

NetApp

NS0-155 Exam

NetApp Certified 7-Mode Data Administrator

Questions & Answers Demo

Version: 8.2

Question: 1

How can you “throttle” SnapValue updates and baseline transfers so that the primary or secondary is not transmitting data as it can?

- A. Use the -k option in the snapvault start or snapshot modify commands.
- B. SnapVault does not support throttling of network throughout.
- C. Use the snapvault throttle command.
- D. Use the -k option in the snapvault initialize command.

Answer: A

Explanation:

http://backdrift.org/man/netapp/man1/na_snapvault.1.html

The snapvault subcommands are:

start [-r] [-k n] [-t n] [-w] [-p {inet | inet6 unspec}] [-o options]

[-S [primary_filer:]primary_path] secondary_qtree

modify [-k n] [-t n] [-p {inet | inet6 | unspec}] [-o options] [-S

primary_filer:primary_path] sec_ondary_qtree

The -k option sets the maximum speed at which data is transferred in kilobytes per second. It is used to throttle disk, CPU, and network usage. If this option is not set, the filer transmits data as fast as it can. The setting applies to the initial transfer as well as subsequent update transfers from the primary.

Question: 2

Node 1 in a clustered pair detects that it has lost connectivity to one of its disk shelves. Node 1 is still up, but it cannot see one of its disk shelves. However, the partner node, Node 2, can see all of the Node 1’s disk shelves. Which feature will cause Node 2 to monitor this error condition for a period of three minutes by default, and then forcibly take over Node 1 if the error condition persists?

- A. Auto enable of giveback
- B. Negotiated Fail Over
- C. Takeover on panic
- D. Cf.quickloop.enable

Answer: B

Explanation:

<http://www.wafl.co.uk/cf/>

<http://wafl.co.uk/options/>

Negotiated failover is a general facility which supports negotiated failover on the basis of decisions made by various modules.

cf.takeover.on_disk_shelf_mismatch

This option allows negotiated takeover to be enabled when the cluster nodes detect a mismatch in disk shelf count. By default, this option is set to off.

This option is available only when cluster is licensed and changing the value on one filer automatically changes the value on the partner filer.

Question: 3

In Data ONTAP, the root user is exempt from those two quotas: _____. (Choose two)

- A. User quotas
- B. Tree quotas
- C. Root quotas
- D. Group quotas
- E. File quotas

Answer: A, D

Explanation:

http://backdrift.org/man/netapp/man1/na_quota.1.html

User and group quotas do not apply to the root user or to the Windows Administrator account; tree quotas, however, do apply even to root and the Windows Administrator account.

Question: 4

Which two Volume SnapMirror (VSM) relationships are supported? (Choose two)

- A. Data ONTAP 8.0.2 64-bit --> Data ONTAP 8.1 64-bit
- B. Data ONTAP 8.0.2 32-bit --> Data ONTAP 8.0.2 64-bit
- C. Data ONTAP 7.3.2 32-bit --> Data ONTAP 8.1 64-bit
- D. Data ONTAP 7.3.2 32-bit --> Data ONTAP 8.0.2 64-bit

Answer: A, C

Volume SnapMirror Interoperability Matrix.

		Destination Volume					
Data ONTAP		7.3.x	8.0.x		8.1		
Source Volume	Aggregate Type	32-bit	32-bit	64-bit	32-bit	64-bit	
	7.3.x	32-bit	Yes	Yes ¹	No	Yes ¹	Yes ¹
	8.0.x	32-bit	No	Yes	No	Yes ¹	Yes ¹
		64-bit	No	No	Yes	Yes ¹²	Yes ¹
	8.1	32-bit	No	No	No	Yes	Yes
		64-bit	No	No	No	Yes ²	Yes

Reference: SnapMirror Async Overview and Best Practices guide.

Volume SnapMirror Source	Volume SnapMirror Destination	Support?
Data ONTAP 7.2*	Data ONTAP 7.2*, 7.3*, 8.0*	Yes
Data ONTAP 7.3*	Data ONTAP 7.3*, 8.0*	Yes
Data ONTAP 7.3*	Data ONTAP 7.2*	No

Question: 5

An aggregate is composed of twelve 36-Gigabyte disks. A drive fails and only 72-Gigabyte spare disks are available. Data ONTAP will then perform what action.

- A. Chooses a 72-Gigabyte disk and use it as is.
- B. Chooses a 72-Gigabyte disk and right-size it.
- C. Halts after 24 hour of running in degraded mode.
- D. Alerts you that there are no 36-Gigabyte disks and wait for one to be inserted.

Answer: B

Explanation:

<https://library.netapp.com/ecmdocs/ECMP1196986/html/GUID-8A91F6A2-61B8-42C3->

AEAB7D5AA9BBE9CF.html

If a disk fails and no hot spare disk that exactly matches the failed disk is available, Data ONTAP uses the best available spare. Understanding how Data ONTAP chooses an appropriate spare when there is no matching spare enables you to optimize your spare allocation for your environment.

Data ONTAP picks a non-matching hot spare based on the following criteria (and more; in the link above):

If the available hot spares are not the correct size, Data ONTAP uses one that is the next size up, if there is one.

The replacement disk is downsized (aka right sized) to match the size of the disk it is replacing; the extra capacity is not available.

Question: 6

Which statement is true about expanding an aggregate from 32-bit to 64-bit in place?

- A. All aggregates are automatically converted from 32-bit to 64-bit with the Data ONTAP 8.1 upgrade.
- B. The expansion is triggered by an aggr convert command.
- C. The expansion is triggered by adding disks to exceed 16 TB.
- D. The 32-bit aggregates are degraded and must be Volume SnapMirrored to a new 64-bit aggregates with Data ONTAP 8.1 upgrade.

Answer: C

Explanation:

<http://www.ntapeek.com/2011/12/how-in-place-expansion-works.html>

To upgrade an aggregate in-place, the only available method is to add disks to expand the aggregate to >16TB.

Question: 7

What utility on the storage system will allow you to capture network packet information?

- A. Snoop
- B. Netstats
- C. Pktt
- D. Traceroute

Answer: C

Explanation:

The pktt command controls a simple on-filer packet tracing facility. Packets can be captured into a trace buffer then dumped to a file, or the captured data can be logged to a file. The data is stored in "tcpdump" format, and can be directly viewed with tcpdump, ethereal, and perhaps other viewers. The output can also be converted using the editcap(1) program to a variety of other formats, including Sniffer, NetMon, and snoop.

Question: 8

The root admin on the UNIX box receives an "Access Denied" message when he attempts to access a newly mounted qtree. What's the most likely cause of this error?

- A. The qtree is missing from the /etc/hosts file.
- B. NFS is turned off on the storage system.
- C. The qtree is set to ntfs security style.
- D. The qtree has not been exported.

Answer: C

Explanation:

<http://hd.kvsconsulting.us/netappdoc/801docs/html/ontap/smg/GUID-866DD79D-D2A6-4FBA-B0BDA63560FECEFF.html>

<https://library.netapp.com/ecmdocs/ECMP1196891/html/GUID-F6304830-B124-40A5-95D6-C646AA85E9E0.html>

Only hostnames and IPs are referenced in /etc/hosts, so A is incorrect. The question states that the qtree was mounted, and B and D would preclude this so they are incorrect.

You need to set a qtree to mixed or unix security style to support UNIX users.

Question: 9

Which two modes support using SnapMirror over multiple network paths?

- A. Standalone
- B. Partner
- C. Multi
- D. Failover
- E. Give back

Answer: C, D

Explanation:

<https://library.netapp.com/ecmdocs/ECMP1196991/html/GUID-75BD4540-A103-4E2B-8B0A-E6BBC0D3F43F.html>

http://backdrift.org/man/netapp/man5/na_snapmirror.conf.5.html

Two multipath methods are supported:

The multiplexing mode causes snapmirror to use both paths at the same time. If one should fail, it will switch to use the remaining path only and use both again should the failing path be repaired.

Failover mode causes snapmirror to use the first path as the desired path and only use the second path should problems arise with the first path.

Question: 10

Identify the two commands that could be entered on the SnapMirror destination storage system. (Choose two)

- A. Options snapmirror.access on
- B. Options snapmirror.resync on
- C. Snapmirror initialize
- D. Snap mirror resync

Answer: C, D

Explanation:

https://library.netapp.com/ecmdocs/ECMP1196979/html/man1/na_snapmirror.1.html

snapmirror initialize [-S source] [-k kilobytes] [-s src_snap] [-c create_dest_snap] [-w] destination Starts an initial transfer over the network. An initial transfer is required before update transfers can take place.

The initialize command must be issued on the destination node. If the destination is a volume, it must be restricted (see na_vol(1) for information on how to examine and restrict volumes). If the destination is a qtree, it must not already exist (see na_qtree(1) for information on how to list qtrees). If a qtree already exists, it must be renamed or removed (using an NFS or CIFS client), or snapmirror initialize to that qtree will not work.

snapmirror resync [-n] [-f] [-S source] [-k kilobytes] [-s src_snap]

[-c create_dest_snap] [-w] destination Resynchronizes a broken-off destination to its former source, putting the destination in the snapmirrored state and making it ready for update transfers. The resync command must be issued on the destination node.

Question: 11

Which two operations can be performed with the SnapDrive for windows graphical user interface?

- A. Create volume
- B. Create Snapshot copies
- C. Create File
- D. Create Disk

Answer: B, D

Explanation:

https://library.netapp.com/ecm/ecm_download_file/ECMM1278920

What SnapDrive does

SnapDrive software integrates with the Windows Volume Manager so that NetApp filers can serve as virtual storage devices for application data in Windows 2000 Server and Windows Server 2003 environments.

SnapDrive is dependent on the virtual disk service. The virtual disk service must be started on the host prior to installing SnapDrive. SnapDrive manages virtual disks (LUNs) on a NetApp filer, making these virtual disks available as local disks on Windows hosts. This allows Windows hosts to interact with the virtual disks just as if they belonged to a directly attached redundant array of independent disks (RAID).

SnapDrive provides the following additional features:

It enables online storage configuration, virtual disk expansion, and streamlined management.

It integrates NetApp Snapshot technology, which creates point-in-time images of data stored on virtual disks.

It works in conjunction with SnapMirror® software to facilitate disaster recovery from either asynchronously or synchronously mirrored destination volumes.

Question: 12

Which storage system command would display the WWPNs of hosts that have logged into storage system using a Fibre Channel connection?

- A. fcp config
- B. fcp initiator show
- C. fcp show -i
- D. fcp show initiator

Answer: D

Explanation:

https://library.netapp.com/ecmdocs/ECMP1196979/html/man1/na_fcp.1.html

fcp show initiator [-v] [adapter]

If no adapter name is given, information about all initiators connected to all adapters are shown.

The command displays the portname of initiators that are currently logged in with the Fibre Channel target adapters. If the portname is in an initiator group setup through the igroup command, then the group name is also displayed. Similarly, all aliases set with the fcp wwpn-alias command for the portname are displayed as well.

If the -v flag is given, the command displays the Fibre Channel host address and the nodename/portname of the initiators as well.

Question: 13

In a Fiber Channel configuration, the host's HBA is referred to as the _____, and the storage system's HBA to as the _____.

- A. Target, initiator
- B. Primary, secondary
- C. Initiator, target
- D. Secondary, primary

Answer: C

Question: 14

Which action will cause a currently in-sync SnapMirror relationship to fail out of sync?

- A. Running snapmirroe update on the source storage system.
- B. Running snapmirror release on the source storage system.
- C. Modifying the /etc/snapmirror.conf file for the relationship on the source storage system
- D. Modifying the /etc/snapmirror.conf file for the relationship on the destination storage system.

Answer: D

Explanation:

https://library.netapp.com/ecmdocs/ECMP1196979/html/man1/na_snapmirror.1.html

https://library.netapp.com/ecmdocs/ECMP1196979/html/man5/na_snapmirror.conf.5.html

snapmirror update [-S source] [-k kilobytes] [-s src_snap] [-c create_dest_snap] [-w] destination

For asynchronous mirrors, an update is immediately started from the source to the destination to update the mirror with the contents of the source.

For synchronous mirrors, a snapshot is created on the source volume which becomes visible to clients of the destination volume.

The snapmirror update command must be issued on the destination filer.

```
snapmirror release source { filer:volume | filer:qtree }
```

Tell SnapMirror that a certain direct mirror is no longer going to request updates.

If a certain destination is no longer going to request updates, you must tell SnapMirror so that it will no longer retain a snapshot for that destination. This command will remove snapshots that are no longer needed for replication to that destination, and can be used to clean up SnapMirror-created snapshots after snapmirror break is issued on the destination side.

The source argument is the source volume or qtree that the destination is to be released from. The destination argument should be either the destination filer and destination volume name or the destination filer and destination qtree path. You can use a line from the output of the snapmirror destinations command as the set of arguments to this command. The /etc/snapmirror.conf file exists on the node containing the mirror used for SnapMirror.

Question: 15

You are trying to do a single file SnapRestore for a file, but you are receiving an error message that the directory structure no longer exists. Which is the most likely explanation?

- A. Once the directory structure has been deleted, you cannot restore the file using single file SnapRestore. You must now SnapRestore the volume.
- B. Snapshot copies have been created since the original directory structure was deleted.
- C. You must recreate the directory structure before trying to restore the file.
- D. You cannot restore a file to an alternate location.

Answer: C

Explanation:

<https://library.netapp.com/ecmdocs/ECMP1196991/html/GUID-35C9A4BA-02BA-4965-A366-698C1299E29D.html>

You can restore a file only to the volume where it was originally located. The directory structure to which a file is to be restored must be the same as that specified in the path. If this directory structure does not exist, you must create it before restoring the file.