

2V0-620^{Q&As}

vSphere 6 Foundations Beta

Pass VMware 2V0-620 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.leads4pass.com/2v0-620.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by VMware
Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers



QUESTION 1

An administrator is cloning and configuring five new web server virtual machines.

What would be the benefit of configuring resource shares for the new VMs?

- A. To prioritize access to a resource during contention.
- B. To guarantee access to a resource during contention.
- C. To prioritize access to a resource before contention occurs.
- D. To guarantee access to a resource before contention occurs.

Correct Answer: A

Configuring resource shares for the new Vms can prioritize access to a resource during contention Reference: <http://www.settlersoman.com/vmware-vsphere-basics-understanding-resource-pools/>

QUESTION 2

Which two statements are true regarding VMware vSphere Flash Read Cache (vFRC)? (Choose two.)

- A. Cache fills and cache evictions happen in the granularity of a cache block size.
- B. vFRC caches data from both read and write I/Os, but write I/Os are always serviced by the underlying storage.
- C. vFRC caches data from both read and write I/Os, but write I/Os are always serviced by the underlying cache data.
- D. Cache fills and cache evictions happen in the granularity of the disk block size.

Correct Answer: AB

Cache fills and cache evictions happen in the granularity of a cache block size. This value ranges from 4KB to 1MB to enable you to best configure your cache block size based on the I/O size of workloads. Even though cache fills and cache evictions happen in the granularity of a cache block size, actual read I/O serviced by the cache can be smaller than the cache block size. For example, if the cache block size is 64KB, and a 4KB read I/O request is issued by the guest virtual machine, and if the data is not available in the cache, a 4KB read is issued to the VMDK. When populating the cache, the vFRC algorithm looks for a 64KB region to place the new 4KB data. If no free space is available, a 64KB region is evicted and the space is used to hold the new 4KB data. The remaining 60KB region in the 64KB cache block is marked as invalid. The cache block size parameter therefore has profound effects on performance

Reference: <https://www.vmware.com/files/pdf/techpaper/vfrc-perf-vsphere55.pdf>

QUESTION 3

An administrator has a virtual machine that requires four times the compute resources than other virtual machines on the same ESXi 6.x host.

How should the administrator configure the virtual machine settings, in order to be prepared for any resource contention?

- A. Set the shares of the priority virtual machine to High and the rest to Low.
- B. Set the shares of the priority virtual machine to High.
- C. Set the shares of the priority virtual machine to High and the rest to Normal.
- D. Set the shares of the priority virtual machine to Normal and the rest to Low.

Correct Answer: A

For the virtual machine that requires more resources, you can set the shares of the priority VM to high and set the rest to medium or low.

Reference: <http://www.globalknowledge.net/mea-shared-content/documents/645372/645389/>

WP_VM_Resource_Pools.pdf (page 3)

QUESTION 4

An administrator wants to add a web server to an existing multi-tier application consisting of three virtual machines: A web server A database server An application server

The web server should be added to the application when the primary web server reaches: 70% vCPU utilization 55% active memory

Which option will achieve this result?

- A. Create a virtual machine alarm with an action to run a script that starts a new instance of the web server.
- B. Create a host cpu and memory alarm with an action to run a script that starts a new instance of the webserver.
- C. Configure High Availability application monitoring for the web server and set it to trigger deployment of a new instance of the web server.
- D. Configure Fault Tolerance on the virtual machine and leave the secondary machine disabled until needed.

Correct Answer: A

You have to create a new VM alarm with an action to run a script that starts with a new instance of the web server.

QUESTION 5

How are ports scaled on vSphere Standard Switches (vSS)?

- A. Ports on a vSS can be dynamically scaled up and down.
- B. Ports on a vSS can only be statically scaled up or down.
- C. Ports on a vSS can only be dynamically scaled down.
- D. Ports on a vSS can only be dynamically scaled up.

Correct Answer: A

To ensure efficient use of host resources on hosts running ESXi 5.5 and later, the number of ports of standard switches are dynamically scaled up and down. A standard switch on such a host can expand up to the maximum number of ports supported on the host. Reference: <https://pubs.vmware.com/vsphere-60/index.jsp?topic=%2Fcom.vmware.vsphere.networking.doc%2FGUID-350344DE-483A-42ED-B0E2-C811EE927D59.html>

[2V0-620 PDF Dumps](#)

[2V0-620 VCE Dumps](#)

[2V0-620 Study Guide](#)