



VMware

VCAP5-DCD

VMware Certified Advanced Professional 5 - Data Center Design

QUESTION: 262

A multitier website has an RTO of 4 hours and an RPO of 15 minutes. The database crashes at 5 p.m. and causes the website to be unavailable. Which of the following is true? (Select all that apply.)

- A. The website must be restored by 5:15 p.m.
- B. The website must be restored by 9 p.m.
- C. At restore time, the data must be correct from at least from 4:45 p.m.
- D. At restore time, the data must be from at least 1 p.m.

Answer: B, C

QUESTION: 263

You are the VMware consultant at a large enterprise. A senior manager is in charge of a project that needs new infrastructure for its application. The application will be used by 3,200 users. Three virtual machines are requested via the help desk. One of the three is a virtual machine with 32 vCPUs allocated for processing. The notes of the request say that the large number of vCPUs (the company template is 2 vCPUs) is justified by the senior manager. The notes also reference the following details from the application manual:

For every 100 users in the workload, the processing tier of the application requires one agent handler service.

Each agent handler requires one single-core CPU.

Which of the following should the VMware consultant do next?

- A. Configure the 32-vCPU box as requested and talk to the users during the rest of the project to understand the application as much as possible for ease of support and to understand the impact.
- B. Create a proof of concept server with 2 vCPUs and engage with the user/application vendor to understand actual resources required for the company workload and support requirements. Suggest other deployment strategies, such as scaling out the processing tier with a lower number of vCPUs per guest virtual machine but more virtual machines within the tier.
- C. Deny the virtual machine request and suggest that the processing tier be physical.

Answer: B

QUESTION: 264

According to HA best practices, which of the following admission control policies is most appropriate for a cluster of hosts with virtual machines of very different memory and CPU reservations?

- A. CPU and Memory % Configuration
- B. Number of Hosts Tolerate
- C. Failover Host

Answer: A

QUESTION: 265

A requirement in a server consolidation project states: “The consolidation ratio must be as high as possible.” This has resulted in a lower number of larger-capacity hosts in the vSphere HA cluster. Which of the following is true?

- A. The restart time of a VM could be greater than when using smaller hosts.
- B. The restart time of a VM could be lower than when using smaller hosts.

Answer: A

QUESTION: 266

A host has three dual-port network cards. What is the maximum number of vSwitches that can be integrated with the physical network, assuming that network redundancy is not a requirement?

- A. 3
- B. 6
- C. 256
- D. 512

Answer: B

QUESTION: 267

As a platform designer, you have been asked to create a design for a vSphere 5 platform. Based on existing capacity state analysis, the platform will require 30 hosts and 250 guest VMs. Due to an existing vendor relationship with the enterprise, the preferred

database for all datacenter systems is Microsoft SQL 2008. Which of the following vCenter deployments can be used based on this constraint?

- A. VCSA with a remote Microsoft SQL 2008 database
- B. Windows-based vCenter with a local Microsoft SQL 2008 database
- C. VCSA with an embedded Microsoft database
- D. Any of the above would meet the requirements.

Answer: B

QUESTION: 268

RTO and RPO requirements can have a direct impact on the size and number of virtual datastores used in a vSphere platform.

- A. True
- B. False

Answer: A

QUESTION: 269

You are a platform designer constructing a physical design from an existing approved logical design. Out of the vendor proposals, there are two proposed solutions that could be used. Which of the following options is the most important factor when making a decision?

- A. Community and vendor-based best practices
- B. Existing vendor relationships
- C. Project requirements
- D. Project budget

Answer: C

QUESTION: 270

You have been asked to design a virtual machine for a new application developed in-house. The application will be load tested before it is finally made production ready. The application developers have created documentation based on data flow and entity dependency. The functionality of the application is critical to the business for a defined

project with a client base of a maximum of 12,000 users. Which of the following scaling approaches would best suit this scenario?

- A. Predictive
- B. Adaptive

Answer: A

QUESTION: 271

You are a virtualization consultant at a company that develops its own inhouse website trading systems. The new website will be deployed in a vSphere 5 platform, which consists of two clusters and four ESXi hosts (one for each project—the SDLC hosts—for all levels of testing and development prior to production deployment).

The production cluster has been configured for fully automated DRS and VMware H

A. During the SDLC, the website platform was developed on a single development ESXi host inside a single VLAN. The system was successfully tested within the SDLC ESXi hosts and then deployed to the production cluster in a single vApp and a single VLAN.

Following a successful midweek, the new application starts to show signs of latency. Suddenly several applications across the cluster are being impacted by network latency issues. Metrics are collected and show that the amount of network traffic between the database and application servers of the new website platform are at 96% utilization on the physical switch links between ESXi hosts in the production cluster. What should you do?

A. Recommend that the website platform be fixed and removed from the production system.

B. Create a DRS affinity rule to ensure that the database and applications servers stay on the same

ESXi host and do not traverse the physical network when communicating.

C. Configure Network I/O Control to ensure that the network latency is alleviated between the database and applications servers.

Answer: B

QUESTION: 272

You are working on a project to deploy a critical multitier web application. The application is as follows:

Number of Servers	Server Role	Memory Allocation	CPU Allocation	Disk Allocation
20	Web	8 GB	2 vCPU	80 GB system, 150 GB data
5	Processing	2 GB	2 vCPU	80 GB system, 50 GB data
2	Database	32 GB	8 vCPU	80 GB system, 300 GB logs, 500 GB data

In the scenario described, the SAN space is limited. Which of the following options would help save space while still providing the infrastructure shown above?

- A. Confirm that 80 GB is actually used on the servers in each tier and configure thin provisioning to save space.
- B. Reduce the amount of memory in each web server by 2 GB to reduce the amount of swap space created.
- C. Create a resource pool hierarchy for the web, processing, and database tiers in the application and configure a reservation for the total amount in each tier. This would reduce the amount of space used within the datastore.

Answer: C

QUESTION: 273

You are working on a project to deploy a critical multitier web application. The application is as follows:

Number of Servers	Server Role	Memory Allocation	CPU Allocation	Disk Allocation
20	Web	8 GB	2 vCPU	80 GB system, 150 GB data
5	Processing	2 GB	2 vCPU	80 GB system, 50 GB data
2	Database	32 GB	8 vCPU	80 GB system, 300 GB logs, 500 GB data

In the scenario described, which of the following would be a good design choice for memory management in the vSphere platform?

- A. Configure each web server in the application to have a limit of 4 GB.

- B. Create a business-as-usual task to monitor the usage and increase it as required.
 C. Configure a DRS affinity rule to split 50% of the web servers between two hosts. D.
 Configure web servers to have high share values.

Answer: B

QUESTION: 274

You are working on a project to deploy a critical multitier web application. The application is as follows:

Number of Servers	Server Role	Memory Allocation	CPU Allocation	Disk Allocation
20	Web	8 GB	2 vCPU	80 GB system, 150 GB data
5	Processing	2 GB	2 vCPU	80 GB system, 50 GB data
2	Database	32 GB	8 vCPU	80 GB system, 300 GB logs, 500 GB data

In the scenario described, which of the following would be a good alternate design choice to discuss with the vendor/web development team?

- A. Consider consolidating the database servers into one large server for ease of management, backup, and licensing costs.
 B. Consider creating a file server with UNC file shares presenting the read-only data for web servers. This would reduce the amount of space required from approximately 3 TB to approximately 150 GB. C. Reduce the specifications of each web server and increase the number of servers to distribute the load more widely across the cluster.

Answer: B

QUESTION: 275

Entry criteria should be included in all areas of documentation designed for implementation and validation.

- A. True
 B. False

Answer: A

QUESTION: 276

A validation plan should be based on estimated workloads and user behavior.

- A. True
- B. False

Answer: B

QUESTION: 277

A validation plan is used to do which of the following? (Select all that apply.)

- A. Verify the design
- B. Verify that the system is functional
- C. Verify that the system meets requirements
- D. Meet current best practices

Answer: A, B, C

QUESTION: 278

vSphere cluster functionality such as DRS and HA does not need to be part of a validation plan. This kind of documentation is only concerned with guest VM workload validation, not platform-wide validation.

- A. True
- B. False

Answer: B

QUESTION: 279

The technical approach and order of delivery can have an impact on which of the following? (Select all that apply.)

- A. Overall delivery time

- B. People resource contention (right person, right time, right task)
- C. Validation phase duration

Answer: A, B, C

Download Full Version From <https://www.certkillers.net>



DON'T KNOW
OR NO PREFERENCE

Pass your exam at First Attempt....Guaranteed!