DUMPSBOSS.

Advanced Cloud Architecture

SOA C90.05

Version Demo

Total Demo Questions: 5

Total Premium Questions: 50

Buy Premium PDF

https://dumpsboss.co

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

QUESTION NO: 1

Which of the following statements about the system that is established by the application of the Storage Workload Management pattern are? (Choose all that apply)

A. The system can distribute workloads across multiple cloud storage devices at runtime.

B. The system supports the grouping of cloud storage devices to improve the balanced distribution of data.

C. The system can use the automated scaling listener to equalize the workload across cloud storage devices.

D. The system can incorporate the use of the LUN migration program to distribute the runtime workload.

ANSWER: A B C D

Explanation:

The Storage Workload Management pattern establishes a storage-management architecture that monitors storage demand and actively distributes I/O activity and data placement among available cloud storage devices. Consequently, the system can distribute workloads across multiple cloud storage devices at runtime, rather than binding all activity to one storage resource. It also supports grouping cloud storage devices into logical pools, which provides the scope needed to balance data and workload placement across comparable resources.

An automated scaling listener can monitor storage-related utilization or workload conditions and initiate scaling or redistribution actions. Used with the workload-management capability, this enables the system to equalize demand across cloud storage devices as conditions change. LUN migration is likewise an applicable mechanism because it moves logical storage volumes between physical storage resources while preserving the logical unit presented to consumers. This permits runtime redistribution of storage workloads and data without requiring applications to be redesigned around a particular physical device. These capabilities align with cloud computing's resource-pooling and elasticity characteristics described by NIST SP 800-145 and with storage virtualization practices, including logical-unit migration, documented by IBM Storage Virtualize documentation.

QUESTION NO: 2

In a cloud architecture resulting from the application of the Cloud Bursting compound pattern, the _____ mechanism is used to maintain synchronizing between a cloud consumer's on-premise and cloud-based IT resources.

A. hypervisor

B. resource pooling

C. automated scaling listener

D. resource replication

ANSWER: D

Explanation:

Resource replication is correct because the Cloud Bursting compound pattern expands an on-premise environment into cloud-based resources when demand exceeds local capacity. For this hybrid arrangement to work reliably, relevant IT resources, application components, configuration data, and/or operational data need to be kept aligned across the on-premise and cloud environments. The resource replication mechanism provides this synchronization by maintaining corresponding copies of resources and propagating updates between the primary on-premise environment and the cloud-based extension. This allows cloud-hosted capacity to operate as a coordinated continuation of the consumer's existing environment during a burst, rather than as an isolated deployment. The exact replication approach can be synchronous or asynchronous depending on workload requirements, particularly the acceptable recovery-point objective and latency tolerance. In general cloud architecture practice, replication supports availability, continuity, and hybrid workload operation by ensuring that the alternate or extended environment has the required current state. For background on cloud bursting as a hybrid-cloud scaling approach, see IBM Think: Cloud Bursting. For guidance on replication as a mechanism for maintaining data consistency and resiliency across environments, see AWS: Data Replication for Disaster Recovery.

QUESTION NO: 3

The following statements describe how the automated scaling listener mechanism can be used as part of the system that is established by the application of the Elastic Resource Capacity pattern.

Which of these statements are false? (Choose Two)

A. The automated scaling listener notifies the intelligent automation engine to execute a script that instructs the hypervisor to allocate additional resources.

B. The automated scaling listener assigns additional CPUs to the virtual server, based on the workload processing requirements.

C. The automated scaling listener monitors message traffic between cloud service consumers and the cloud services being hosted by the virtual server to which the pattern has been applied.

D. The automated scaling listener withdraws previously assigned CPUs when they are no longer needed and returns them to the shared resource pool.

ANSWER: B D

Explanation:

The false statements are “The automated scaling listener assigns additional CPUs to the virtual server, based on the workload processing requirements.” and “The automated scaling listener withdraws previously assigned CPUs when they are no longer needed and returns them to the shared resource pool.” In the Elastic Resource Capacity pattern, the automated scaling listener has an observation and notification role. It detects workload conditions and threshold events, then communicates those events to the intelligent automation engine. The intelligent automation engine is the component that applies scaling policy, selects the required action, and invokes the relevant scripts or management interfaces. Those scripts direct the hypervisor or cloud resource-management platform to allocate or reclaim processing capacity. Keeping monitoring separate from decision-making and infrastructure control supports loose coupling, clearer operational responsibilities, and reusable automation logic. This separation is also reflected in production autoscaling platforms: monitoring signals and rules trigger an autoscale action, while the cloud platform performs the resource change. See Microsoft Azure Autoscale overview and AWS EC2 Auto Scaling documentation. Therefore, directly assigning or withdrawing CPUs is not a responsibility of the automated scaling listener itself.

QUESTION NO: 4

The system that results from the application of the Elastic Disk Provisioning pattern uses the thin-provisioning program to create virtual server disks.

A. True

B. False

ANSWER: A

Explanation:

True is correct. The Elastic Disk Provisioning pattern addresses the need to provide virtual servers with storage that can be allocated efficiently and expanded as demand grows, rather than reserving the entire maximum disk capacity at initial deployment. It uses a thin-provisioning program to create virtual server disks with a logical capacity that may be greater than the physical storage initially consumed. Physical blocks are allocated as data is written, allowing the cloud environment to conserve storage capacity while still presenting the required disk size to the virtual server.

This approach supports elasticity because storage can be provisioned and extended with less up-front allocation and reduced operational interruption. Thin provisioning is therefore a core enabling mechanism for the pattern’s goal of on-demand, scalable virtual-disk capacity. VMware similarly describes thin provisioning as allocating physical storage space to a virtual disk only as the disk consumes it, rather than allocating all configured capacity immediately. See VMware vSphere Storage documentation and Microsoft Disk Management overview.

QUESTION NO: 5

The cloud architecture that results from the application of the _____ pattern relies on the _____mechanism to synchronize cloud storage devices.

A. Synchronized Operating State. resource cluster

B. Redundant Storage, resource replication

C. Zero Downtime, hypervisor

D. Service Load Balancing, failover system

ANSWER: B

Explanation:

Redundant Storage, resource replication is correct because the Redundant Storage pattern establishes multiple copies of cloud-resident data on separate storage devices or locations to improve availability and resilience. Merely maintaining multiple storage devices is insufficient if their contents can diverge; the copies must be kept consistent as data is created, changed, or deleted. The resource replication mechanism performs this synchronization by propagating data and updates between the primary storage resource and its redundant counterpart(s). This allows a replicated copy to remain usable if a storage device, availability zone, or other underlying storage component becomes unavailable.

Replication is therefore the operational mechanism that realizes the pattern's redundancy objective. In practical cloud platforms, replication can be synchronous or asynchronous, and implementation details determine the resulting consistency and recovery characteristics. For example, Amazon S3 Replication automatically and asynchronously copies objects and associated metadata to designated destination buckets, illustrating how replication maintains redundant cloud storage resources. The pattern's emphasis is on preserving accessible copies of data through synchronized replication, rather than simply deploying additional storage capacity. See Amazon S3 Replication documentation and Arcitura Cloud Computing Patterns.