

## [2017-Jul.-New2V0-621D VCE Dumps from Braindump2go Guarantee 100% Pass 2V0-621D VMware Exam[Q61-Q70]

2017/July VMware 2V0-621D Exam Dumps with PDF and VCE New Updated in [www.Braindump2go.com](http://www.Braindump2go.com) Today!100% 2V0-621D Real Exam Questions! 100% 2V0-621D Exam Pass Guaranteed! 1.|2017 New 2V0-621D Exam Dumps (PDF & VCE) 256Q&As Download:<https://www.braindump2go.com/2v0-621d.html> 2.|2017 New 2V0-621D Exam Questions & Answers Download:<https://drive.google.com/drive/folders/0B75b5xYLjSSNN0M4cTJyOXZySk0?usp=sharing> QUESTION 61An administrator created a six node Virtual SAN cluster, created a fault domain, and moved three of the six nodes into that domain.A node that is a member of the fault domain fails.What is the expected result? A. The remaining two fault domain members are treated as failed.B. The remaining two fault domain members stay protected by the domain.C. One of the non-member nodes will be automatically added to the fault domain.D. VMware High Availability will restart virtual machines on remaining nodes in the domain. Answer: AExplanation:When the member of the fault domain fails, the remaining two fault domain members are treated as failed too. QUESTION 62Where is a Virtual SAN Fault Domain configured? A. VMware Virtual SAN Cluster configurationB. VMware High Availability Cluster configurationC. Distributed Resource Scheduler configurationD. Datacenter Advanced Settings configuration Answer: AExplanation:If your Virtual SAN cluster spans across multiple racks or blade server chassis in a data center and you want to make sure that your hosts are protected against rack or chassis failure, you can create fault domains and add one or more hosts to it.A fault domain consists of one or more Virtual SAN hosts grouped together according to their physical location in the data center. When configured, fault domains enable Virtual SAN to tolerate failures of entire physical rack as well as failures of a single host, capacity device, network link or a network switch dedicated to fault domains. Fault domains cannot be configured for stretched or metro clusters.The number of failures your cluster can tolerate depends on the number of failures a virtual machine is provisioned to tolerate. For example, when a virtual machine is configured with Number of failures to tolerate=1 and using multiple fault domains, Virtual SAN can tolerate a single failure of any kind and of any component in a fault domain, including the failure of an entire rack.When you configure fault domains on a rack and provision a new virtual machine, Virtual SAN ensures that protection objects, such as replicas and witnesses are placed on different fault domains. If, for example, a virtual machine's storage policy is Number of failures to tolerate=n, Virtual SAN requires a minimum of 2\*n+1 fault domains in the cluster. When virtual machines are provisioned in a cluster with fault domains using this policy, the copies of the associated virtual machine objects are stored across separate racks.Reference:

<http://pubs.vmware.com/vsphere-60/index.jsp?topic=%2Fcom.vmware.vsphere.virtualsan.doc%2FGUID-8491C4B0-6F94-4023-8C7A-FD7B40D0368D.html> QUESTION 63Which statement is true for the Path Selection Plug-In VMW\_PSP\_MRU? A.

VMW\_PSP\_MRU is default for a majority of active-active and active-passive arrays.B. VMW\_PSP\_MRU will remain on the selected path even if the state were to change.C. VMW\_PSP\_MRU is recommended for Virtual SAN.D. VMW\_PSP\_MRU will have no preferred path setting for the Plug-In. Answer: DExplanation:The host selects the path that it used most recently. When the path becomes unavailable, the host selects an alternative path. The host does not revert back to the original path when that path becomes available again. There is no preferred path setting with the MRU policy. MRU is the default policy for active-passive storage devices.Reference:

[http://pubs.vmware.com/vsphere-60/index.jsp?topic=%2Fcom.vmware.vcli.examples.doc%2Fcli\\_manage\\_storage.6.7.html](http://pubs.vmware.com/vsphere-60/index.jsp?topic=%2Fcom.vmware.vcli.examples.doc%2Fcli_manage_storage.6.7.html)

QUESTION 64Which two tasks does the Pluggable Storage Architecture (PSA) perform? (Choose two.) A. Handles I/O queueing to the logical devices.B. Handles physical path discovery, but is not involved in the removal.C. Handles physical path discovery and removal.D. Handles I/O queueing to FC storage HBAs. Answer: ACExplanation:When coordinating the VMware NMP and any installed third-party MPPs, the PSA performs the following tasks:Loads and unloads multipathing plug-ins.Hides virtual machine specifics from a particular plug-in. Routes I/O requests for a specific logical device to the MPP managing that device. Handles I/O queueing to the logical devices.Implements logical device bandwidth sharing between virtual machines. Handles I/O queueing to the physical storage HBAs.Handles physical path discovery and removal.Provides logical device and physical path I/O statistics.Reference:

<http://pubs.vmware.com/vsphere-51/index.jsp?topic=%2Fcom.vmware.vsphere.storage.doc%2FGUID-C1C4A725-8BE4-4875-919E-693812961366.html> QUESTION 65Which two statements are true regarding Storage Multipathing Plug-Ins? (Choose two.) A.

The default Path Selection Policy is VMW\_PSP\_MRU for iSCSI or FC devices.B. The default Path Selection Policy is VMW\_PSP\_FIXED for iSCSI or FC devices.C. VMW\_PSP\_MRU is typically selected for ALUA arrays by default.D. VMW\_PSP\_FIXED is typically selected for ALUA arrays by default. Answer: BCReference:

[https://pubs.vmware.com/vsphere-51/index.jsp?topic=%2Fcom.vmware.vcli.examples.doc%2Fcli\\_manage\\_storage.6.5.html](https://pubs.vmware.com/vsphere-51/index.jsp?topic=%2Fcom.vmware.vcli.examples.doc%2Fcli_manage_storage.6.5.html)

QUESTION 66 What is the command to list multipathing modules on an ESXi 6.x host? A. esxcli storage core list plugin --plugin-class=MPB. esxcli storage core list plugin --class-plugin=MPC. esxcli storage core plugin list --plugin-class=MPD. esxcli storage core plugin list --class-plugin=MP Answer: C Reference:

<http://darrylcauldwell.com/advanced-vsphere-5-x-storage-masking-multipathing-filtering/> QUESTION 67 Which two solutions require Physical Mode Raw Device Mapping (RDM)? (Choose two.) A. Direct access to the storage array device B. Virtual Machine Snapshots C. Hardware Acceleration D. Guest Clustering across ESXi hosts Answer: AD Explanation: Normally, Direct access to storage array device and guest clustering across ESXi hosts required the use of Physical Mode raw device mapping.

QUESTION 68 A device's vStorage API for Array Integration (VAAI) support status command line output shows: naa.500253825002a865 VAAI Plugin Name: ATS Status: unsupported Clone Status: unsupported Zero Status: supported Delete Status: unsupported What is the corresponding VAAI support status in the vSphere Web Client? A. Unknown B. Supported C. Not supported D. Unsupported Answer: A Explanation: The VAAI support status will be unknown. Reference:

<https://pubs.vmware.com/vsphere-55/index.jsp?topic=%2Fcom.vmware.vsphere.storage.doc%2FGUID-0520FD37-D7AD-4FBA-9A2E-E5F8211FCBBB.html> QUESTION 69 Refer to the Exhibit. What will be the result of selecting the highlighted device? A.

Datastore will grow up to 200.01GB using the remaining free space on the device. B. Datastore will add 200.01GB by adding the device as a second extent. C. The device size can be expanded to be larger than 200.01 GB in size. D. The device is not suitable for this operation. Answer: A Explanation: The datastore will use up 200.01 GB of free space on the device. QUESTION 70 An administrator observes that virtual machine storage activity on an ESXi 6.x host is negatively affecting virtual machine storage activity on another host that is accessing the same VMFS Datastore. Which action would mitigate the issue? A. Enable Storage I/O Control. B. Configure Storage DRS. C. Enable the Dynamic Queue Depth Throttling option. D. Configure the Disk.SchedNumReqOutstanding parameter. Answer: A Explanation: SIOC is extremely powerful, it can increase your consolidation ratios on the storage side, allowing more VM's per datastore. Which leads to lower storage costs and less administrative overhead. So how does it work? At a basic level SIOC is monitoring the end to end latency of a datastore. When there is congestion (the latency is higher than the configured value) SIOC reduces the latency by throttling back VM's who are using excessive I/O. Now you might say, I need that VM to have all of those I/O's, which in many cases is true, you simply need to give the VMDK(s) of that VM a higher share value. SIOC will use the share values assigned to the VM's VMDK's to prioritize access to the datastore. Just simply turning SIOC on will guarantee each VMDK has equal access to the datastore, shares fine tune that giving you the ability to give VMDK's more or less priority during times of contention. Reference:

<http://blogs.vmware.com/vsphere/2014/05/enabling-monitoring-storage-io-control.html> !!!RECOMMEND!!! 1. | 2017 New 2V0-621D Exam Dumps (PDF & VCE) 256Q&As Download: <https://www.braindump2go.com/2v0-621d.html> 2. | 2017 New 2V0-621D Study Guide Video: YouTube Video: [YouTube.com/watch?v=EeODpwZixxo](https://www.youtube.com/watch?v=EeODpwZixxo)