Free Questions for C90.06 by certsdeals

Shared by Lowe on 24-05-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

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.

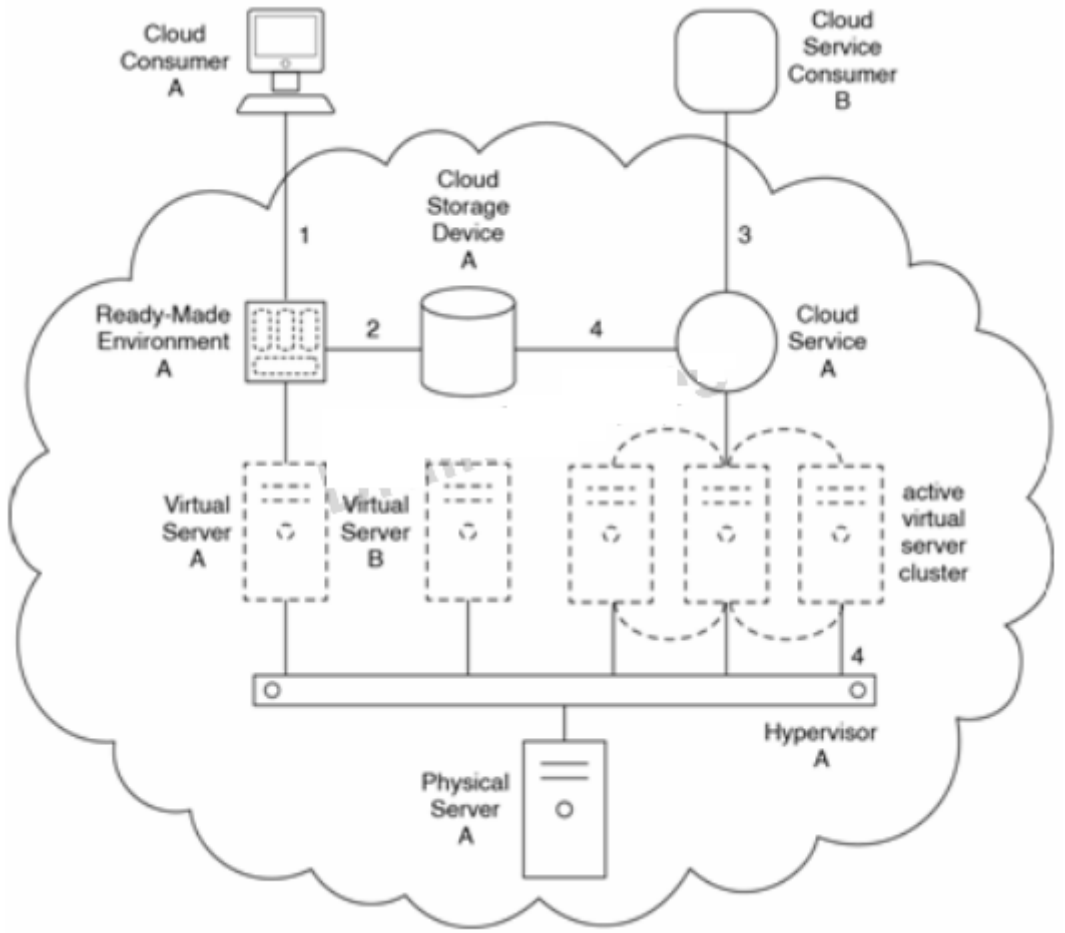

Developers access Ready-Made Environment A to work on the development of a new solution (1).

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?

Options:

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 2

Question Type: MultipleChoice

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).

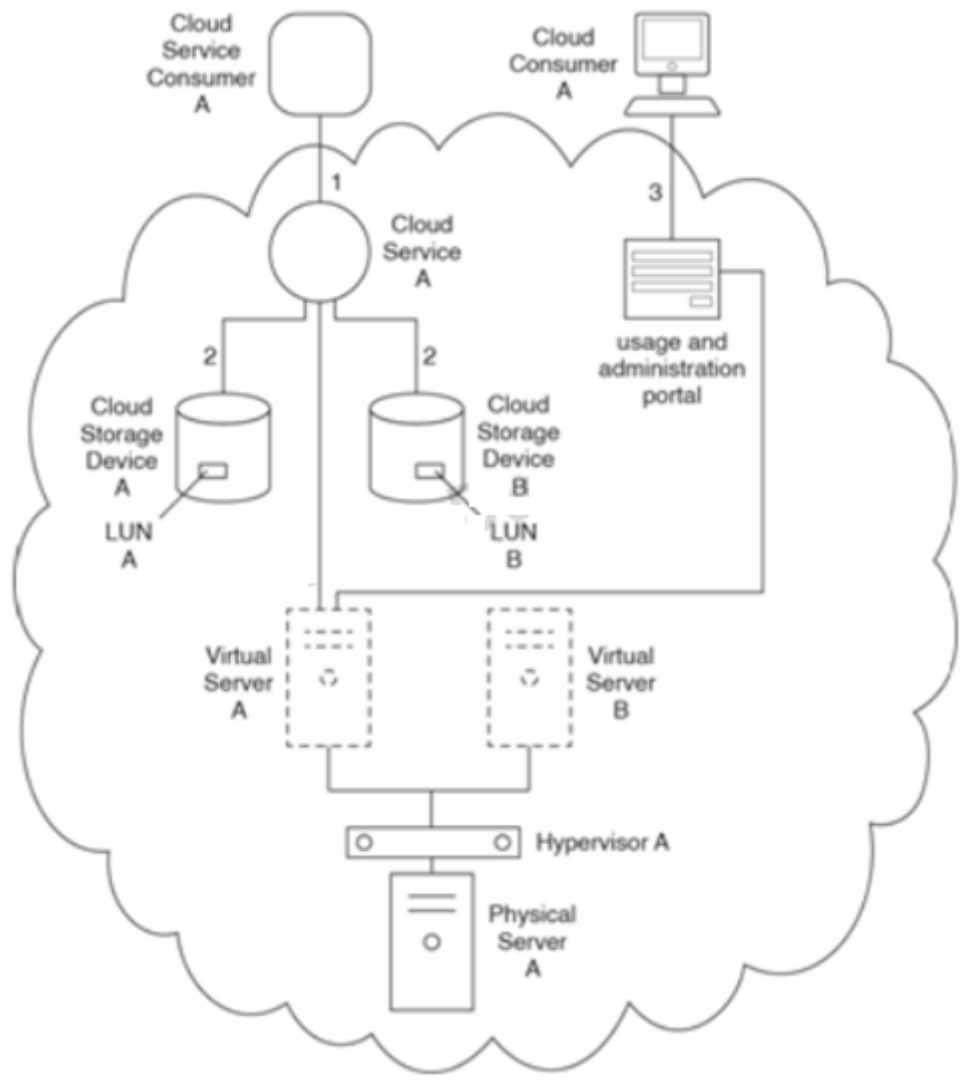

Cloud Service Consumer A and Cloud Consumer A belong to Organization A, which is leasing an laaS environment from the cloud provider.

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?

Options:

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 3

Question Type: **MultipleChoice**

Cloud Service A is hosted by Virtual Server A, which is hosted by Hypervisor A that resides on Physical Server A. Cloud Storage Device A is used to store media library data that is continuously replicated with a redundant, secondary implementation of Cloud Storage A (not shown). Access to Cloud Service A is monitored by Pay-Per-Use Monitor A. Access to Cloud Storage Device A is monitored by Pay-Per-Use Monitor B. Pay-Per-Use Monitors A and B capture billing-related usage data that is forwarded to a billing management system that

is hosted by Physical Server B.

Cloud Service Consumer A accesses Cloud Service A and the usage data is captured by Pay-Per-Use Monitor A (1). Cloud Consumer B accesses Cloud Storage Device A via a usage and administration portal that it uses to upload media data (2). This usage is captured by Pay-Per-Use Monitor B (3). Pay-Per-Use Monitors A and B store collected usage data in the billing management system (4), which is later used by the cloud provider to bill for the usage of Cloud Service A and Cloud Storage Device A.

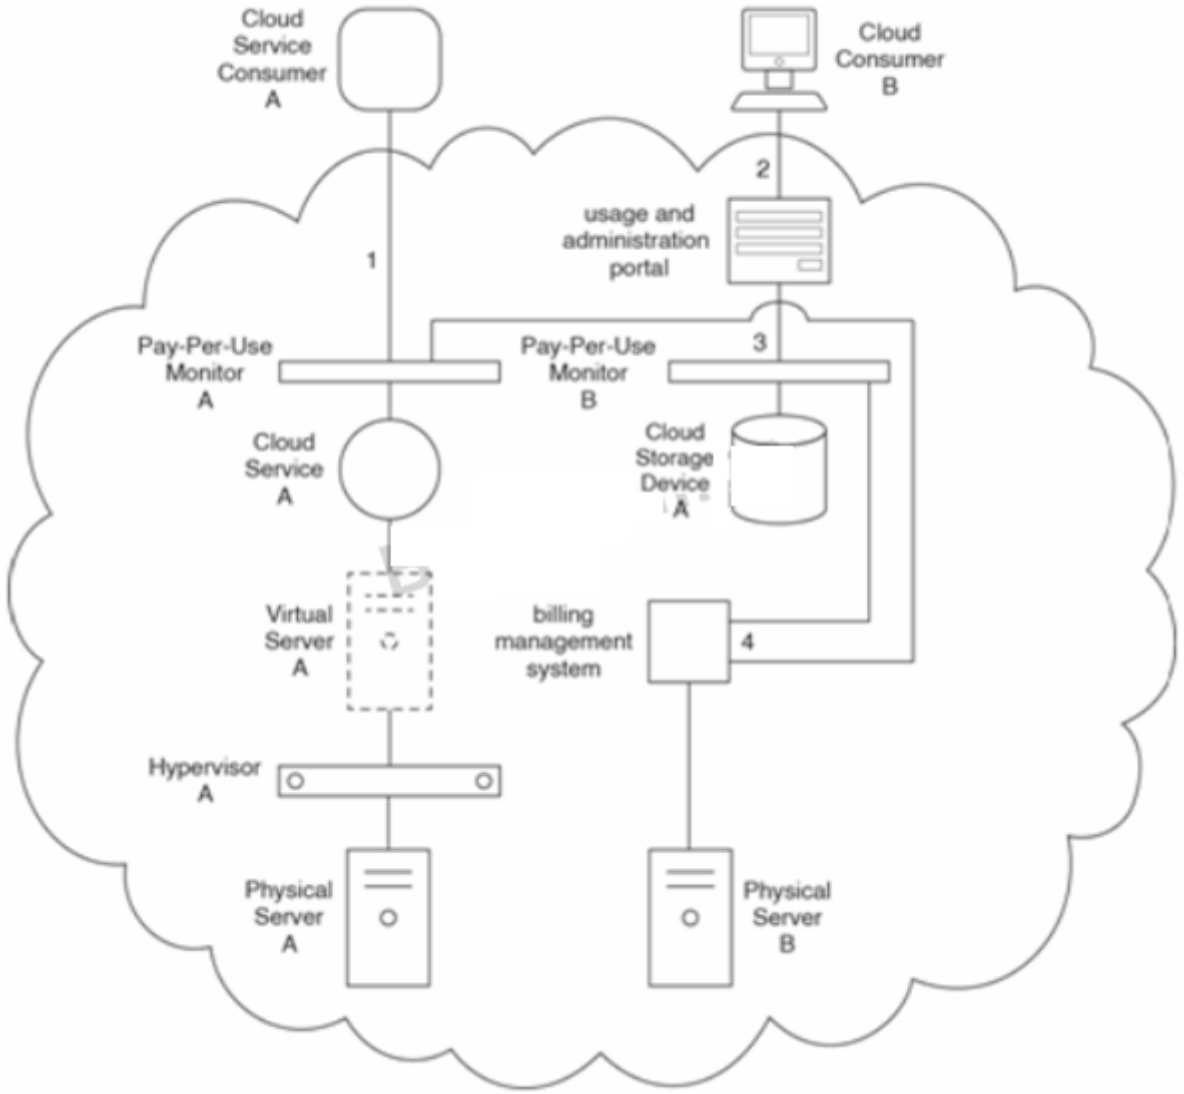
Cloud
Service
Consumer
A
Cloud
Consumer
B
2
usage and
administration
portal
1
Pay-Per-Use
Monitor
A
Pay-Per-Use
Monitor
B
3
Cloud
Service
A
Cloud
Storage
Device
A
Virtual
Server
A
billing
management
system
4
Hypervisor
A
Physical
Server
A
Physical
Server
B

Each service instance of Cloud Service A requires a virtual server with 2 virtual CPUs and 4 GBs of RAM at a package price of $2.00 for each initial invocation and an additional $0.50 for each consecutive 60 seconds of usage. Cloud Service Consumer A accesses Cloud Service A twice in one day. The two exchanges with Cloud Service A last 60 seconds and 120 seconds. For that one day, the organization that owns Cloud Service Consumer A is billed $6.50, which it determines is incorrect. After complaining to the cloud provider, it is discovered that the rapid provisioning system responsible for provisioning instances of Cloud Service A is not de-provisioning Cloud Service A when Cloud Service Consumer A indicates it has completed an exchange. Instead, Cloud Service A is de-provisioned after a 60 second timeout that occurs after Cloud Service Consumer A is completed with an exchange.

Storage space on Cloud Storage Device A can only be purchased in units of terabytes (TBs), with each TB costing $1 per day. Cloud Consumer B purchases 5 TBs of storage space on day 1 and stores 5 TBs of data on days 6 and 7. Cloud Consumer B was expecting to be billed $10.00, but is billed $35. After raising a complaint, Cloud Consumer B is informed by the cloud provider that cloud consumers are billed based on the allocation of storage space, regardless of how much storage space they actually use.

Which of the following statements describes a solution that can update the cloud architecture to avoid these billing-related problems and discrepancies?

Options:

A- The Pay-as-You-Go pattern can be applied together with the Usage Monitoring pattern to establish a monitoring and billing system capable of de-provisioning Cloud Service A instances when they are no longer required. The Dynamic Data Normalization pattern can be applied to eliminate any redundant data stored by Cloud Consumer A so that the amount of required storage space is minimized.

B- The Platform Provisioning pattern can be applied to create an intelligent automation script capable of immediately de-provisioning cloud service instances. The Redundant Storage pattern can be applied to introduce a secondary cloud storage device for which storage space can be billed based on actual usage.

C- The Self-Provisioning pattern can be applied to enable the organization that owns Cloud Service Consumer B to configure how and to what extent Cloud Service A instances need to be provisioned. The Resource Management pattern can be applied to establish a storage system that bills cloud consumers for actual storage space usage only.

D- None of the above.

Answer:

D

Question 4

Question Type: MultipleChoice

Cloud Service A and Cloud Service B perform different functions but both share access to Cloud Storage Device A when fulfilling requests from cloud service consumers that require data access.

Cloud Services A and B are hosted by Virtual Server A, which is hosted by Hypervisor A on Physical Server A.

Cloud Service Consumer A accesses Cloud Service A to issue a request for data (1). Cloud Service A queries a database in Cloud Storage Device A to retrieve the data (2). Upon receiving the requested data, Cloud Service Consumer A combines it with additional data to form a new collection of data. Cloud Service Consumer A then accesses Cloud Service B and provides it with the new data (3). Cloud Service B accesses a different database in Cloud Storage Device A to store the new data (4). Cloud Consumer B accesses the usage and administration portal to upload new data (5). The data is uploaded to Cloud Storage Device B (6).

Cloud Service Consumer A belongs to Organization A. Cloud Consumer B belongs to Organization B.

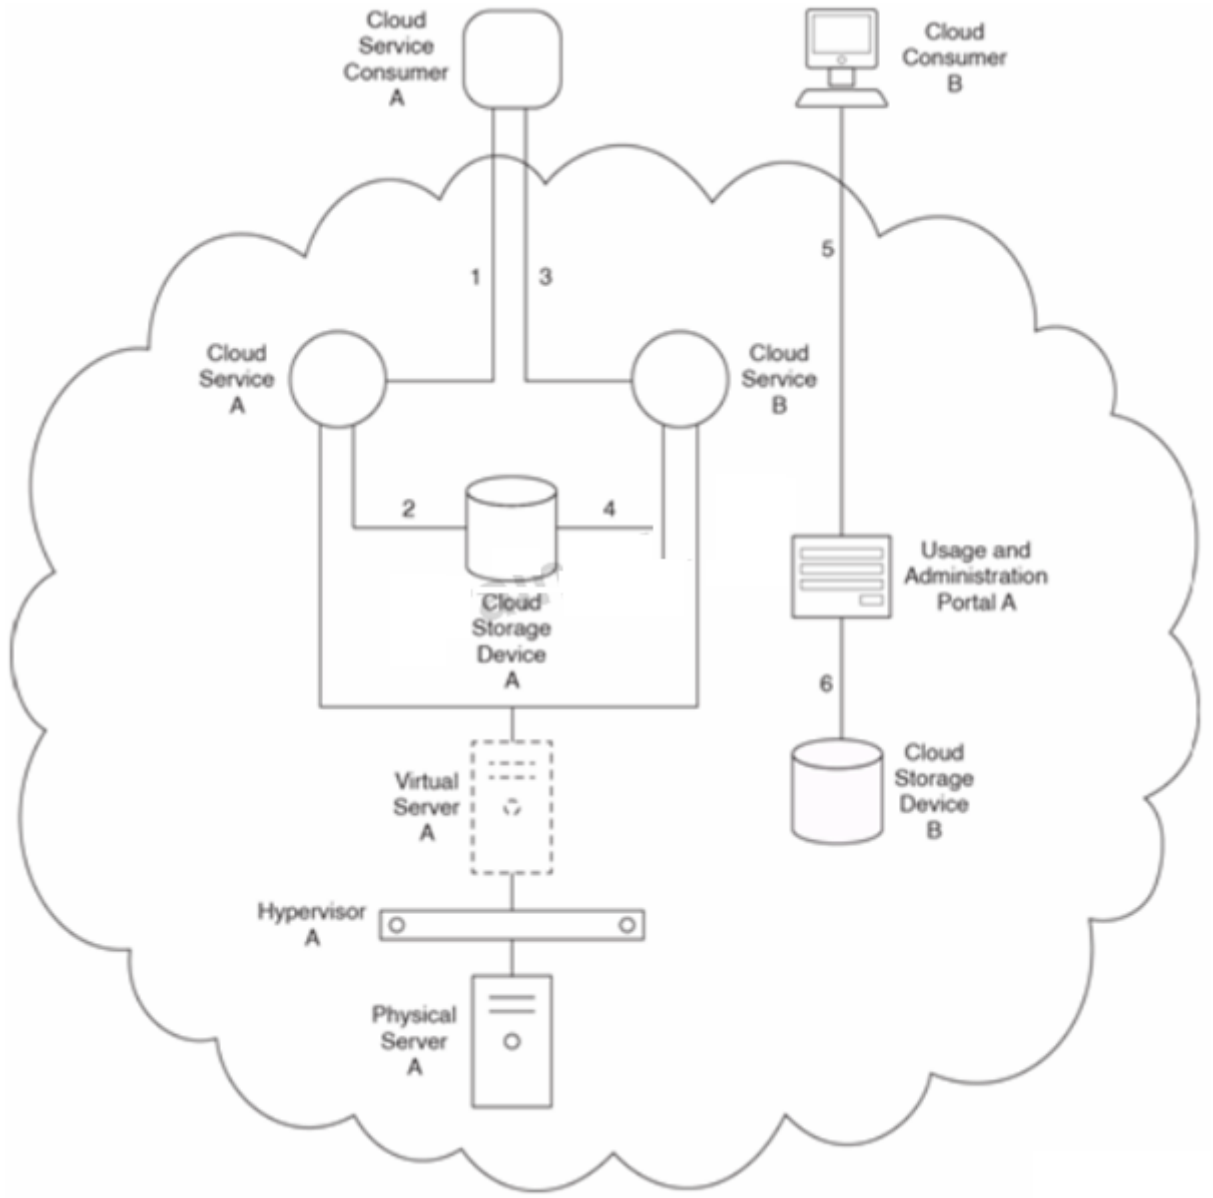
Cloud Service Consumer A
Cloud Consumer B
1
3
5
Cloud Service A
Cloud Service B
2
4
Cloud Storage Device A
Usage and Administration Portal A
6
Cloud Storage Device B
Virtual Server A
Hypervisor A
Physical Server A

Cloud Service A is a SaaS product offered by the cloud provider to the general public, and is therefore used by numerous cloud consumers from different organizations at different times. Cloud Service B is also a SaaS product as part of the same overall solution as Cloud Service A.

However, because a given cloud service consumer only needs to access Cloud Service B when the data it receives from Cloud Service A meets certain criteria, it is not used nearly as much as Cloud Service A. Cloud Service A currently has a hard threshold allowing no more than 10 concurrent instances of it to exist at once. One day, Cloud Service Consumer A attempts to access Cloud Service A as the eleventh cloud service consumer, and is predictably rejected.

Cloud Service Consumer A belongs to Organization A, one of the cloud provider's most important customers. Therefore, when Organization A complains about not being able to access Cloud Service A during peak usage times, the cloud provider agrees to provide a solution.

As a result of a natural disaster, the cloud provider's data center that houses Physical Server A becomes unexpectedly unavailable. Physical Server A subsequently becomes unavailable for nearly two days. This outage exceeds what the cloud provider guaranteed in its original SLA and the cloud provider agrees to not charge Organization for usage fees for an entire month as compensation. However, the unavailability of Physical Server A had a significant impact on Organization As business, resulting in financial loss and loss of confidence of its clients.

Organization A informs the cloud provider that it cannot continue working with this cloud unless the cloud provider can guarantee that the availability of Physical Server A will no longer be dependent on a single data center or a single geographic region.

Organization B receives its latest monthly invoice from the cloud provider and discovers that the charges are identical to the invoice it received last month, even though the usage and administration portal shows that its data usage is a third less. They bring this issue to the attention of the cloud provider and are informed that they are currently subscribed to a fixed-allocation plan.

The cloud provider explains that in order to get them on a plan whereby they are charged only for actual data usage, Cloud Storage Device B would need to be upgraded and a system capable of tracking runtime usage would need to be established. Organization B asks the cloud provider to make these changes.

Which of the following statements provides a solution that can address Organization A's and Organization B's issues?

Options:

A- The Resource Pooling pattern can be applied to pool the IT resources used by Cloud Service A so that requests from Cloud Service Consumer A can utilize these resources during peak usage times. The Non-Disruptive Service Relocation pattern can be applied so that, in the event of failure, Cloud Storage Device B can be migrated at runtime to a cloud in another geographic region. The Elastic Disk Provisioning pattern can be applied so that Organization B is only charged for the amount of data storage it uses.

B- The Service Load Balancing pattern can be applied to balance the workloads across multiple Cloud Service A implementations. The Synchronized Operating State pattern can be applied so that Virtual Server A is automatically synchronized to a secondary implementation of Physical Server A located in a different geographic region. The Usage Monitoring pattern can be applied so Cloud Consumer B's resource consumption is tracked and logged at runtime.

C- The Service Load Balancing pattern can be applied to create redundant service implementations of Cloud Service A, so that a load-balancing system can distribute workloads across the service implementations dynamically. The Zero Downtime pattern can be applied to establish a secondary deployment of Physical Server A in a different data center located in a different geographic region. The Elastic Disk Provisioning pattern is applied together with the Payas-You-Go pattern to establish systems that ensure that Cloud Consumer B is charged only for the amount of data and resources it consumes.

D- None of the above.

Answer:

C

Question 5

Question Type: MultipleChoice

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 fulfil the request (7).

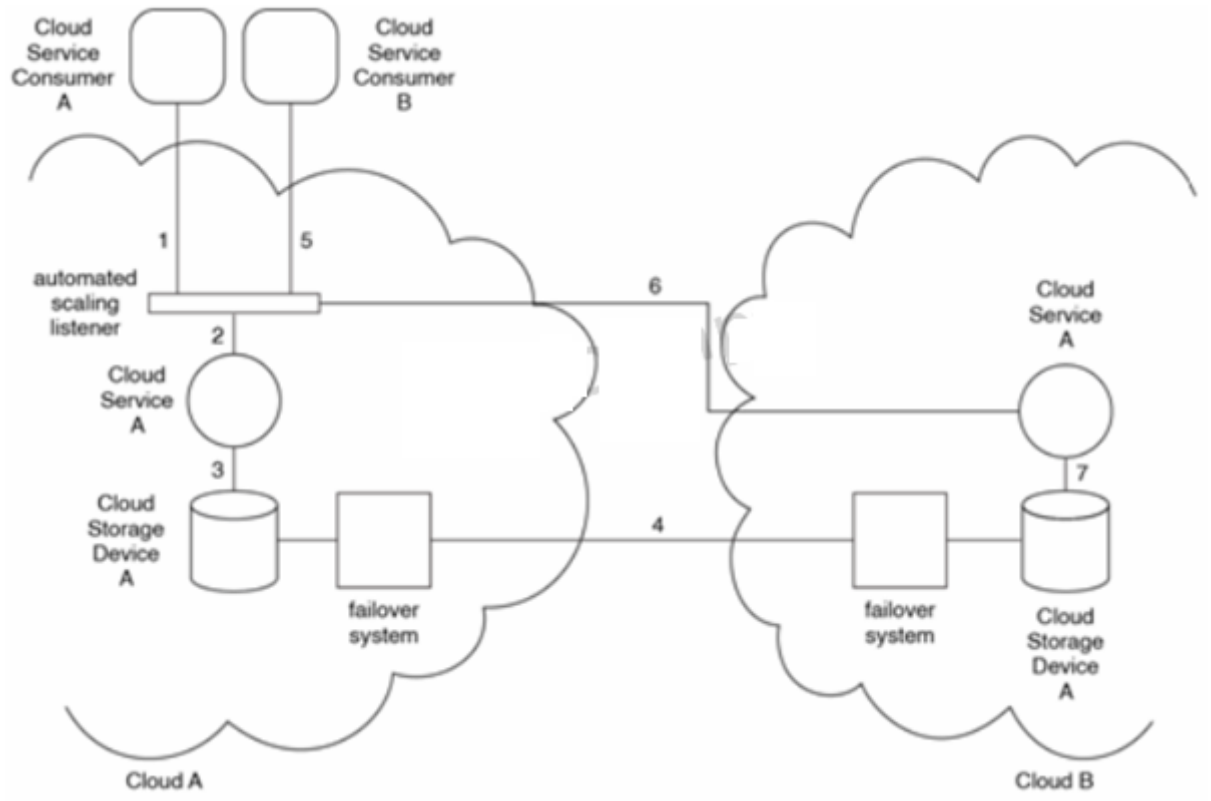

An unexpected outage occurs in Cloud A, making Cloud Service A unavailable. Organization A notices that its cloud resource administrator can continue accessing data in Cloud Storage Device A via a usage and administration portal. Cloud Service Consumer A is unable to access data in Cloud Storage Device 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?

Options:

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 6

Question Type: **MultipleChoice**

Cloud Service A accesses LUN Aon Cloud Storage Device A when it receives requests to process data from cloud consumers. Cloud Service A is hosted by Virtual Server A. The usage and administration portal can be used to access and manage the data in Cloud Storage Device B, which is also hosted by Virtual Server A. Virtual Server A is further hosted by Hypervisor A, which resides on Physical Server A. Virtual Server B is part of a virtual server cluster hosted by Hypervisor B. which resides on Physical Server B. Physical Server C is not in use and does not yet have an operating system installed.

Cloud Service Consumer A sends a request to Cloud Service A (1), which accesses data in LUN Aon Cloud Storage Device A (2). Cloud Consumer B uses the usage and administration portal to upload new data (3). The data is placed in LUN B on Cloud Storage Device B (4).

Cloud Service Consumer A and Cloud Consumer B belong to Organization A, which is leasing Virtual Server A and Virtual Server B from the cloud provider. Organization A also proceeds to lease Physical Server C as part of a new laaS agreement it signs with the cloud provider.

Organization A wants to provision Physical Server C with a number of legacy systems that cannot be deployed on virtual servers. However, when it attempts to do so, it realizes that its laaS package only provides Physical Server C as an out-of-the-box hardware server without anything installed on it. In order to deploy its legacy systems Organization A requires that Physical Server C first has an operating system installed, but it has no means of remotely provisioning Physical Server C with an operating system.

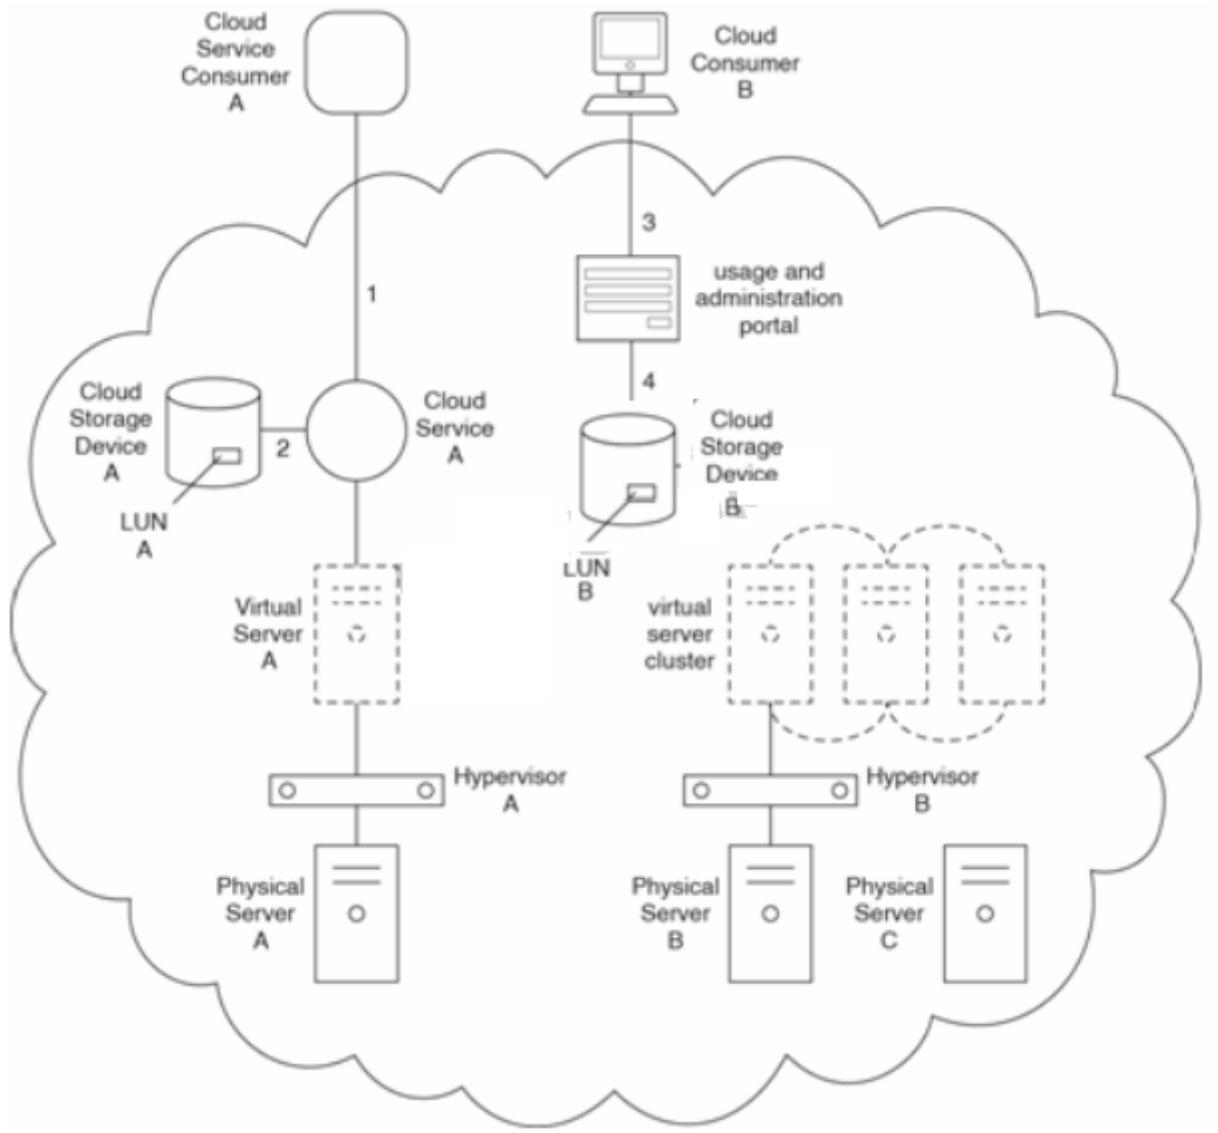
Cloud Service Consumer A
Cloud Consumer B
1
3
usage and administration portal
4
Cloud Storage Device A
2
Cloud Service A
Cloud Storage Device B
LUN A
LUN B
Virtual Server A
virtual server cluster
Hypervisor A
Hypervisor B
Physical Server A
Physical Server B
Physical Server C

Organization A would like to deploy two of its legacy systems on Virtual Server A and to further extend Cloud Service A's functions so that it can be used as an external interface for cloud service consumers to access legacy system features. Additionally, Organization A would like to deploy three of its mission-critical legacy systems on Virtual Server B in order to take advantage of the additional performance and failover benefits provided by the virtual server cluster that Virtual Server B is part of. Each of the five legacy systems is comprised of dozens of components that need to be installed individually. Instead of manually installing each component of each legacy system, Organization A would like to customize workflows that can automate these deployment tasks.

During the first few months of working with its cloud-based legacy systems. Organization A receives a number of complaints from users that the cloud-based legacy systems are at times behaving erratically. However, when cloud resource administrators with Organization A review the cloud provider's reports that log usage, downtime and other runtime characteristics, they do not find any indication of erratic behavior or any other comparable problems. After some further investigation, the cloud resource administrators determine that the nature of the erratic behavior is specific to proprietary features of the legacy systems and is therefore not monitored or logged by the cloud provider's standard audit monitor, pay-per-use monitor or automated scaling listener.

The cloud resource administrators recommend that a new service agent be developed with features customized to monitor the legacy systems.

Which of the following statements provides a solution that can address Organization A's requirements?

Options:

A- The Bare-Metal Provisioning pattern can be applied to remotely provision Physical Server C with the operating system required to deploy the legacy systems. The Automated Administration pattern can be applied to enable Organization A to create custom scripts that can carry out the deployment of the legacy system components via the use of an intelligent automation engine. To provide Organization X with the tools to monitor IT resource usage and collect usage data so that security breaches and other impacts do not occur, the Usage Monitoring pattern can be applied to establish the required custom monitoring functionality.

B- The Bare-Metal Provisioning pattern can be applied to enable Organization A to provisioning Physical Server C with legacy systems after the operating system has been installed. The Synchronized Operating State pattern can be applied to consolidate Organization A's legacy systems via a centralized administration portal from which it can then automate their deployment.
The Automated Administration pattern can be applied to establish a series of workflow scripts customized to monitor and log proprietary legacy system behavior.

C- The Rapid Provisioning pattern can be applied to enable Physical Server C to be remotely provisioned with the operating system and legacy systems. The Centralized Remote Administration pattern can be applied to enable Organization A's employees to remotely manage and administer legacy system deployment. The Pay-as-You-Go pattern can be applied to establish the custom monitoring functionality required by Organization A's legacy systems.

D- None of the above.

Answer:

A

Question 7

Question Type: MultipleChoice

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.

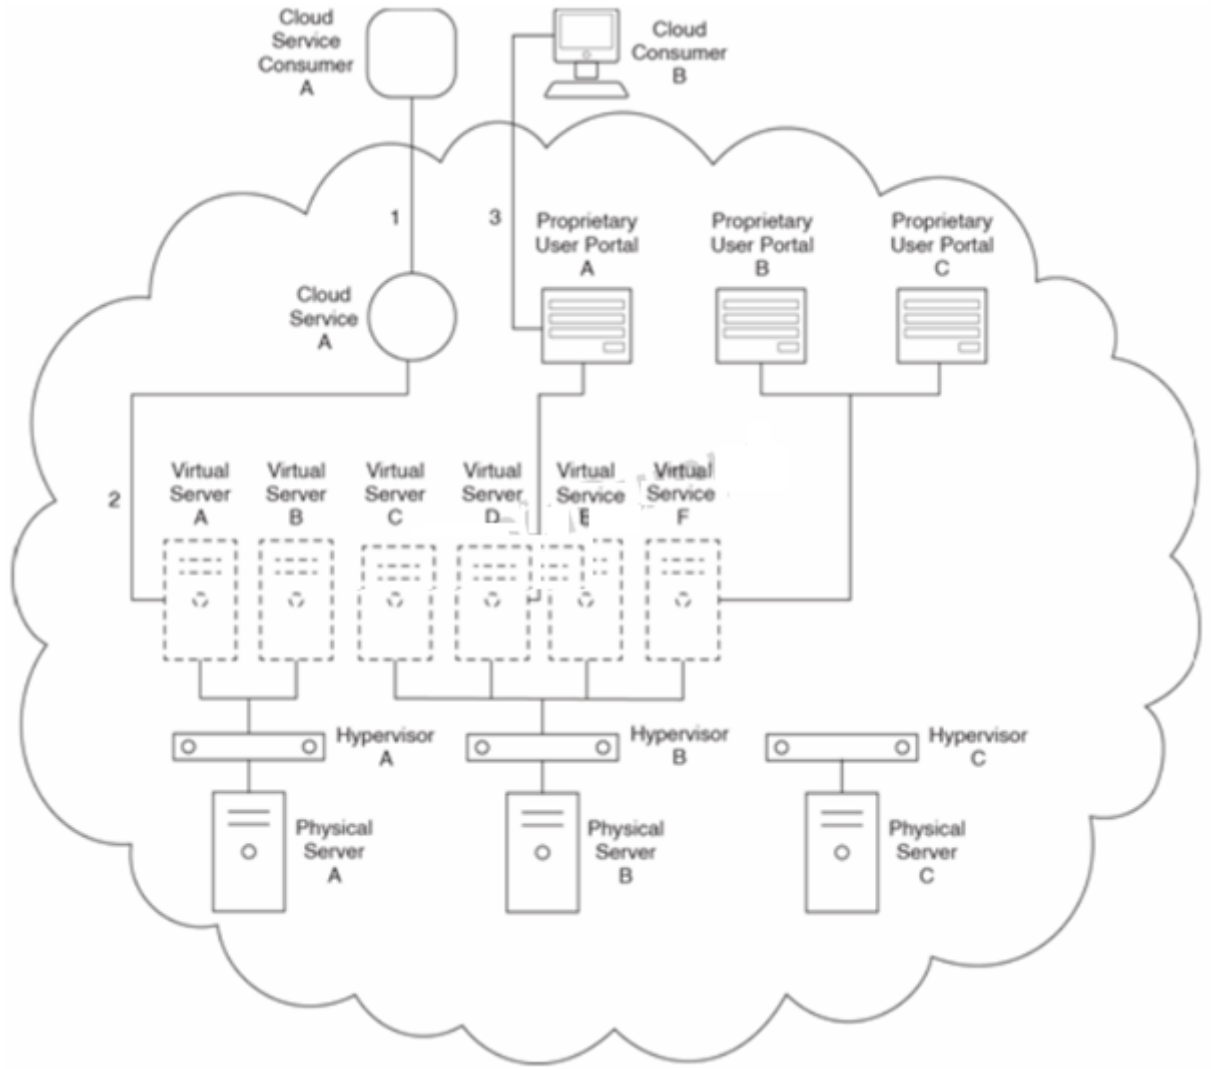

The cloud shown in the figure is a private cloud. Department A and Department B share IT resources within the private cloud and are part of the same organization. Cloud Service Consumer A belongs to Department A and Cloud Consumer B belongs to Department B.

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?

Options:

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 A. 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 8

Question Type: **MultipleChoice**

Cloud Service A is hosted by Virtual Server A. Cloud Storage Device A contains LUN A. Cloud Storage Device A is a multi-tiered cloud storage device with different types of disk groups that perform at different levels. LUN A is located in the disk group with the highest

performance level.

Cloud Service B is hosted by Virtual Server B. Virtual Servers A and B are hosted by HypervisorA,

which is installed on a physical server (not shown) that resides in Cloud A. A redundant implementation of LUN A is replicated synchronously to Cloud Storage Device C. Cloud Storage

Device C does not support multiple types of disk groups and resides in Cloud B, which is located in a different geographic region than Cloud A. Requests that cloud service consumers send to Cloud Services A and B are intercepted by an automated scaling listener responsible for initiating scaling activities.

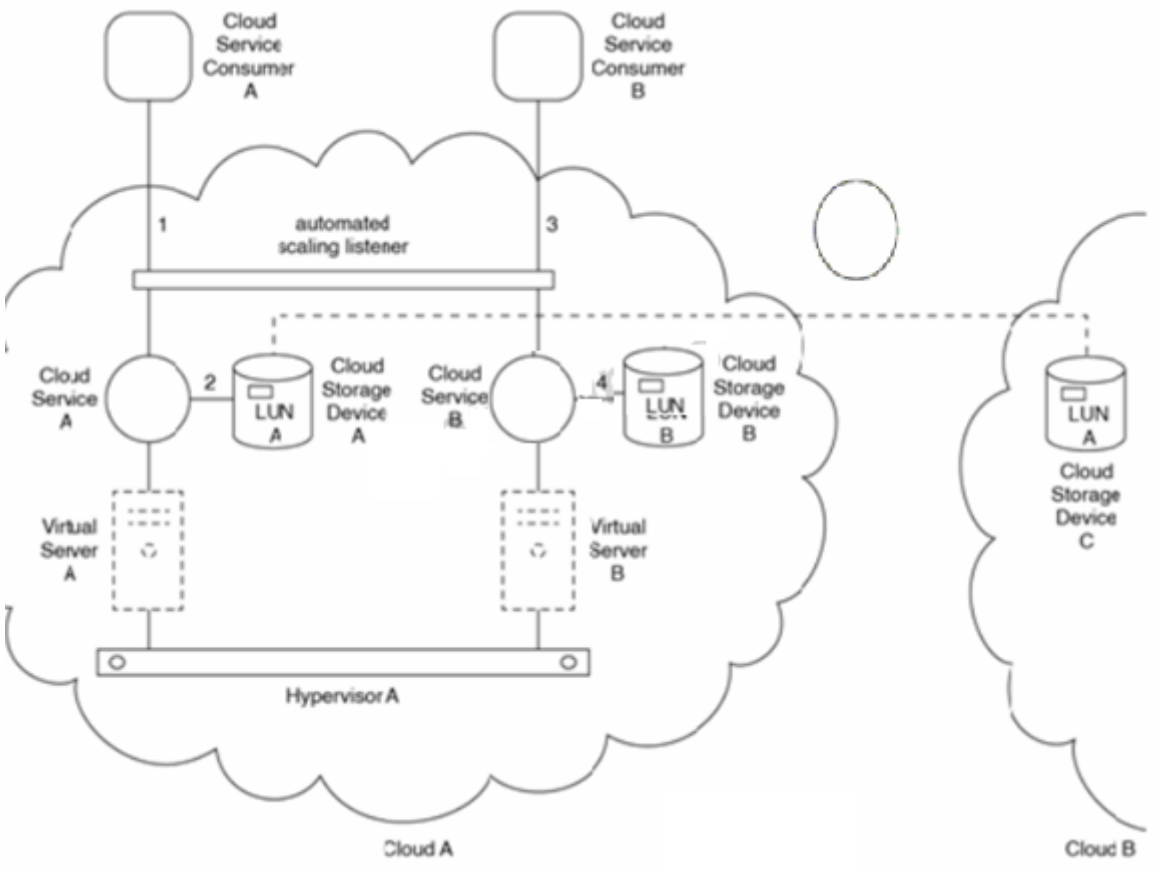

Cloud Service Consumer A issues a request to Cloud Service A (1). To process the request, Cloud Service A accesses LUN Aon Cloud Storage Device A (2). Cloud Service Consumer B issues a request to Cloud Service B (3). To process the request, Cloud Service B accesses LUN B on Cloud Storage Device B (4).

When Cloud Service Consumer A accesses Cloud Service A, there is usually no noticeable performance fluctuation, even during peak usage periods. However, recently, Cloud Storage Device A became unexpectedly unavailable, requiring that Cloud Service A access LUN A on Cloud Storage Device C instead. During the following outage period for Cloud Storage Device A,

Cloud Service Consumer A encounters inconsistent performance from Cloud Service A, including

unusual delays that occur whenever the data requested by Cloud Consumer A isn't cached and

Cloud Service A is required to retrieve the data from LUN A.

Which of the following statements describes a solution that can address this problem?

Options:

A- The Storage Maintenance Window pattern can be applied so that future outages of Cloud Storage Device A do not occur unexpectedly. The Resource Pooling and Resource Reservation patterns can be further applied to establish a resource pool on Cloud A that has resources reserved specifically for Cloud Service A. This will prevent other cloud service consumers, such as Cloud Service Consumer B, from competing for Cloud Service A's resources.

B- The Shared Resources pattern can be applied to prevent Cloud Service A from encountering performance issues when IT resources hosted by Hypervisor A are accessed by other cloud service consumers. The Cross-Storage Device Vertical Tiering pattern can be applied to enable Cloud Storage Device A to scale to a higher performance disk type when an outage occurs.

C- The Cloud Balancing pattern can be applied to enable Cloud Service A to switch over to Cloud Storage Device C if Cloud Storage Device A becomes unavailable. The Dynamic Data Normalization pattern can be further applied to streamline and reduce the quantity of the data being stored by LUN A within Cloud Storage Device A, so as to correspondingly reduce the performance impacts during high usage volumes.

D- None of the above.

Answer:

D

Question 9

Question Type: MultipleChoice

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).

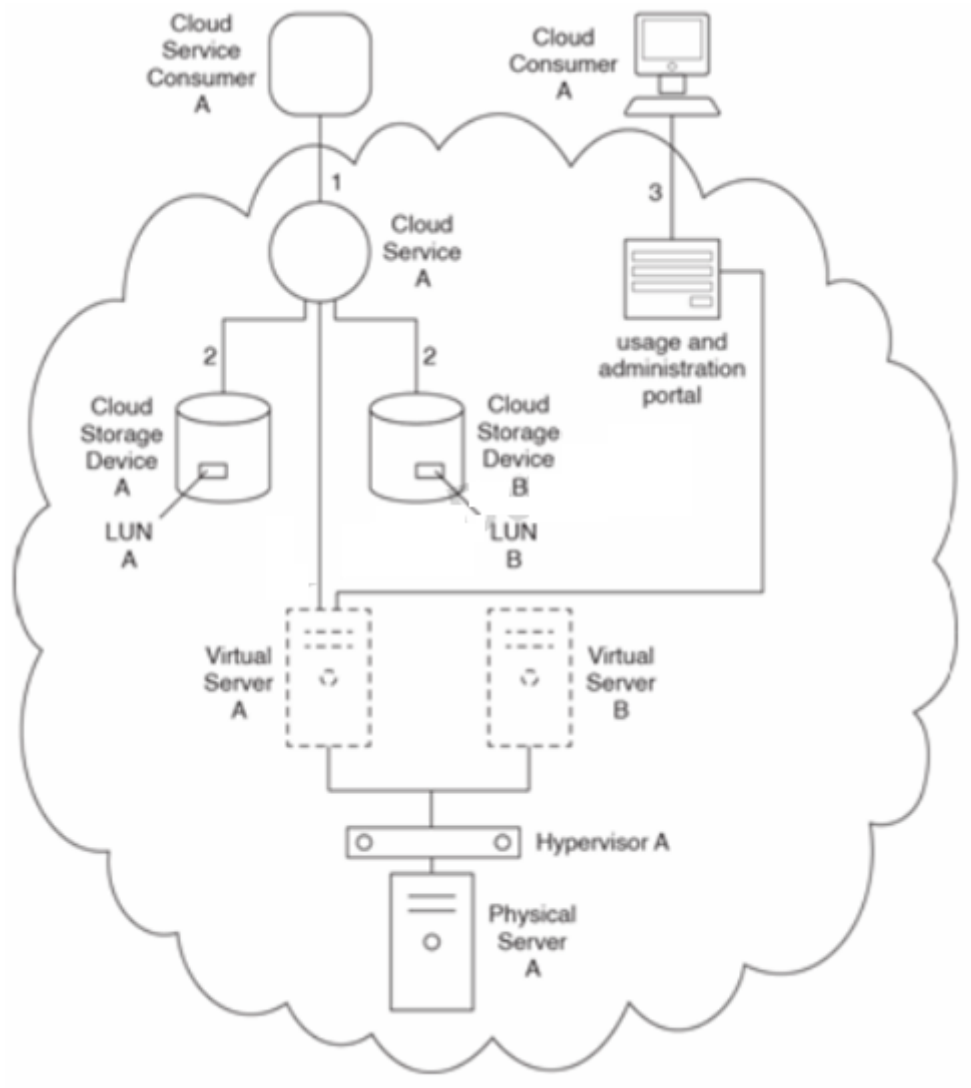

Cloud Service Consumer A and Cloud Consumer A belong to Organization A, which is leasing an laaS environment from the cloud provider.

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?

Options:

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 10

Question Type: **MultipleChoice**

Virtual Server A is hosted by Hypervisor A, which resides on Physical Server A. Virtual Server A

hosts Cloud Services A and B. Virtual Server B is hosted by Hypervisor B on Physical Server B.

Physical Server C is currently not being used.

Cloud Service Consumer A sends a request to Cloud Service A that is intercepted by Pay-Per-Use

Monitor A

(1), which collects billing-related usage data that is later forwarded to the billing management system

(2). Cloud Service A receives and processes the request

(3). Cloud Consumer B accesses the usage and administration portal

(4) to access data on Cloud Storage Device B. Pay-Per-Use Monitor B intercepts the data access to collect and forward billing-related usage data to the billing management system

(5). Cloud Storage Device B processes the data access request from Cloud Consumer B

(6).

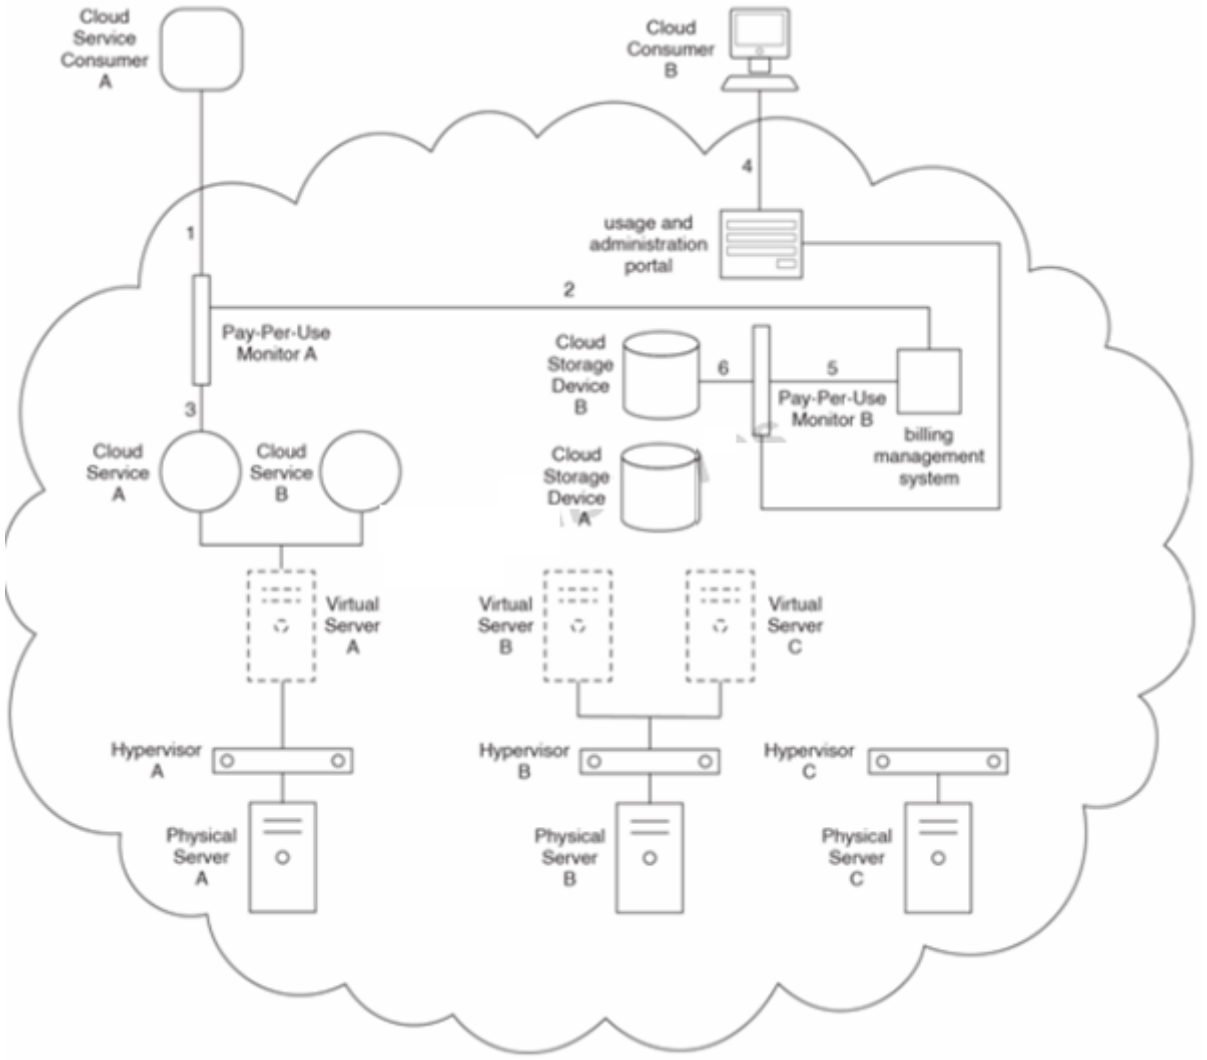

Cloud Service Consumer A and Cloud Consumer B belong to Organization A

Cloud Storage Device B is accessed on a regular basis by Cloud Consumer B. However,

managers at Organization A receive reports from their cloud resource administrator that Cloud Storage Device B is unavailable for longer periods and more frequently than what they expected, based on the SLA availability guarantee they were provided by the cloud provider. This results in wasted time when the cloud resource administrator attempts to upload or access data and then discovers that Cloud Storage Device B is unavailable. The cloud resource administrator requires a means of checking for the availability of Cloud Storage Device B prior to attempting access.

As the workload increases on Physical Server B, Cloud Consumer B begins to receive runtime exceptions and degraded data access performance from Cloud Storage Device B. It is determined that the cause of the deteriorating performance is a network bottleneck that has formed on Physical Server B due to its bandwidth capacity being reached, primarily because of other cloud consumer organizations also sharing its hosted IT resources.

Organization A receives a monthly billing statement that shows the charges for the total usage of Cloud Service A during that period. However, Organization A requires a more detailed breakdown of the types of usage and their associated costs. For example, Cloud Service Consumer A can issue requests for information by employees within Organization A and on behalf of clients of Organization A. Organization A requires a breakdown of the usage costs incurred on behalf of clients so that it can bill the clients for this usage accordingly. The cloud provider informs Organization A that it has no existing monitor that can collect and log this detailed usage information and suggests that Organization A customize its own monitor.

Which of the following statements lists the patterns that can be applied to solve these three problems?

Options:

A- Real-time Resource Availability, Elastic Network Capacity, Usage Monitoring

B- Persistent Virtual Network Configuration, Elastic Network Capacity, Load Balanced Virtual Server Instances

C- Load Balanced Virtual Switches, Elastic Resource Capacity, Automated Administration

D- None of the above.

Answer:

A

Question 11

Question Type: **MultipleChoice**

Cloud Service A is installed on Virtual Server A and the database it accesses is located on Cloud Storage Device A. Both Virtual Servers A and B are hosted by Hypervisor A. Requests from cloud

service consumers are intercepted by an automated scaling listener that automatically routes subsequent requests to additional instances of Cloud Service A whenever the given usage of an instance exceeds two concurrent requests.

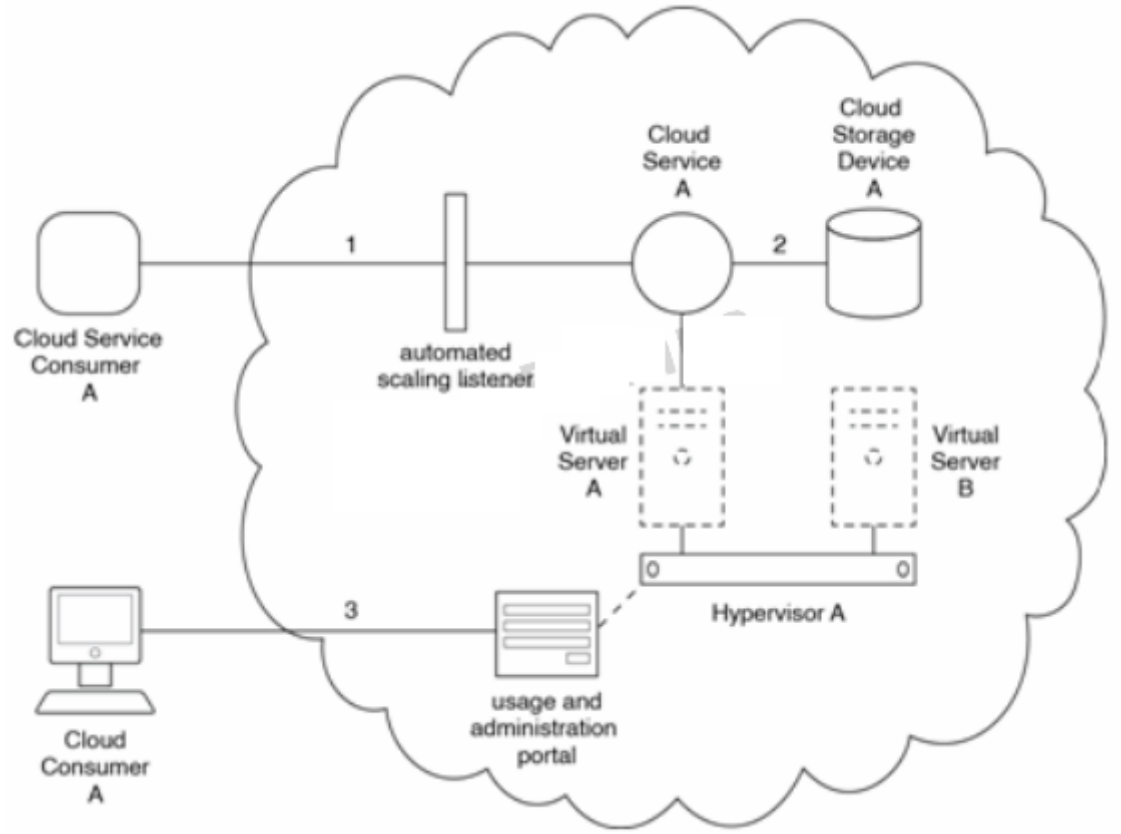

Cloud Service Consumer A accesses Cloud Service A (1), which either sends a query or a read/write request to a database on Cloud Storage Device A (2). A usage and administration portal is available, enabling Cloud Consumer A to view the billing and usage history of Virtual Servers A and B (3).

Cloud Service Consumer A and Cloud Consumer A are owned by Organization A. which performs several tests on the cloud architecture that produce the following results:

A stress test is performed to generate workloads on Virtual Servers A and B to gauge their load capacity. This test reveals that both virtual servers have firm workload thresholds. If the workload capacity on either virtual server reaches its threshold, further processing requests are rejected.

An availability test shows that Cloud Service A becomes unavailable whenever Hypervisor A crashes.

A security test is carried out during which the cloud architecture is accessed by a malicious cloud consumer that disables the path used by Cloud Service A to access Cloud Storage Device A, thereby causing all subsequent cloud service consumer requests to be replied to with data access errors.

Which of the following statements describes a solution that addresses the concerns raised by the three tests?

Options:

A- The Resource Reservation pattern can be applied to ensure that Virtual Servers A and B are not accessed by any cloud consumers other than Organization A, thereby enabling their respective capacity to be maximized. A second hypervisor can be implemented and the Synchronized Operating State pattern can be applied to emulate the usage of the resource cluster mechanism with the two hypervisors. This will prevent Cloud Service A from being affected if one of the hypervisors fails. The Service State Management pattern can be applied to establish a secondary cloud storage device that can be accessed by Cloud Service A whenever Cloud Storage Device A becomes inaccessible.

B- The Elastic Resource Capacity pattern can be applied to enable resources to be assigned to the virtual servers dynamically. The Hypervisor Clustering pattern can be applied to avoid jeopardizing the availability of Cloud Service A when its underlying hypervisor fails. The Multipath Resource Access pattern can be applied to establish an alternative path to Cloud Storage Device A. Cloud Service A can then be designed to access Cloud Storage Device A via the alternative path whenever access via the original path fails.

C- The Elastic Resource Capacity pattern can be applied to enable resources to be assigned to the virtual servers dynamically. The Resource Pooling pattern can be applied to allow Hypervisor A to be part of a larger hypervisor pool. The Cross-Storage Device Vertical Tiering pattern can be applied to allow Cloud Service A to access Cloud Storage Device A via different tiers.

D- None of the above.

Answer:

B

Question 12

Question Type: MultipleChoice

Cloud Service A is hosted by Virtual Server A, which is hosted by Hypervisor A on Physical Server A. Virtual Server B is hosted by Hypervisor B on Physical Server B.

Cloud Service Consumer A accesses Cloud Service A and the request is intercepted by an SLA monitor (1). Cloud Service A receives the request (2) and accesses Cloud Storage Device A and Cloud Storage Device B (3).

Cloud Service Consumer A belongs to Organization A, which is leasing all of the IT resources shown in the figure as part of an laaS environment.

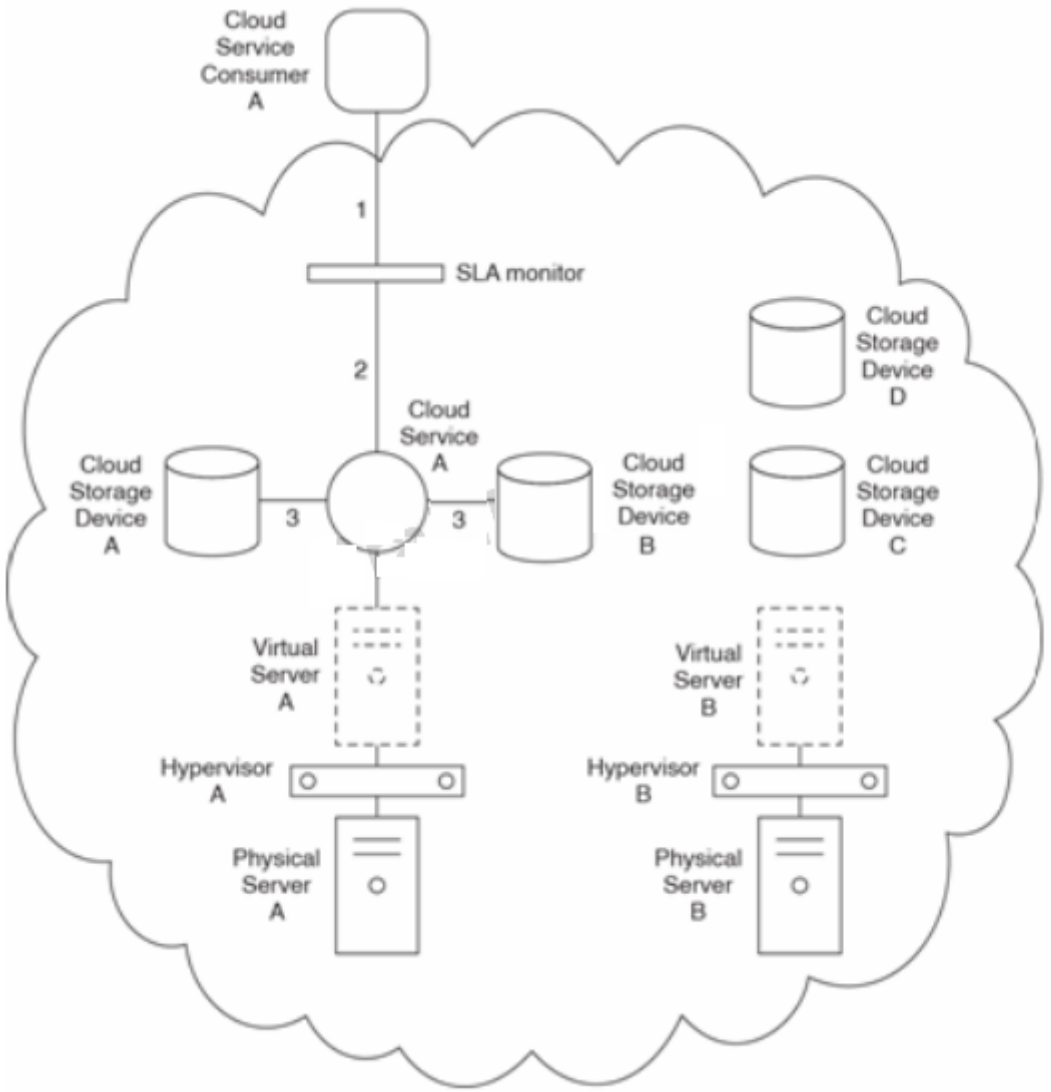
Cloud
Service
Consumer
A
1
SLA monitor
2
Cloud
Service
A
Cloud
Storage
Device
D
Cloud
Storage
Device
A
3
3
Cloud
Storage
Device
B
Cloud
Storage
Device
C
Virtual
Server
A
Virtual
Server
B
Hypervisor
A
Hypervisor
B
Physical
Server
A
Physical
Server
B

Cloud Storage Device B has a higher performance capacity than Cloud Storage Device A. Cloud Storage Device C has a higher performance capacity than Cloud Storage Device B. The requests being received by Cloud Service A from Cloud Service Consumer A have recently increased in both quantity and in the amount of data being queried, written and read from Cloud Storage Device A. As a result, Cloud Storage Device A's capacity is frequently reached and it has become unstable at times, timing out with some requests and rejecting other requests.

Cloud Storage Device C is used by Organization A to store backup data on a daily basis. One day, a hardware failure within Cloud Storage Device C results in the permanent loss of data.

Organization A requires a system that will prevent this type of failure from resulting in data loss.

The cloud provider is planning to implement a routine maintenance schedule for Cloud Storage Devices A, B, and C and issues a notice stating that the new schedule will start next week. An outage of 30 minutes every Thursday and Sunday at 8:00 PM is needed for the maintenance tasks. Upon hearing this, Organization A complains that they cannot afford to have Cloud Storage Devices A and B become inoperable, especially not during the weekdays.

Which of the following statements describes a solution that can address Organization A's issues?

Options:

A- The Intra-Storage Device Vertical Data Tiering pattern can be applied to enable dynamic scaling between Cloud Storage Devices A, B and C. The Dynamic Failure Detection and Recovery pattern can be applied to establish a resilient watchdog system that is able to respond dynamically to prevent data loss. The Service State Management pattern can be applied to keep a copy of the data in Cloud Storage Devices A, B and C during the maintenance outages.

B- The Cross-Storage Device Vertical Tiering pattern can be applied to enable dynamic scaling between Cloud Storage Devices A, B and C. The Redundant Storage pattern can be applied by designating Cloud Storage Device D as the secondary storage to which

Organization A's data can be replicated. In order to prevent planned or unplanned outages from affecting Organization A's data access, the Storage Maintenance Window pattern can be applied to replicate the data in Cloud Storage Device D for retrieval before the outages begin.

C- The Load Balanced Virtual Switches pattern can be applied to increase the bandwidth of Physical Server A so that data processing problems within Cloud Storage Device A can be prevented. The Non-Disruptive Service Relocation pattern can be applied to automatically relocate Cloud Storage Device A to Physical Server B so that data access is not interrupted. The Storage Maintenance Window pattern can be applied to replicate the data in Cloud Storage Device D for retrieval before the outages begin.

D- None of the above.

Answer:

B