

IT 认证电子书



质 量 更 高 服 务 更 好

半年免费升级服务

<http://www.itrenzheng.com>

Exam : **VCAP5-DCD**

Title : VMware Certified Advanced
Professional 5 - Data Center
Design (VCAP5-DCD)

Version : Demo

1. When designing a storage platform, which of the following should be considered as part of the overall design?

- A. Capacity
- B. I/O requirements of the applications to be supported
- C. Disk latency tolerance
- D. Growth rate
- E. All of the above

Answer: E

Explanation:

A storage platform logical design requires in-depth analysis of factors that can affect applications. In the case of storage, aspects that relate to the physical layer - such as the amount of usable space required for services; the size, number, and speed of disks; and how fast the data is being produced - could have a substantial impact to the storage platform. Chapter 4, Storage Design - The Logical & Physical Approaches

2. You are the technical designer for a vSphere platform transformation project. After conducting SME interviews and using various platform information-gathering methods, you have created a high-level design document.

This document specifies the following:

Requirements:

- R1. The solution must not have a single point of failure.
- R2. Production applications must not have an outage of more than 10 seconds.
- R3. Data must be based in the UK.
- R4. There is a 7-year retention policy for contracts.
- R5. Applications should support existing and developing workloads for the next 3 years' growth.

Spec of servers: Web 1 vCPU, 2 GB RAM, 100 GB storage

App 1 vCPU, 4 GB RAM, 100 Gb storage

DB 2 vCPU, 16 GB RAM, 750 Gb storage

At a late stage in the software development life cycle of a production application developed inhouse, an unfortunate issue was identified when the application was deployed to the production vSphere cluster. The production core stacked switch has capacity issues, and this is having a serious impact on all applications for which the switch is providing network services. Within the test system, the application works as intended in the single test VLAN and with a single-host configuration.

Which of the following could be tried to help in this situation, from a vSphere perspective? (Choose two)

- A. Redevelop the application for a virtual platform
- B. Place the application into a single-vApp network
- C. Add DRS rules to keep network traffic within the same host, where possible
- D. Configure network I/O control

Answer: B,D

Explanation:

If you limit the application traffic to a specific dedicated network (that is, a separate VLAN) and using enforced DRS affinity rules, the application traffic will not traverse the ESXi host's physical network interfaces. This will ensure that the impact of the application is minimized, while also ensuring that the application itself is not limited. The application servers already have vCPU settings. This suggests that

the system has already being virtualized in both types of environments. Network I/O control could be useful in the event of contention, but the role of a designer would be to plan to prevent contention where possible; other options would be more beneficial. In addition, Requirement 5 specifies that the system should work with workloads over the next 3 years. Network I/O control would suggest contention very early in the platform history. Chapter 4, Network - Logical and Physical Design to allow applications to flow

3. Based on the information in the high-level design extract from question 2, what storage protocol is unsuitable?

- A. NFS
- B. iSCSI
- C. FC

Answer: A

Explanation:

The requirements state that the application uptime requirements would not be met if using vSphere HA alone. vSphere HA would invoke a restart of guest virtual machines after at least 10 seconds. VMware fault tolerance would meet the uptime requirements; a failover would result in zero downtime of the application. This technology can be used only with VMFS; therefore, NFS cannot be used in this design. Chapter 4, Network - Logical and Physical Design to allow applications to flow

4. Based on the information in the high-level design extract for question 2, which type of data store would be required?

- A. VMFS
- B. NFS

Answer: A

Explanation:

Because fault tolerance would be the only VMware technology that would meet the technical requirements, VMFS is the only choice here. Chapter 4, Storage Design - The Logical & Physical Approaches

5. Which of the following vSphere cluster technologies would meet the application requirements specified in the high-level design extract for question 2?

- A. FT
- B. HA

Answer: A

Explanation:

vSphere HA would need to wait at least 10 seconds before a restart would be possible. This would not meet requirements. In the event of a restart, it could be possible to be without the application feed database for about 15 minutes. VMware FT would enable the service to be provided without additional redevelopment. Chapter 4, Storage Design - The Logical & Physical Approaches

6. A project requirement _____

- A. must be achieved
- B. can be set as aspirational but may never be achieved

Answer: A

Explanation:

A project requirement is a specific deliverable that the business has said must be provided at the end of a project. It must always be achieved. Chapter 4, Storage Design - The Logical & Physical Approaches

7. A logical diagram shows vendor-specific information.

- A. True
- B. False

Answer: B

Explanation:

A logical diagram shows the high-level components of a design or platform. It may include descriptions such as data flow and entity relationships. However, it does not show specific vendors (such as Dell or HP). Chapter 2, Creating the design

8. Storage DRS is enabled by default.

- A. True
- B. False

Answer: B

Explanation:

Chapter 4, Storage Design - The Logical & Physical Approaches

9. Storage DRS can be enabled and can balance I/O levels based on _____

- A. real-time stats
- B. historical stats

Answer: B

Explanation:

Storage DRS can balance I/O by using storage vMotion. This would be based on historical values for a period of time. Storage DRS would make decisions based on these figures. If the system were using real-time stats, the machines could potentially migrate a lot more. This could have an effect on the latency figures for some systems. The larger historical sampling in storage DRS allows the system to make a decision with limited risk of impacting I/O - unlike moving a virtual machine over and over again. Chapter 4, Storage Design - The Logical & Physical Approaches

10. The vCenter database is powered off ungracefully, along with a three cluster nodes in a separate cluster. Will the failed existing ESXi host nodes restart correctly using Auto Deploy?

- A. Yes
- B. No

Answer: A

Explanation:

Auto Deploy enables an untrained IT professional to increase capacity while ensuring quality and consistency within a cluster. However, there is a dependency on the vCenter server and Auto Deploy server roles at different times. In this case, the ESXi hosts have already booted successfully. This means that all subsequent reboots will be controlled using configuration from the Auto Deploy server, not vCenter. If the vCenter database is impacted, the ESXi hosts will automatically restore. Even in the event of no vCenter availability, hosts will use the last-known copy of a distributed switch or the configured vSphere switch. Chapter X, Putting it all together

11. A best practice can change over time.

- A. True
- B. False

Answer: A

Explanation:

A best practice is based on experience with delivering a solution. A new product will have limited best-practice information except for vendor- or vendor-partner-based recommendations. The longer a technology is used in a community, the more potential success; therefore, best practices can change over time. Chapter 6, Validation, Thoughts, and Processes

12. Only a software vendor develops a standard process or configuration.

- A. True
- B. False

Answer: B

Explanation:

A best practice is based on experience with delivering a solution. A new product will have limited best-practice information except for vendor- or vendor-partner-based recommendations. The longer a technology is used in a community, the more potential success; therefore, best practices may be recommended by non-vendor-related people. It is up to a designer to validate this information rather than ignore it straightaway. Chapter 6, Validation, Thoughts, and Processes

13. A logical design can contain which of the following? (Choose three)

- A. BIOS settings
- B. Data flow
- C. Key component relationships
- D. High-level diagrams
- E. Server host names

Answer: B,C,D

Explanation:

A logical diagram can be used to show service components and how they relate to each other. A logical diagram is high level in nature and is independent of the hardware used for a solution. Chapter 2, presenting the Data Gathered

14. You are a VMware architect employed to design and build application environments and the underlying platform for a software development company. The company is using a “devops” approach and has engaged with you at the beginning of the first planning meeting/pre-sprint.

The software has yet to be written. However; you have some high-level specifications that are subject to change at each of the early meetings.

The _____ is more logical for virtual machine sizing.

- A. predictive sizing approach based on current software
- B. adaptive sizing approach

Answer: B

Explanation:

There are two types of scaling approaches in virtual platforms: adaptive and predictive. Both approaches are usable; however, a designer must work out which is the best approach for each project. The adaptive approach to guest virtual machine design is a major advantage that virtualization technology brings to a datacenter. It allows a guest virtual machine to be created initially on a few metrics or high-level requirements (such as a limited budget) and scaled (that is, grown, not shrunk) through the virtual machine lifecycle. In many cases, this scaling can be performed hot/live without outage/downtime to the running applications. The predictive approach, on the other hand, requires more planning and research, and you must know the metrics before deployment. The advantage here is that once the virtual machine has been deployed, only application maintenance is required. The application requirements are known. Chapter 5, Design Approaches

15. You are working on a hybrid cloud project, where production applications (yet to be fully developed) are to be deployed.

Which of the following is a project requirement?

- A. The production data must be in the UK at all times.
- B. The hosting partner provides sufficient resources without overcommitment to support application load.
- C. The hosting provider meets uptime expectations.
- D. The development team provides the software on time.

Answer: A

Explanation:

A requirement is a definable and measurable project item. It must be part of the delivery, and its removal cannot be justified. Chapter 2, Definitions of Key Terms

16. In the project life cycle, the _____ defines the vision.

- A. IT architect
- B. software vendor
- C. business

Answer: C

Explanation:

By reviewing the current configuration of a system and mapping the final status of components as described in the vision, the business identifies the gap between the two. A gap state analysis is a review of the items or tasks that are required to progress between the stages. Such an analysis is useful in creating milestone plans, implementation orders, and so on. Chapter 2, Gathering Information and Spotting the Gaps

17. A(n) _____ is an item that is taken to be true in the design phases but has not been tested prior to execution.

- A. requirement
- B. constraint
- C. assumption
- D. risk

Answer: C

Explanation:

An assumption could impact your design. It is good to ensure that project team and wider business are

aware of assumptions and that they are agreed and validated as part of the design consideration process.
Chapter 2, Gathering Information and Spotting the Gaps

18. You are a virtualization consultant working on a DR project. You have proposed a solution that uses SAN replication technology to replicate production virtual machine files. This meets the cold standby requirement.

Which of the following could be a design constraint?

- A. The hardware currently being used in the datacenter is out of support.
- B. The company is undecided about the choice of centralized storage to be used in the enterprise.
- C. The company is at the end of year 1 of a 3-year contract for the point-to-point link between Site A and Site
- D. This link is currently 10 MB.

Answer: C

Explanation:

An item from the current state analysis that would restrict or change a design choice is potentially a constraint. Such factors can limit a design and in some cases prevent the realization of a project vision.
Chapter 2, Gathering Information and Spotting the Gaps

19. You are a technical consultant designing a solution for an online retail company. The project vision is to deploy a hybrid cloud, with the internal team developing the website on internal infrastructure and migrating production-ready applications to a hosting provider. The project is expected to ease deployment and require less infrastructure capital expenditure - without lowering application quality.

Which of the following is a risk to the project?

- A. The solution must adhere to ISO 27001.
- B. Change control of the hosting vCloud platform is not under full control of the internal business.
- C. The hosting provider outsources the platform support to the platform vendor.
- D. The applications to be deployed on the production hosts are not fully developed, although a beta exists.

Answer: B

Explanation:

The technical design process enables the platform designer to work in a methodical way and with others who have similar objectives. A risk is that a project item could potentially prevent the vision or aspects of it (requirements) from being created. Chapter 1, The Technical Design Process

20. A DR project requires a recovery time objective (RTO) of 4 hours and has a recovery point objective (RPO) of 2 hours. The system fails at 8 p. m. on a Sunday evening. At what time is the application expected to be available to the end users?

- A. Midnight
- B. 10 p. m.
- C. 2 a. m.
- D. 11 p. m.

Answer: A

Explanation:

The recovery time objective (RTO) - also known as the return to operation - is the amount of time it takes

to restore a service after a failure has taken place. The recovery point objective (RPO) is the point in time to which the system needs to be restored following a failure. Chapter 3, RTO, RPO, and All That Stuff?