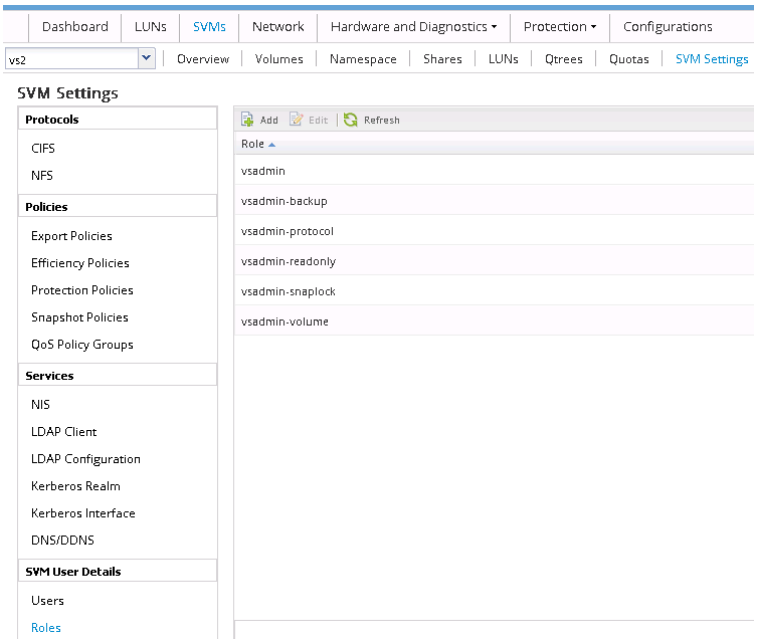

Step 4.

Select a Vserver-scoped administrative role and scroll through the **Role Attributes** pane to see the commands that are available to a user with this role.

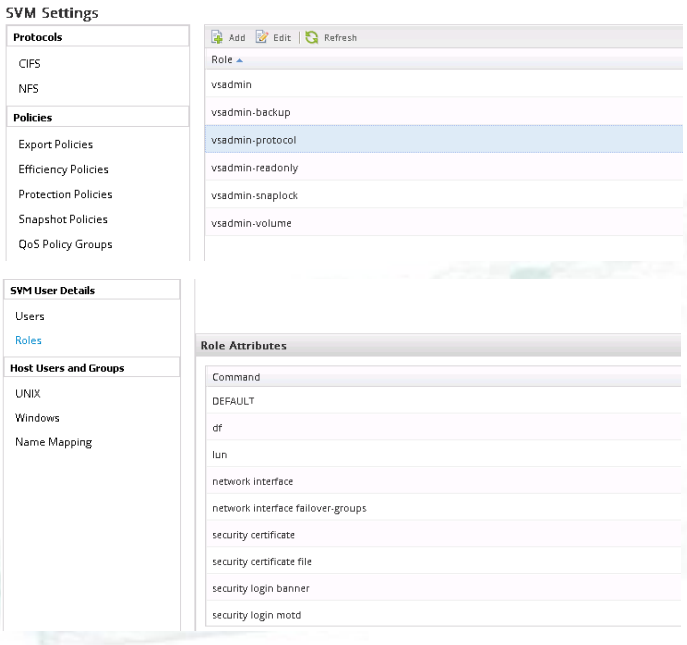

Step 5.

Examine the other roles and compare the commands that are available to users who are assigned to the roles.

Step 6.

Select **SVMs > vs2 > SVM Settings > Users** to display the users for your Vserver, vs2.

Step 7.

Click **Add** to create a user.

Step 8.

Name the user **vs2admin**, choose a password, and click the **Add** button on the right side of the window.

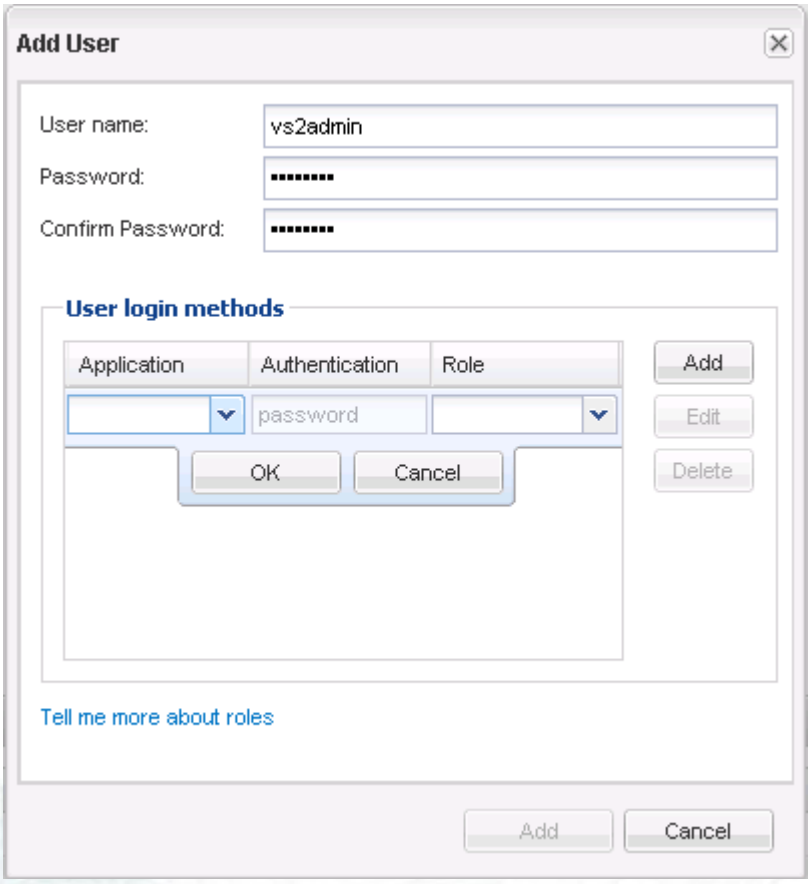

Step 9.

Select the application **ssh** and the role **vsadmin**, and click **OK**.

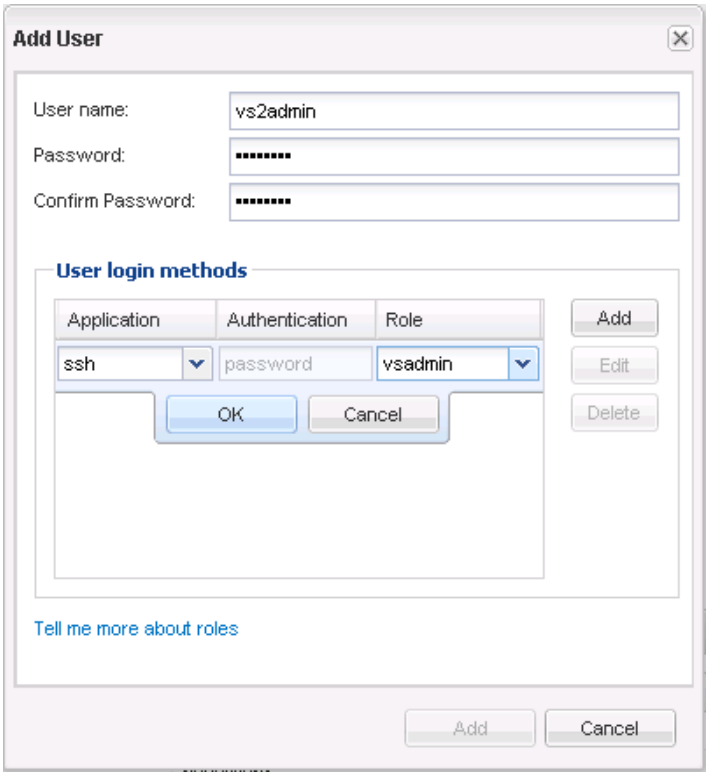

Step 10.

Click **Add** at the bottom of the window to finish creating the user vs2admin.

Step 11.

Add another user with the following properties:

- User name: **vs2intern**
- Application: **ssh**
- Role: **vsadmin-readonly**

Step 12.

Use PuTTY to start an SSH session to the data LIF IP (10.34.x1.40 or 41) address that was management-enabled in the first step of this task, and log in as vs2admin.

Step 13.

Display the status of the cluster.

vs2::> cluster show

The command fails. Why?

Step 14.

Display all available volumes and observe the SVM that are represented in the output.

vs2::> volume show

Step 15.

Modify a volume.

vs2::> volume modify -volume vs2_vol01 -comment "modified by vs2admin"

Step 16.

Verify the change.

vs2::> volume show -volume vs2_vol01 -fields comment

Step 17.

Use PuTTY to start another SSH session to the same data LIF, and log in as **vs2intern**.

Step 18.

Display network ports.

vs2::> network port show

The command fails.

Step 19.

Display network interfaces, examine the LIFs that are displayed, and compare the list to a list of LIFs that is displayed for the cluster admin user.

vs2::> network interface show

Step 20.

Modify a volume.

vs2::> volume modify -volume vs2_vol01 -comment "modified by vs2intern"

The command fails. Why?

Step 21.

Close both of the vs2 administrative PuTTY sessions.

END OF EXERCISE

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B

IČO 357 034 66
IČ DPH SK2020341158

LAB 15: Best practices – optional (if LAB 12 was not done)

TASK 1. SVM root volume protection

To protect the Storage Virtual Machine (SVM) namespace root volume, you can create a load sharing mirror volume on every node in the cluster, including the node in which the root volume is located. Then you create a mirror relationship to each load-sharing mirror volume and initialize the set of load-sharing mirror volumes.

On every node, create a load sharing mirror volume. For protecting the root volume of a Storage Virtual Machine (SVM), you must create a FlexVol volume on every node of the cluster and designate it as a load-sharing mirror destination. A group of load-sharing mirror destination volumes that replicate from the same source volume is referred to as a load-sharing mirror set.

Step 1.

Identify the data aggregates in the nodes of the cluster by using the **aggr show** command.

cluster1::> aggr show -root false

Aggregate Size Available Used% State #Vols Nodes RAID Status

--- --- --- --- --- --- --- ---

n01_aggr1 4.39GB 1.70GB 61% online 4 cluster1-01 raid_dp, normal

n01_fp1 7.03GB 6.27GB 11% online 1 cluster1-01 raid_dp, hybrid, normal

n02_aggr1 7.03GB 3.97GB 44% online 5 cluster1-02 raid_dp, normal

3 entries were displayed.

Step 2.

On each node, create a load-sharing mirror volume by using the volume create command with the -type parameter set to DP. The destination volume that you create must be the same size or greater than the SVM root volume.

cluster1::> volume create -vserver vs2 -volume vs2_m1 -aggregate n01_aggr1 -size 20mb -state online -type DP

cluster1::> volume create -vserver vs2 -volume vs2_m2 -aggregate n02_aggr1 -size 20mb -state online -type DP

Step 3.

Verify that the volume type is set to DP for the load-sharing mirror volumes by using the **volume show** command.

cluster1::> vol show -vserver vs2

Vserver Volume Aggregate State Type Size Available Used%

--- --- --- --- --- --- --- ---

vs2 vs2_m1 n01_aggr1 online DP 20MB 19.89MB 0%

vs2 vs2_m2 n02_aggr1 online DP 20MB 19.89MB 0%

vs2 vs2_root n02_aggr1 online RW 20MB 18.86MB 5%

vs2 vs2_vol01 n01_aggr1 online RW 400MB 379.8MB 5%

4 entries were displayed.

Step 4.

Create a load-sharing mirror relationship between the SVM root volume and each of the load sharing mirrors created on every node of the cluster by using the **snapmirror create** command with the -type parameter set to LS.

cluster1::> snapmirror create -source-path vs2:vs2_root -destination-path vs2:vs2_m1 -type LS -schedule hourly

cluster1::> snapmirror create -source-path vs2:vs2_root -destination-path vs2:vs2_m2 -type LS -schedule hourly

Step 5.

The type attribute of the load-sharing mirror volumes changes from DP to LS. Verify that the load-sharing mirror relationships are created and their mirror state is **Uninitialized** by using the **snapmirror show** command.

cluster1::> snapmirror show -source-volume vs2_root

Source Destination Mirror Relationship Total Last

Path Type Path State Status Progress Healthy Updated

--- --- --- --- --- --- --- ---

cluster1://vs2/vs2_root

LS cluster1://vs2/vs2_m1

Uninitialized Idle - - -

cluster1://vs2/vs2_m2

Uninitialized Idle - - -

2 entries were displayed.

Step 6.

Initialize all the load-sharing mirrors in the set by using the **snapmirror initialize-ls-set** command.

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B

IČO 357 034 66
IČ DPH SK2020341158

cluster1::> snapmirror initialize-ls-set -source-path vs2:vs2_root

[Job 61] Job is queued: snapmirror initialize-ls-set for source "cluster1://vs2/vs2_root".

Step 7.

Verify that the load-sharing mirror relationships are in the Snapmirrored state by using the **snapmirror show** command.

cluster1::> snapmirror show -source-volume vs2_root

Source Destination Mirror Relationship Total Last

Path Type Path State Status Progress Healthy Updated

---

cluster1://vs2/vs2_root

LS cluster1://vs2/vs2_m1

Snapmirrored Idle - true -

cluster1://vs2/vs2_m2

Snapmirrored Idle - true -

2 entries were displayed.

Step 8.

You should update the set of load-sharing mirror volumes to make the changes in the root volume visible to all the clients before the next scheduled update. For example, when a new volume is mounted on the root volume of the SVM, you should update the set of load-sharing mirror volumes. Update all the load-sharing mirror volumes in the set by using the **snapmirror update-lsset** command.

cluster1::> snapmirror update-ls-set -source-path vs2:vs2_root

[Job 62] Job is queued: snapmirror update-ls-set for source "cluster1://vs2/vs2_root".

Restore procedure:

If the Storage Virtual Machine (SVM) root volume becomes unavailable and you have protected it with a set of load-sharing mirrors, you can promote one of the mirrored volumes and then rename it to take the place of the original SVM source volume. For more information read the “SVM Root Volume Protection Express Guide”

END OF EXERCISE.

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B
IČO 357 034 66
IČ DPH SK2020341158

LAB SOLUTIONS

LAB1

TASK1: Put a “P” or “L” beside each item in the following list of Data ONTAP cluster-mode concepts to indicate whether it is physical or logical.

__P__ node

__P__ disk

__P__ aggregate

__L__ virtual server (Vserver)

__L__ cluster

__P__ network port

__L__ flexible volume

__L__ Snapshot copy

__L__ SnapMirror copy

__P__ host bus adapter (HBA)

__L__ LIF

TASK2: Put an “N,” “A,” or “D” (or a combination of those letters) beside each item to indicate whether it has a node Vserver, administration Vserver, or data Vserver scope.

__N__ disk

__D__ namespace

__D__ data LIF

__N__ network port

__A__ cluster management LIF

__D__ flexible volume

__N__ aggregate

__D__ Snapshot copy

__N__ host bus adapter (HBA)

__D__ LUN

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B

IČO 357 034 66
IČ DPH SK2020341158

LAB4

TASK 1: Identify kernel and user-space processes

Step 1.

Using the following list of cluster components, show the path of a network-attached storage (NAS) write request by labeling each component with a number, starting with 1 for the initiator of a write request and proceeding in order to the destination of the request.

_**4**__Cluster session manager (CSM)

_**5**__D-blade

_**2**__Data network port

_**1**__NFS or CIFS client

_**3**__N-blade

_**7**__Disks

_**6**__Nonvolatile RAM (NVRAM)

TASK 2: Explain RDB and quorum concepts

Step 1. List the names of the RDB units.

Answer: Management, VifMgr, VLDB, BCOM

Step 2. Which RDB unit keeps track of the data that is used to operate the clustershell?

Answer: Management

Step 3. Which RDB unit keeps track of the volumes and aggregates and which volumes are on which aggregates?

Answer: VLDB

Step 4. How many healthy nodes are needed in a 16-node cluster to maintain a quorum?

Answer: 9, or 8 + Epsilon

TASK 3: vol0 and Vserver root volumes uses and limitations

Step 1. What is the minimum number of vol0 volumes in a 20-node cluster? What is the maximum?

Answer: 20. 20.

Step 2. What is stored on a vol0 volume?

Answer: Log files, RDB databases.

ALEF Distribution SK, s.r.o.,
Galvaniho Business Centrum IV, Galvaniho 17/C, 821 04 Bratislava, Slovakia
T +421 (2) 4920 3888, F +421 (2) 4920 3899, E-mail sk-reception@alef.com
Obchodný register Okresného súdu Bratislava 1, Oddiel: Sro, vložka číslo 12187/B

IČO 357 034 66
IČ DPH SK2020341158

Step 3. For the following characteristics, write **vol0**, **vsroot**, or **both** to match the volumes with their characteristics.

a) Is a flexible volume __**both**___

b) Can have Snapshot copies __ **both** ___

c) Can have mirrors __**vsroot**___

d) Can be accessed by NFS or CIFS clients __ **vsroot** ___

e) Can be backed up to tape __ **vsroot** ___

f) Can be copied __ **vsroot** ___

g) Can be accessed by the systemshell __**vol0**___

h) Can have junctions to other volumes __ **vsroot** ___

i) Cannot be moved to another aggregate __**vol0**___