

[2017-NewValid Braindump2go NS0-158 Exam PDF Dump 121q Offer[23-33

2017 May New NS0-158 Exam Dumps with VCE and PDF Updated in www.Braindump2go.com Today!100% Real Exam Questions! 100% Exam Pass Guaranteed! 1.|2017 New Version NS0-158 Exam Dumps (PDF & VCE) 121Q&As Download: <http://www.braindump2go.com/ns0-158.html> 2.|2017 New Version NS0-158 Exam Questions & Answers Download: <https://drive.google.com/drive/folders/0B75b5xYLjSSNUWNTQ3dmbXc1d2s?usp=sharing> QUESTION 23Which statement is true about CFO and SFO in clustered ONTAP? A. CFO aggregates are given back first, then SFO aggregates are given back one at a time.B. SFO aggregates are given back first, then CFO aggregates are given back one at a time.C. All aggregates are given at the same time.D. None, clustered ONTAP supports SFO aggregates only. Answer: A QUESTION 24When creating a LIF for NFS or CIFS, what is the default failover policy? A. PriorityB. NextavailC. Round robinD. Disabled Answer: B QUESTION 25 What is considered a node in a clustered ONTAP? A. An HA storage controller pairB. A single storage controllerC. A data VserverD. A data port Answer: B QUESTION 26A block from a Flash Pool read operation is initially marked as _____. A. hotB. warmC. neutralD. coldE. evict Answer: C QUESTION 27What is the maximum number of nodes supported for a cluster serving only NAS protocols? A. 1B. 2C. 4D. 8E. 16F. 24 Answer: F QUESTION 28A customer has an aggregate, aggr1_01 with three RAID groups: rg0, rg1, and rg2. Each RAID group has 15 disks, and the aggregate's maximum RAID group size is set to 18 disks. The customer adds three additional disks to the aggregate using the following command:storage aggregate add-disks -aggregate aggr1_01 -diskcount 3How are the RAID groups affected by this command? A. The disks are added to a new RAID group rg3.B. Each RAID group receives one additional disk.C. The three disks are added to RAID group rg2.D. The three disks are added to RAID group rg0. Answer: CExplanation: Tested with a 4 node clustered Data ONTAP 8.3 cluster Cluster::*> storage aggregate add-disks -aggregate aggr2 -diskcount 3Cluster::*> run -node node03 sysconfig -Vvolume aggr2 (3 RAID groups): group 2: 18 disks group 1: 15 disks group 0: 15 disks QUESTION 29A customer has a 1000 GB volume that is at 25% capacity and has no existing Snapshot copies. The customer creates a Snapshot copy of the volume, creates a new 250 GB file, and deletes an existing 100 GB file. How much space is free in the volume? A. 350 GB.B. 400 GB.C. 500 GB.D. 600 GB Answer: C QUESTION 30A user is reporting trouble using a SVM management LIF. You have verified host configuration, network connectivity, and that the LIF responds to pings. What should you do to continue troubleshooting this problem? A. Ensure the SVM management LIF's failover policy is correct.B. Ensure the SVM management LIF's administrative status is up.C. Ensure the SVM management LIF is on its home port.D. Ensure the SVM management LIF's firewall policy is set to management. Answer: D QUESTION 31You notice your indirect I/O performance between nodes in your 4-node cluster is slowing down from prior week. Each node has two 10 GB connections to the cluster interconnect, one to each switch. You are using four out of 12 available ports on each cluster interconnect switch. Your cluster network utilization is over 80%. What should you do to increase cluster interconnect performance? A. Change the SFP+ from 10 Gb to 16 Gb on each node and switch.B. Add two more cluster interconnect connections per node.C. Add additional ISLs between the two cluster interconnect switches.D. Change the cable from TwinAx to optical with SFP+ for better throughput. Answer: B QUESTION 32A customer wants to use FlashCache but does not want to lose the information stored in the FlashCache card when a takeover occurs in a high-availability solution. What should the customer do to make sure this feature is enabled? A. Confirm that the aggregates using the FlashCache card are configured appropriately.B. Set the option to rewarm the FlashCache in both nodes of the HA pair.C. Verify that the FlashCache card is located in the appropriate slot in the NetApp controller.D. Verify that the FlashCache card has been configured only on the node where the feature is required. Answer: BExplanation: The system does not serve data from a Flash Cache or Flash Cache 2 module when a node is shutdown. However, the WAFL external cache preserves the cache during a graceful shutdown and can serve "warm" data after giveback. The WAFL external cache can preserve the cache in Flash Cache modules during a graceful shutdown. It preserves the cache through a process called "cache rewarming," which helps to maintain system performance after a graceful shutdown. For example, you might shut down a system to add hardware or upgrade software. Cache rewarming is enabled by default if you have a Flash Cache or Flash Cache 2 module installed. Cache rewarming is available when both nodes in an HA pair are running Data ONTAP 8.1 or later.

<https://library.netapp.com/ecmdocs/ECMP1196798/html/GUID-AEB76252-545D-4886-ABE9-B91452FDD3AF.html> QUESTION 33A customer has configured SnapDrive for Windows to assist in creating Snapshot copies of a LUN. In SnapDrive, the transport protocol configuration uses the vsadminuser account for authentication. SnapDrive fails to connect to the SVM. Why did this failure occur? A. Transport protocol requires a domain Windows account.B. The vsadmin account is locked or does not have a password.C. SnapDrive requires a cluster-wide privileged account.D. Transport protocol requires a local Windows account. Answer: C !!!RECOMMEND!!! 1.|2017 New Version NS0-158 Exam Dumps (PDF & VCE) 121Q&As Download:

<http://www.braindump2go.com/ns0-158.html> 2.|2017 New Version NS0-158 Study Guide Video: YouTube Video:

[YouTube.com/watch?v=5Dhu18qCw4Q](https://www.youtube.com/watch?v=5Dhu18qCw4Q)