PassQueen
a.k.a dumpsqueen.com

LSI SVM5 Sales Consultant

LSI L50-503

Version Demo

Total Demo Questions: 12

Total Premium Questions: 129

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Volume A	61
Topic 2, Volume B	48
Total	129

QUESTION NO: 1

An SVM prospect just called requesting that you virtualize their data center using VMware. What are three ways that SVM's volume management capabilities will help this customer?

(Choose three.)

A. SVM allows users to quickly and efficiently create and manage volumes for each of the virtual machines.

B. SVM eliminates the need to deal with ongoing LUN management at the array level.

C. SVM will prevent the virtual servers from running out of space.

D. SVM allows storage be provisioned and reallocated after use with the same simplicity as the virtual servers.

E. SVM will size the volumes to the actual data size.

ANSWER: A B D

Explanation:

SVM volume management complements VMware virtualization by applying the same abstraction and operational flexibility to storage that VMware applies to compute resources. “SVM allows users to quickly and efficiently create and manage volumes for each of the virtual machines.” is correct because storage can be presented and managed in logical units that align with virtual-machine and application requirements, rather than requiring a separate, rigid physical-storage workflow for every deployment.

“SVM eliminates the need to deal with ongoing LUN management at the array level.” is correct because SVM provides a virtualization layer for storage administration. This centralizes the logical management of capacity and reduces dependence on the underlying array’s individual LUN configuration as virtualized workloads change.

“SVM allows storage be provisioned and reallocated after use with the same simplicity as the virtual servers.” is also correct. Rapid provisioning and reassignment of logical capacity supports the dynamic lifecycle of VMware workloads, including new deployments, migrations, and changing application demand. VMware similarly describes storage virtualization as abstracting physical storage resources into logical storage presented to virtual machines. See VMware vSphere and vSphere Storage documentation.

QUESTION NO: 2

Your customer explains that previously they used a host-based mirroring product. After their most recent failure, the remote site was unusable due to corruption. How would multiMirror handle this problem?

A. multiMirror will not corrupt data based on its internal ECC correction code.

B. Host-based mirroring uses IP which can be less reliable than FC.

C. multiMirror uses the reliability of the RAID devices to accomplish a much more reliable copy.

D. Snapshot-enabled mirroring allows multiMirror to retain several recovery points at the remote site for rollback purposes.

ANSWER: D

Explanation:

Snapshot-enabled mirroring allows multiMirror to retain several recovery points at the remote site for rollback purposes. This directly addresses the customer’s concern: a mirror can faithfully replicate application-level corruption, deletion, or other bad changes as well as valid writes. By maintaining multiple point-in-time snapshots at the recovery location, multiMirror provides a recovery history rather than only the most current mirrored image. Following discovery of corruption, an administrator can identify a snapshot created before the bad data was replicated and roll back or recover data from that known-good recovery point. This separates replication from recoverability: the remote copy remains current for normal disaster-recovery needs, while retained snapshots provide an additional layer of protection against logical failures. Snapshot-based recovery points are a recognized storage-data-protection approach because they preserve a point-in-time view that can be used to restore data after an unwanted change. The SNIA definition of a snapshot and related storage terminology is available in the SNIA

Online Dictionary. For broader storage and data-protection product information from the LSI successor organization, see Broadcom Storage.

QUESTION NO: 3

How can multiCopy help when trialing performance requirements of applications running on new servers and/or storage?

A. multiCopy allows applications to be copied and made available to the same server with no application downtime and allows the performance increases to the server to be measured.

B. multiCopy asynchronously mirrors data between arrays of heterogeneous storage, and testing can then be carried out on one side of the mirror.

C. multiCopy allows applications to be copied between different tiers of heterogeneous storage and then made available to test servers for performance testing.

D. multiCopy allows servers to read and write data from and to anywhere in the SAN allowing servers to write to the fastest disk available.

ANSWER: C

Explanation:

multiCopy allows applications to be copied between different tiers of heterogeneous storage and then made available to test servers for performance testing. This capability is valuable in a proof-of-concept or hardware-refresh trial because it creates a separate, usable copy of the application data on the target storage tier or array. The copied data can be presented to a new server, allowing the organization to run realistic workload and performance tests against the proposed infrastructure without requiring the production application to be moved permanently first. Since SANsymphony virtualizes heterogeneous storage, the source and destination can be different storage systems or tiers, enabling a meaningful comparison of application behavior, throughput, and latency on the candidate platform. The test can therefore use representative data and application conditions while preserving the production environment for normal business operations. This aligns with DataCore SANsymphony's storage-virtualization approach, which centrally manages and provisions capacity across mixed storage hardware, and with its data-services positioning for non-disruptive data mobility and copies. See DataCore SANsymphony and DataCore resources.

QUESTION NO: 5

Which three actions will help mitigate budget cuts? (Choose three.)

A. Migrate data to less expensive tiers.

B. Deploy rapid application recovery.

C. Avoid Tier 1 storage where possible.

D. Stripe every volume.

E. Implement server and storage virtualization.

ANSWER: A C E

Explanation:

Migrating data to less expensive tiers, avoiding Tier 1 storage where possible, and implementing server and storage virtualization are complementary ways to reduce infrastructure spending without treating all workloads as though they require the same performance and availability level. Data tiering places infrequently accessed or lower-value information on lower-cost capacity while reserving premium media for workloads with genuine performance requirements. This improves the value obtained from existing storage investments and can defer new high-end capacity purchases. IBM describes storage tiering as moving data between storage classes according to access and performance needs: IBM: Storage tiering.

Avoiding Tier 1 storage where it is not justified applies the same principle at the purchasing and provisioning stage: expensive, highly resilient storage is aligned to business-critical applications, while suitable workloads use more economical tiers. Server and storage virtualization further mitigates budget pressure by consolidating resources, increasing utilization, simplifying allocation, and enabling capacity to be assigned or moved more efficiently across workloads. Virtualization can reduce hardware footprint and operating overhead while preserving flexibility. VMware explains that virtualization abstracts physical infrastructure resources so they can be pooled and managed more efficiently: VMware: What is virtualization?.

QUESTION NO: 6

Your customer calls to discuss their DPMs and their consolidation project's impact on their DPM port count. They are moving from a pair of dual controllers with four ports to one dual controller with eight ports. Which statement accurately describes what they should do?

A. They will need four more ports on each DPM.

B. They will need four fewer ports on each DPM.

C. They will need eight more ports on each DPM.

D. They will not need to change their DPM port count.

ANSWER: D

Explanation:

They will not need to change their DPM port count. The relevant calculation is the aggregate number of controller ports that the DPMs must support, rather than the number of controller enclosures or controller pairs. Before consolidation, the customer has two dual-controller systems with four ports per controller: four controllers multiplied by four ports equals sixteen ports. After consolidation, the customer has one dual-controller system with eight ports per controller: two controllers multiplied by eight ports also equals sixteen ports. Since the total port requirement remains sixteen, the existing DPM port count remains appropriately sized. This illustrates an important consolidation principle: a change in physical system count does not by itself change the required managed-port capacity; capacity should be reassessed using the total ports presented by all controllers in the resulting configuration. Product documentation and legacy LSI resources are available through Broadcom's Support Portal and Download Search.

QUESTION NO: 7

Capacity trapped in allocation is a result of which two administrative behaviors? (Choose two.)

A. over allocation to increase the number of operational requests

B. over allocation of volumes to minimize application disruptions

C. over allocation of volumes for rapid application recovery

D. over allocation to minimize frequency of file system size expansions

ANSWER: B D

Explanation:

Capacity becomes trapped in allocation when administrators assign more storage to a volume than the application is actively consuming, leaving capacity unavailable for other workloads even though it is not storing useful data. **over allocation of volumes to minimize application disruptions** is a common cause: teams often provision generous headroom in advance to avoid an urgent storage-change request, a maintenance window, or any risk that an application will run out of space. That reserved capacity is administratively committed even when it remains unused.

over allocation to minimize frequency of file system size expansions similarly creates stranded capacity. File-system expansions involve planning, approvals, and operational effort, so administrators may allocate substantial anticipated growth at the start. Although this can simplify short-term administration, it can leave large amounts of physical capacity assigned to

volumes or file systems that have not yet grown into it. Storage virtualization and thin-provisioning practices help organizations present logical capacity while consuming physical capacity as data is actually written, improving utilization and reducing this allocation gap. IBM's overview of thin-provisioned volumes and capacity allocation concepts is available in the IBM Storage Virtualize documentation; related storage-efficiency guidance is available from IBM Storage FlashSystem.

QUESTION NO: 8

Department A has run out of storage space while Department B has excess space while sharing the same SAN. In which two ways does SVM solve this typical customer problem? (Choose two.)

A. centralized storage management through virtualization of storage and thin provisioning

B. migration of storage

C. centralized storage management through pooling of storage and dynamic allocation

D. mirroring of storage

ANSWER: A C

Explanation:

Centralized storage management through virtualization of storage and thin provisioning addresses stranded capacity by presenting physical storage resources through a common logical management layer. Virtualization decouples the storage volumes delivered to departments from the particular physical arrays on which capacity resides, allowing available capacity to be used more flexibly. Thin provisioning further helps by allocating physical capacity as data is actually written rather than reserving all requested capacity in advance. This improves utilization and can defer unnecessary storage purchases.

Centralized storage management through pooling of storage and dynamic allocation solves the same business problem at the capacity-management level. A shared pool aggregates available capacity, and dynamic allocation enables storage administrators to assign or expand usable capacity for the department that needs it without treating each department's originally assigned physical storage as an isolated, fixed resource. Together, these capabilities make excess capacity available where demand exists while maintaining centralized control of the SAN environment. These concepts align with the storage-virtualization and thin-provisioning definitions maintained by the Storage Networking Industry Association and with Broadcom's enterprise storage portfolio information: SNIA: What Is Storage Virtualization? and Broadcom Storage Products.

QUESTION NO: 9

With intelligent scripts (App-Pack) running on the application server, which type of snapshots are created using multiView?

A. crash consistent

B. application consistent

C. user consistent

D. snapshot consistent

ANSWER: B

Explanation:

application consistent is correct. App-Pack uses intelligent, application-aware scripts on the application server to coordinate snapshot creation with the state of the protected workload. Before multiView creates the storage snapshot, these scripts can prepare or quiesce the application and its data, helping ensure that pending writes are handled and that application data is captured in a recoverable state. The resulting snapshot is therefore suitable for application-aware recovery, rather than merely reflecting whatever data happened to be on disk at an arbitrary instant.

This distinction is especially important for transactional workloads, where data files, logs, and in-memory activity must be coordinated for a reliable restore. By integrating application-side processing with the multiView snapshot operation, App-Pack adds the consistency controls required for recovery of supported applications. This aligns with the general storage-

industry principle that application-aware snapshot workflows coordinate the application and storage layers to provide recoverable copies. See Broadcom's storage product information and Microsoft's guidance on Volume Shadow Copy Service for background on application-coordinated snapshot consistency.

QUESTION NO: 10

Which software is used to manage EMC Invista?

A. EMC Control Center (ECC)

B. Navisphere

C. Element Manager

D. Replication Manager

ANSWER: C

Explanation:

Element Manager is the dedicated management application for EMC Invista. Invista was EMC's network-based storage-virtualization platform, designed to present and control heterogeneous back-end storage through virtualization engines. Administrators used Element Manager to perform the platform-specific operational tasks required to configure and administer an Invista environment, including discovering and managing virtualization components, creating and presenting virtual volumes, configuring paths, and reviewing system status. In other words, Element Manager provided the central administration interface for the Invista infrastructure rather than merely managing an individual array or a separate data-protection workflow.

This distinction matters in a sales-consultant context because Invista's value proposition depended on abstraction and control across supported storage resources. The management tool therefore had to understand the Invista virtualization layer, its engines, and its virtual-volume configuration. Element Manager was built for precisely that role. Dell Technologies maintains support and product-documentation access through its Support portal; historical EMC product material, including legacy platform documentation, can also be located through the Dell Technologies knowledge base.

QUESTION NO: 11

Your customer is considering purchasing SVM, but has a question about backups. You explain that SVM allows for zero-impact backups. What are three benefits of zero-impact backups that you can describe to your customer? (Choose three.)

A. SVM zero-impact backups create compressed files so that operations resume after an unexpected event quickly and non-disruptive to normal operations.

B. SVM servers coordinate backup operations so that applications experience only a few seconds of interruption during normal backup procedures.

C. multiView snapshots as part of a backup solution maximize uptime by reducing RPO and RTO of key business applications.

D. SVM zero-impact backups improve business continuity by enabling rapid resumption of operations after unexpected events.

E. SVM zero-impact backups eliminate the backup window, so backups can occur at any time while applications remain online.

ANSWER: C D E

Explanation:

Zero-impact backups are designed to protect data without requiring production applications to be taken offline or placed into a conventional backup window. "multiView snapshots as part of a backup solution maximize uptime by reducing RPO and RTO of key business applications" is a benefit because point-in-time snapshots provide recoverable versions of application

data. This supports a lower recovery-point objective by preserving recent data states and a lower recovery-time objective by enabling fast restoration from a local snapshot or replicated copy. “SVM zero-impact backups improve business continuity by enabling rapid resumption of operations after unexpected events” is also correct: recoverable snapshot copies are a core business-continuity capability, allowing operations to resume quickly following data loss, corruption, or a service disruption. Finally, “SVM zero-impact backups eliminate the backup window, so backups can occur at any time while applications remain online” captures the operational value of non-disruptive protection. Administrators can schedule protection around business and recovery requirements rather than around periods when applications must be unavailable. Snapshot-based data protection is commonly used to retain consistent recovery points while minimizing impact on running workloads. See NetApp’s overview of Snapshot policies and its guidance on RPO planning for local data protection.

QUESTION NO: 12

What are three cost structures affected by SVM? (Choose three.)

actual data

B. copies for protection

C. RAID protection

D. planned downtime

E. allocated but unused capacity

Answer: B,D,E

Explanation:

A. copies for protection

B. RAID protection

C. planned downtime

D. allocated but unused capacity

ANSWER: A C D

Explanation:

SVM affects the cost of **copies for protection**, **planned downtime**, and **allocated but unused capacity** because storage virtualization centralizes capacity management and makes storage resources more flexible across the environment. Virtualized storage can reduce the number of separate protection copies required by enabling more efficient use of shared storage services and policy-based data protection. It also supports nondisruptive data movement and maintenance activities, helping organizations avoid the operational and business costs associated with scheduled outages. In addition, pooling capacity improves utilization: capacity that would otherwise be reserved on an individual array or application basis can be made available to workloads that need it. These effects address both capital expense, through deferred capacity purchases and reduced duplicate storage, and operational expense, through lower administration and downtime costs. This is consistent with the storage-virtualization value proposition of improving utilization, availability, and operational efficiency. See Broadcom’s information on its enterprise storage portfolio at Broadcom Storage and its overview of storage software and management capabilities at Broadcom Storage Software.