

DUMPSBOSS.

Data ONTAP 7-Mode Administrator

Network Appliance NS0-155

Version Demo

Total Demo Questions: 10

Total Premium Questions: 189

Buy Premium PDF

<https://dumpsboss.co>

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, Volume A	89
Topic 2, Volume B	100
Total	189

QUESTION NO: 1

What command would create a LUN clone named lun0.clone from the LUN lun0 on the volume vol2?

- A. lun clone start create /vol/vol2/lun0.clone -b /vol/vol2/lun0 snap_recent
- B. lun clone create /vol/vol2/lun0.clone -b /vol/vol2/lun0 snap_recent
- C. lun clone create /vol/vol2/lun0 -b /vol/vol2/lun0.clone snap_recent
- D. lun clone create /vol/vol2/lun0.clone -b /vol/vol2/lun0 -s snap_recent

ANSWER: B

Explanation:

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

```
lun clone create clone_lunpath [ -o noreserve ] -b parent_lunpath parent_snap
```

QUESTION NO: 2

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

- A. file quotas
- B. group quotas
- C. root quotas
- D. tree quotas
- E. user quotas

ANSWER: B E

QUESTION NO: 3

Which three statements apply to the SnapMirror visibility_interval argument? (Choose three.)

- A. The default interval value is five minutes.
- B. The smallest interval value supported is 30 seconds.
- C. This variable controls the view of the data on the destination system.
- D. By setting a value, you specify the amount of time before an automatic snapshot is created on the source volume that is synchronously mirrored.

ANSWER: B C D

Explanation:

<http://www.netapp.com/us/media/tr-3326.pdf>

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

<http://hd.kvsconsulting.us/netappdoc/733docs/html/ontap/onlinebk/GUID-47724B21-7C7E-4AC9-A497-9505B42C27AF.html>

With SnapMirror Sync and SnapMirror Semi-Sync, changes to the source volume do not show up immediately on the destination volume, even though the changes have been replicated. The changes are first shown after the source system creates an automatic Snapshot copy of the source volume.

Snapshot copies created on the source volume are immediately propagated to the destination volume. The automatic Snapshot copies are created every three minutes by default.

To change the interval for automatic Snapshot copies, change the `visibility_interval` in the

`snapmirror.conf` file; however, performance can degrade if set to a smaller interval because frequent

Snapshot copies cause additional processing such as cache consistency checking on the destination system.

There is also a disadvantage in setting this to a large value. When the connection between the systems has an outage and SnapMirror goes into asynchronous mode, SnapMirror uses the last common Snapshot copy to transition into synchronous mode. This means that all data from the last valid common Snapshot copy will need to be replicated from the source to the destination storage system. If the `visibility_interval` is set to a large value, a large amount of data might have to be transferred and it might take longer time for SnapMirror to transition into synchronous mode. For these reasons, keep the default value of three minutes.

`visibility_interval`

The value for this argument is a number optionally followed by the suffixes: s (seconds), m (minutes) or h (hours). If a suffix is not specified, value is interpreted as seconds. This argument controls the amount of time before an automatic snapshot is created on the source volume that is synchronously mirrored. The value is the number of seconds between automatically created snapshots. The default value is 3 minutes. A small number here can negatively affect the performance of the mirror. This argument is ignored for asynchronous mirrors.

`visibility_interval={xs | xm | xh}`

Determines the amount of time before an automatic Snapshot copy is created on the source volume that is being replicated using synchronous or semi-synchronous SnapMirror. When replicating synchronously or semisynchronously, using SnapMirror, changes to the source volume do not show immediately on the destination volume, even though the changes have been replicated. The changes are shown only after the source system takes an automatic Snapshot copy of the source volume. This event occurs every three minutes by default.

You can change the interval for automatic Snapshot copies, but performance can degrade if you set smaller intervals because Snapshot copies are taken more often. The smallest interval you can set is 30 seconds.

Note: The s, m, and h suffixes specify seconds, minutes, and hours, respectively.

QUESTION NO: 4

To collect per client NFS statistics on a storage system, the option `nfs.per_client_stats.enable` must be set on.

A. True

B. False

ANSWER: A

Explanation:

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

nfs.per_client_stats.enable

Enables/disables the collection and display of perclient NFS statistics, as described in na_nfsstat(1). Valid values for this option are on or off. The default value for this option is off.

QUESTION NO: 5

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_miscompare

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 NO: 6

In an active/active configuration in takeover mode, the surviving controller interfaces will reflect the identity of the _____.

- A. local interfaces only

- B. failed partner interfaces only
- C. local and failed partner interfaces
- D. interfaces as defined in the /etc/rc file

ANSWER: D

Explanation:

In takeover mode, this network interface assumes the identity of the network interface on the partner

<http://www.wafl.co.uk/tag/ifconfig/#sthash.3rsjaDZD.dpuf>

For the partner interface configuration to be persistent across reboots, you must include the ifconfig command in the /etc/rc file

<http://hd.kvsconsulting.us/netappdoc/801docs/html/ontap/nag/GUID-525FCF4E-109D-4E6F-9256-C6E56438CA9A.html>

QUESTION NO: 7

Which two are Data ONTAP commands for displaying performance information? (Choose two.)

- A. nfsstat
- B. perfstat
- C. sio_ontap_soll
- D. sysstat

ANSWER: A D

Explanation:

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

https://library.netapp.com/ecmdocs/ECMP1196890/html/man1/na_nfsstat.1.html

nfsstat and sysstat are OnTap commands used for performance information.

sio_ontap_soll is not an OnTap command.

Perfstat is a GUI performance tool, not an OnTap command. Think of PerfMon for Windows to remind you that PerfStat is a Windows tool. Note how many of the ONTAP performance commands end in stat.

QUESTION NO: 8

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

- A. snapmirror resync
- B. snapmirror release

- C. snapmirror initialize
- D. options snapmirror.access on

ANSWER: A C

QUESTION NO: 9

A snapshot is a frozen, read-only image of the entire Data ONTAP active file system that reflects the state of the _____ at the time the snapshot was created.

- A. qtrees
- B. volume
- C. directory
- D. storage system

ANSWER: B

Explanation:

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

Snapshots can be made of volumes or aggregates; most commonly the former. The question refers to a file system, or more specifically a volume.

You cannot snapshot directories or qtrees alone.

QUESTION NO: 10

Which two statements apply to an entry in the /etc/exports file regarding exporting directories/qtrees and volumes? (Choose two.)

- A. Each volume must be exported separately.
- B. Exports are accessible only to hosts with names.
- C. User names used in an export must be resolvable.
- D. Exports can access multiple paths on the same volume.

ANSWER: A D

Explanation:

https://library.netapp.com/ecmdocs/ECMP1196890/html/man5/na_exports.5.html

IP addresses and wildcards can be used when specifying host targets in exports.

User names are not valid exports options, save for the anon option which utilizes a local system UID or user name (i.e. no user name mapping or resolution is performed.)

The highest level in the filesystem that can be specified for an export is the volume level. Implicitly this means that each volume must be exported separately.

An export provides access to paths below the export, implying that multiple paths on the same volume can be exported from an export.