DUMPSQUEEN

Cloud Architecture Lab

SOA C90.06

Version Demo

Total Demo Questions: 5

Total Premium Questions: 15

Buy Premium PDF

https://dumpsqueen.com

support@dumpsqueen.com

QUESTION NO: 1

A cloud provider has two cloud environments (Cloud A and Cloud B) that are in different geographical regions. Cloud Service A resides in Cloud A. A redundant implementation of Cloud Service A resides in Cloud B. An automated scaling listener is used in Cloud A to automatically balance the workload of requests for Cloud Service A across the two redundant implementations. Cloud Service A is required to access Cloud Storage Device A, which also resides in Cloud A. A redundant implementation of Cloud Storage Device A is located in Cloud B. A failover system is used to ensure that if the Cloud Storage Device A implementation in Cloud A fails, the Cloud Storage Device A implementation in Cloud B takes its place.

Cloud Service Consumer A is owned by Organization A. Cloud Service Consumer A sends a request to Cloud Service A (1). The automated scaling listener intercepts the request and directs it to the Cloud Service A implementation in Cloud A (2). This Cloud Service A implementation attempts to access Cloud Storage Device A in Cloud A, but Cloud Storage Device A fails (3). The failover system redirects the request to Cloud Storage Device A in Cloud B (4). Cloud Service Consumer B sends a request to Cloud Service A (5). The automated scaling listener intercepts the request and directs it to the Cloud Service A implementation in Cloud B (6). This Cloud Service A implementation accesses Cloud Storage Device A in Cloud B to fulfill the request (7).

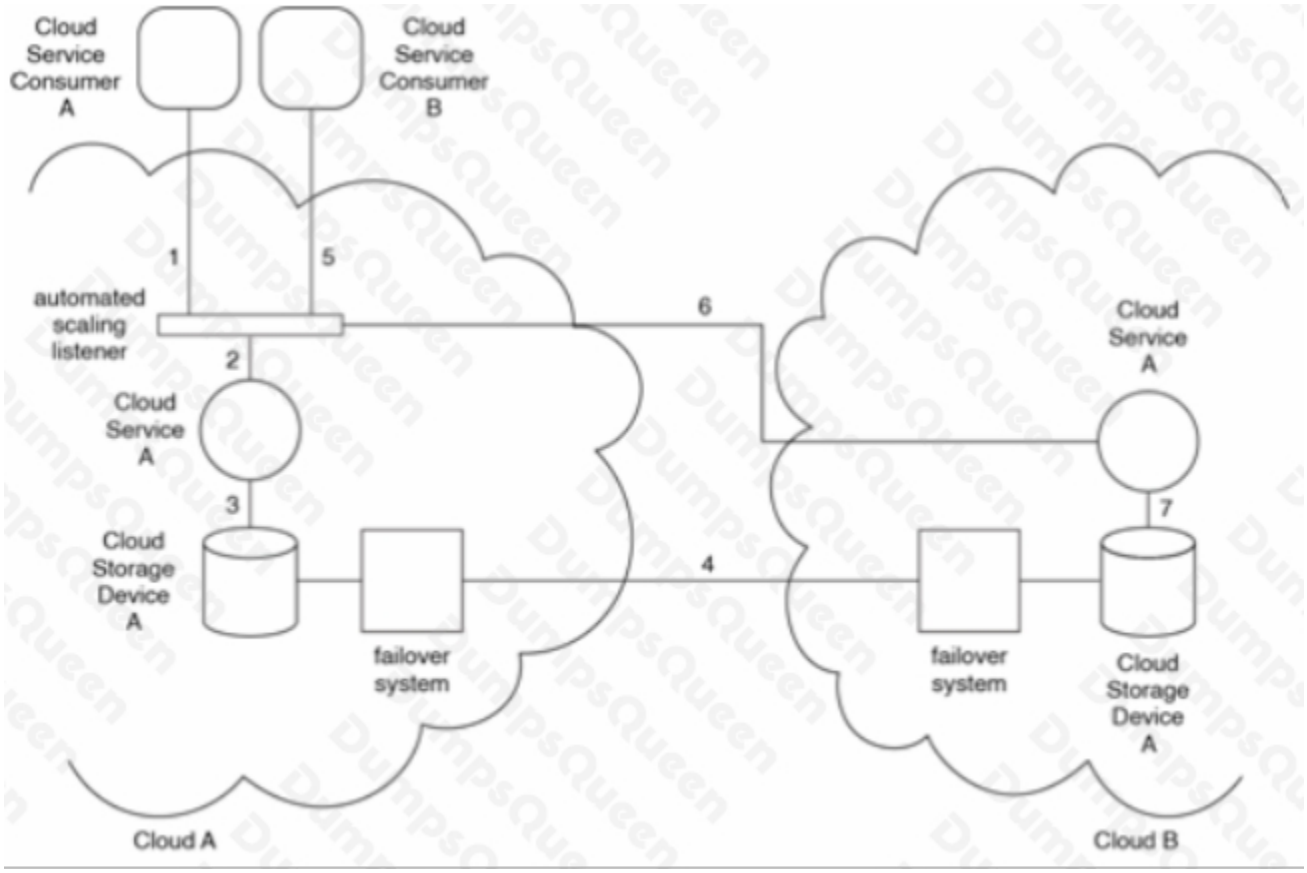

An unexpected outage occurs in Cloud A, making Cloud Service A unavailable. Organization A notices that its cloud resource administrator can continue accessing data in Cloud Storage Device A via a usage and administration portal. Cloud Service Consumer A is unable to access data in Cloud Storage Device A via Cloud Service A during the outage. The cloud resource administrator manually restarts Cloud Service A and it continues to function normally.

Organization A needs to change the cloud architecture so that when Cloud Service A fails, three automated attempts are made to recover it before a manual restart is required.

Due to data storage regulations, Organization A is prohibited from storing some types of data across more than one cloud storage device. A large amount of the data in Cloud Storage Device A is subject to these regulations. Because of an increase in usage, the capacity of Cloud Storage Device A has reached its limit, resulting in regular delays and lag time when processing data access requests during peak usage times.

A management change by another cloud consumer organization inadvertently reconfigures settings in the failover system used in Cloud A for Cloud Storage Device A. Organization A complains to the cloud provider who promises to take the steps required to prevent management tasks performed by other cloud consumer organizations from affecting IT resources being used by Organization A.

Which of the following statements describes a solution that can resolve all of these issues?

A. The Zero Downtime pattern can be applied to establish a cross-cloud failover system for the two Cloud Service A implementations. The Cross-Storage Device Vertical Tiering pattern can be applied to vertically scale data in Cloud Storage Device A across multiple other cloud storage devices dynamically. The Centralized Remote Administration pattern can be applied to establish a logical network perimeter around Organization A's IT resources, thereby protecting them from other cloud consumer organizations.

B. The Dynamic Failure Detection and Recovery pattern can be applied so that if Cloud Service A in Cloud A fails, a watchdog system attempts to automatically recover Cloud Service A. Assuming Cloud Storage Device A has support for multiple disk types, the Intra-Storage Device Vertical Data Tiering pattern can be applied so that Cloud Storage Device A is equipped with dynamic vertical scaling. The Resource Management pattern can be applied to allow cloud consumer organizations to perform management tasks on IT resources without impacting IT resources being used by other cloud consumer organizations.

C. The Load Balanced Virtual Server Instances pattern can be applied to balance the virtual servers hosting Cloud Service A implementations across the two cloud environments. The Storage Workload Management pattern can be applied to balance workloads across both Cloud Storage Device A implementations. The Resource Reservation pattern can be applied to establish a physical network boundary around Organization A's IT resources, thereby protecting them from other cloud consumer organizations.

D. None of the above.

ANSWER: B

QUESTION NO: 2

Physical Server A hosts Hypervisor A which hosts Virtual Server A, Virtual Server B and an active cluster comprised of three virtual servers. Virtual Server A hosts Ready-Made Environment A. Ready-Made Environment A uses Cloud Storage Device A to store data related to the ready-made environment and its users. Cloud Service A is hosted by a high-availability (HA) virtual server cluster. Hypervisor A is a cluster-enabled hypervisor.

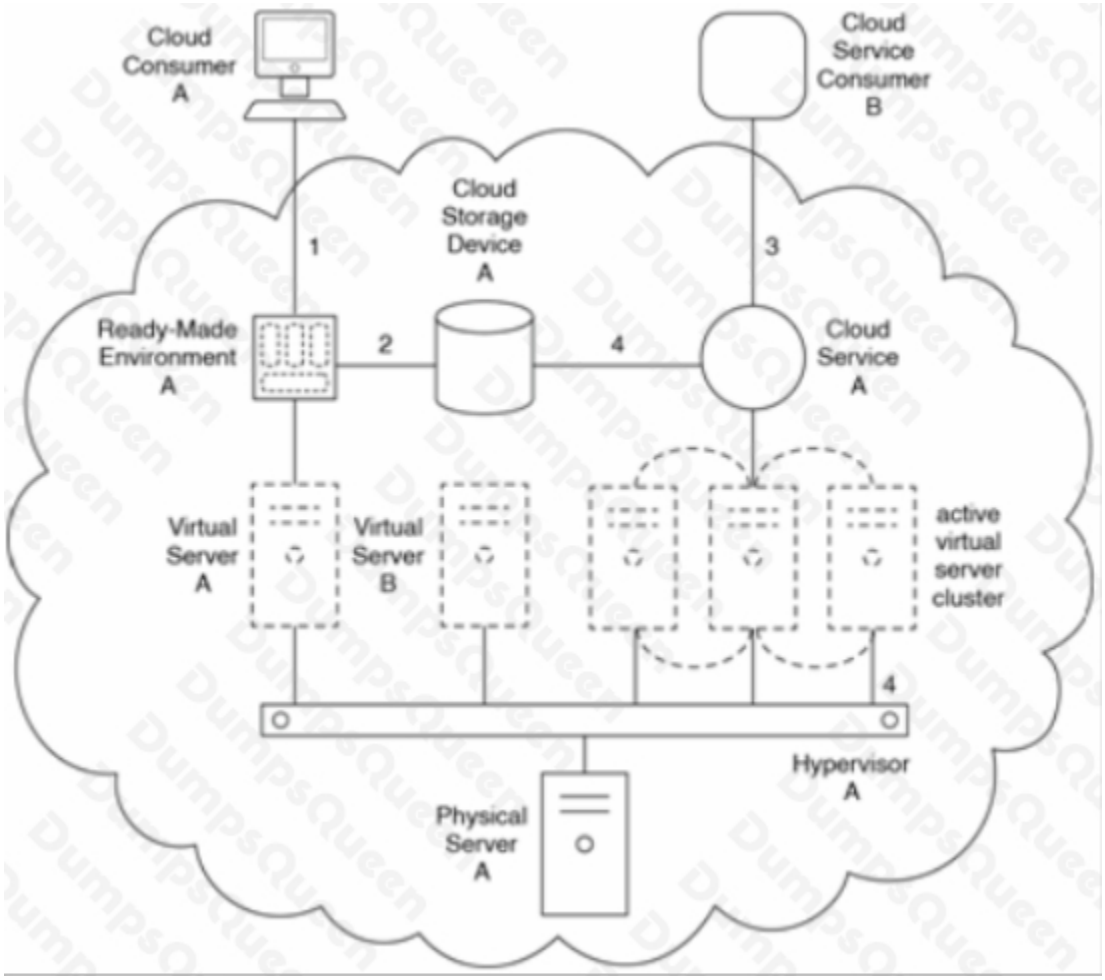

Developers access Ready-Made Environment A to work on the development of a new solution (1). During this usage. Ready-Made Environment A regularly reads and writes data to Cloud Storage Device A (2). Cloud Service Consumer B accesses Cloud Service A (3). Cloud Service A queries data residing in Cloud Storage Device A in response to processing requests from Cloud Service Consumer B (4).

Hypervisor A is made part of a cluster of hypetvisors. Ready-Made Environment A, which is still hosted by Virtual Server Aon Hypervisor A, subsequently becomes unexpectedly unavailable. It then takes twenty minutes to pass before Virtual Server A and Ready-Made Environment A become available again on Hypervisor B (a hypervisor that is also part of the hypervisor cluster). This delay is considered unacceptable by Cloud Consumer A. Furthermore, after being relocated to Hypervisor B, Virtual Server A is unable to connect to the network. By the time the cloud provider rectifies this second problem, Cloud Consumer A experiences a total of two hours of downtime.

Which of the following statements describes a solution that can minimize or entirely avoid a delay for the runtime relocation of Ready-Made Environment A?

A. The Load Balanced Virtual Server Instances pattern can be applied in combination with the Elastic Network Capacity pattern in order to establish a system whereby Ready-Made Environment A can be smoothly transitioned between hypervisors in the same cluster, while its underlying virtual server maintains the network connection.

B. The Hypervisor Clustering pattern was incorrectly applied and therefore needs to be re-applied correctly in order to establish a native system capable of instantly relocating virtual servers between hypervisors within the same cluster. The Direct I/O Access pattern can then also be applied so that the virtual servers retain their network configurations regardless of which hypervisor within the cluster they reside on.

C. The Non-Disruptive Service Relocation pattern can be applied to place a secondary copy of Ready-Made Environment A on Hypervisor B. The Persistent Virtual Network Configuration pattern can be applied so that virtual servers retain network configurations when moving to other hypervisors.

D. The Load Balanced Virtual Server Instances and Persistent Virtual Network Configuration patterns can be applied together to ensure that the virtual servers retain their network configurations when moving to another hypervisor. The Redundant Storage pattern can further be applied to move the ready-made environment to another hypervisor without service impact.

ANSWER: C

QUESTION NO: 3

Virtual Server A and Virtual Server B are hosted by Hypervisor A, which resides on Physical Server A. Virtual Server A hosts Cloud Service A. Virtual Server C. Virtual Server D, Virtual Server E and Virtual Server F are hosted by Hypervisor B on Physical Server B. Physical Server C, which hosts Hypervisor C, is currently not being used.

Cloud Service Consumer A accesses Cloud Service A (1), which accesses files stored in a folder on Virtual Server A (2). Cloud Consumer B uses Proprietary User Portal A to administer legacy software (not shown) installed on Virtual Server D (3). Proprietary User Portal B and Proprietary User Portal C are also available for accessing additional legacy systems located on Virtual Server F; however, they are not often used.

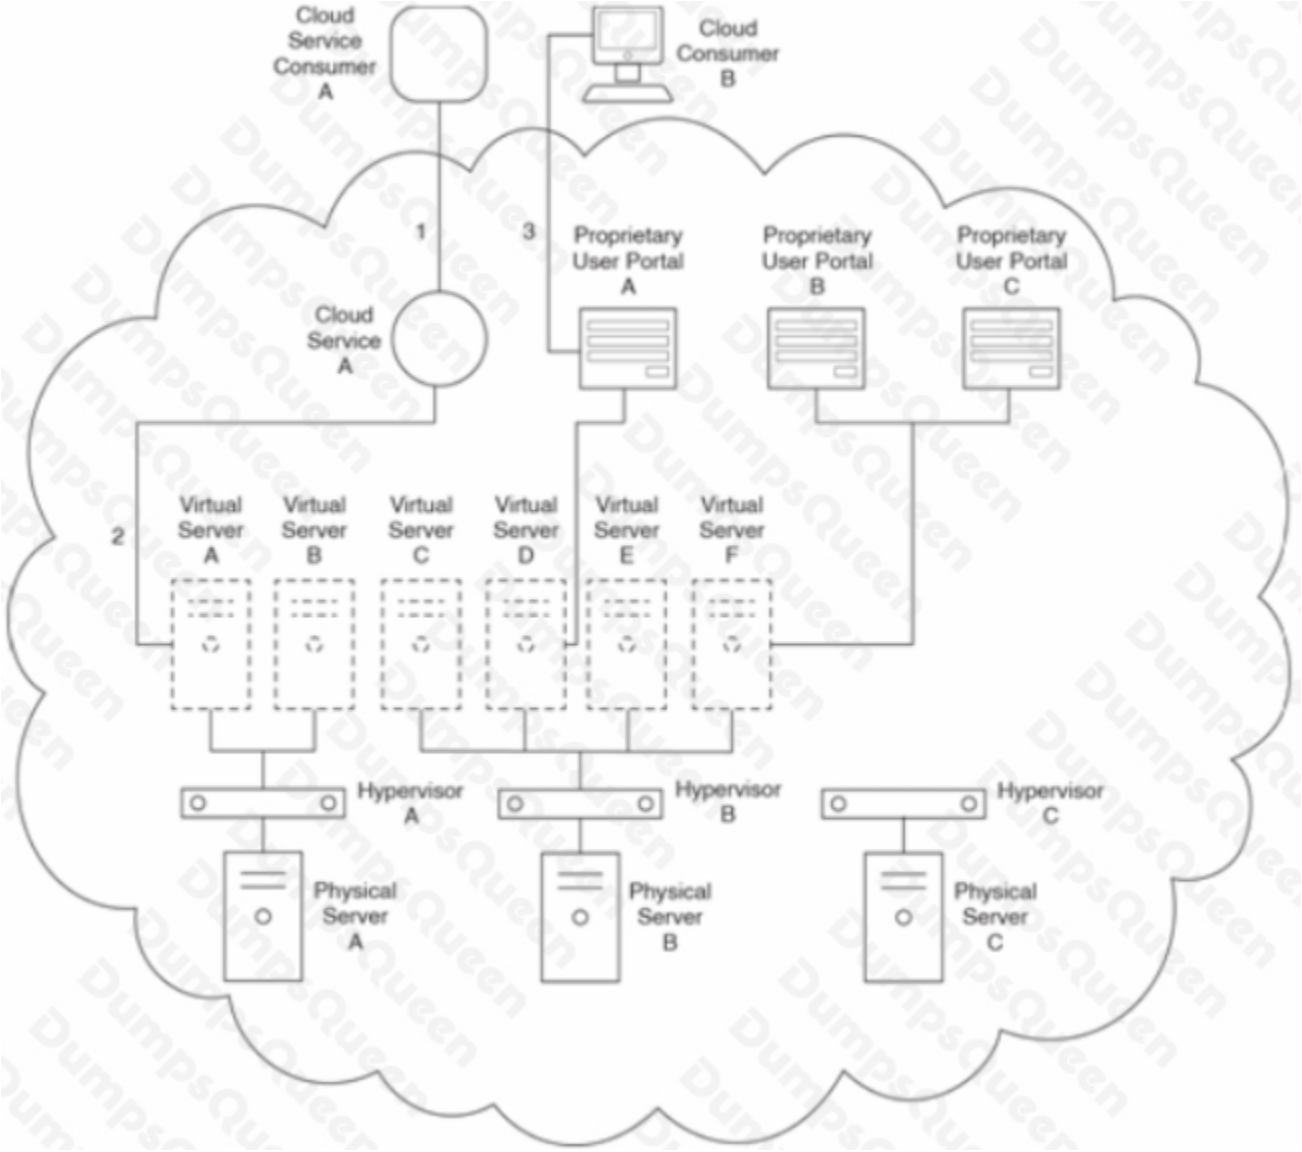

The cloud shown in the figure is a private cloud. Department A and Department B share IT resources within the private cloud and are part of the same organization. Cloud Service Consumer A belongs to Department A and Cloud Consumer B belongs to Department B.

During routine access of Cloud Service A by Cloud Service Consumer A, the Department A cloud resource administrator is notified that a hardware fault is occurring within Physical Server A that will soon cause it to fail. The cloud resource administrator scrambles to arrange for Cloud Service A to be relocated but is unable to do so before Physical Server A does fail. It takes several more hours of downtime until, with the cooperation of the cloud provider, the Cloud Service A implementation is successfully moved to Physical Server C and made live again. Managers at Department A demand that a system be put in place to avoid this scenario in the future.

Cloud Service A was initially developed specifically for Department A's Cloud Service Consumer A. However, recently Department B has indicated that it will be developing its own cloud service consumer that will also need to regularly access Cloud Service A. After this new cloud service consumer is deployed, both Department A and Department B experience occasional runtime errors when their cloud service consumers attempt to access Cloud Service A at the same time.

Cloud Service A accesses a legacy system on Virtual Server A that requires regular updates and patches to stay current. Each time the legacy system is updated, Cloud Service A needs to undergo an update as well, during which it needs to be

temporarily unavailable. Department A managers ask the cloud provider to extend the cloud architecture so that a duplicate, secondary implementation of Cloud Service A can be made available while the primary implementation undergoes a maintenance update.

Which of the following statements provide a solution that can adequately resolve all of Departments A and B's issues?

A. The Resource Reservation pattern can be applied to protect the Cloud Service A implementation via the use of a logical network perimeter. The Workload Distribution pattern can be applied to introduce a load balancing system for Cloud Service The Zero Downtime pattern can be applied to establish a system that allows Cloud Service A to be constantly available, even during maintenance outages.

B. The Resource Pooling pattern can be applied to pool together Physical Server A, B and C, thereby enabling the Cloud Service A implementation to be migrated to a different physical server when its hosting physical server fails. The Dynamic Scalability pattern can be applied to establish a system whereby multiple instances of Cloud Service A can be created and an automated scaling listener can be used to redirect concurrent requests to the Cloud Service A instances. The Load Balanced Virtual Server Instances pattern can be applied to establish a system that distributes instances of Cloud Service A to Virtual Server B.

C. The Non-Disruptive Service Relocation pattern can be applied to establish a system that uses live VM migration to move the virtual server hosting Cloud Service A to a new physical server without allowing any downtime. The Dynamic Scalability pattern can be applied to establish a system whereby multiple instances of Cloud Service A can be created and an automated scaling listener can be used to redirect concurrent requests to the Cloud Service A instances. The Non-Disruptive Service Relocation pattern can be applied to establish a system whereby cloud service consumer requests to Cloud Service A can be temporarily redirected to a duplicate implementation of Cloud Service A while the original implementation undergoes a maintenance outage.

D. None of the above.

ANSWER: C

QUESTION NO: 4

Cloud Service A requires access to Cloud Storage Device A and Cloud Storage Device B. Cloud Service A is hosted by Virtual Server A. Virtual Server A and Virtual Server B are hosted by Hypervisor A, which resides on Physical Server A.

Cloud Service Consumer A sends a request to access Cloud Service A (1). Cloud Service A retrieves data from Cloud Storage Device A and Cloud Storage Device B (2). Cloud Consumer A uses the usage and administration portal to access resource usage reports for Cloud Service A (3).

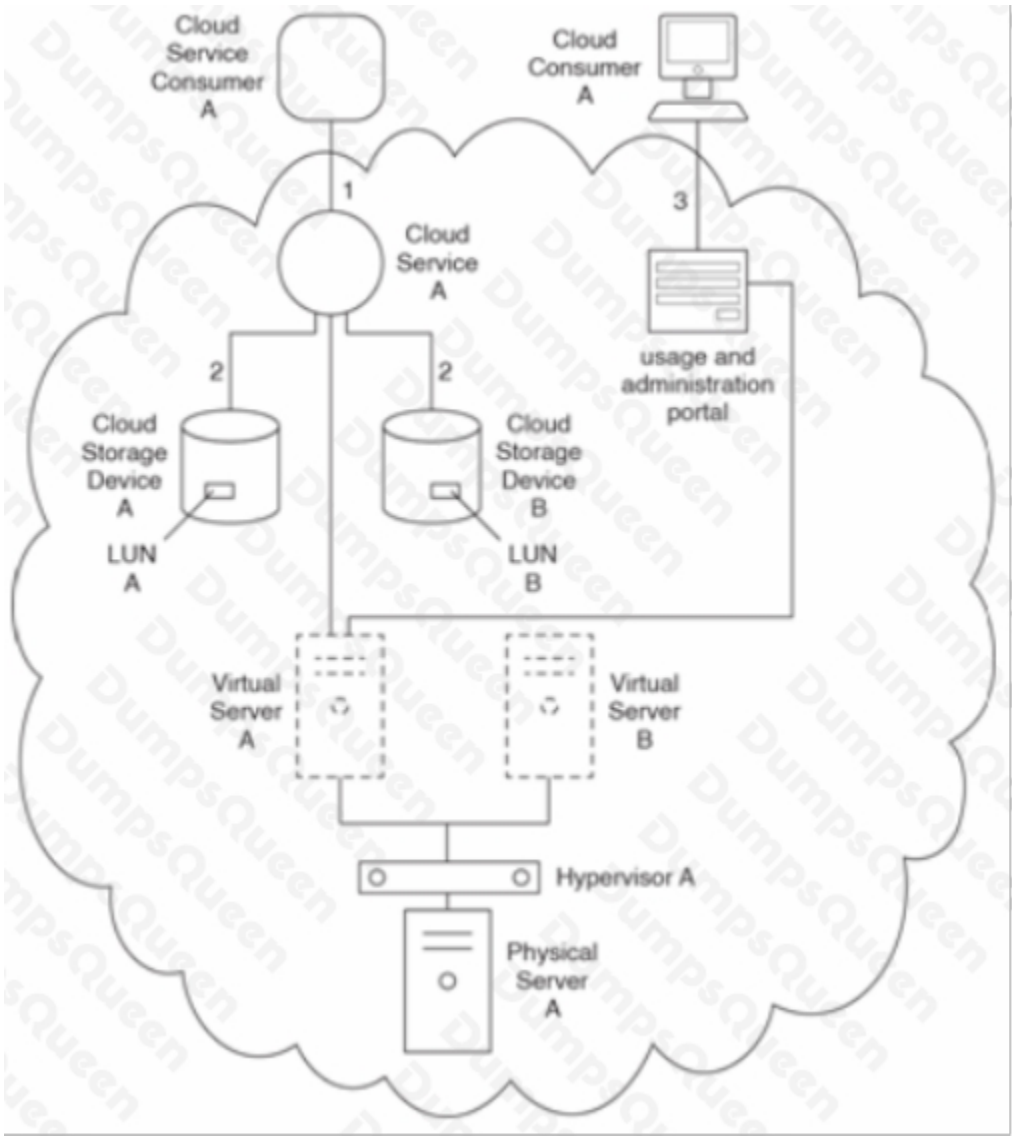

Cloud Service Consumer A and Cloud Consumer A belong to Organization A, which is leasing an IaaS environment from the cloud provider.

The cloud provider makes Cloud Service A available to several new cloud service consumers. Additionally, new LUNs are created on Cloud Storage Devices A and B for new cloud consumers to perform regular data access functions. This increase in workload causes Virtual Server A to fail during peak usage periods. Organization A and the new cloud consumer organizations request that the cloud provider find a way to dynamically support the higher usage workloads.

Organization A keeps its master files and data in LUN B in Cloud Storage Device B. One day, a cloud resource administrator accidentally changes the path used to access LUN B. The original path cannot be retrieved. The cloud resource administrator informs Organization A's IT department that it must change any systems or tools it uses to access LUN B to the new path. This causes significant challenges, as well as a costly period of disruption. Organization A asks the cloud provider to create a system that would help avoid disruption in access to LUN B, if this was to ever happen again.

The cloud provider has made Cloud Storage Device A part of a resource pool of synchronized cloud storage devices. Organization A is sharing Cloud Storage Device A with another cloud consumer organization. When cloud consumers from both organizations access Cloud Storage Device A at the same time, they encounter a resource constraint condition that causes Cloud Storage Device A to fail. Organization A requests that the cloud provider extend the existing cloud architecture to prevent this situation from happening again.

Which of the following statements provides a solution that can address all of these problems?

A. The Elastic Network Capacity pattern can be applied to implement a system that dynamically assigns network ports to Virtual Server A before its processing capacity thresholds are reached. The Redundant Physical Connection for Virtual Servers pattern can be applied to create an alternative path to LUN B in Cloud Storage Device B. The Resource Pooling pattern can be applied to synchronize Cloud Storage Device A with other cloud storage devices.

B. The Resource Reservation pattern can be applied to dynamically provision resources to Virtual Server A whenever its processing thresholds are being reached. The Persistent Virtual Network Configuration pattern can be applied to establish a persistent hyperlink to LUN B over the virtual network that cannot be lost. The Elastic Resource Capacity pattern can be applied to prevent Cloud Storage Device A from encountering resource constraints.

C. The Elastic Resource Capacity pattern can be applied to establish a system that can dynamically allocate resources to Virtual Server A. The Multipath Resource Access pattern can be applied to establish a multipathing system that can provide an alternative path to LUN B in Cloud Storage Device B. The Resource Reservation pattern can be applied to establish a system that enables Organization A to have exclusive access to pre-defined resources on Cloud Storage Device A for a given period of time.

D. None of the above.

ANSWER: C

QUESTION NO: 5

Cloud Storage Device A contains LUN A and can be accessed by external cloud consumers via a proprietary user portal that is used primarily by cloud consumers to upload and manage data for backup purposes. Pay-Per-Use Monitor A tracks the amount of data being uploaded and forwards this information to a billing management system. Cloud Storage Device B is a secondary cloud storage device. Data from Cloud Storage Device A is replicated synchronously to Cloud Storage Device B using a storage replication program (not shown). Cloud Storage Device A is further administered by a cloud resource administrator that works for the cloud provider, who accesses the cloud storage device via an internal usage and administration portal.

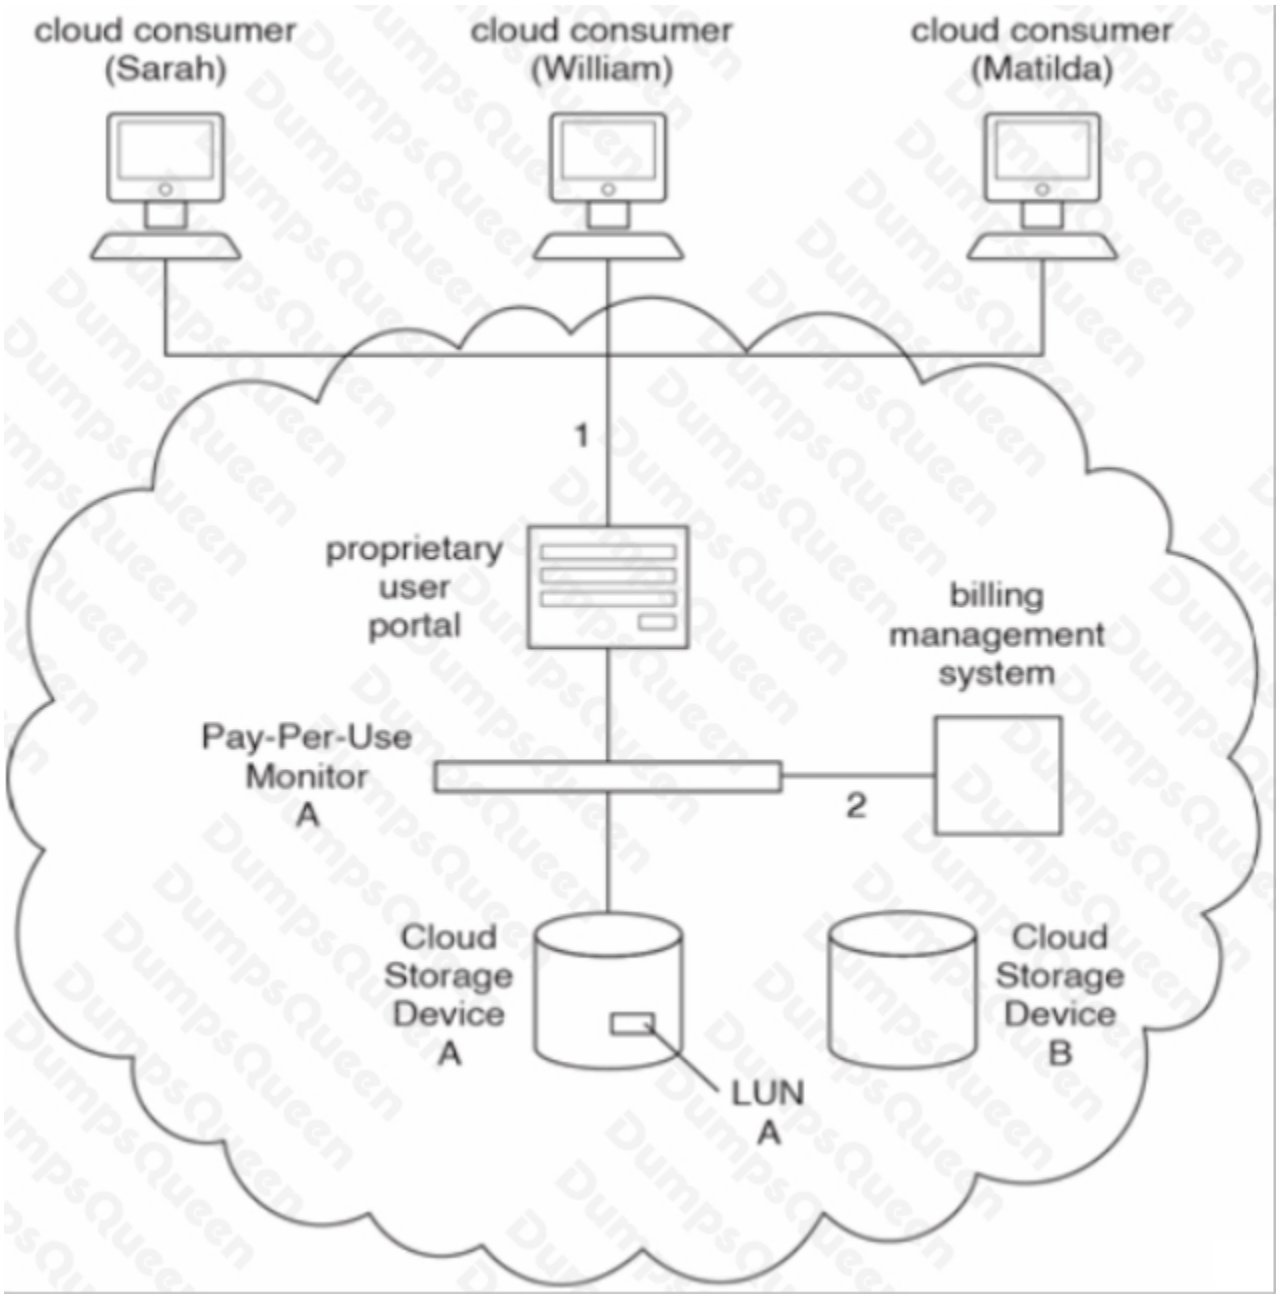

Three different cloud consumers (Sarah. William, Matilda) access Cloud Storage Device A to upload data for backup purposes (1). These three cloud consumers belong to different

departments within the same organization, and therefore all three share LUN A. Pay-Per-Use Monitor A collects the storage space data and forwards it to the billing management system (2).

The cloud provider offers premium and discount storage plans. With the premium plan, the data stored on Cloud Storage Device A is also replicated to Cloud Storage Device B. With the discount plan, the data stored on Cloud Storage Device A is not replicated. Both plans are based on fixed-disk storage allocation. The cost of the premium plan is $5 per week for every GB of storage space and the cost of the discount plan is $2 per week for every GB of storage space. The SLA from the cloud provider guarantees availability of 97% for access to Cloud Storage Device A.

The three cloud consumers use Cloud Storage Device A as follows:

Which of the following statements lists the patterns that can be applied to address the three issues raised by the three cloud consumers?

A. Storage Workload Management, Elastic Disk Provisioning, Centralized Remote Administration

B. Elastic Disk Provisioning, Dynamic Data Normalization, Zero Downtime

C. Storage Maintenance Window, Dynamic Failure Detection and Recovery, Broad Access

D. None of the above.

ANSWER: B