Free Questions for C90.06 by ebraindumps

Shared by Bowman on 06-06-2022

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

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.

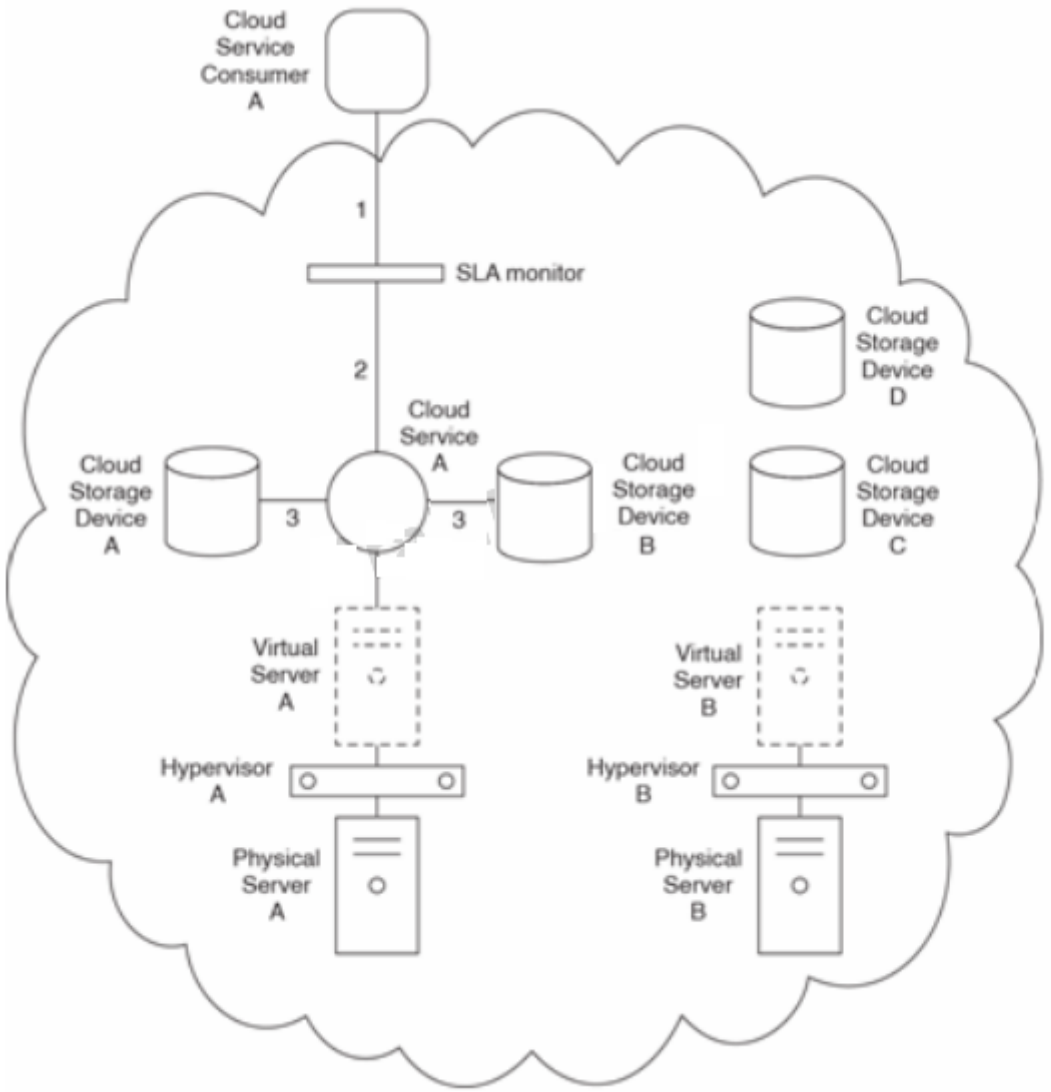
Cloud
Service
Consumer
A
1
SLA monitor
2
Cloud
Service
A
Cloud
Storage
Device
D
Cloud
Storage
Device
A
3
3
Cloud
Storage
Device
B
Cloud
Storage
Device
C
Virtual
Server
A
Virtual
Server
B
Hypervisor
A
Hypervisor
B
Physical
Server
A
Physical
Server
B

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

Question 2

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.

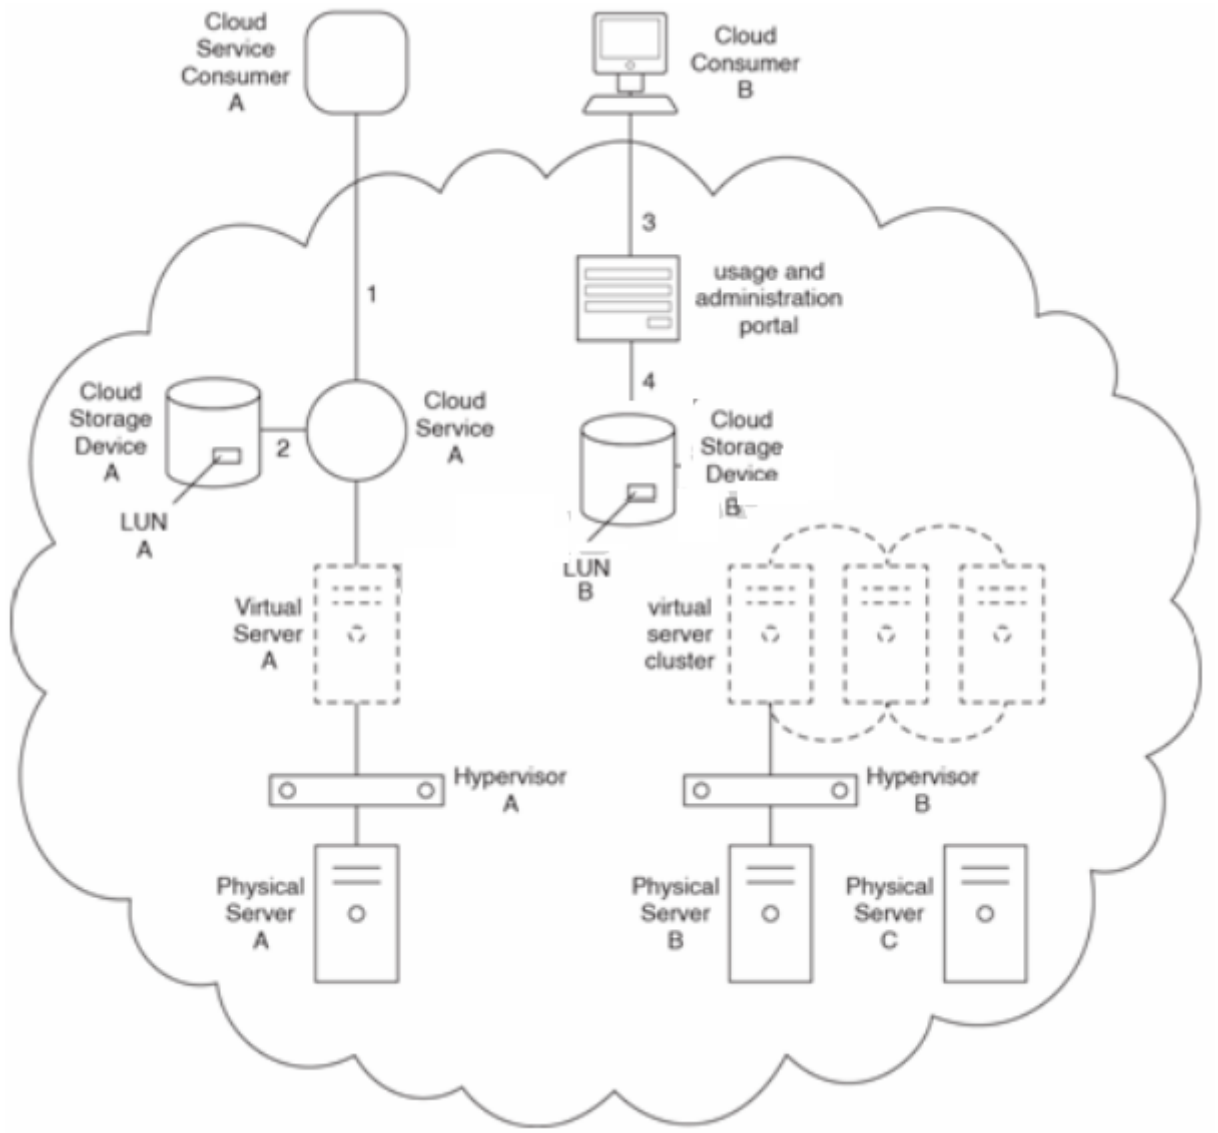
Cloud Service Consumer A
Cloud Consumer B
1
3
usage and administration portal
4
Cloud Storage Device A
2
Cloud Service A
Cloud Storage Device B
LUN A
LUN B
Virtual Server A
virtual server cluster
Hypervisor A
Hypervisor B
Physical Server A
Physical Server B
Physical Server C

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 3

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.

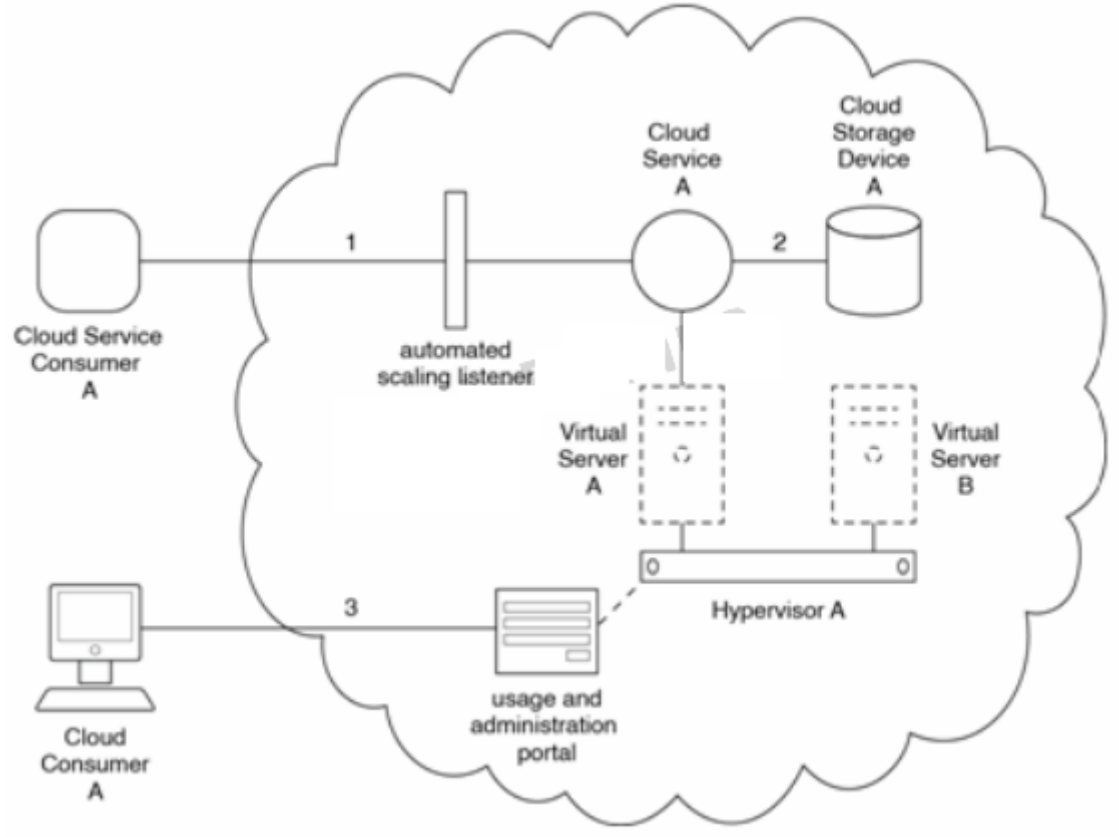

Cloud Service Consumer A accesses Cloud Service A (1), which either sends a query or a read/write request to a database on Cloud Storage Device A (2). A usage and administration portal is available, enabling Cloud Consumer A to view the billing and usage history of Virtual Servers A and B (3).

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 4

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.

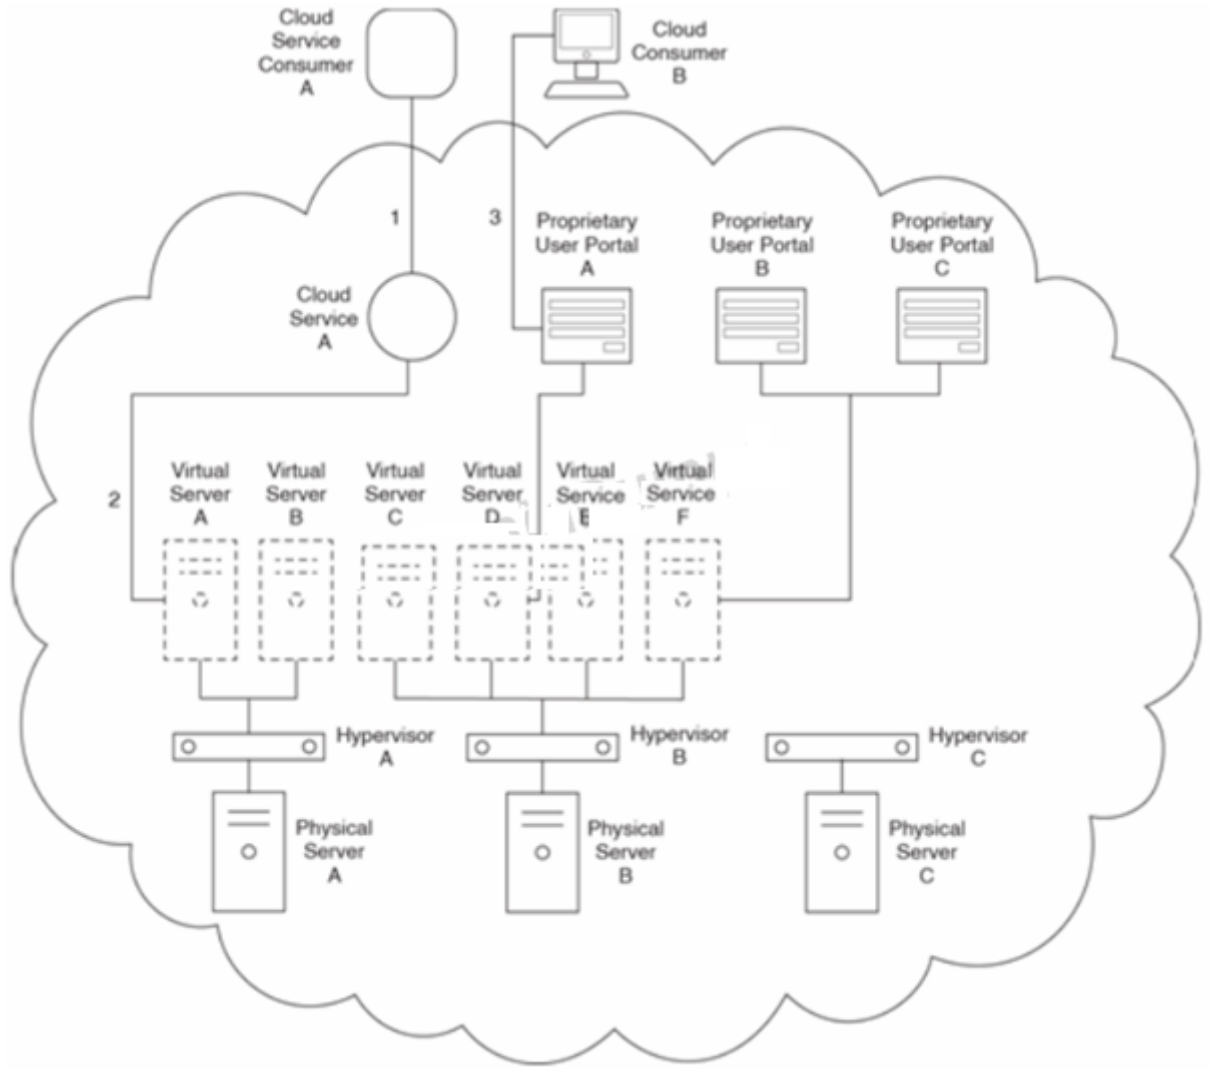

The cloud shown in the figure is a private cloud. Department A and Department B share IT resources within the private cloud and are part of the same organization. Cloud Service Consumer A belongs to Department A and Cloud Consumer B belongs to Department B.

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 5

Question Type: **MultipleChoice**

Ready-Made Environment A is hosted by Virtual Server A and Ready-Made Environments is hosted by Virtual Server B. Virtual Servers A and B are hosted by Hypervisor A, which is part of a hypervisor cluster. An automated scaling listener intercepts cloud consumer

requests and automatically invokes the on-demand generation of additional instances of ready-made environments, as required.

A self-service portal and a usage and administration portal are also available to cloud consumers.

The self-service portal can be used to request the provisioning of a new ready-made environment.

Any cloud consumer that has already had a ready-made environment provisioned can configure and view information about that ready-made environment via the usage and administration portal.

Cloud Consumer A accesses Ready-Made Environment A to work on the development of a new cloud service (1). Cloud Consumer B accesses Ready-Made Environment B to test a recently completed application comprised of three cloud services (2). Cloud Consumer C accesses the self-service portal to request the creation of a new ready-made environment (3).

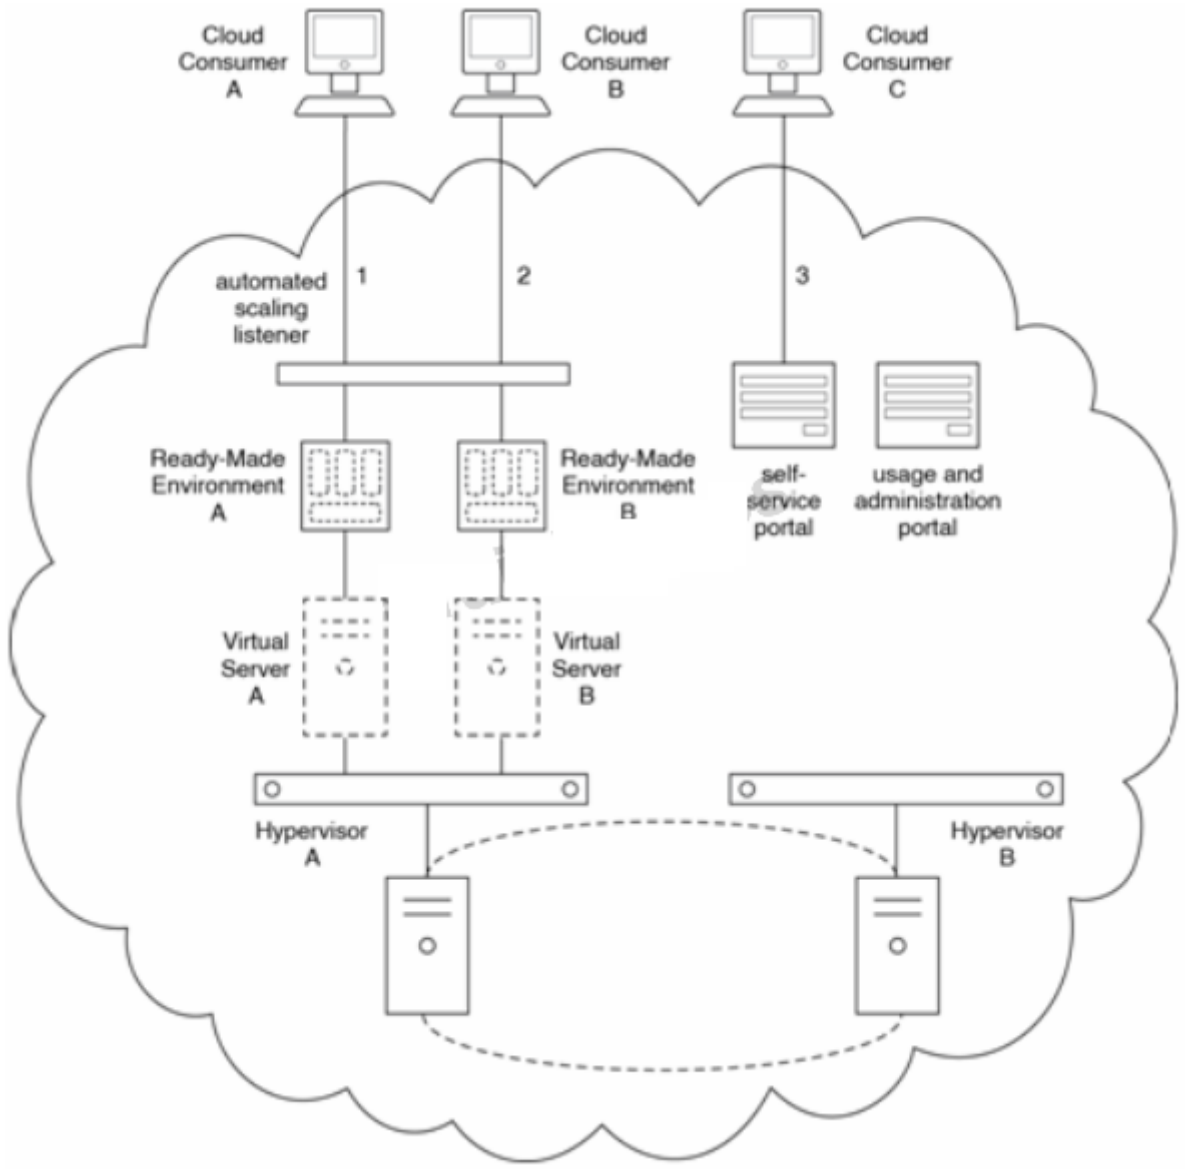
Cloud Consumer A
Cloud Consumer B
Cloud Consumer C
automated scaling listener
1
2
3
Ready-Made Environment A
Ready-Made Environment B
self-service portal
usage and administration portal
Virtual Server A
Virtual Server B
Hypervisor A
Hypervisor B

The cloud provider is required to perform an emergency maintenance outage on a cloud storage device used by all ready-made environments. The unplanned outage takes two hours. During this period, Cloud Consumers A and B are unable to access Ready-Made Environments A and B and Cloud Consumer C receives an error when submitting a request to create a new ready-made environment.

After the maintenance outage is over, Cloud Consumers A and B encounter the following problems:

Cloud Consumer A is unable to recover session data that was kept in memory for an extended period, prior to the time of the outage.

Cloud Consumer B has no access to Virtual Server B, which was moved to Hypervisor B during the maintenance outage. When Cloud Consumer B attempts to ping Virtual Server B, the request times out.

Even though Cloud Consumer C is able to log into the usage and administration portal to confirm that its ready-made environment was successfully provisioned, the unexpected outage has raised concerns about the stability of the ready-made environment's underlying infrastructure. Cloud Consumer C informs the cloud provider that it cannot proceed with its lease of the ready-made environment if there are future occurrences of this type of maintenance outage.

Which of the following statements can help address the problems and concerns of the three cloud consumers?

Options:

A- A combination of the Load Balanced Virtual Server Instances and Synchronized Operating State patterns can be applied to establish a system capable of deferring state across multiple cloud storage devices, each located on a different virtual server. The Elastic Disk Provisioning pattern can be applied to persist virtual server configuration data across hypervisors so that connectivity is preserved whenever a virtual server is relocated to a different hypervisor. The Zero Downtime pattern can be applied to ensure that none of the ready-made environments or virtual servers are subject to a maintenance outage in the future.

B- The Elastic Disk Provisioning and Cross-Storage Device Vertical Tiering patterns can be applied to establish a cloud architecture that supports a set of cloud storage devices, each with different tiers that cloud consumers can choose to scale to The Synchronized Operating State pattern can be applied in combination with the Hypervisor Clustering pattern to avoid further virtual server and ready-made environment connectivity problems. The Redundant Storage pattern can be applied so that if a cloud storage device fails, a secondary implementation of the cloud storage device automatically takes over processing tasks, thereby avoiding outages.

C- The Service State Management pattern can be applied to establish a system that can persist session data in a database. The Persistent Virtual Network Configuration pattern can be applied to centralize the configuration data necessary for virtual servers to remain accessible after they have been relocated to different hypervisors. The Storage Maintenance Window pattern can be applied to establish a system that allows cloud storage devices to be maintained without causing outages.

D- None of the above.

Answer:

C

Question 6

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).

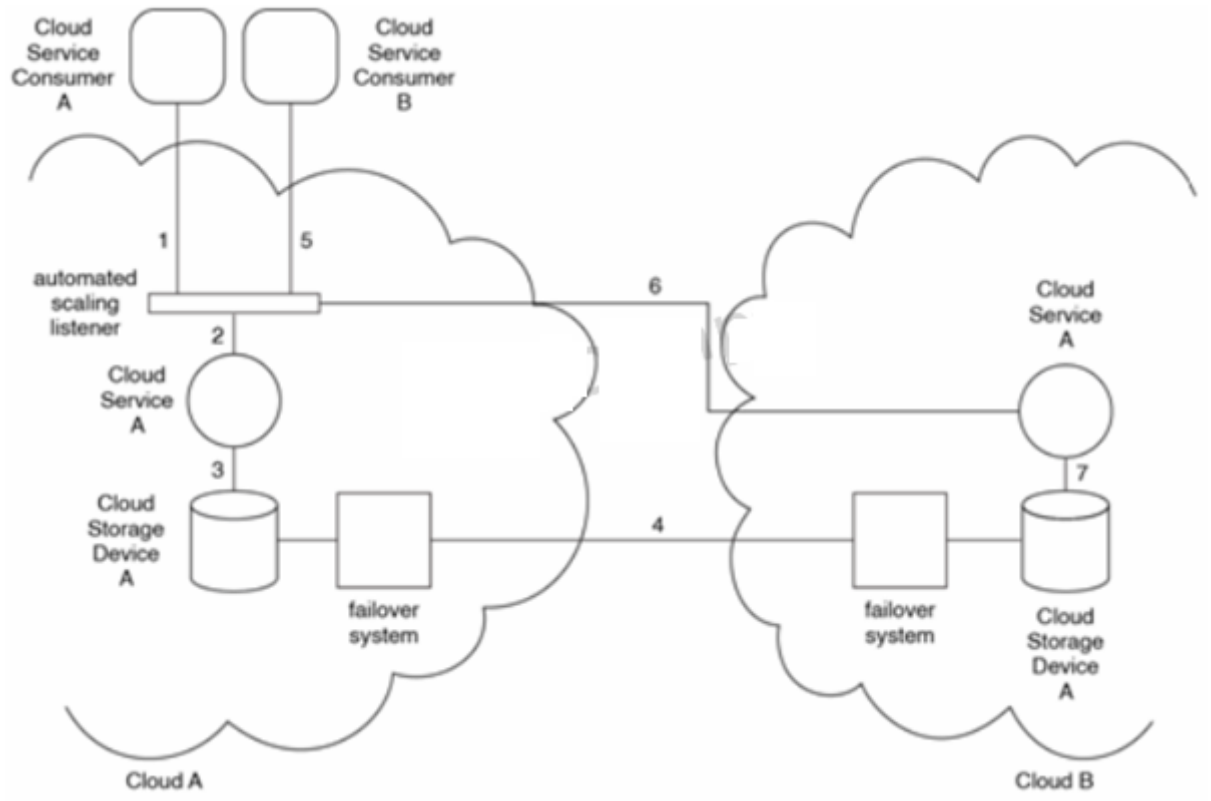

An unexpected outage occurs in Cloud A, making Cloud Service A unavailable. Organization A notices that its cloud resource administrator can continue accessing data in Cloud Storage Device A via a usage and administration portal. Cloud Service Consumer A is unable to access data in Cloud Storage Device 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 7

Question Type: **MultipleChoice**

Cloud Sen/ice A is hosted by Virtual Server A, which is hosted by Hypervisor A on Physical Server A. Cloud Service B is hosted by Virtual Server B. Virtual Server C hosts Cloud Services C and D.

Virtual Server B and Virtual Server C are hosted by Hypervisor B on Physical Server B.

Cloud Service Consumer A accesses Cloud Service A (1). Cloud Service Consumer B accesses Cloud Service A (2). Cloud Service Consumer C accesses Cloud Service A (3) and then accesses Cloud Service B (4).

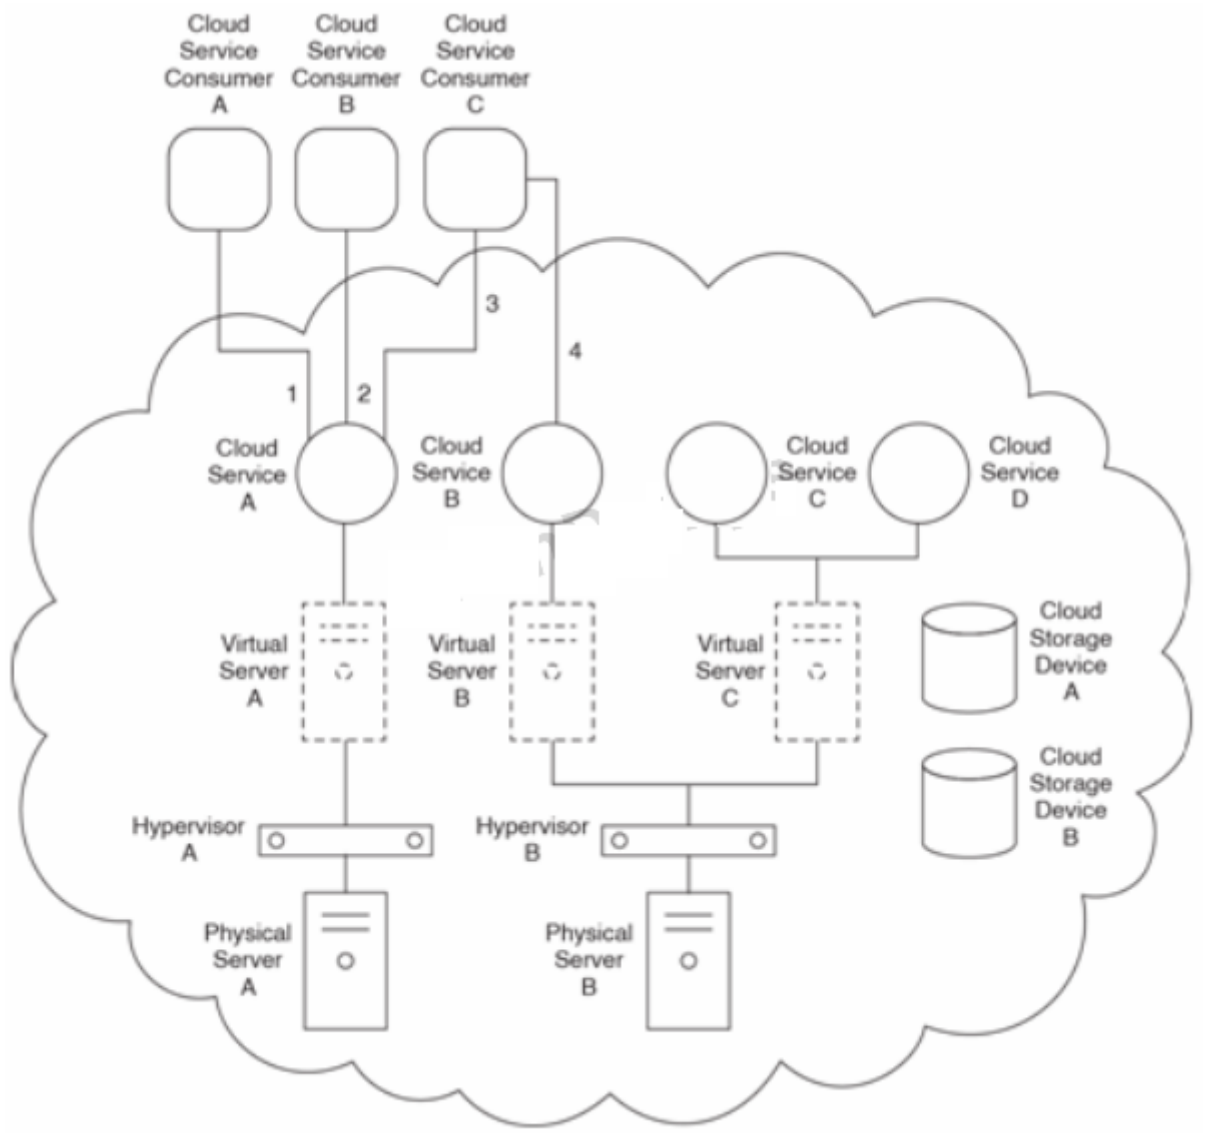
Cloud Service Consumer A
Cloud Service Consumer B
Cloud Service Consumer C
1
2
3
4
Cloud Service A
Cloud Service B
Cloud Service C
Cloud Service D
Virtual Server A
Virtual Server B
Virtual Server C
Cloud Storage Device A
Cloud Storage Device B
Hypervisor A
Hypervisor B
Physical Server A
Physical Server B

Cloud Service Consumers A, B and C simultaneously access Cloud Service A. Cloud Service Consumer C receives a runtime exception and its request for access is rejected. It is determined that Cloud Service Consumer C attempted to upload a large amount of input data for Cloud Service A, which exceeded the bandwidth threshold of the virtual network. The cloud architecture needs to be improved to avoid this from happening again.

Cloud Service Consumer C's repeated access of Cloud Service B imposes workloads that are large and highly unpredictable. After some time, Cloud Service B begins to delay its responses and sometimes times out entirely. The cloud resource administrator discovers that Virtual Server B is unstable and close to failure primarily because its CPU and memory resources are being used to their maximum capacity.

Cloud Services C and D are being positioned as SaaS products for use by a range of cloud consumer organizations. After their initial release, they begin to quickly use up the available memory in Virtual Server C, primarily because of the large amounts of state and session data they need to place into memory for extended periods.

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

Options:

A- Elastic Network Capacity, Load Balanced Virtual Server Instances, Service State Management

B- Elastic Resource Capacity, Service Load Balancing, Synchronized Operating State

C- Persistent Virtual Network Configuration. Load Balanced Virtual Switches, Service State Management

D- None of the above.

Answer:

A

Question 8

Question Type: MultipleChoice

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.

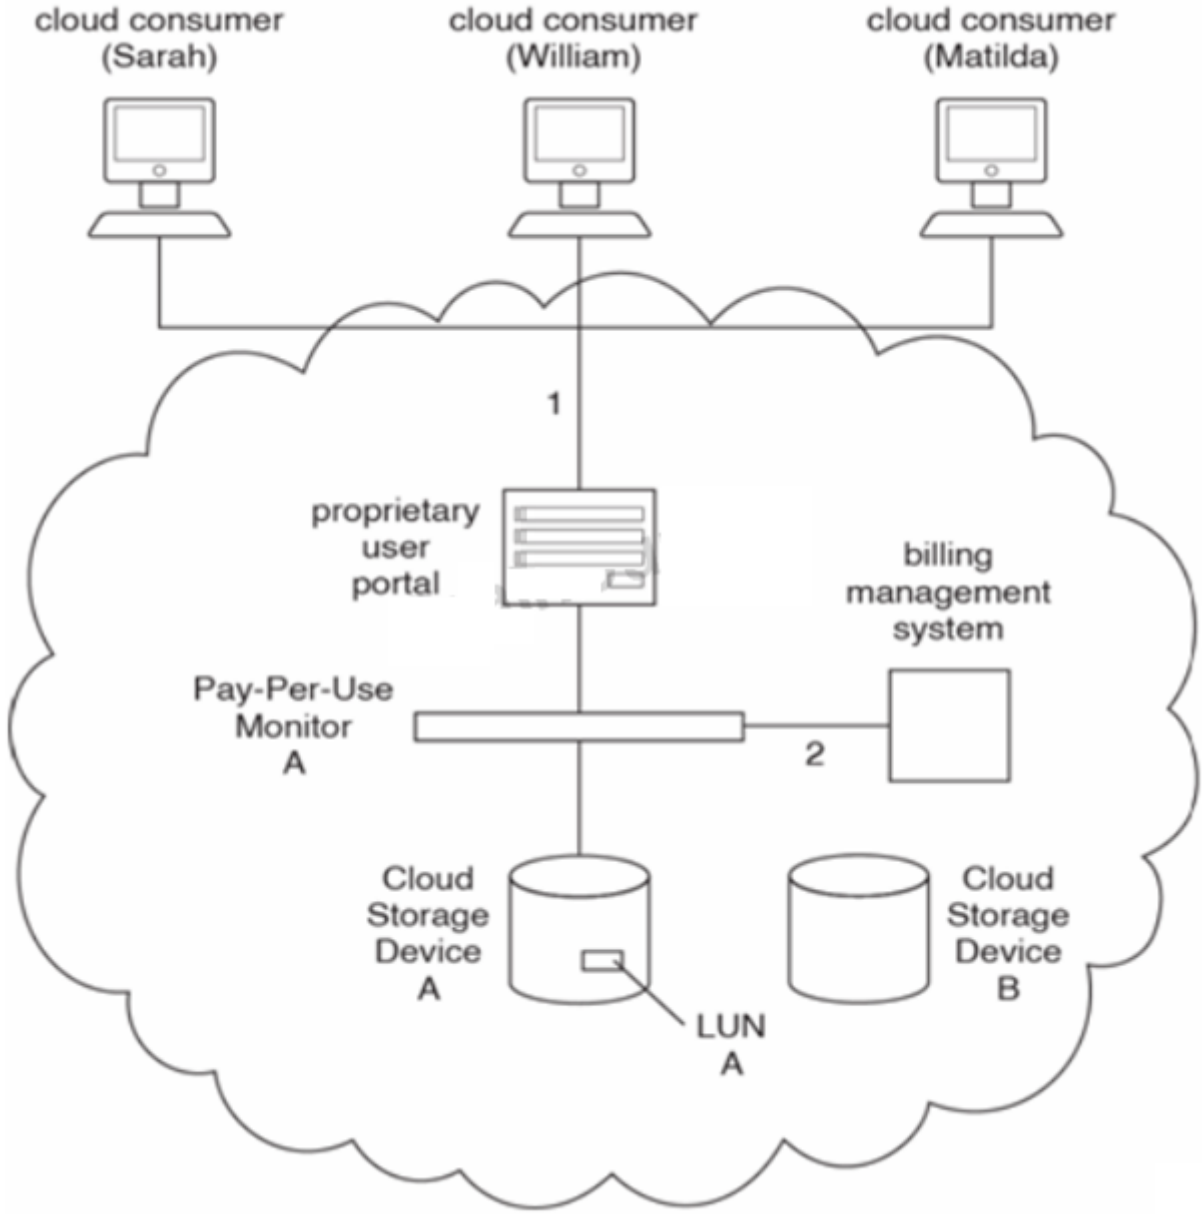
cloud consumer
(Sarah)
cloud consumer
(William)
cloud consumer
(Matilda)
1
proprietary
user
portal
billing
management
system
Pay-Per-Use
Monitor
A
2
Cloud
Storage
Device
A
Cloud
Storage
Device
B
LUN
A

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:

Sara signs up for the discount backup plan and is allocated 50 GBs of storage space. A week later, she uploads 10 GBs of backup data. A week after that, she uploads another 35 GBs. She later finds out that for those two weeks her organization was billed $200, which is more than she was expecting. After she complains to the cloud provider, she learns about how fixed-disk storage allocation is billed. She asks the cloud provider to change her account to a different storage plan where she is only billed for the actual amount of storage space she uses at any given time. The cloud provider assures her that a new system will be set up to accommodate this request.

William is on the premium backup plan. He uploads 1 to 3 GBs of important business data every day. After two weeks of daily uploads, he realizes that the centralized nature of the backup data on the cloud makes it more convenient for reporting purposes than the distributed nature of the same data in his on-premise environment. He uses an analysis tool to run queries against the cloud-based data. However, due to the large quantity of redundant data, the queries end up being ineffective and take too long to run. He asks the cloud provider to find a solution that can streamline the cloud-based data by reducing redundancy. By reducing the overall quantity of the data, the analysis queries will run faster.

Matilda is on the discount backup plan and uploads backup data on a daily basis. Over the course of two weeks she uploads over 200 GBs of data. She performs a daily backup at the end of each day by identifying the data to back up and then using the proprietary user portal to run the cloud backup procedure. This procedure can take 5 to 45 minutes, depending on the amount of data she is uploading. Matilda performs this as her last task of the day and therefore initiates the procedure before she leaves, but does not wait for it to complete. One day, she receives an e-mail from the cloud provider explaining that the backup procedure from the previous day had failed due to an unexpected hardware failure that occurred on Cloud Storage Device A. The notification e-mail goes on to state that this type of failure falls within the 97% availability guarantee of her organization's SLA, and is therefore in compliance with the current provisioning agreement. Had a disaster occurred that night, the on-premise data could have been lost and Matilda would be held accountable. Matilda contacts the cloud provider to demand that the provisioning agreement be amended to upgrade their existing SLA to the maximum possible availability (which, for this cloud provider, is 99.999%). The cloud provider agrees to establish a system to accommodate this request.

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

Options:

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

Question 9

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).

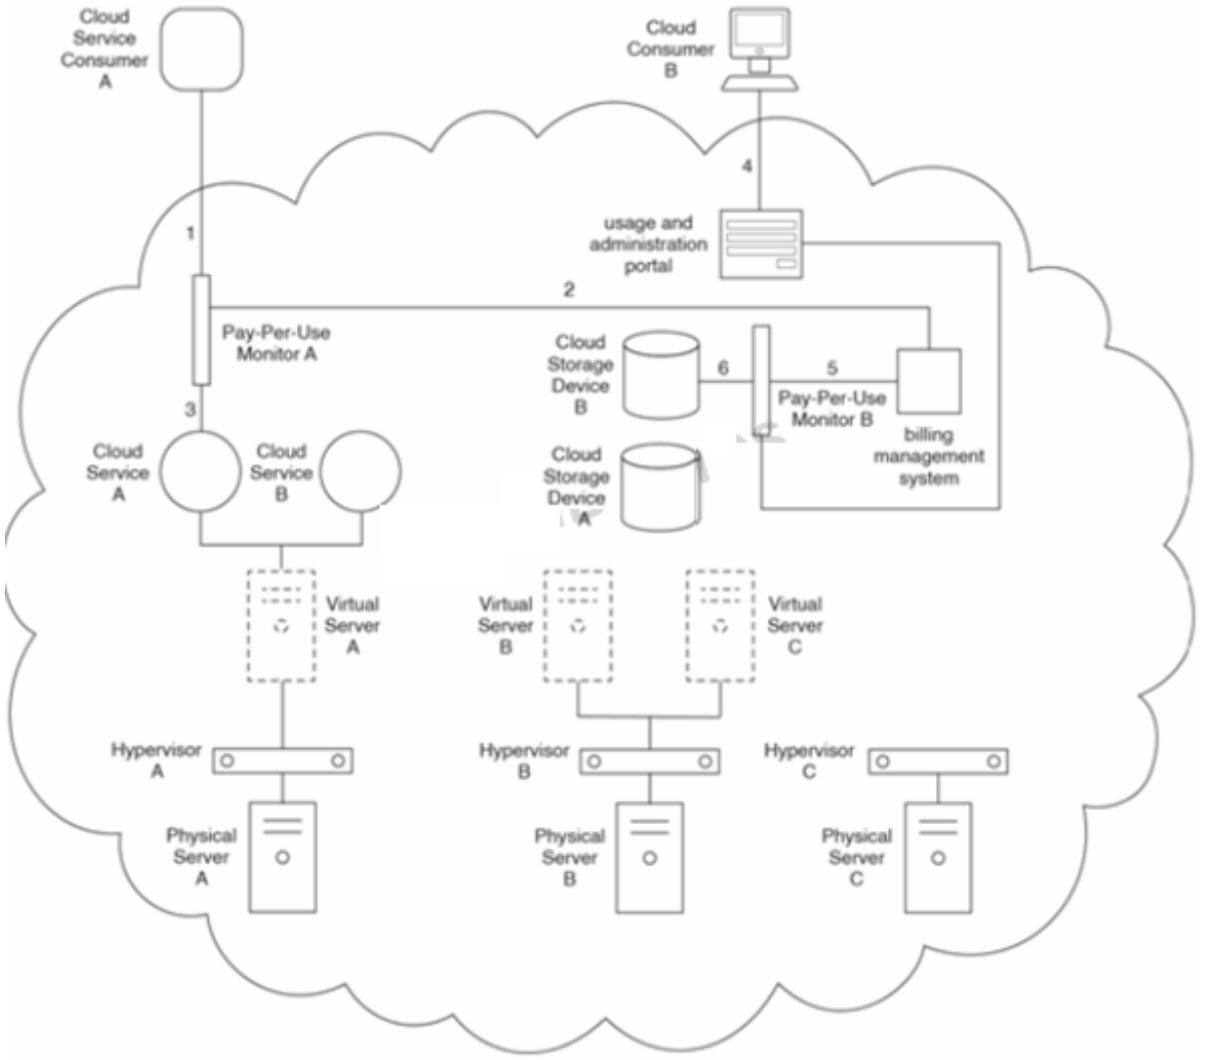

Cloud Service Consumer A and Cloud Consumer B belong to Organization A

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 10

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.

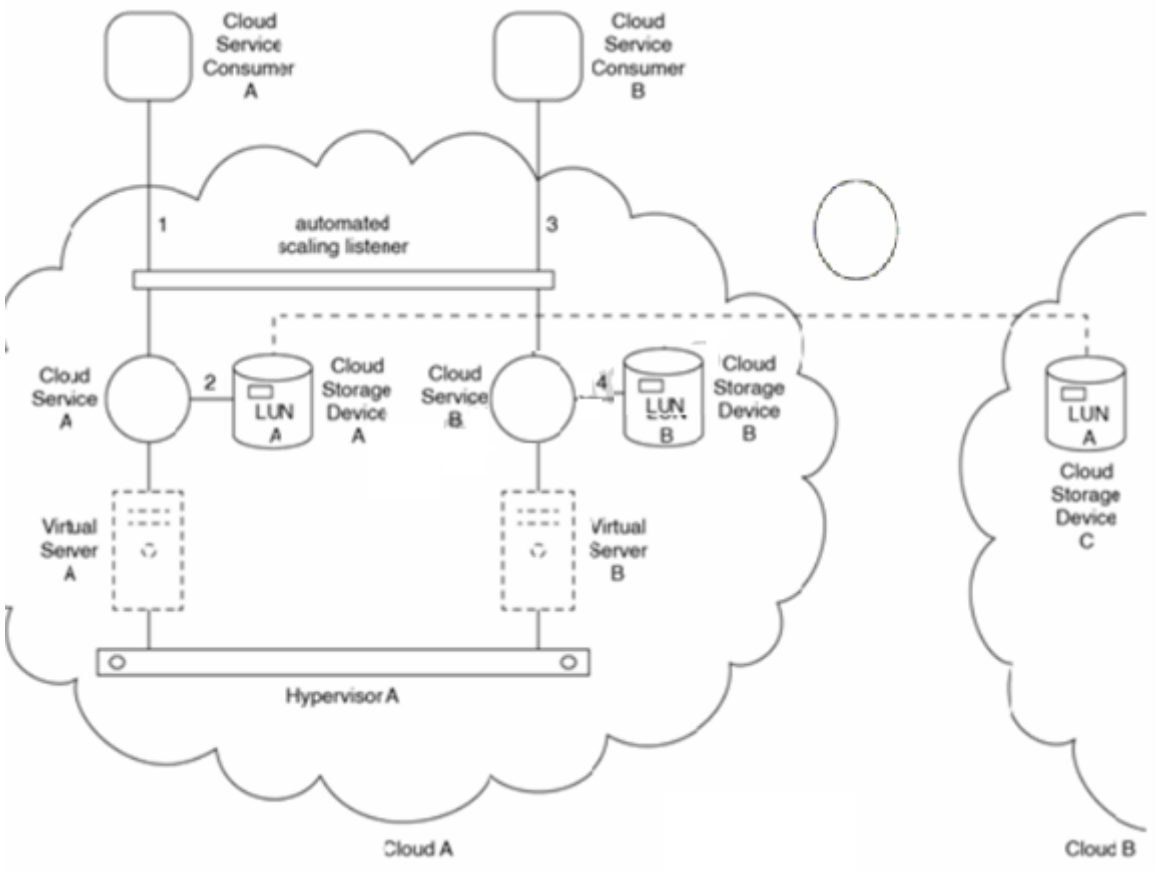

Cloud Service Consumer A issues a request to Cloud Service A (1). To process the request, Cloud Service A accesses LUN Aon Cloud Storage Device A (2). Cloud Service Consumer B issues a request to Cloud Service B (3). To process the request, Cloud Service B accesses LUN B on Cloud Storage Device B (4).

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 11

Question Type: MultipleChoice

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

Virtual Server B hosts Cloud Service B and Cloud Service C.

Cloud Service Consumer A accesses Cloud Service A (1), which then accesses LUN A or B on Cloud Storage Device A (2). After receiving the requested data from Cloud Service A, Cloud Service Consumer A forwards the data to Cloud Service B (3), which then writes it to Cloud Storage Device B (4).

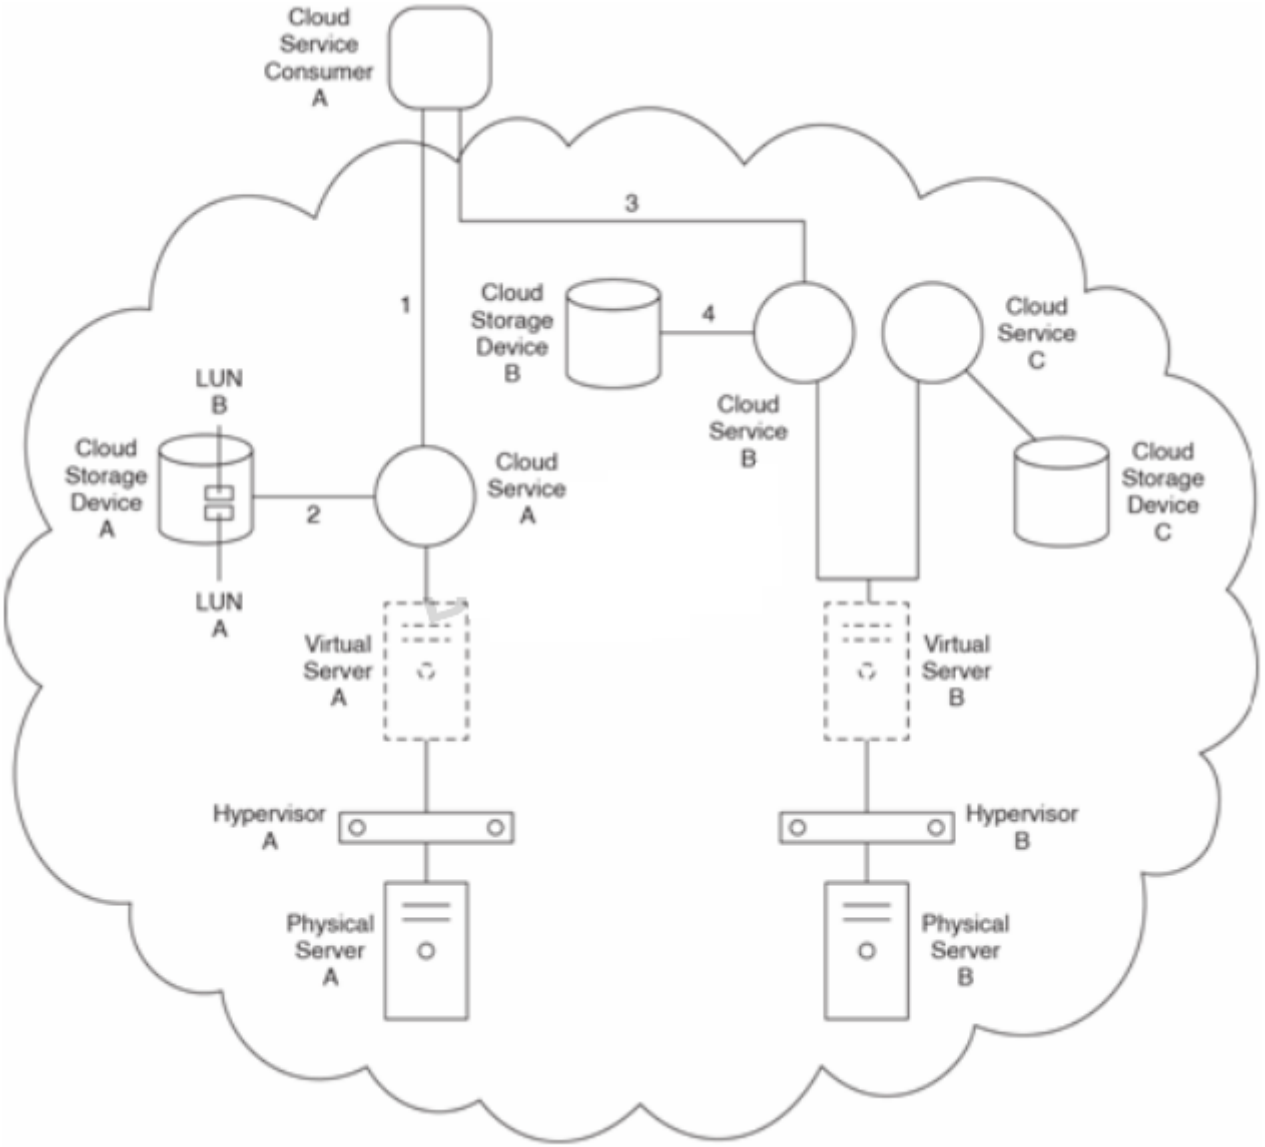
Cloud
Service
Consumer
A
1
3
Cloud
Storage
Device
B
4
Cloud
Service
C
LUN
B
Cloud
Service
B
Cloud
Storage
Device
A
Cloud
Service
A
2
Cloud
Storage
Device
C
LUN
A
Virtual
Server
A
Virtual
Server
B
Hypervisor
A
Hypervisor
B
Physical
Server
A
Physical
Server
B

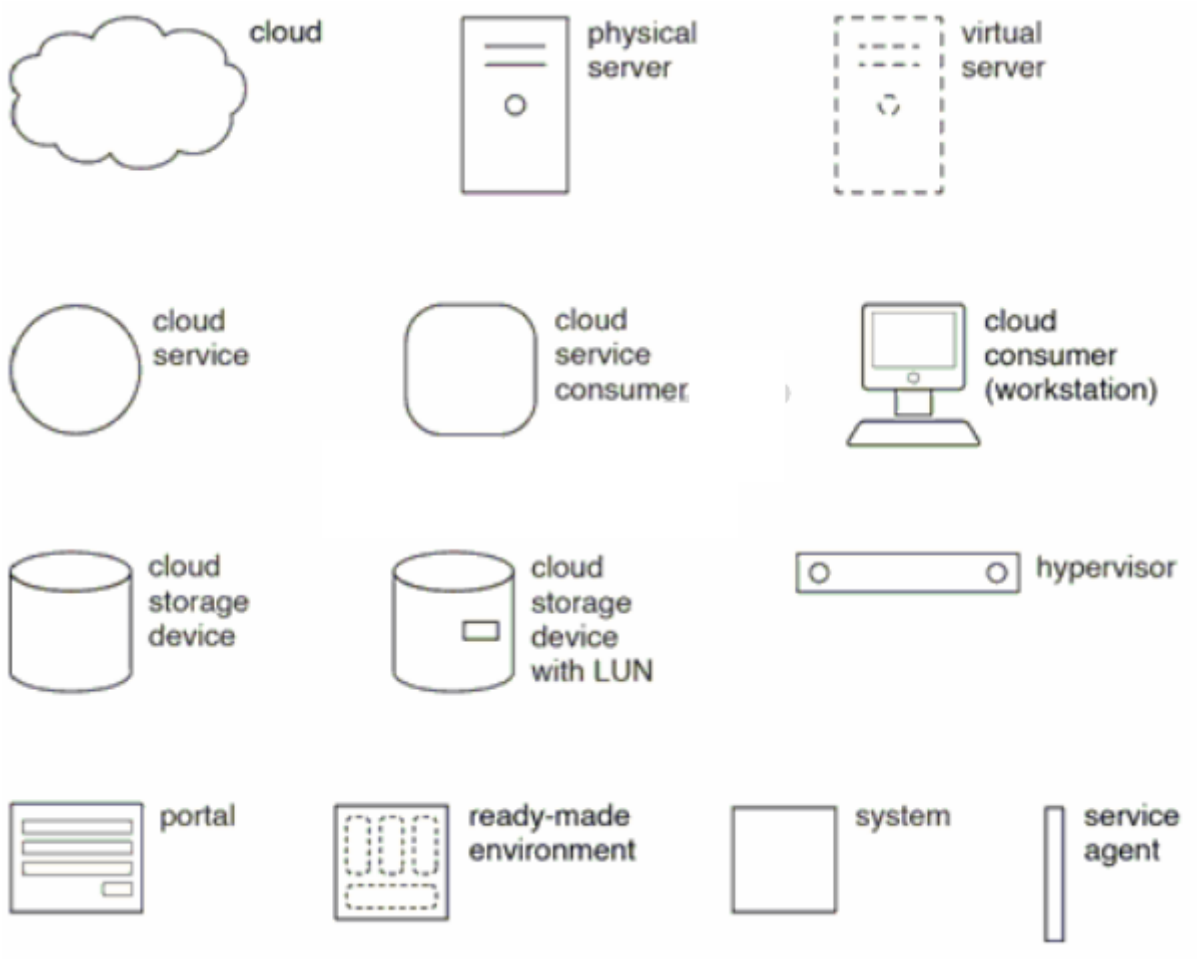

Cloud Service Consumer A belongs to Organization A,

Organization A uses LUN A and LUN B on Cloud Storage Device A to store their important client account data. Cloud Storage Device A is a low-performance cloud storage device, which begins to cause performance issues as more data is added to LUNs A and B and as Cloud Service Consumer A performs data access requests more frequently. Organization A asks that its cloud architecture be upgraded

to process increased quantities of data and higher volumes of data requests.

Organization A has been leasing a PaaS environment that it used to build Cloud Service A, which it would like to make available to the general public. Organization A needs to establish a system capable of monitoring usage of Cloud Service A for billing purposes.

The cloud provider is using a usage data collection and reporting system that gathers information on Organization A's hosted IT resources approximately ten hours after the time of usage. One day, Organization A attempts to retrieve information on whether Virtual Server B has available Cloud Service C instances. They discover that they are unable to obtain the current status of Virtual Server B. Organization A demands a system that provides instant availability reporting.

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

Options:

A- Cross-Storage Device Vertical Tiering, Pay-as-You-Go. Self-Provisioning

B- Service Load Balancing, Pay-as-You-Go, Multipath Resource Access

C- Intra-Storage Device Vertical Data Tiering, Usage Monitoring, Centralized Remote Administration

D- None of the above.

Answer:

D