[2018-Feb-NewFree NetApp 195Q NS0-155 Dumps PDF Braindump2go Offers[100-110

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 100After upgrading to DOT 8.07-mode, you can convert existing 32-bit aggregates into 64-bit aggregates.A. TrueB. FalseAnswer: BExplanation: http://www.ntapgeek.com/2011/12/how-in-place-expansion-works.htmlTo upgrade an aggregate in-place, the only available method is to add disks to expand the aggregate to >16TB. There is an obscure method using diag mode but that is not supported or mentioned in the exam content. The answer is always add disks to grow past 16TB, or migrate.QUESTION 101With 64-bit aggregates, the number of FlexVols that can be created on a storage controller is _____.A. 500B. 1000C. 6400D. 64000Answer: A Explanation:http://www.netapp.com/us/products/storage-systems/fas3200/fas3200-tech-specs.aspx There is no easily remembered reason why its 500. Some people like to remember that there is one other question on the exam that requires you to remember an arbitrary number and the answer is also 500 (distance of a fiber interconnect at 2Gbps).QUESTION 102The largest volume that can be created on a 64-bit aggregate is ____.A. 10 TBB. 16 TBC. 64 TBD. Equal to size of the aggregateAnswer: DExplanation: http://www.netapp.com/us/system/pdf-reader.aspx?m=tr-3786.pdf&cc=usTrick question.... Depends on the array and the maximum shelves and disks and aggregate size, which can change over time. But a volume can never be greater than the aggregate - save for a new feature in 8.2 called Infinite Volume that is not covered by this exam.QUESTION 103Which command will show the FCP target(s) on a storage system? A. fcp show adapter B. fcp show all C. fcp show -tD. fcp status Answer: AExplanation: https://library.netapp.com/ecmdocs/ECMP1196979/html/man1/na fcp.1.htmlfcp show adapter [-v] [adapter]If no adapter name is given, information about all adapters are shown. This command displays information such as nodename/portname and link state about the adapter. If the -v flag is given, this command displays additional information about the adapters.QUESTION 104You are having problems accessing three recently-created LUNs. Given the output below, the reason is because all of the LUNs are using the same LUN ID of 0.lun show -mLUN pathMapped to LUNID-----/vol/vol1/lun0 solaris-igroup 0/vol/vol1/lun1 windows-igroup 0/vol/vol1/qtree1/lun2 aix-igroup 0/vol/vol1/qtree1/lun3 linux-igroup 0A. TrueB. FalseAnswer: BExplanation: The 'lun' in the path is irrelevant. The 0 at the end of each statement denotes the LUN ID.LUN IDs must be unique per igroup.QUESTION 105Using the output below, a co-worker determined that these are type "Solaris" LUNs. What would your determination be? A. These are type "Solaris" LUNs.B. Not enough information is given here. Run the lun map command to get the information requested.C. Not enough information is given here. Run the lun show -v command to get the information requested.D. Not enough information is given here. Run the lun status command to get the information requested. Answer: CExplanation: https://library.netapp.com/ecmdocs/ECMP1196979/html/man1/na lun.1.htmllun show [-v | -m | -c [all | mapped | offline | online | unmapped | staging | -g initiator_group | -n node | -l vol_name | lun_path] Displays the status (lun_path, size, online/offline state, shared state) of the given LUN or class of LUNs. With the -v option supplied, additional information (comment string, serial number, LUN mapping, HA Pair Shared Volume Information) is also displayed. With the -m option supplied, information about lun_path to initiator_group mappings is displayed in a table format. With the -c option supplied, information about LUN cloning status is displayed. A specific LUN can be indicated by supplying its lun_path. When an initiator_group is specified, status is reported for all LUNs that are mapped to the initiator group. When a node is specified, status is reported for all LUNs that are mapped to initiator groups which contain that node. When staging is specified, information about the temporary LUNs preserved in the staging area is reported. When vol name is specified, status is reported for all the LUNs in that volume. Mapped LUNs are ones with at least one map definition. A LUN is online if it has not been explicitly made offline using the lun offline command. QUESTION 106When 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 commandB. Take the aggregate offline and use the destroy -f commandC. 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.htmlaggr 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 107When replicating data between 32-bit and 64-bit aggregates, which two are allowed? (Choose two)A. NDMPcopyB. QSMC. vol copyD. VSMAnswer: ABExplanation: http://www.netapp.com/us/system/pdf-reader.aspx?m=tr-3786.pdf&cc=ushttp://wafl.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. QUESTION 108What types of migration are supported from 32-bit to 64-bit?A. migration of qtrees via QSMB. migration of volumes via ndmpcopyC. migration of volumes via SnapMirrorD. Both A & BE. Both B & CF. Both A & CAnswer: DExplanation: 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.ndmpcopy is similarly operating with data already read from the volume and so the destination volume type is not a consideration.QUESTION 109How do you migrate LUNs under volumes from 32-bit to 64-bit?A. lrepB. ndmpcopyC. rsyncD. snapvaultAnswer: BExplanation: http://wafl.co.uk/ndmpcopy/Think of ndmpcopy 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 110Which 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 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: CDEExplanation: Section: Data ProtectionNetApp 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 gtree 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!!!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