



- Expert Verified, Online, **Free**.



CERTIFICATION TEST

- CertificationTest.net - Cheap & Quality Resources With Best Support

A customer wants to virtualize an Oracle database with vSphere 6.5, but is concerned about its performance. Which three design elements will ensure optimum performance? (Choose three.)

- A. Share as much memory as possible with the balloon driver.
- B. Use VMXNET3 for the network adapter.
- C. Create affinity rules for the virtual machine to a single physical socket.
- D. Use VMware Paravirtual SCSI adapters for data and log vDisk.
- E. Enable Hyper-Threading.



Suggested Answer: BDE

  **xpl0ziv** Highly Voted 5 years, 4 months ago

BDE.

<https://blogs.vmware.com/performance/2017/05/oracle-database-performance-vsphere-6-5-monster-virtual-machines.html>

upvoted 11 times

  **Ilstrixll** 5 years, 4 months ago

Thank you - the document makes it crisp and clear that the answer should be BDE.

upvoted 3 times

  **VCAPito** Highly Voted 5 years, 7 months ago

I like this website as students can share their ideas and opinions!



Avoid at all cost sites like prepaway, they're ripping you off with fake and wrong Q&As

upvoted 6 times

  **chafik** Most Recent 4 years, 2 months ago

Correct Answer: BDE

upvoted 2 times

  **Ophone** 4 years, 5 months ago

BDE is correct

upvoted 1 times

  **lotso** 4 years, 9 months ago

BDE without any doubt

upvoted 3 times

  **Marcossss** 4 years, 11 months ago

BDE.

correct answer

upvoted 3 times

  **vmwarelive** 5 years, 4 months ago

Correct Answer: BCD

upvoted 4 times

  **mariogg85** 5 years ago

IMHO I think that the correct answer may be BDE. The answer C says that you configure affinity rules for uses one physical socket. Ok, but why reason?

What do we gain from it?

upvoted 1 times

  **mariogg85** 5 years ago

IMHO the answer D it is more reasonable but with scsi paravirtual you can get a great performance.

upvoted 2 times

A development team must provide layer 2 network isolation between virtual machines that are in the same VLAN. The solutions architect must provide additional security between the virtual machines on the same subnet.

How can this be done without consuming more VLANs?

- A. Use Virtual Switch Tagging.
- B. Use Private VLANs.
- C. Use Virtual Guest Tagging.
- D. Use External Switch Tagging.

Suggested Answer: B

  **dax**  5 years, 6 months ago

A PVLAN is a VLAN with configuration for Layer 2 isolation from other ports within the same broadcast domain or subnet

upvoted 8 times

  **chafik**  4 years, 2 months ago

Correct Answer: B

upvoted 1 times

  **Ranjith_24284** 5 years, 6 months ago

A is the Answer

upvoted 2 times

  **dax** 5 years, 6 months ago

reqmt: additional security between the virtual machines on the same subnet.

vlan tagging wont provide any difference as they are still in same subnet.

correct answer is B

upvoted 11 times

DRAG DROP -

A company is outsourcing its support operations to an external service provide and plans to complete the project by April 1.

- ⇒ The external support engineers must have the ability to power cycle, create, and edit virtual machines settings within their assigned vSphere site.
- ⇒ The company maintains three vCenter servers in Enhanced Linked Mode that are run as virtual machines in the supported infrastructure.
- ⇒ The vCenter servers will be supported by the external service provider.
- ⇒ Each vCenter server is connected to its own local Platform Services Controller and MSSQL database server.
- ⇒ The company will provide escalation support and physical access on a per request basis.
- ⇒ 99.9% ESXi host uptime is required in this environment, but no SLA has been specified for the hosted applications.

Drag each statement to its appropriate concept.

Select and Place:

Statement	Concept
The SLA allows for only 8 hours of downtime per year.	Risk
Remote support engineers are able to edit virtual machines settings.	Constraint
Missing the April 1 deadline will result in additional costs to the company.	Assumption
The company is appropriately staffed to support escalation and physical access.	Requierement

Suggested Answer:

Statement	Concept
	Missing the April 1 deadline will result in additional costs to the company.
	The company is appropriately staffed to support escalation and physical access.
	Remote support engineers are able to edit virtual machines settings.
	The SLA allows for only 8 hours of downtime per year.

Ilstrixll Highly Voted 5 years, 4 months ago

SLA is 8h -> Requirement

Remote support -> Constraint

Missing deadline -> Risk

Company is staffed -> Assumption

upvoted 22 times

SuperTed Highly Voted 5 years, 7 months ago

I would say

The SLA allows for ... -> Constraint


Remote Support engineers are able .. -> Requirement

Missing the April 1 .. -> Risk

The company is appropriately .. -> Assumption



But I am not quite sure. The design constraints are factors that restrict the options the architect can use to satisfy the design requirements. Assumptions are made by the architect and have not yet been validated. Assumptions are not accepted as a fact until they have been validated.

upvoted 20 times

  **dax** 5 years, 5 months ago

agree with your answer

upvoted 3 times

  **Kof_81** 4 years, 3 months ago

Agree with @SuperTed, SLA is a non-functional requirement (HOW a system should behave), thus it is a Constraint (at least according to the vSphere Design Cook Book).

upvoted 1 times

  **chafik**  4 years, 2 months ago



The SLA allows for ... -> Constraint

Remote Support engineers are able .. -> Requirement

Missing the April 1 .. -> Risk

The company is appropriately .. -> Assumption

upvoted 3 times

  **Chrisog** 4 years, 2 months ago

Risk: Missing deadline will result in additional cost. That's a risk to a company

Constraints: 99.9% ESXi host uptime is required in this environment, but no SLA has been specified...SLA allowing for only 8 hours of downtime per year is a constraint in meeting your 99.9% uptime

Assumption: "The company will provide escalation support and physical access on a per request basis" This assumes that "The company is appropriately staffed to support escalations and Physical assess.

Requirement: "External engineers must have the ability to power cycle, create, and edit virtual machines settings within their assigned vSphere site."

So the statement already Requires that "Remote Support Engineers are able to editVirtual Machines settings

upvoted 1 times



  **adrian_des** 4 years, 8 months ago

SLA - Constraint - by def a limitation

Remote support - Requirement - "external support engineers must have the ability"

agree with what others have for Risk and Assumption

upvoted 2 times

  **Ophone** 4 years, 5 months ago

"**99.9% ESXi host uptime is required in this environment, but no SLA has been specified for the hosted applications."

Constraint must something existed and limit the design

upvoted 1 times

  **mo_rada** 4 years, 9 months ago

As always, the answer is in the question...

**99.9% ESXi host uptime is required in this environment, but no SLA has been specified for the hosted applications.

upvoted 1 times

  **VCAPito** 5 years, 7 months ago

SLA is 8h (99.9% ESXi host uptime is required) -> Requirement

Remote support (they will have permission, but will they be skilled enough?) -> Constraint/Assumption (Have to rely on/wait for them for change and assume they are able to (which they should, it's their job after all))

Missing deadline -> Risk

Company is staffed -> Assumption/Constraints (More or less same as per Remote support)

I have doubts regarding which one is constraints and which one is assumption. As Ted said, assumptions require validation, and from the info provided there aren't enough info to validate them.

If anyone can shed more light, we'd be grateful

upvoted 4 times

A customer has requested a vSphere 6.5 deployment design that utilizes vCenter Server and the use of VMware-recommended best practices for securing vCenter Server.

Which three actions would satisfy these requirements? (Choose three.)

- A. Utilizing vSphere CLI and vSphere SDK for Perl scripts.
- B. Restricting vCenter Server access to only the management network.
- C. Assigning the default Administrator role to all administrator users.
- D. Synchronizing time in the vCenter Server with a NTP source.
- E. Removing expired and revoked certificates from vCenter Server system.

Suggested Answer: ABD

  **diegof1**  5 years, 8 months ago

BDE. Leaving expired or revoked certificates or leaving vCenter Server installation logs for failed installation on your vCenter Server system can compromise your environment. <https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.security.doc/GUID-F583EF9D-49A0-438F-8A8E-DD6E0A11186E.html>

upvoted 14 times

  **VCAPito**  5 years, 7 months ago



Agreed with BDE - <https://docs.vmware.com/en/VMware-vSphere/6.5/vsphere-esxi-vcenter-server-65-security-guide.pdf>

upvoted 8 times

  **chafik**  4 years, 2 months ago

Correct Answer: BDE

upvoted 1 times



  **Chrisog** 4 years, 2 months ago

Correct answers are BDE.

Answer A is wrong: CLI is risky, in fact disabled by default on ESXi servers. Should only enabled and used when necessary and disabled afterwards


C: is wrong also. Obviously assigning default administrator role to everyone is not the way to go

upvoted 1 times

  **Ophone** 4 years, 5 months ago

BDE is correct

upvoted 1 times

  **Lasitus** 4 years, 6 months ago

BDE. Information from <https://docs.vmware.com/en/VMware-vSphere/6.5/vsphere-esxi-vcenter-server-65-security-guide.pdf>

B - page 100

D page 101

E - page 90

upvoted 3 times

  **MistaPablo** 4 years, 5 months ago

Agree, E - page 99

upvoted 1 times

  **vmwarelive** 5 years, 4 months ago

Correct Answer: BDE

upvoted 7 times

A database administrator is operating a virtual machine (VM) configured with 16 vCPU and 64GB of RAM. A recent performance audit has indicated that this virtual machine is oversized and is using less than 60% of its configured CPU and memory capacity.

- ⇒ The ESXi host that contains this VM has 2 physical processors with 10 cores per processor, and 128GB of RAM.
- ⇒ This physical host's architecture is split into two equal NUMA nodes.

Which vCPU and RAM configuration for the VM allows for the most resources, but also provides the performance benefit of local NUMA access?

- A. 16 vCPU and 32GB RAM
- B. 4 vCPU and 16GB RAM
- C. 10 vCPU and 64GB RAM
- D. 12 vCPU and 64GB RAM

Suggested Answer: B

🗲️ 👤 **diegof1** Highly Voted 👍 5 years, 8 months ago

C is right. Because audit indicates that the VM is using less than 60% (9.6 vCPU and 38.4 GB of Memory), then the best configuration that takes into account the NUMA architecture is 10 vCPU with 64GB so that it is not generated Traffic between memory banks of physical processors. (10 cores because it is greater than 60% which is 9.6), in addition this configuration would be executed in single NUMA Node).

<https://blogs.vmware.com/performance/2017/03/virtual-machine-vcpu-and-vnuma-rightsizing-rules-of-thumb.html>
upvoted 8 times

🗲️ 👤 **duke_srg** Highly Voted 👍 5 years, 8 months ago

Answer C is the right one in terms of NUMA locality and VM resources consumed
upvoted 5 times

🗲️ 👤 **chafik** Most Recent 🕒 4 years, 2 months ago

Correct Answer: C
upvoted 1 times

🗲️ 👤 **mariogg85** 5 years ago

IMHO the answer would be C. The VM has 16 vCPU and 64 GB RAM like it says on the statement the vm is using less of 60%. With this point we can deduce that this VM could be resized with 9 or 10 vCPU and 38-39GB RAM. It also indicates that they want to benefit from the benefits of the NUMA node (2 Nodes with 10 Cores and 64 GB each NUMA node). With this information the answer C is the right it should be resized to 10vCPU and 64 GB.
upvoted 2 times

🗲️ 👤 **GodMan114** 5 years, 3 months ago

I too like answer C. For a moment I was worried about ESXi overhead not being taken into consideration, but still seems to be the most logical answer.
upvoted 1 times

🗲️ 👤 **vmwarelive** 5 years, 4 months ago

Correct Answer: C
upvoted 1 times

🗲️ 👤 **Grande** 5 years, 7 months ago

I agree with Deof1's analysis. A vCPU lower than 9.6 will mean VM is under-sized. C is correct
upvoted 1 times

When implementing update policies for the vSphere environment, which would be the VMware-recommended way to update the vCenter Server Appliance (VCSA) when an underlying operating system (OS) patch is released?

- A. Introduce a policy that requires a system administrator to check if a new appliance update (which might include an OS update) is available from the downloads section of MyVMware portal, and follow the VCSA documentation to apply the update.
- B. Do nothing-the VCSA applies all OS updates automatically without any human interaction.
- C. Introduce a policy that requires a system administrator to go online and check with the OS vendor to see if a new version is available. If it is, download it manually, log in to the VCSA with the root credentials, and proceed with the OS update.
- D. Configure VMware Update Manager to download the OS update and apply it on a scheduled basis.

Suggested Answer: A



  **vCat**  5 years, 3 months ago

A is the correct answer.



Currently there is no automatism to update/upgrade a VCSA
upvoted 5 times

  **chafik**  4 years, 2 months ago

Correct Answer: A
upvoted 1 times

  **Ilstrixll** 5 years, 3 months ago

Yes A is the correct answer
upvoted 1 times

  **hendri** 5 years, 4 months ago

i think A is correct answer.
upvoted 2 times

DRAG DROP -

Sort the traffic by whether it can be encrypted natively by vSphere.

Select and Place:

Traffic	Can be encrypted
vMotion traffic	Place here
VM traffic	Place here
SNMPv3	Place here
vSphere Client	

Cannot be encrypted
Place here

Suggested Answer:

Traffic	Can be encrypted
	vMotion traffic
	VM traffic
	SNMPv3

Cannot be encrypted
vSphere Client

Grande Highly Voted 5 years, 7 months ago

vSphere by default uses certificates to communicate between various infrastructure components - and by that traffic is encrypted already - see
See the "Server Configuration Guide", Chapter 11, page 227
"Encryption and Security Certificates for ESX Server."

VM traffic requires an external KMS server where encryption keys are stored. an external component
<https://www.starwindsoftware.com/blog/encryption-of-vmware-vsphere-6-5-virtual-machines-and-vmotion-migrations-and-their-performance>
upvoted 12 times

duke_srg Highly Voted 5 years, 7 months ago

VM traffic can't be encrypted natively by vSphere
upvoted 6 times

chafik Most Recent 4 years, 2 months ago

Encrypted natively:

vMotion traffic

vSphere client

SNMPv3

Not encrypted natively:

VM traffic

upvoted 1 times

🗨️ 👤 **BilboRowbear** 4 years, 3 months ago

Even though VMs can be encrypted on storage using a KMS, this does NOT encrypt the VM network traffic. It only encrypts the files on the data store. Also, the vSphere Client traffic has always been encrypted using HTTPS. Look at any Hands On Lab at the vSphere Client client URL for proof. HTTP is unencrypted and HTTPS is encrypted.

Both VMotion & SNMPv3 traffic can be encrypted.

upvoted 1 times

🗨️ 👤 **Lasitus** 4 years, 6 months ago

Encrypted natively:

VMotion traffic

vSphere client

SNMPv3

Not encrypted natively:

VM traffic

upvoted 5 times

🗨️ 👤 **admin007** 4 years, 7 months ago

is vsphere client traffic natively encrypted ?

upvoted 1 times

🗨️ 👤 **alex3** 4 years, 7 months ago

no it is not

upvoted 2 times

🗨️ 👤 **SuperTed** 5 years, 7 months ago

I agree with duke, VM traffic can't be encrypted natively by vSphere

upvoted 4 times

A solution architect has finished conducting interviews and gathering requirements for a company, and has determined that the logical requirements are:

- ⇒ Two data centers for high availability
- ⇒ Synchronous replication to meet the zero minute RPO
- ⇒ Separating management workloads from application workloads
- ⇒ Dedicated 10Gb uplink for each low latency server
- Single management point for the entire environment
-

Which two actions would meet the design requirements? (Choose two.)

- A. Configure 1 Port Group with a dedicated 10Gb Uplink for low latency servers.
- B. Deploy two clusters, one for management workloads and one for application workloads.
- C. Build 2 Port Groups, one for management servers and one for application servers.
- D. Install two vCenter Servers in Enhanced Link Mode.


Suggested Answer: AB

  **diegof1** Highly Voted 5 years, 8 months ago

Answer is BD.

Installing two vCenters connected with Enhanced Linked Mode meets the requirement of a single administration point for the entire environment. To separate the Administration charges from the Application loads, two clusters must be deployed. On the other hand, 1 dedicated Port Group for 10GB is a good option, but NOT with a single dedicated 10GB Uplink as it would generate a single point of failure for low-latency servers.

upvoted 12 times

  **VCAPito** Highly Voted 5 years, 7 months ago

Agreed with BD. 2 clusters to separate workloads and vCenter in ELM can be managed from one single point.


The other answers are wrong as the servers require one 10Gbps link EACH, not shared.

upvoted 6 times

  **chafik** Most Recent 4 years, 2 months ago

Correct Answer: BD

upvoted 1 times

  **yfirman** 4 years, 11 months ago

yes correct answer is BD

upvoted 3 times

  **vmwarelive** 5 years, 4 months ago

Correct Answer: BD

upvoted 4 times

A solution architect has made the following design decisions:

- ⇒ Leverage existing hardware that is certified on earlier versions of vSphere but is NOT on HCL for ESXi 6.5.
- ⇒ Upgrade vCenter Server to version 6.5.
- ⇒ Configure separate clusters based on ESXi versions 5.5, 6.0, and 6.5 for newly purchased, certified hardware.
- ⇒ The underlying CPU family is compatible.
- ⇒ There is enough resources available to vMotion virtual machines (VMs).

Given this scenario, what is the correct statement about the ability to vMotion virtual machines between versions of ESXi?

- A. VMs created in vSphere 5.x must be upgraded first to newer virtual hardware and then be vMotioned to vSphere 6.x.
- B. VMs created in vSphere 6.x environment with default settings can be moved to ESXi 5.x.
- C. VMs can be vMotioned to the same or newer version of ESXi.
- D. VMs that are created after the vCenter Server 6.5 upgrade can be vMotioned between any supported versions of ESXi.

Suggested Answer: A

🗲️ 👤 **diegof1** Highly Voted 👍 5 years, 8 months ago

C is the correct answer.

Is not necessary to upgrade Virtual Hardware to migrate between ESXi versions.

On the other hand, you can migrate VMs between hosts with same ESXi version or newer versions, without restriction.

upvoted 19 times

🗲️ 👤 **chafik** Most Recent ⌚ 4 years, 2 months ago

Correct Answer: C

upvoted 1 times

🗲️ 👤 **vmwarelive** 5 years, 4 months ago

Correct Answer: C

upvoted 2 times

🗲️ 👤 **dax** 5 years, 5 months ago

We have esxi upgrade from 5.5 to 6.5. vms are unable to migrate back to esxi 5.5. but the rest of the options are more far off... So still correct is C

upvoted 1 times

🗲️ 👤 **Grande** 5 years, 7 months ago

I agree with diegof

https://docs.vmware.com/en/VMware-vSphere/6.5/com.vmware.vsphere.vm_admin.doc/GUID-64D4B1C9-CD5D-4C68-8B50-585F6A87EBA0.html

upvoted 2 times

DRAG DROP -

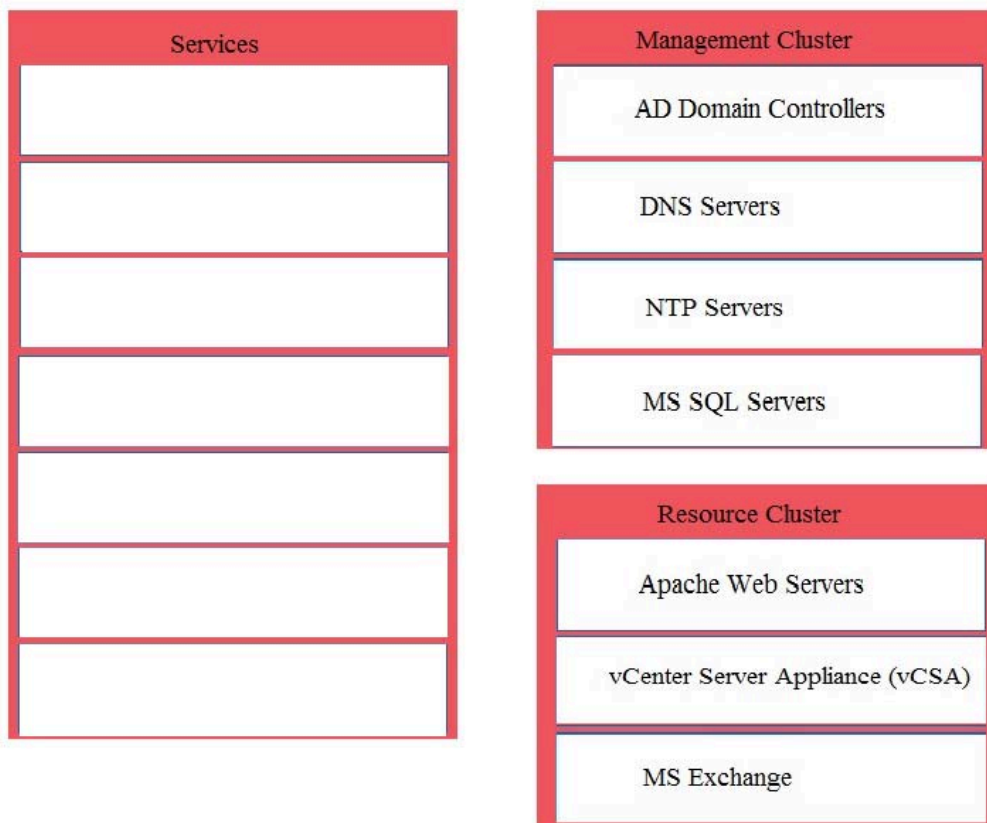
According to VMware-recommended best practices, on which cluster should each of the services be placed?

Select and Place:

Services	Management Cluster
MS Exchange	Place here
DNS Servers	Place here
NTP Servers	Place here
MS SQL Servers	Place here
Apache Web Servers	
AD Domain Controllers	
vCenter Server Appliance (vCSA)	

Resource Cluster
Place here
Place here
Place here

Suggested Answer:



diegof1 Highly Voted 5 years, 8 months ago

Correct Answer:

Management Cluster = DNS, NTP, AD, VCSA

Resource Cluster = MS Ex, MS SQL, Apache

upvoted 23 times

VCAPito Highly Voted 5 years, 7 months ago

Can't find the documentation supporting this, but I agree with Diego answer, as vCenter depends on NTP, DNS and AD:

Management cluster: DNS, NTP, AD DC, VCSA

Resource cluster: MS Exchange, MS SQL Server, Apache

upvoted 5 times

chafik Most Recent 4 years, 2 months ago

Correct Answer:

Management Cluster = DNS, NTP, AD, VCSA

Resource Cluster = MS Ex, MS SQL, Apache

upvoted 1 times

lotso 4 years, 9 months ago

I agree.

Management: AD, DNS, NTP and VCSA

Resource cluster: MS Exchange, SQL, Apache.

Consider that exchange, sql and apache are applications and your infrastructure can exist without them, but not without NTP, DNS, AD, etc

upvoted 2 times

neo20011 4 years, 9 months ago

basically VCSA and any others it needs are in the mgmt cluster: VCSA, AD, NTP, DNS. Anything outside of that (specifically resources dedicated to end-user w into the resource cluster. ref:

http://download3.vmware.com/vcat/vcat31_documentation_center/index.html#page/Architecting%20a%20vCloud/3a%20Architecting%20a%20VMware%20v

upvoted 1 times

  **mariogg85** 5 years ago

I agree with both.

upvoted 1 times

DRAG DROP -

A company would like to utilize its current infrastructure but wants to adopt virtualization to consolidate its environment. Currently, the infrastructure contains:

- ⇒ Server with 2 x 8 cores CPUs and 98GB of memory
- ⇒ Backup LAN with a single physical switch
- ⇒ Production LAN
- ⇒ Sufficient datastore performance as determined by the storage team.

Match the existing infrastructure component to its appropriate concept.

Select and Place:

Component	Concept
production LAN	Risk
backup LAN has a single physical switch	Constraint
server with 2 x 8 cores CPUs and 96GB of memory	Assumption
sufficient datastore performance as determined by the storage team	Requirement

Suggested Answer:

Component	Concept
	sufficient datastore performance as determined by the storage team
	backup LAN has a single physical switch
	production LAN
	server with 2 x 8 cores CPUs and 96GB of memory

charithabuddhika Highly Voted 5 years, 8 months ago

Risk <--> backup LAN has a single physical switch

Constraint <--> server with 2 x 8 cores CPUs and 96GB of memory

Assumption <--> sufficient datastore performance as determined by the storage team

Requirement <--> production LAN

upvoted 47 times

micus 4 years, 11 months ago

i agree with you :)

upvoted 1 times

mo_rada 4 years, 9 months ago

Also agree. Although the term "requirement" doesn't seem to make any sense to me in this scenario nor apply, "Production LAN" would be the answer by process of elimination.

upvoted 3 times

mo_rada 4 years, 9 months ago

Agreed - Requirement = something not yet existing or implemented = Production LAN

Constraint is something that already exists and can't be changed, require Validation = The \$hitty server specs
upvoted 1 times

  **vmwarelive** Highly Voted 5 years, 4 months ago

Risk <--> backup LAN has a single physical switch

Constraint <--> server with 2 x 8 cores CPUs and 96GB of memory

Assumption <--> sufficient datastore performance as determined by the storage team

Requirement <--> production LAN

upvoted 13 times

  **chafik** Most Recent 4 years, 2 months ago

Risk <--> backup LAN has a single physical switch

Constraint <--> server with 2 x 8 cores CPUs and 96GB of memory

Assumption <--> sufficient datastore performance as determined by the storage team

Requirement <--> production LAN

upvoted 1 times

  **MohamedFouad** 5 years, 6 months ago

Agreed on chaith



Risk <--> backup LAN has a single physical switch

Constraint <--> server with 2 x 8 cores CPUs and 96GB of memory

Assumption <--> sufficient datastore performance as determined by the storage team

Requirement <--> production LAN

upvoted 5 times

  **diegof1** 5 years, 8 months ago



Risk <--> sufficient datastore performance as determined by the storage team

Constraint <--> backup LAN with a single physical switch

Assumption <--> server with 2 x 8 cores CPUs and 96GB of memory



Requirement <--> production LAN

upvoted 1 times

  **hoo** 5 years, 6 months ago

Wrong answer

upvoted 5 times

  **Ilstrixll** 5 years, 4 months ago

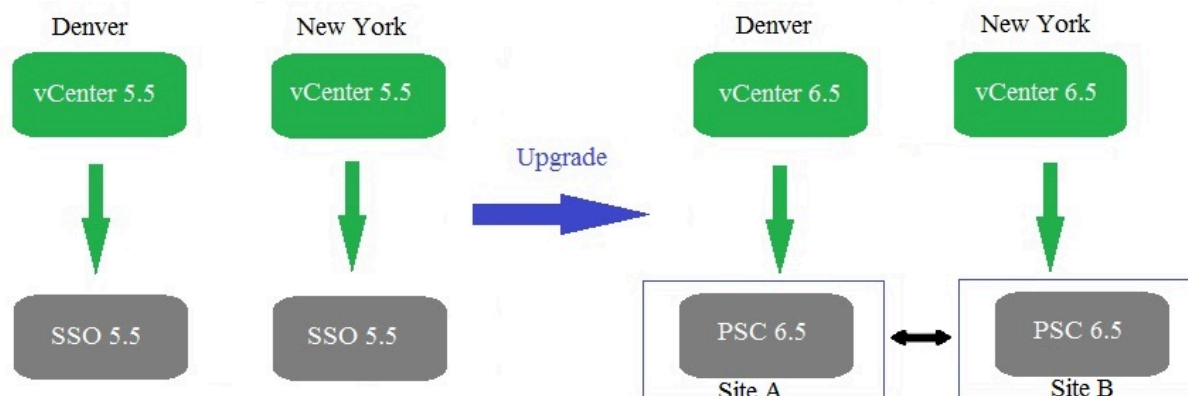
I have to disagree - this is not correct

upvoted 1 times

The system administrator team is planning to upgrade its vCenter Server 5.5 environments to version 6.5.

- ⇒ Each vCenter 5.5 is pointing to a Single Sign On (SSO) server that has a dedicated virtual machine
- ⇒ The SSO servers are currently in independent SSO domains.
- ⇒ During the upgrade process, the administrators would like to combine their two SSO domains into a single one.

View the exhibit.



Referring to the exhibit, which upgrade scenario would accomplish this?

- A. 1. Upgrade the Denver SSP server to a 6.5 PSC. 2. Upgrade the Denver vCenter Server 5.5 to version 6.5. 3. Use the migration utility to upgrade the New York vCenter Server to 6.5. 4. Choose to join it to the Denver PSC.
- B. 1. Upgrade the Denver SSP server to a 6.5 PSC. 2. Use the migration utility to upgrade the New York SSO server. 3. Choose to join the existing SSO domain during the second upgrade. 4. Upgrade both of the vCenter Servers to 6.5.
- C. 1. Upgrade both of the SSO Servers to 6.5 PSCs. 2. Upgrade both of the vCenter Servers to 6.5. 3. Install a new 6.5 PSC in the same SSO domain as the Denver 6.5 PSC. 4. Repoint the New York vCenter Server to the newly-installed PSC.
- D. 1. Install a New York SSO 5.5 server in the same SSO domain as the Denver SSO server. 2. Repoint the New York vCenter Server to the newly-installed SSO server. 3. Upgrade both of the SSO Servers to 6.5 PSCs. 4. Upgrade both of the vCenter Servers to 6.5.

Suggested Answer: B

VCAPito Highly Voted 5 years, 7 months ago

Agree with D, SSOs need to be linked before the upgrade. See also <https://kb.vmware.com/s/article/2130433>
upvoted 10 times

abo2020 4 years, 5 months ago

I always trust VCAPito
upvoted 4 times

Grande Highly Voted 5 years, 7 months ago

I believe this question is referring to consolidating sso domain and upgrading

<http://cloudvm.es/consolidate-sso-domain-before-vsphere-6-5-upgrade> also see

<https://virtualinsanity.com/index.php/2017/02/25/vcsa-consolidating-sso-domains-and-ensuring-a-successful-migration/>

D is correct

upvoted 6 times

chafik Most Recent 4 years, 2 months ago

Correct Answer: D
upvoted 1 times

Kay02 4 years, 5 months ago

Agree With D;

After som searching finally found:

<https://vmarena.com/wp-content/uploads/2018/04/vSphere-6.5-Upgrade.pdf>

1. vSphere SSO Domain Consolidation: Here the upgrade process starts with vSphere SSO domain consolidation to enable single sign-on functionality. Keep in mind this can only be done on vSphere 5.5; once the first node is updated to vSphere 6.5 the ability to consolidate is no longer available. For scenarios that don't require vSphere SSO domain consolidation, begin with Step 2.

upvoted 1 times

🗨️ 👤 **admin007** 4 years, 7 months ago

C Seems to be correct answer

upvoted 1 times

🗨️ 👤 **labi** 4 years, 7 months ago

why? in 6.5 u can't repoint to other PSC which is in another vSphere domain .. I think, the correct answer is D.

upvoted 2 times

🗨️ 👤 **admin007** 4 years, 7 months ago

Thanks for pointing out , I see this now.

"In vSphere 6.5 the ability to repoint vCenters between sites is no longer available. When you deploy or upgrade to vSphere 6.5, your vCenter Servers are stuck to whatever PSC Site it currently resides."

upvoted 1 times

🗨️ 👤 **oud** 5 years, 2 months ago

I think c

upvoted 1 times

🗨️ 👤 **vmwarelive** 5 years, 4 months ago

Correct Answer:D

upvoted 3 times

An organization's security policy requires a design where the ESXi hosts will be manageable only through vCenter Server. Which two security configurations will help meet this requirement? (Choose two.)

- A. enable lockdown mode strict
- B. disable DCUI access
- C. enable lockdown mode normal
- D. disable shell access

Suggested Answer: BC

 **diegof1** Highly Voted 5 years, 8 months ago


Correct Answer:

AD

In the Lockdown mode strict mode the DCUI service is disabled, while in normal lockdown mode users of the DCUI.Access exception list can enter. Now, if the Shell is disabled, not even users in the exception list will be able to access this service.

<https://docs.vmware.com/en/VMware-vSphere/6.5/com.vmware.vsphere.security.doc/GUID-F8F105F7-CF93-46DF-9319-F8991839D265.html#GUID-F8F105F7-CF93-46DF-9319-F8991839D265>

upvoted 17 times

 **MohamedFouad** Highly Voted 5 years, 6 months ago

Correct Answer:

AD

lockdown Strict

Disable shell

upvoted 5 times

 **chafik** Most Recent 4 years, 2 months ago

Correct Answer: AD

upvoted 1 times

 **lotso** 4 years, 9 months ago

AD

<https://kb.vmware.com/s/article/1008077#:~:text=Strict%20Lockdown%20mode%3A,and%20Exception%20Users%20are%20defined.>

upvoted 1 times

 **oud** 5 years, 2 months ago

I think A C

When the host is running, available services depend on whether lockdown mode is enabled, and on the type of lockdown mode.

In strict and normal lockdown mode, privileged users can access the host through vCenter Server, either from the vSphere Web Client or by using the vSphere Web Services SDK.

Direct Console Interface behavior differs for strict lockdown mode and normal lockdown mode.

In strict lockdown mode, the Direct Console User Interface (DCUI) service is disabled.

In normal lockdown mode, accounts on the Exception User list can access the DCUI if they have administrator privileges. In addition, all users who are specified in the DCUI.Access advanced system setting can access the DCUI.

If the ESXi Shell or SSH is enabled and the host is placed in lockdown mode, accounts on the Exception Users list who have administrator privileges can use these services. For all other users, ESXi Shell or SSH access is disabled. Starting with vSphere 6.0, ESXi or SSH sessions for users who do not have administrator privileges are terminated.

upvoted 2 times

 **charithabuddhika** 5 years, 8 months ago

Answer is AD

upvoted 4 times

A company is in the process of deploying a modern video-streaming application.

- ⇒ The application is able to scale (expand and collapse) its streaming nodes in the form of CentOS 7.x 64bit virtual machines, based on demand.
- ⇒ This IO-Intensive application has a high CPU demand and generates a significant number of disk operations (IOPS).
- ⇒ To host the application, the company decided to implement a brand-new VMware cluster with vSphere 6.x.
- ⇒ The company would like a significant reduction in CPU utilization as well as possible increase in throughput.

Which virtual disk adapter should be recommended for the company's physical design?

- A. LSI Logical Parallel
- B. VMware Paravirtual
- C. BusLogic Parallel
- D. LSI Logic SAS

Suggested Answer: B

Reference -

<http://virtuallyhyper.com/2012/09/vcap5-dcd-objective-3-5-determine-virtual-machine-configuration-for-a-vsphere-5-physical-design/>

🗨️ 👤 **diegof1** Highly Voted 👍 5 years, 8 months ago
B is correct.

According to documentation, PVSCSI adapters are high performance that can provide high throughput and low CPU consumption.

<https://pubs.vmware.com/workstation-9/index.jsp?topic=%2Fcom.vmware.ws.using.doc%2FGUID-A0438F6C-6651-4A38-853A-0A7A494E23DF.html>
upvoted 10 times

🗨️ 👤 **chafik** Most Recent 🔍 4 years, 2 months ago
Correct Answer:B
upvoted 1 times

🗨️ 👤 **AdhamELHariry** 4 years, 3 months ago
VMware Paravirtual is correct.
upvoted 1 times

🗨️ 👤 **lotso** 4 years, 9 months ago
B.

VMware Paravirtual (aka PVSCSI) – this vSCSI controller is virtualization aware and was been designed to support very high throughput with minimal processing cost and is therefore the most efficient driver. In the past, there were issues if it was used with virtual machines that didn't do a lot of IOPS, but that was resolved in vSphere 4.1.

<https://blogs.vmware.com/vsphere/2014/02/vscsi-controller-choose-performance.html>
upvoted 1 times

A solution architect has been tasked with designing a new environment for a company's growing needs, and has obtained this information:

- ⇒ Uptime is critical during regular business hours when 95% of the transactions occur. Application uptime must be 99.9% during those hours.
- ⇒ In a true disaster, the business can withstand a day of data loss and half a day of downtime.
- ⇒ The company is one year into a 5-year contract with the co-lo data center.
- ⇒ The building that is currently occupied no longer has any floor space available, but the company has 3 empty racks of space. The co-lo can provide up to 11KVA of power per rack.
- ⇒ There are current contains with Dell to provide servers and with Cisco to provide the network components
- ⇒ The network team has standardized on an end-to-end 10Gb network.

Based on this information, what are two requirements for the new design? (Choose two.)

- A. RTO of 24 hours.
- B. RTO of 12 hours.
- C. The application must be available 99,9% during business hours.
- D. 11 KVA of power is available per rack.

Suggested Answer: AC

🗲️ 👤 **diegof1** Highly Voted 5 years, 8 months ago

Correct Answer:

BC

- Uptime is critical during regular business hours when 95% of the transactions occur. Application uptime must be 99.9% during those hours.

- In a true Disaster, the business can withstand a day of data loss 24H RPO (Recovery Point Objective) and half a day of downtime 12 H RTO (Recovery Time Objective).

upvoted 32 times

🗲️ 👤 **chafik** Most Recent 4 years, 2 months ago

Correct Answer:BC

upvoted 1 times

🗲️ 👤 **mo_rada** 4 years, 9 months ago

This is another obvious incorrect one that needs to be fixed.

Recovery Time is 12 hours. RPO = 24 hours

upvoted 1 times

A customer has storage arrays from two different storage vendors at two different sites. The customer wants to restore operations at the secondary site in the event of a disaster.

Which VMware technology must be used to meet this requirement?

- A. vSphere replication
- B. vSphere Data Protection
- C. array-based replication
- D. vSphere Fault Tolerance

Suggested Answer: A

  **VCAPito** Highly Voted 5 years, 7 months ago

Correct A, vSphere Replication is independent of the underlying storage <https://www.vmware.com/uk/products/vsphere/replication.html>
upvoted 10 times

  **chafik** Most Recent 4 years, 2 months ago

Correct Answer:A
upvoted 1 times

  **QualifiedExpert** 4 years, 5 months ago

Correct, the [B.] vSphere Data Protection is already discontinued.
upvoted 1 times

The ability to live-migrate all virtual machines between two clusters is a requirement in the customer's design. Which two clusters and EVC configurations will accomplish this? (Choose two.)

Cluster 1

- ESXi 6.0
- Intel Skylake CPUs
- EVC Enabled: AMD Opteron™ “Steamroller” Generation

Cluster 2

- ESXi 6.5
- AMD Steamroller CPUs
- EVC Enabled: AMD Opteron™ “Steamroller” Generation

A.

Cluster 1

- ESXi 6.5
- Intel® Broadwell CPUs
- EVC Disabled

Cluster 2

- ESXi 6.0
- Intel® Broadwell CPUs
- EVC Disabled

B.

Cluster 1

- ESXi 5.5
- AMD Piledriver CPUs
- EVC Enabled: AMD Opteron™ “Piledriver” Generation

Cluster 2

- ESXi 6.5
- AMD Steamroller CPUs
- EVC Enabled: AMD Opteron™ “Piledriver” Generation

C.

Cluster 1

- ESXi 6.5
- Intel Broadwell CPUs
- EVC Enabled: Intel® “Broadwell” Generation

Cluster 2

- ESXi 6.5
- Intel Sandy Bridge CPUs
- EVC Enabled: Intel® “Sandy Bridge” Generation

D.

Suggested Answer: AD

 **diegof1**  5 years, 8 months ago

BC is correct answer

B. Cluster 1

ESXi 6.5

Intel Broadwell CPUs

EVC Disabled

Cluster 2

ESXi 6.0

Intel Broadwell CPUs

EVC Disabled

C. Cluster 1

ESXi 5.5

AMD Piledriver CPUs

EVC Enabled: AMD Opteron™;Piledrive"; Generation

Cluster 2

ESXi 6.5
AMD Steamroller CPUs
EVC Enabled: AMD Opteron™ "Piledriver"; Generation

For Live-Migration processors must be from the same Intel or AMD vendor and have EVC enabled at the level supported by the two clusters together.

<https://docs.vmware.com/en/VMware-vSphere/6.5/com.vmware.vsphere.vcenterhost.doc/GUID-FEC87C0B-7276-4152-8EAA-915305E64FED.html>

<https://kb.vmware.com/s/article/1003212?lang=es>
upvoted 21 times

  **VCAPito**  5 years, 7 months ago

Agreed B & C



A) No - EVC doesn't work with different CPU vendors

B) Yes, but officially not possible - same CPU and family, EVC disabled would make it a Yes. However, Broadwell is not officially supported on 6.0 as per VCG

C) Yes - same CPU, different family but EVC enabled to the lowest family on both clusters (<https://kb.vmware.com/s/article/1005764> for EVC support)



D) No - same CPU, different family and EVC enabled on a newer family (Broadwell) on one cluster (so can't migrate from Broadwell to Sandy Bridge as some features won't be available)

upvoted 10 times

  **dax** 5 years, 5 months ago

not sure about B. vms can migrate from esxi ver 6 to 6.5 but not vice versa

upvoted 1 times

  **MoD** 5 years, 4 months ago

As long as you keep the VM's on the same hardware version (HW Version 11), you can live migrate (vMotion) from 6.0 to 6.5 and vice versa

upvoted 5 times

  **mo_rada** 4 years, 9 months ago

Correct

upvoted 1 times

  **chafik**  4 years, 2 months ago

BC is correct answer

B. Cluster 1

ESXi 6.5

Intel Broadwell CPUs

EVC Disabled

Cluster 2

ESXi 6.0

Intel Broadwell CPUs

EVC Disabled

C. Cluster 1

ESXi 5.5

AMD Piledriver CPUs

EVC Enabled: AMD Opteron™ "Piledriver"; Generation

Cluster 2

ESXi 6.5

AMD Steamroller CPUs

EVC Enabled: AMD Opteron™ "Piledriver"; Generation

upvoted 1 times

  **stillface** 4 years, 6 months ago

Where is D answer?

upvoted 1 times

  **DineshKu07** 4 years, 6 months ago

is there on top of D

upvoted 1 times

  **mo_rada** 4 years, 9 months ago

Confirmed. This is another question that needs to be corrected
upvoted 1 times

Which two types of workloads are efficiently consolidated when virtualize? (Choose two.)

- A. Workloads that do NOT require user input and are constantly processing large amounts of batched data.
- B. Workloads that will consume all available assigned resources.
- C. Workloads that are NOT CPU bound; most of their time is spent waiting for external events such as user interaction.
- D. Workloads that do NOT require access to specific physical resources such as a hardware dongle or graphics card.

Suggested Answer: AC

🗲️ 👤 **Grande** Highly Voted 5 years, 7 months ago

types of workloads are efficiently consolidated

C,D are most appropriate

upvoted 18 times

🗲️ 👤 **chafik** Most Recent 4 years, 2 months ago

Correct Answer: CD

upvoted 2 times

🗲️ 👤 **MonditoVCAPoSSJ4** 4 years, 8 months ago

Agree with C and D

upvoted 2 times

🗲️ 👤 **lotso** 4 years, 9 months ago

It seems C and D.

"For applications that are not CPU-bound, CPU virtualization likely translates into an increase in CPU use. If spare CPU capacity is available to absorb the overhead, it can still deliver comparable performance in terms of overall throughput."

<https://docs.vmware.com/en/VMware-vSphere/6.5/com.vmware.vsphere.resmgmt.doc/GUID-9986F393-4D36-48AB-B839-8FF25ECEE94.html>

upvoted 2 times

🗲️ 👤 **ddanilidis** 4 years, 11 months ago

I think A and D

upvoted 1 times

🗲️ 👤 **nobt** 4 years, 12 months ago

A,C are opposite to each other.

so, They will be efficiently consolidate.

upvoted 1 times

🗲️ 👤 **vmwarelive** 5 years, 4 months ago

Correct Answer: CD

upvoted 3 times

A company provides critical financial and statistical data for several major banks:

- ⇒ The company ensures that the bank's customer data is secure and that analytics data is available when needed.
- ⇒ Customers rely on this data before making crucial business and financial decisions.
- ⇒ Just a few minutes of downtime can result in loss of revenue and trust.
- ⇒ To meet high-availability requirements, the company's IT infrastructure components must be redundant.
- ⇒ The company established three data centers across the globe and interconnected them with high-speed WAN links.
- ⇒ Due to the rapid growth of its customers and their increasing demands, the compute, network, and storage were procured and managed by the company's enterprise system administrators group.

What are its two key challenges? (Choose two.)

- A. Data centers across the globe possess manageability problems.
- B. Availability of business applications must be ensured.
- C. Regulatory requirements must be met.
- D. Hardware-defined data centers have limitations.

Suggested Answer: BC

🗳️ 👤 **khanti** Highly Voted 5 years, 4 months ago

BC.

A - Manageability is a challenge, but not the key.

B - Availability is the primary goal and hence the key challenge

C - Supplying customer data globally necessarily dictates strict regulatory requirements

D - No mention in the question and too generic

upvoted 23 times

🗳️ 👤 **VCAPito** Highly Voted 5 years, 7 months ago

I would say A and D. B and C are requirements.

I've found that in most cases the word "must" indicates requirements, while "problems" and "limitations" are something that you have to find your way around, therefore challenges.

upvoted 7 times

🗳️ 👤 **mo_rada** 4 years, 9 months ago

This is another really objective question, but this logic makes sense to me.

upvoted 1 times

🗳️ 👤 **chafik** Most Recent 4 years, 2 months ago

Correct Answer: AD

upvoted 1 times

🗳️ 👤 **hotdogwater** 4 years, 4 months ago

A and D. C is invalid as it isn't even mentioned there are regulatory params. D is more of an opinion.

upvoted 1 times

🗳️ 👤 **some1saiddiegozawesome** 5 years, 4 months ago

CD

D >> Agreed with everyone.

C>> Because it is a challenge to comply to all regulatory compliance due to the fact that the datacenter is located in 3 different global locations. just a thought!

upvoted 3 times

🗳️ 👤 **vmwarelive** 5 years, 4 months ago

Correct Answer: AD

upvoted 2 times

🗳️ 👤 **MohamedFouad** 5 years, 6 months ago

AD correct ones

upvoted 2 times

A company has requested assistance with a new cross-site failover design to support business-critical applications.

- ⇒ It has two sites which are very well-connected, and latency is less than 5ms round trip.
- ⇒ The customer requires that its applications be restarted even in the event of a total site failure.
- ⇒ The applications must be kept online even when migrated during maintenance.

Storage arrays at either site support both synchronous and asynchronous replication.

▪

Which two options are accurate application requirements for this scenario? (Choose two.)

- A. The design must ensure continuous application uptime even during a total site failure.
- B. The design must prioritize application availability.
- C. The design must ensure application recoverability at the second site.
- D. The applications are latency-sensitive.

Suggested Answer: *CD*

🗲️ 👤 **diegof1** Highly Voted 5 years, 8 months ago

The Correct Answer is BC

In the event of a total failure of the site, inevitably there will be a restart of the VMs to the other site. Then you must prioritize the availability of the application and ensure the recovery of the application at the second site, when the design is done.

upvoted 21 times

🗲️ 👤 **VCAPito** Highly Voted 5 years, 7 months ago

Agree with B and C. Customer requires apps to be restarted, thus implying that a short downtime is tolerated

upvoted 7 times

🗲️ 👤 **chafik** Most Recent 4 years, 2 months ago

Correct Answer: BC

upvoted 1 times

🗲️ 👤 **vmwarelive** 5 years, 4 months ago

Correct Answer: BC

upvoted 5 times

A virtualization administrator has been tasked with migrating several business applications from physical to virtual. The administrator must also migrate the virtual machines from VMware Workstation to vSphere 6.5, using vCenter Converter Standalone 6.1. In this scenario, which two source types are supported? (Choose two.)

- A. powered-off Windows Server 2008 physical machine
- B. powered-on Windows Server 2000 Workstation virtual machine
- C. powered-off Windows Server 2008 Workstation virtual machine
- D. powered-on Windows Server 2008 physical machine

Suggested Answer: *CD*

  **VCAPito** Highly Voted 5 years, 7 months ago

C and D - https://www.vmware.com/pdf/convsa_61_guide.pdf
upvoted 13 times

  **chafik** Most Recent 4 years, 2 months ago

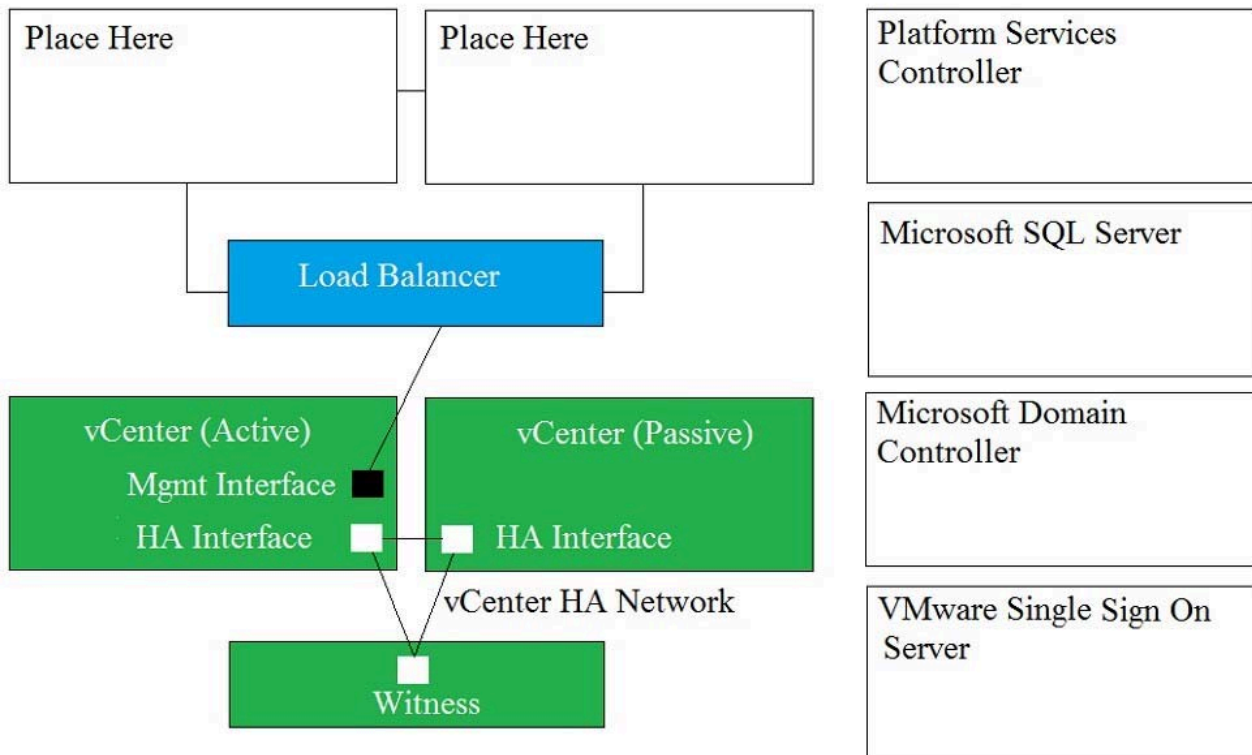
Correct Answer: CD
upvoted 1 times

DRAG DROP -

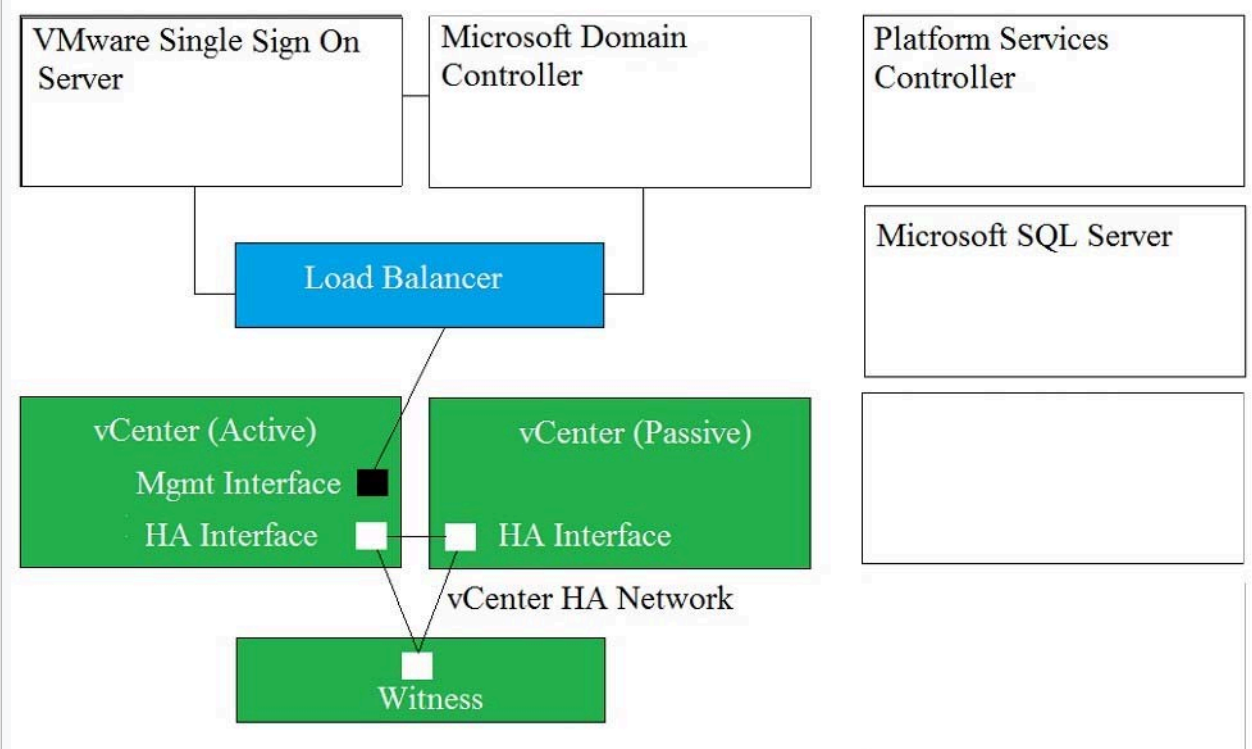
In the vCenter HA configuration below, drag the two correct components to the blank boxes in this diagram. The same component may be used more twice.

(Choose two.)

Select and Place:



Suggested Answer:



The Answer is Platform Services Controller en both boxes.

<https://kb.vmware.com/s/article/2147672>

upvoted 23 times

  **chafik** Most Recent 4 years, 2 months ago

The Answer is Platform Services Controller en both boxes.

upvoted 1 times

  **mdinisoae** 5 years, 3 months ago

vCenter HA was introduced with VMware Infrastructure 6.5 => the right answer is Platform Service Controller (see section "vCenter HA with an External Platform Services Controller")

<https://docs.vmware.com/en/VMware-vSphere/6.5/com.vmware.vsphere.avail.doc/GUID-F08F7501-FC06-4A83-806D-89CA489FA3AF.html>

VMware Single Sign On is part of VMware Infrastructure 5.x and doesn't have support for vCenter High Availability

upvoted 2 times


  **gskar2000** 5 years, 3 months ago

https://virtualizationreview.com/articles/2017/03/31/vcenter-high-availability-in-vsphere-6_5-tutorial.aspx

Answer is Platform Services Controller en both boxes.

upvoted 4 times

View the exhibit.



Referring to the exhibit, which appliance or device belongs in the square with the question mark?

- A. Firewall Appliance
- B. Load Balancer
- C. Platform Services Controller
- D. vCenter Server Appliance

Suggested Answer: B

hoo 5 years, 6 months ago

B is correct, same as question 22!

upvoted 7 times

DRAG DROP -

Match the business statement to its appropriate concept.

Select and Place:

Business statement	Concept
The operations must be automated and scable	Security
The transactions must be services in under 2ms.	Performance
The maximum recovery time objective is two hours.	Availability
The integrity of transactions cannot be compromised.	Manageability
The services cannot be interrupted for more than 5 minutes per year.	Recoverability

Suggested Answer:

Business statement	Concept
	The integrity of transactions cannot be compromised.
	The transactions must be services in under 2ms.
	The services cannot be interrupted for more than 5 minutes per year.
	The operations must be automated and scable
	The maximum recovery time objective is two hours.

 **VCAPito** Highly Voted 5 years, 7 months ago

Agree with provided answer:

Operation must scale -> Manageability

Transactions <2ms -> Performance

Max RTO 2h -> Recoverability

Transaction integrity -> Security

Service cannot be interrupted -> Availability

upvoted 21 times

 **MonditoVCAPoSSJ4** Most Recent 4 years, 8 months ago

Agree With the provided solution

upvoted 4 times

A developer is tasked with building an application to process shipping requests. The developer is consulting the vSphere team to determine failover options and performance best practices.

- ⇒ The development team is providing three physical ESXi hosts with 8 CPU cores and 256GB of RAM per host.
- ⇒ The developer does NOT know how many virtual machines they will require.

Which virtual machine (VM) sizing strategy will provide the highest level of uptime, individual VM performance, and failover capacity?

- A. A few large 8 vCPU VMs per host protected by vSphere HA.
- B. Many small 1 vCPU VMs participating in an OS level clustered application protected by vSphere HA.
- C. A few large 8 vCPU VMs per host protected by vSphere Fault Tolerance.
- D. Many small 1 vCPU VMs participating in an OS level clustered application protected by vSphere Fault Tolerance.

Suggested Answer: B

🗳️ 👤 **VCAPito** Highly Voted 5 years, 7 months ago

- A) Matching vCPU with the physical CPU cores provide best performance. However, we don't know how many VMs are required, therefore there's a risk of a high CPU overcommitment (especially in case of failover), which will impact on performance. Also, if the application doesn't need/use 8 cores, their use will be scheduled anyway therefore wasting resources (we don't know that, but if 1 vCPU is enough for the application then it would be counterproductive to configure it with 8)
- B) It'd be good practice to start a VM with lower vCPU possible and then increase if needed. This solution lower the risk of over-commitment, and guest clustering is supported for HA. This provides both guest and hardware failover. Might or might not have lower performance, requirements of the application (and host network/storage specifications) are not mentioned.
- C) No, FT supports up to 4 vCPU
- d) No, FT doesn't support OS (Microsoft at least) failover clustering: <https://docs.vmware.com/en/VMware-vSphere/6.5/vsphere-esxi-vcenter-server-651-setup-mscs.pdf> p.13

A bit tricky not knowing more details about application and OS, but I'd say B (assuming HW supports it), as it provides failover for both OS and HW failures and won't cause high CPU overcommitment (hence answering all 3 questions).

upvoted 21 times

🗳️ 👤 **dan80** 5 years ago

i will go with A - working with OS level clustered application is canceling option for best performance, there is no need for 100% uptime.

upvoted 2 times

🗳️ 👤 **abo2020** 4 years, 5 months ago

I always trust VCAPito but i will have to go with B this time my friend, No need for 100% uptime so no need for the OS level Clustered APP.

upvoted 1 times

🗳️ 👤 **duke_srg** Highly Voted 5 years, 7 months ago

B is the right one:

- guest level clustering + vSphere HA provides all the best result for 3 requirements
- can be scaled both up and out later when resource requirements will be clear

upvoted 12 times

🗳️ 👤 **ianol** 4 years, 5 months ago

The answer has to be either A or B.

B seems to make the most sense as you say.

It is not D -> 6.5 FT is limited to 4vCPU . This exam is for 6.5 only.

<https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/certification/vmw-vcap65-dcv-design-3v0-624-guide.pdf>

It is not C since it's not supported.

upvoted 1 times

🗳️ 👤 **chafik** Most Recent 4 years, 2 months ago

correct answer:B

upvoted 1 times

🗨️ 👤 **adrian_des** 4 years, 8 months ago

I would have to say "D" as there is no mention of shared storage therefore vSphere HA is out. The question says "VM sizing strategy...highest level of uptime, individual VM performance, and failover capacity" - all that can only be achieved with option "D". "A" won't work as there is no shared storage. "B" can work but it doesn't provide the "highest level of uptime". "C" is not possible with the exam version of vSphere 6.5 U1.

upvoted 1 times

🗨️ 👤 **labi** 4 years, 7 months ago

FT require shared storage as well: "Fault Tolerance avoids "split-brain" situations, which can lead to two active copies of a virtual machine after recovery from a failure. Atomic file locking on shared storage is used to coordinate failover so that only one side continues running as the Primary VM and a new Secondary VM is respawned automatically." <https://docs.vmware.com/en/VMware-vSphere/6.5/com.vmware.vsphere.avail.doc/GUID-623812E6-D253-4FBC-B3E1-6FBFDF82ED21.html>

upvoted 1 times

🗨️ 👤 **mdinisoae** 5 years, 3 months ago

In addition to the FT limitation (4/8 vCPU per VM and 8 vCPU per host), there is a statement on Course "VMware vSphere - Design Workshop", Module 8 - Virtual Machine Design:

Rightsize your VMs: Start with minimum guest resource requirements, then increase resources to improve application performance:

- Use vRealize Operations Manager to help you rightsize your VMs and make efficient use of host resource capacity
- Where possible, decrease resources to reduce wasted capacity.

upvoted 5 times

🗨️ 👤 **Shiv20** 5 years, 3 months ago

B is the right answer

upvoted 3 times

🗨️ 👤 **hoo** 5 years, 6 months ago

Reading the question my first answer was B. VCAPito also has a very good explanation which was my also close to my thought process while reviewing the question.

upvoted 3 times

🗨️ 👤 **diegof1** 5 years, 8 months ago

Answer is C (if VCAP 6.5 considers maximum and minimum of vSphere 6.5)

For vSphere 6.5 the maximum of vCPU per VM for FT is 4vCPU, maximum 8vCPU protected on a host and only 4 VMs protected. While in vSphere 6.5 U1 and 6.7 it supports up to 8vCPU per VM, but the same 8vCPU protected on a host and the same 4 VMs protected. It makes no sense to have an OS cluster that offers high availability and spend twice as many resources to protect them with FT in a 3 ESXi cluster with 8 cores. Since it is enough to protect them with HA.

Considering that the statement says the highest level of Uptime and Individual Performance then this covers it with FT and 8vCPU per VM. This way you could protect: A single VM per host (in vSphere 6.5U1 or higher) located on host1, the secondary copy would be on the other host2 and host3 would remain for high availability. For this reason, the answer is C (Assuming that VCAP 6.5 also takes into account the maximum values of 6.5 U1). If, on the other hand, it only takes into account only the maximum values of vSphere 6.5 U1 then the Answer would be D, since it could NOT protect 8vCPU (but if 4 VM of 1 vCPU), and the highest level of Uptime would achieve it with OS Level cluster and FT.

upvoted 1 times

🗨️ 👤 **VCAPito** 5 years, 7 months ago

It can't be C, as on 6.5 FT is limited to 4 vCPU with Ent+ (8 vCPU is on 6.7, exam refers to 6.5 GA, not even U3)

<https://docs.vmware.com/en/VMware-vSphere/6.5/com.vmware.vsphere.avail.doc/GUID-57929CF0-DA9B-407A-BF2E-E7B72708D825.html> .

upvoted 7 times

A company has requested a new vSphere 6.5 design that will allow it to finally break the 80 virtualization barrier by virtualizing its resource-intensive application.

- ⇒ The application is highly available by design and includes application-aware clustering software capable of operating as a fully distributed system.
- ⇒ The company's Application Version 2.0 consists of 386 small applications and middleware with non-persistent storage and 24 database virtual machines at each data center.
- ⇒ When coupled with a proper load balancing solution, this application can continue operating even with the loss of an entire data center, but the small applications and middleware tiers within a data center must exist within the same broadcast domain.
- ⇒ The database tier is tightly controlled with a firewall policy that only allows middleware tier access, and is replicated to other sites using a dedicated circuit.

Which two application requirements apply to this scenario? (Choose two.)

- A. The application will require the configuration of an IGMP stub and helper.
- B. Shared storage is required by the application clustering software.
- C. The application will require one large subnet.
- D. The application will require a method of balancing and recovering sessions between sites.

Suggested Answer: BD

🗨️ 👤 **VCAPito** Highly Voted 5 years, 7 months ago

Initially I thought B and D, with a doubt on C, but after I re-evaluated it, I think the correct answer is:

- A) No, as apps need to be on the same broadcast domain, not multicast
- B) No, applications have non-persistent storage and work as a fully distributed system (therefore replicating data rather than sharing it), and DB tier is replicated to the other site
- C) Yes, the 386 small applications and 24 middleware DB VMs need to be on the same broadcast domain (so at least a /23 is required, which is not that large, but more than the default /24). A dedicated circuit is used for replication, but applications must access the DB on the same subnet.
- D) Yes, LB required in the event of a data center loss

Some links: <https://www.freecodecamp.org/news/a-thorough-introduction-to-distributed-systems-3b91562c9b3c/> ,

https://en.wikipedia.org/wiki/Clustered_file_system

upvoted 20 times

🗨️ 👤 **Ophone** 4 years, 6 months ago

Fantastic

upvoted 1 times

🗨️ 👤 **vmwarelive** Highly Voted 5 years, 4 months ago

Correct Answer: CD

upvoted 7 times

🗨️ 👤 **chafik** Most Recent 4 years, 2 months ago

Correct Answer: CD

upvoted 1 times

DRAG DROP -

A company is a leading for an online travel booking system with over a \$1,000,000 turnover each day. The company wants to leverage VMware cloud solutions to consolidate, scale, and ensure high availability for all of its data centers.

Match each business requirement to its appropriate design concept.

Select and Place:

Business Requirement	Design Concept
Architecture must support data center-level resiliency.	Performance
Design should require less than 5 minutes to restore after failure.	Recoverability
WAN link between the data centers must be private.	Security
Data centers should be administered by a centralized console.	Manageability
Infrastructure must be properly provisioned to ensure application response time.	Availability

Suggested Answer:

Business Requirement	Design Concept
	Infrastructure must be properly provisioned to ensure application response time.
	Architecture must support data center-level resiliency.
	WAN link between the data centers must be private.
	Data centers should be administered by a centralized console.
	Design should require less than 5 minutes to restore after failure.

SuperTed Highly Voted 5 years, 8 months ago

Architecture must support datacenter level resiliency. -> Availability

Design should require less than 5 minutes to restore after failure -> This is a typical RTO requirement: Recoverability

WAN link between the data centers must be private -> Security

Data centers should be administered by a centralized console -> Manageability

Infrastructure must be properly provisioned to ensure application response time. -> Performance

upvoted 48 times

MonditoVCAPOSSJ4 4 years, 9 months ago

Agreed

upvoted 3 times

chafik Most Recent 4 years, 2 months ago

Architecture must support datacenter level resiliency. -> Availability

Design should require less than 5 minutes to restore after failure -> This is a typical RTO requirement: Recoverability

WAN link between the data centers must be private -> Security

Data centers should be administered by a centralized console -> Manageability

Infrastructure must be properly provisioned to ensure application response time. -> Performance
upvoted 1 times

A customer is using a vSphere APIs for Storage Awareness (VASA) compatible storage array. The VASA provider is published as a virtual appliance.

To ensure recoverability, where must the VASA provider and vCenter server virtual machines be stored?

- A. the VASA provider and vCenter Server will be placed in the standard datastore (VMFS, NFS).
- B. the VASA provider and vCenter Server will be placed on the vVol datastore.
- C. the vCenter Server will be placed on the vVol datastore and the VASA provider will be placed on the standard datastore (VMFS, NFS)
- D. the VASA provider will be placed on the vVol datastore and the vCenter Server will be placed on the standard datastore (VMFS, NFS)



Suggested Answer: A

  **VCAPito** Highly Voted 5 years, 7 months ago

Agree with A. VVOLs rely on VASA and vCenter, so putting those on a VVOLs could be catastrophic if any of the components fail
upvoted 21 times

  **dan80** Most Recent 5 years ago

The answer is D - You should not run VASA Provider on a VVOL datastore. Any management operation, including powering on a virtual machine that is on a VVOL, requires that VASA Provider be running. In addition, you would lose access to all VVOLs because VASA Provider would not be able to boot.
<https://library.netapp.com/ecmdocs/ECMP12405937/html/GUID-5B810B73-0233-4F3B-80BE-47A415D2F107.html>
upvoted 1 times

  **dan80** 5 years ago

sorry , answer is A.
upvoted 3 times

A customer wants to make its data available with a RPO of 10 minutes. Replication to the second data center will be done using the network. Which type of storage configuration should be used?

- A. NFS datastore on ESXi 6.0 with vSphere replication appliance 6.0
- B. VMFS datastore on ESXi 6.0 with vSphere replication appliance 6.5
- C. vSAN datastore on ESXi 6.0 with vSphere replication appliance 6.5
- D. VMFS datastore on ESXi 6.0 with vSphere replication appliance 6.0

Suggested Answer: D

🗨️ 👤 **SuperTed** Highly Voted 5 years, 8 months ago
Only C is correct.

The 5 minute RPO requires the source host to be ESXi 6.0 or later for VSAN, and ESXi 6.5 for other supported datastores.

<https://docs.vmware.com/en/vSphere-Replication/6.5/rn/vsphere-replication-65-release-notes.html>
upvoted 20 times

🗨️ 👤 **Darius_Th3D0G** 4 years, 10 months ago
Yes, C is correct.
<https://kb.vmware.com/s/article/2102453>
"In vSphere Replication 6.1.x, you can use the 5 minute Recovery Point Objective (RPO) if the target and the source sites use Virtual SAN storage. If you change the target site to a site that does not use Virtual SAN, you cannot set the RPO value to less than 15 minutes. The 5 minute RPO is limited to 100 VMs per Virtual SAN cluster."
upvoted 2 times

🗨️ 👤 **VCAPito** Highly Voted 5 years, 7 months ago
There is no hints about vCenter (6.5 required for vRep 6.5) or VSAN versions, therefore we must assume those info (as to be 6.5).
So, according to this, only C is valid: "The 5 minute RPO requires the source host to be ESXi 6.0 or later for VSAN, and ESXi 6.5 for other supported datastores" (from <https://docs.vmware.com/en/vSphere-Replication/6.5/rn/vsphere-replication-65-release-notes.html>)
upvoted 8 times

🗨️ 👤 **mariogg85** 5 years ago
Agree with you.
upvoted 1 times

🗨️ 👤 **silvanuwa** Most Recent 4 years, 2 months ago
<https://docs.vmware.com/en/vSphere-Replication/6.5/vsphere-replication-65-user.pdf>
Correct B
upvoted 1 times

🗨️ 👤 **chafik** 4 years, 2 months ago
Correct Answer: C
upvoted 1 times

🗨️ 👤 **vcp555** 4 years, 10 months ago
Answer is C

The 5 minute RPO requires the source host to be ESXi 6.0 or later for vSAN, and ESXi 6.5 for other supported datastores.

<https://docs.vmware.com/en/vSphere-Replication/8.3/com.vmware.vsphere.replication-admin.doc/GUID-9E17D567-A947-49CD-8A84-8EA2D676B55A.html>
upvoted 2 times

🗨️ 👤 **mdinisoae** 5 years, 3 months ago
Correct answer is C

In the past, the 5 min RPO was limited by the presence of vSAN datastore on source & target sites.

<https://kb.vmware.com/s/article/2102453>

"In vSphere Replication 6.1.x, you can use the 5 minute Recovery Point Objective (RPO) if the target and the source sites use Virtual SAN storage. If you change the target site to a site that does not use Virtual SAN, you cannot set the RPO value to less than 15 minutes. "

upvoted 2 times

  **mdinisoae** 5 years, 3 months ago

In the past, the 5 min RPO was limited by the presence of vSAN datastore on source & target sites.

<https://kb.vmware.com/s/article/2102453>

"In vSphere Replication 6.1.x, you can use the 5 minute Recovery Point Objective (RPO) if the target and the source sites use Virtual SAN storage. If you change the target site to a site that does not use Virtual SAN, you cannot set the RPO value to less than 15 minutes. "

upvoted 1 times

  **smirgole** 5 years, 4 months ago

I probably don't agree with most of you, but I think B is still correct (and C as well):

you have mentioned the following from the official doc:

The 5 minute RPO requires the source host to be ESXi 6.0 or later for VSAN, and ESXi 6.5 for other supported datastores.

this is only applicable to VR 8.1 not to VR 6.5, here the statement regarding RPO from the 6.5 doc regarding 5min RPO :

You can use the 5 minute Recovery Point Objective (RPO) if the target and the source sites use VMFS 6.0, VMFS 5.x, NFS 4.1, NFS 3, VVOL, or Virtual SAN 6.0 storage and later.

vSphere Replication 6.5 displays the 5 minute RPO setting when the target and the source site use VMFS 6.0, VMFS 5.x, NFS 4.1, NFS 3, VVOL, or Virtual SAN 6.0 storage and later.

Link to the doc :

<https://docs.vmware.com/en/vSphere-Replication/6.5/com.vmware.vsphere.replication-admin.doc/GUID-9E17D567-A947-49CD-8A84-8EA2D676B55A.html>

upvoted 1 times

  **vmwarelive** 5 years, 4 months ago

Correct Answer: C

upvoted 3 times

  **dax** 5 years, 5 months ago

answer is B. as minimum. no vsan requirements.

In vSphere Replication 6.5.x and 8.x, you can use the 5 minute Recovery Point Objective (RPO) if the target and the source sites use VMFS 6.0, VMFS 5.x, NFS 4.1, NFS 3, VVOL and Virtual SAN 6.0 storage and later. The 5 minute RPO can be applied to a maximum of 100 VMs on VMFS 6.0, VMFS 5.x, NFS 4.1, NFS 3 and Virtual SAN 6.0 storage and later. The maximum for VVOL datastore is 50 VMs.

please disregard my other posts...

upvoted 2 times

  **dax** 5 years, 6 months ago



Should be D. I do replication from source esxi 5.5 goin to vsan cluster here in my work. Using vcenter 6.5.

upvoted 1 times

  **dax** 5 years, 6 months ago

yes C is correct. please disregard above comment

upvoted 2 times

  **diegof1** 5 years, 8 months ago

B and C are correct

The vSphere Replication 6.0 version supports a minimum of 15 minutes of RPO, while version 6.5 supports RPO from 5 minutes under some conditions, which includes support for VMFS6, VMFS 5.x, VSAN among others. However, because VSAN is a more expensive solution, I would leave only option B, as the correct answer.

<https://kb.vmware.com/s/article/2102453>

upvoted 5 times

  **e_amr** 4 years, 5 months ago

only C

upvoted 1 times

DRAG DROP -

A leading steel manufacturer relies on SAP for purchase, sales, and invoice processing.

⇒ It is planning to virtualize its servers to reduce CAPEX and OPEX.

⇒ However, its CIO is concerned about the availability, performance, manageability, recoverability, and security for the SAP database and ERP instance.

Match the business requirement with the appropriate design concept.

Select and Place:

Business Requirement	Design Concept
The application must scale within two hours.	Performance
The application must support clustering between physical and virtual machines.	Recoverability
The application must be restored in 30 minutes in case of disaster.	Security
The application data must be encrypted at rest.	Manageability
The application must generate 200 invoices concurrently.	Availability

Suggested Answer:

Business Requirement	Design Concept
	The application must generate 200 invoices concurrently.
	The application must be restored in 30 minutes in case of disaster.
	The application data must be encrypted at rest.
	The application must support clustering between physical and virtual machines.
	The application must scale within two hours.

 **diegof1** Highly Voted 5 years, 8 months ago

Correct Answer:

The application must scale... --> Manageability


The application must support... --> Availability

The application must be restored... --> Recoverability

The application data must be encrypted... --> Security

The application must generate 200... --> Performance

upvoted 24 times

 **stysonchanvilla** Highly Voted 5 years, 4 months ago

Application must scale simply means making the application serve more users.

So, this should be related to the availability of the application.

Hence,

Correct Answer:

The application must scale... --> Availability


The application must support... --> Manageability

The application must be restored... --> Recoverability

The application data must be encrypted... --> Security

The application must generate 200... --> Performance

upvoted 22 times

  **lotso** 4 years, 9 months ago

I agree

upvoted 2 times

  **ianol** 4 years, 5 months ago

stysonchanvilla's answer is correct.

Availability and scalability are practically synonymous in that scalability enhances availability.

Scalability is the ability of a system to handle the increase in demand without impacting the application's performance or availability. When the demand is too high and there are not enough resources, then it impacts performance

<https://networklessons.com/cisco/evolving-technologies/cloud-performance-scalability-and-high-availability>

upvoted 1 times

  **chafik**  4 years, 2 months ago

Correct Answer:

The application must scale... --> Availability

The application must support... --> Manageability

The application must be restored... --> Recoverability

The application data must be encrypted... --> Security

The application must generate 200... --> Performance

upvoted 1 times

  **admin007** 4 years, 7 months ago

I go with stysonchanvilla.

Availability -> Everything to make sure system is available

Manageability -> What ever is needed to achieve the availability

Application must scale with in 2 hrs so that application is available



By doing Physical and virtual machine clustering we can get that application available in 2 hrs so it's Manageability

upvoted 3 times

  **CristianSalgueiro** 4 years, 8 months ago

VCAPito and Diegof1 --> I agree with both

upvoted 2 times

  **VCAPito** 5 years, 7 months ago

Agreed with Diego. Availability is about having redundancy (cluster), manageability about having a solution easy to deploy, maintain, update and upgrade (scale).

Although in this example the application must scale could also be availability and performance (system could be unavailable/slow if not), but manageability is a better fit

upvoted 4 times

  **MonditoVCAPoSSJ4** 4 years, 9 months ago

Please bear in mind that you are not implementing that clustering it is only asking for it to be supported.

I go with stysonchanvilla's answer

upvoted 1 times