## [2018-Feb-NewValid Braindump2go NS0-155 PDF and VCE 195Q Offer[78-88

2018 Feb New NetApp NS0-155 Exam Dumps with PDF and VCE Free Updated Today! Following are some new NS0-155 Real
Exam Questions: 1.  2018 Latest NS0-155 Exam Dumps (PDF & VCE) 195Q&As Download:
https://www.braindump2go.com/ns0-155.html2. 2018 Latest NS0-155 Exam Questions & Answers Download:
https://drive.google.com/drive/folders/1MxkLZN5KA1zypVTo9BJ6WLvLiN2Rnad8?usp=sharingQUESTION 78When using
SnapRestore to restore a single LUN, the LUN must beA. exported and accessibleB. mapped to a Windows server
C. mounted on a hostD. taken offline or unmappedAnswer: DExplanation:
https://library.netapp.com/ecmdocs/ECMP1196995/html/GUID-577EDC41-FEA0-4987-8F04-55E736962FD1.htmlAlways
unmount the LUN before you run the snap restore command on a volume containing the LUN or before you run a single file
SnapRestore of the LUN. For a single file SnapRestore, you must also take the LUN offline.QUESTION 79For aggregates or
traditional volume SyncMirror to work correctly, the configuration of each plex should be identical with regard toA.
NVRAMB. Mirrored Snapshot copiesC. RAID groups and disk sizesD. NDMP and Snapshot copiesAnswer: CExplanation:
SyncMirror is essentially a RAID-1 mirror; hence the name. The best practice is to mirror storage of equal performance; i.e.
identical RAID groups and disk sizes (number of spindles, IOPS). If you were to mirror two storage devices of different
performance, the slower device would be the limiting factor.QUESTION 80What happens when two aggregates are rejoined?A. It
is not possible to rejoin aggregates.B. Data ONTAP will always retain data from the newest aggregate.C. Both aggregates revert
to the last common snapshot before the split.D. You will have a large new aggregate containing all volumes from both aggregates.
E. One of the aggregates will be overwritten and you will lose any data changed on it after the split. Answer: EExplanation: When
you rejoin split aggregates, Data ONTAP mirrors the data from one aggregate to the other and destroys data that existed on that
aggregate before the rejoin.https://library.netapp.com/ecmdocs/ECMP1196991/html/GUID-B9CB3DA8-E1AF-4BAC-A85B-
64C009010F4A.htmlQUESTION 81You cannot remove the SyncMirror license if one or more mirrored volumes exist.A. TrueB.
FalseAnswer: AExplanation: You cannot disable the SyncMirror license if mirrored aggregates exist and are online.
http://hd.kvsconsulting.us/netappdoc/801docs/html/ontap/onlinebk/GUID-5058FDA7-68AE-4693-
B343-0E9FFEDFC86C.htmlQUESTION 82A mirrored volume has a failed disk, but there are no available spares in the pool from
which the degraded plex was built. Data ONTAP willA. halt all operations to volumes in that poolB. choose spare
disks to reconstruct from the opposite poolC. warn you that there are no available spares from the proper poolD. "halt" after 24
hours if spare disks are not added to the proper poolAnswer: CExplanation:Both `warn you that there are no available spares from
the proper pool' and `halt after 24 hours if spare disk are not added to the proper pool' are correct answers. However, since Data
ONTAP alerts you first, 'warn you that there are no available spares' is the best answer.
https://library.netapp.com/ecm/ecm_download_file/ECMP1196912QUESTION 83The required hardware to add SyncMirror to an
existing system would beA. disk shelves and cluster interconnectB. disk shelves and termination switchesC.
disk shelves, Ethernet adapters, and cablingD. disk shelves, Fibre Channel adapters, and cablingAnswer: DExplanation:
QUESTION~84 Which~two~are~true~when~creating~a~SyncMirror~volume?~(Choose~two.) A.~~The~disks~selected~for~Plex~0~and~Plex~1~~Inches the contraction of the contrac
must come from the same disk shelf.B. The disks selected for Plex 0 and Plex 1 must come from the same disk pools.C. The
number of disks added must be an even number and equally divided between the two plexes.D. The disks are selected first on the
basis of equivalent bytes per sector (bps) size, then on the basis of the size of the disk. If there is no equivalent-sized disk, Data
ONTAP takes a larger capacity disk and right-sizes it. Answer: CDExplanation: The rules for the selection of disks or array LUNs,
for using as mirrored aggregates, are as follows:Disks or array LUNs selected for each plex must be in different pools. The same
number of disks or array LUNs must be in both the plexes. Disks are selected first on the basis of equivalent bytes per sector (bps)
size, then on the basis of the size of the disk. If there is no equivalent-sized disk, Data ONTAP uses a larger-capacity disk, and limits
the size to make it identically sized.Data ONTAP names the plexes of the mirrored aggregate.
https://library.netapp.com/ecmdocs/ECMP1196991/html/GUID-F5F241AF-50C0-4FF3-8C01-96958581C4D9.htmlQUESTION 85
After issuing the aggr split aggr0/plex0 aggrNew commandA. there is a new plex called aggrNewB. there are
two new plexes, plex0 and plexNewC. there is a new mirrored aggregate called aggrNewD. there are two unmirrored aggregates,
aggr0 and aggrNewE. flexible volumes are evenly split between aggr0 and aggrNew Answer: DExplanation:
https://library.netapp.com/ecmdocs/ECMP1196890/html/man1/na aggr.1.html aggr split aggrname/plexname new_aggrname
Removes plexname from a mirrored aggregate and creates a new unmirrored aggregate named new_aggrname that contains the plex.
The original mirrored aggregate becomes unmirrored. The plex to be split from the original aggregate must be functional (not
partial), but it could be inactive, resyncing, or out-ofdate. Aggr split can therefore be used to gain access to a plex that is not up to

date with respect to its partner plex, if its partner plex is currently failed. If the aggregate in which plexname resides is embedded in a traditional volume, aggr split behaves identically to vol split. The new aggregate is embedded in a new traditional volume of the same name. The aggregate in which plexname resides must contain exactly one flexible volume, aggr split will by default rename the flexible volume image in the split-off plex to be the same as the new aggregate. If the original aggregate is restricted at the time of the split, the resulting aggregate will also be restricted. If the restricted aggregate is hosting flexible volumes, they are not renamed at the time of the split. Flexible volumes will be renamed later, when the name conflict is detected while bringing an aggregate online. Flexible volumes in the aggregate that is brought online first keep their names. That aggregate can be either the original aggregate, or the aggregate resulting from the split. When the other aggregate is brought online later, flexible volumes in that aggregate will be renamed. If the plex of an aggregate embedded within a traditional volume is offline at the time of the split, the resulting aggregate will be offline. When splitting a plex from an aggregate that hosts flexible volumes, if that plex is offline, but the aggregate is online, the resulting aggregate will come online, and its flexible volumes will be renamed. It is not allowed to split a plex from an offline aggregate. A split mirror can be joined back together via the -v option to aggr mirror.QUESTION 86Normally, before you split a SyncMirror volume, you ensure that both plexes are \_\_\_\_\_ and \_\_\_\_.A. offline, quiescedB. online, consistentC. online, normalD. restricted, offlineAnswer: CExplanation: Ensure that both plexes of the mirrored aggregate you are splitting are online and operational. https://library.netapp.com/ecmdocs/ECMP1196991/html/GUID-DE08C404-6E9C-47F6-A72E- CDEE2043408C.htmlQUESTION 87You can avoid data corruption when you issue the cf forcetakeover -d command if the remote node is \_\_\_\_\_\_.A. in a giveback modeB. fenced off manuallyC. powered on and accessibleD. powered off and inaccessibleAnswer: DQUESTION 88 You have connected all cables and disk shelves in a fabric-attached MetroCluster. Which command would verify that disks are connected and have dual paths? A. disk showB. sysconfig -pC. vol status -rD. storage show disk -pAnswer: DExplanation: The operator must ensure that the partner filer does not become operational at any time while a filer is in a takeover mode initiated by the use of this command http://backdrift.org/man/netapp/man1/na cf.1.html!!!RECOMMEND!!!1.|2018 Latest NS0-155 Exam Dumps (PDF & VCE) 195Q&As Download:https://www.braindump2go.com/ns0-155.html2.|2018 Latest NS0-155 Study Guide Video: YouTube Video: YouTube.com/watch?v=G74lQEevDwI