

# QUESTIONS & ANSWERS

Kill your exam at first Attempt



**NetworkAppliance**

# NS0-155

*NetApp Certified 7-Mode Data Administrator*



<https://killexams.com/pass4sure/exam-detail/NS0-155>

**QUESTION:** 182

Which SyncMirror command mirrors an existing unmirrored aggregate where Data ONTAP selects the disk drive?

- A. `aggr mirror aggrY`
- B. `aggr mirror aggrY -m`
- C. `aggr mirror aggrT -d 12@144`
- D. `aggr mirror aggrY -m 12@144`

**Answer:** A

**Explanation:**

[https://library.netapp.com/ecmdocs/ECMP1196890/html/man1/na\\_aggr.1.html](https://library.netapp.com/ecmdocs/ECMP1196890/html/man1/na_aggr.1.html) `aggr mirror aggrname`

[ -f ]

[ -n ]

[ -v victim\_aggrname ] [ -d disk1 [ disk2 ... ] ]

Turns an unmirrored aggregate into a mirrored aggregate by adding a plex to it. The plex is either newlyformed from disks chosen from a spare pool, or, if the -v option is specified, is taken from another existing unmirrored aggregate. Aggregate aggrname must currently be unmirrored. Use `aggr create` to make a new, mirrored aggregate from scratch.

Disks may be specified explicitly using -d in the same way as with the `aggr create` and `aggr add` commands.

The number of disks indicated must match the number present on the existing aggregate. The disks specified are not permitted to span disk pools. This behavior can be overridden with the -f option. The -f option, in combination with -d, can also be used to force using disks that have a rotational speed that does not match that of the majority of existing disks in the aggregate.

If disks are not specified explicitly, then disks are automatically selected to match those in the aggregate's existing plex.

The -v option can be used to join victim\_aggrname back into aggrname to form a mirrored aggregate. The result is a mirrored aggregate named aggrname which is otherwise identical to aggrname before the operation.

Victim\_aggrname is effectively destroyed. Victim\_aggrname must have been previously mirrored with aggrname, and then separated via the `aggr split` command.

Victim\_aggrname must be offline. Combined with the -v option, the -f option can be used to join aggrname and victim\_aggrname without prompting the user.

The -n option can be used to display the command that the system will execute without actually making any changes. This is useful for displaying the automatically selected disks, for example.

**QUESTION:** 183

When troubleshooting a double disk failure with SyncMirror, you decide to remove and

destroy a plex from a mirrored aggregate. What steps remove a plex from a mirrored aggregate?

- A. Take the aggregate offline and use the aggr destroy command
- B. Take the aggregate offline and use the destroy -f command
- C. Ensure the plex is online and use the aggr destroy plexname command.
- D. Ensure the plex is offline and use the aggr destroy plexname command.

**Answer:** D

**Explanation:**

[https://library.netapp.com/ecmdocs/ECMP1196890/html/man1/na\\_aggr.1.html](https://library.netapp.com/ecmdocs/ECMP1196890/html/man1/na_aggr.1.html) aggr destroy { aggrname | plexname } [ -f ]

Destroys the aggregate named aggrname, or the plex named plexname. Note that if the specified aggregate is tied to a traditional volume, then the traditional volume itself is destroyed as well.

If an aggregate is specified, all plexes in the aggregate are destroyed. The named aggregate must also not contain any flexible volumes, regardless of their mount state (online, restricted, or offline). If a plex is specified, the plex is destroyed, leaving an unmirrored aggregate or traditional volume containing the remaining plex.

Before destroying the aggregate, traditional volume or plex, the user is prompted to confirm the operation. The -f flag can be used to destroy an aggregate, traditional volume or plex without prompting the user.

The disks originally in the destroyed object become spare disks. Only offline aggregates, traditional volumes and plexes can be destroyed.

**QUESTION:** 184

When replicating data between 32-bit and 64-bit aggregates, which two are allowed? (Choose two)

- A. NDMPcopy
- B. QSM
- C. vol copy
- D. VSM

**Answer:** A, B

**Explanation:**

<http://www.netapp.com/us/system/pdf-reader.aspx?m=tr-3786.pdf&cc=us>  
<http://waf1.co.uk/ndmpcopy/>

Because volume SnapMirror works at the block level, the source and destination must both be FlexVol volumes in the same aggregate type. The source and destination of a volume SnapMirror relationship must either both be FlexVol volumes in 32-bit aggregates or both be FlexVol volumes in 64-bit aggregates.

Therefore you cannot create a volume SnapMirror relationship from a FlexVol volume in a 32-bit aggregate to a FlexVol volume in a 64-bit aggregate or vice versa.

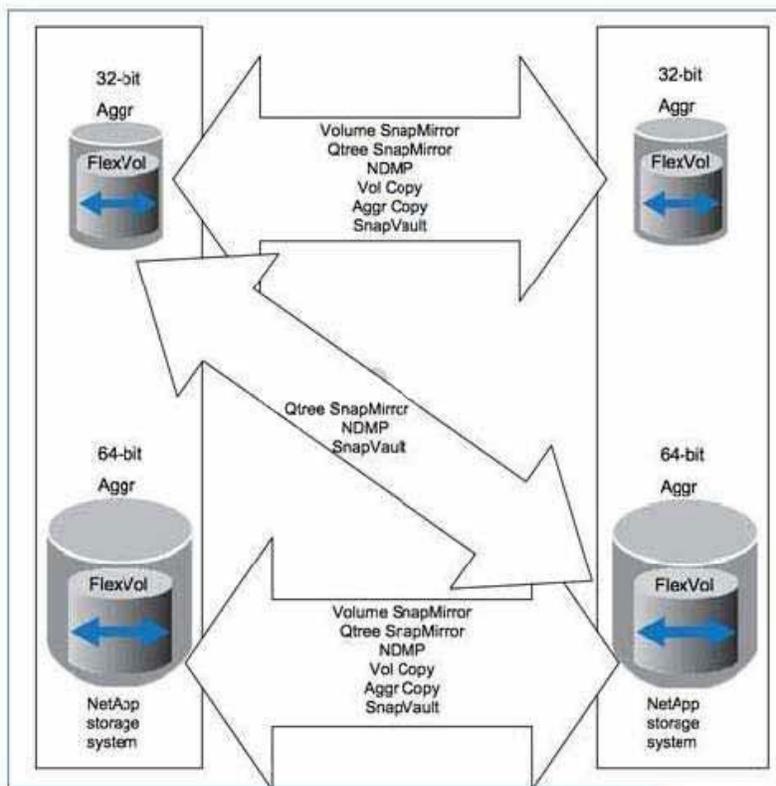


Figure 5) Data replication utilities and their interaction with 32-bit and 64-bit aggregates.

**QUESTION:** 185

What types of migration are supported from 32-bit to 64-bit?

- A. migration of qtrees via QSM
- B. migration of volumes via ndmpcopy
- C. migration of volumes via SnapMirror
- D. Both A & B
- E. Both B & C
- F. Both A & C

**Answer:** D

**Explanation:**

Volume SnapMirror operates at the physical block level. It replicates the contents of an entire volume, including all Snapshot copies, plus all volume attributes verbatim from a source (primary) volume to a target (secondary) volume. Implicitly, the destination volume will be the same as the source volume. Since C is incorrect, A & B must be correct. Qtree snapmirror operates at the logical level, so it is unimportant that the destination system has a different volume configuration.

ndmcopy is similarly operating with data already read from the volume and so the destination volume type is not a consideration.

**QUESTION: 186**

How do you migrate LUNs under volumes from 32-bit to 64-bit?

- A. lrep
- B. ndmcopy
- C. rsync
- D. snapvault

**Answer: B**

**Explanation:**

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

Think of ndmcopy as a sort of ftp for SAN. You can use it to copy files from one volume to another. Because the file is read from one volume and written to another, the constraints of volume types are not an issue.

**QUESTION: 187**

Which three FAS deduplication features exist in ONTAP 8.0 7-Mode?

- A. The deduplication fingerprint and change logs were moved to the root volume /etc/dedup directory.
- B. After breaking the deduplication volume destination mirror, the deduplication (a-sis) process continues uninterrupted.
- C. Qtree SnapMirror is supported on the destination for deduplication by enabling on the source, destination, or both systems.
- D. The fingerprint database and the change logs that the deduplication process uses are located outside the volume, in the aggregate.
- E. The deduplication schedule is not tied to a Qtree SnapMirror update, and can be configured just like the deduplication schedule for any volume.

**Answer:** C, D, E

**Explanation:**

Section: Data Protection

NetApp also recommends that if deduplication is used on the source volume, then it should also be used on the destination volume. However, you are not restricted using deduplication on the source volume only.

The deduplication schedule is not tied to the qtree SnapMirror update. In other words, the deduplication process does not automatically start at the completion of the qtree SnapMirror transfer.

The fingerprint database and the change log files that are used in the deduplication process are located outside of the volume in the aggregate and are therefore not captured in Snapshot copies <http://www.netapp.com/us/media/tr-3505.pdf>

**QUESTION:** 188

To manage SnapVault relationships with an NDMP management application, you must specify a:

- A. directory path
- B. volume and qtree
- C. SnapVault license
- D. username and password

**Answer:** D

**Explanation:**

Add a new backup user to the Backup Operators useradmin group list. Enter the following command:

```
useradmin user add backupuser -g "Backup Operators"
```

Generate an NDMP password for the new user. Enter the following command: ndmp password backupuser

The NDMP password is used to authenticate the NetBackup media server to the storage system. This password is required for setup on the NetBackup management station.

<https://library.netapp.com/ecmdocs/ECMM1278284/html/onlinebk/5snapv37.htm>

**QUESTION:** 189

For OSSV sources, the qtree is the basic unit of SnapVault backups.

- A. True
- B. False

**Answer:** B

**Explanation:**

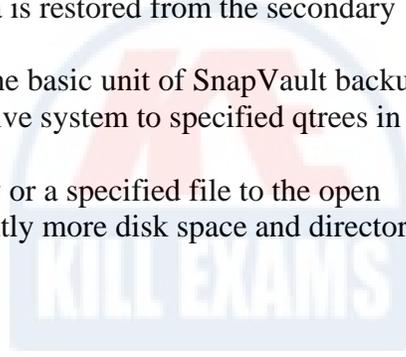
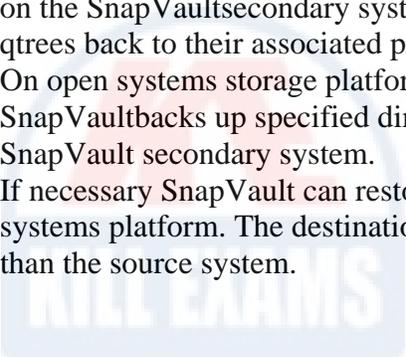
<https://library.netapp.com/ecmdocs/ECMP1196991/html/GUID-06C70D3B-9A06-4D20-B140-EB147B941BDA.html>

The data structures that are backed up and restored through SnapVault depend on the primary system.

On systems running Data ONTAP, the qtree is the basic unit of SnapVault backup and restore. SnapVault backs up specified qtrees on the primary system to associated qtrees on the SnapVaultsecondary system. If necessary, data is restored from the secondary qtrees back to their associated primaryqtrees.

On open systems storage platforms, the directory is the basic unit of SnapVault backup. SnapVaultbacks up specified directories from the native system to specified qtrees in the SnapVault secondary system.

If necessary SnapVault can restore an entire directory or a specified file to the open systems platform. The destination system uses a slightly more disk space and directories than the source system.



For More exams visit <https://killexams.com/vendors-exam-list>



*Kill your exam at First Attempt....Guaranteed!*