

# DUMPS ARENA

## Cloud Architecture Lab

SOA C90.06

Version Demo

Total Demo Questions: 5

Total Premium Questions: 15

Buy Premium PDF

<https://dumpsarena.co>


[sales@dumpsarena.co](mailto:sales@dumpsarena.co)

[sales@dumpsarena.co](mailto:sales@dumpsarena.co)  
[dumpsarena.co](https://dumpsarena.co)

**QUESTION NO: 1**

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.



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?

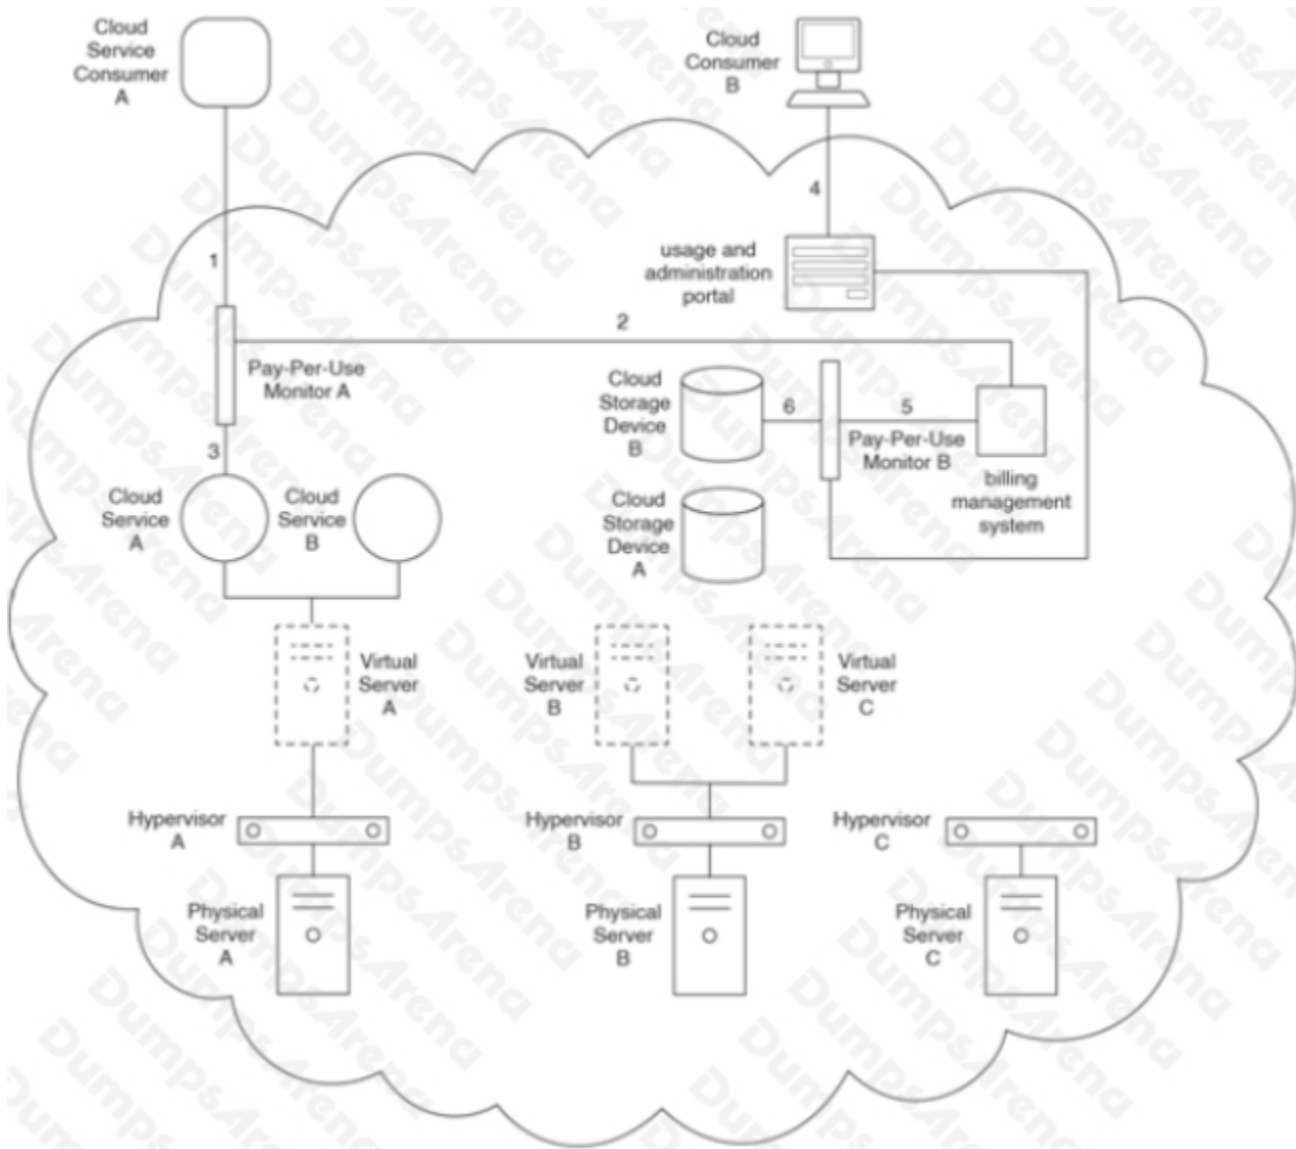
- A.** The Pay-as-You-Go pattern can be applied together with the Usage Monitoring pattern to establish a monitoring and billing system capable of de-provisioning Cloud Service A instances when they are no longer required. The Dynamic Data Normalization pattern can be applied to eliminate any redundant 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 NO: 2

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?

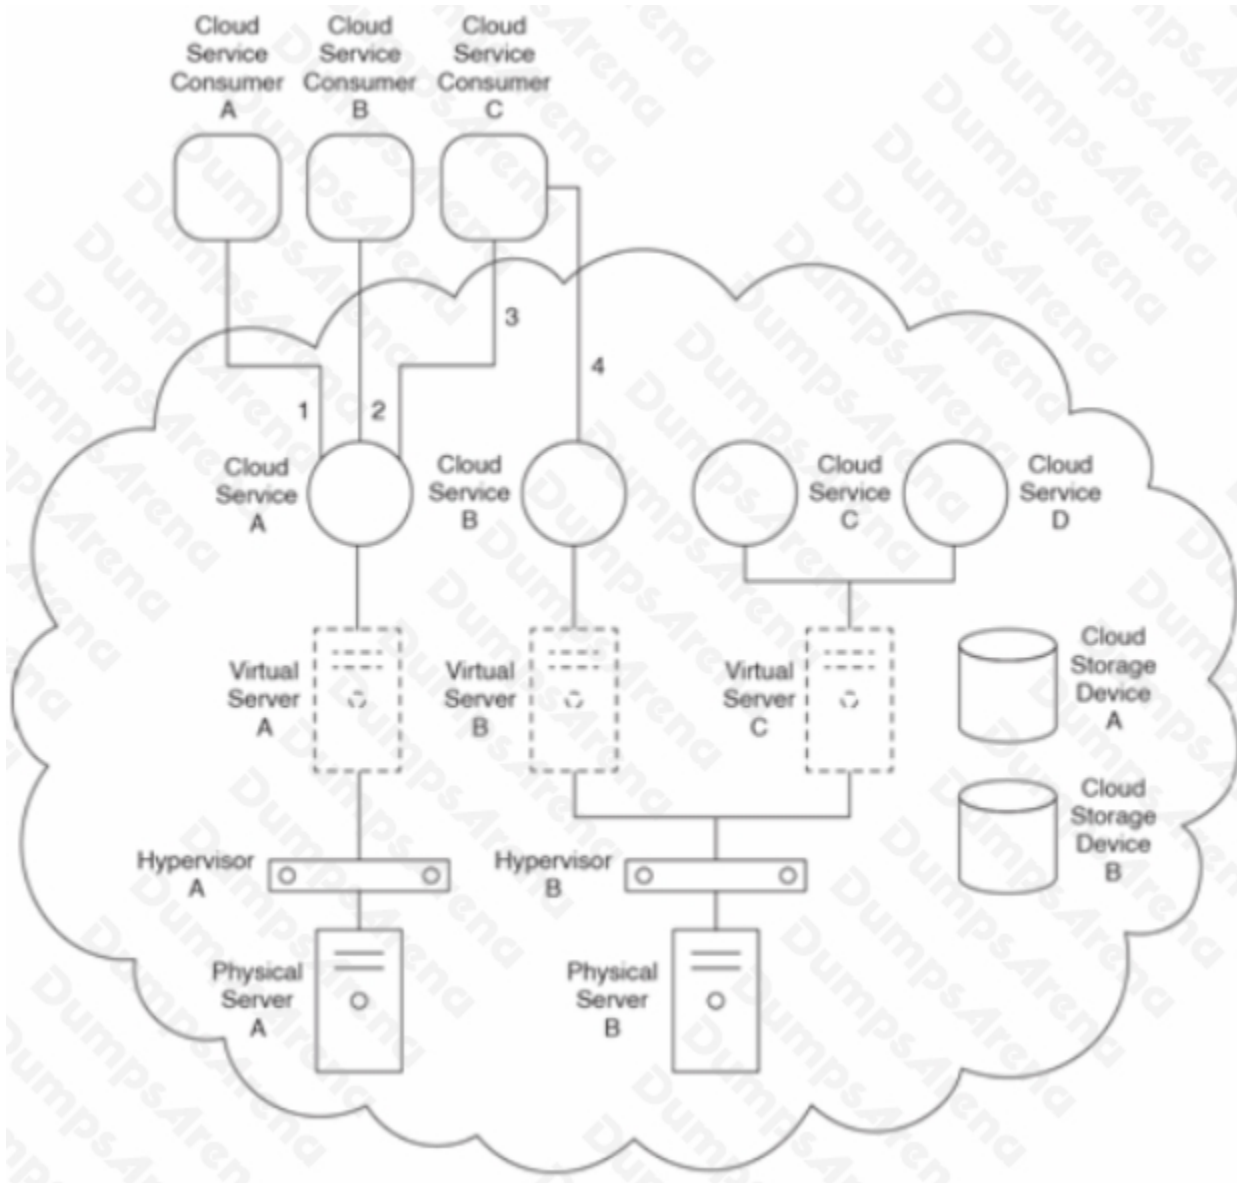
- A. Realtime Resource Availability, Elastic Network Capacity, Usage Monitoring
- B. Persistent Virtual Network Configuration, Elastic Network Capacity, Load Balanced Virtual Server Instances
- C. Load Balanced Virtual Switches, Elastic Resource Capacity, Automated Administration
- D. None of the above.

**ANSWER: A**

### QUESTION NO: 3

Cloud Service 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 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?

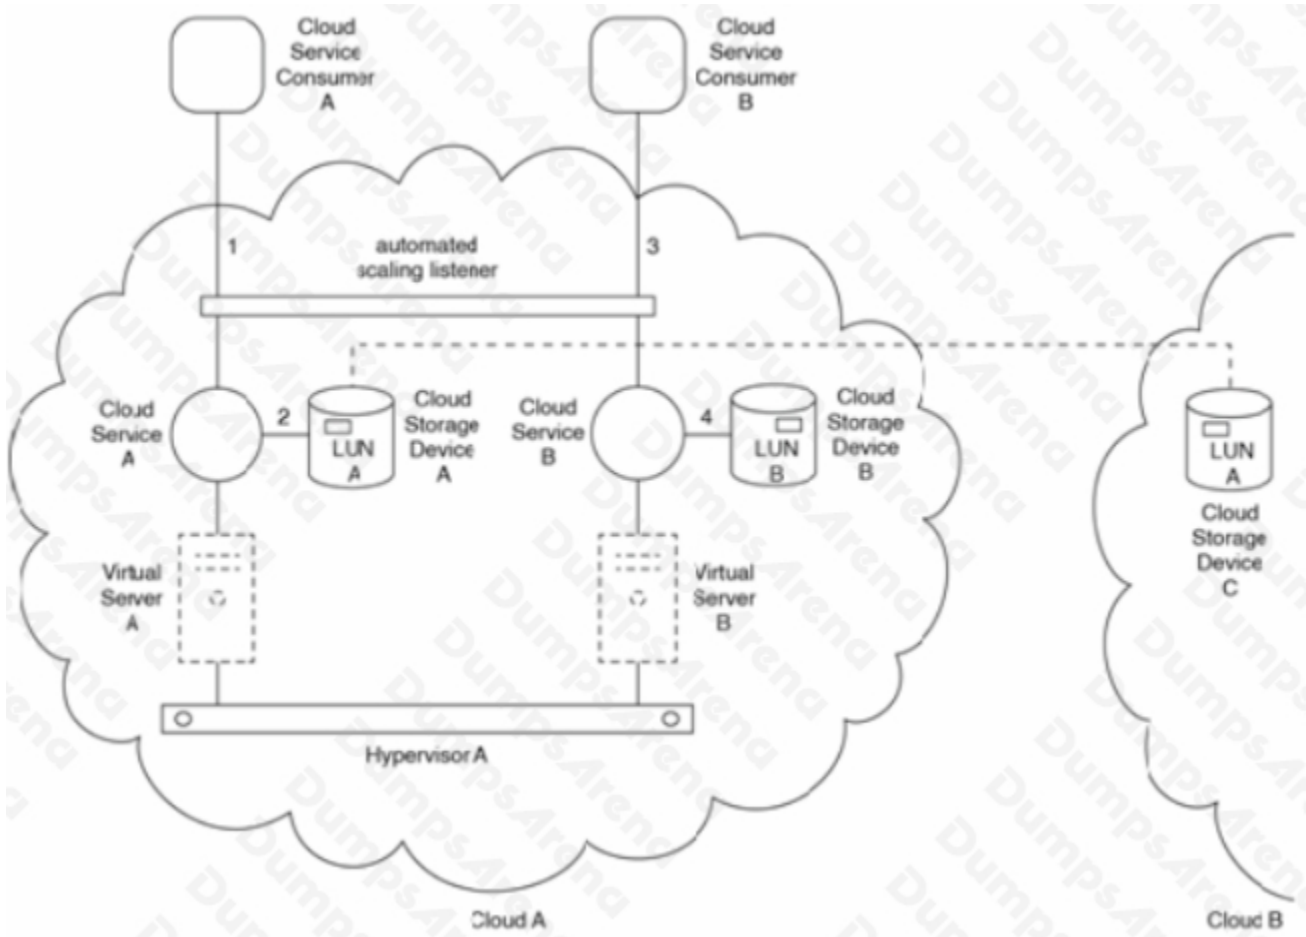
- A. Elastic Network Capacity, Load Balanced Virtual Server Instances, Service State Management

- 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 NO: 4**

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 A on 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?

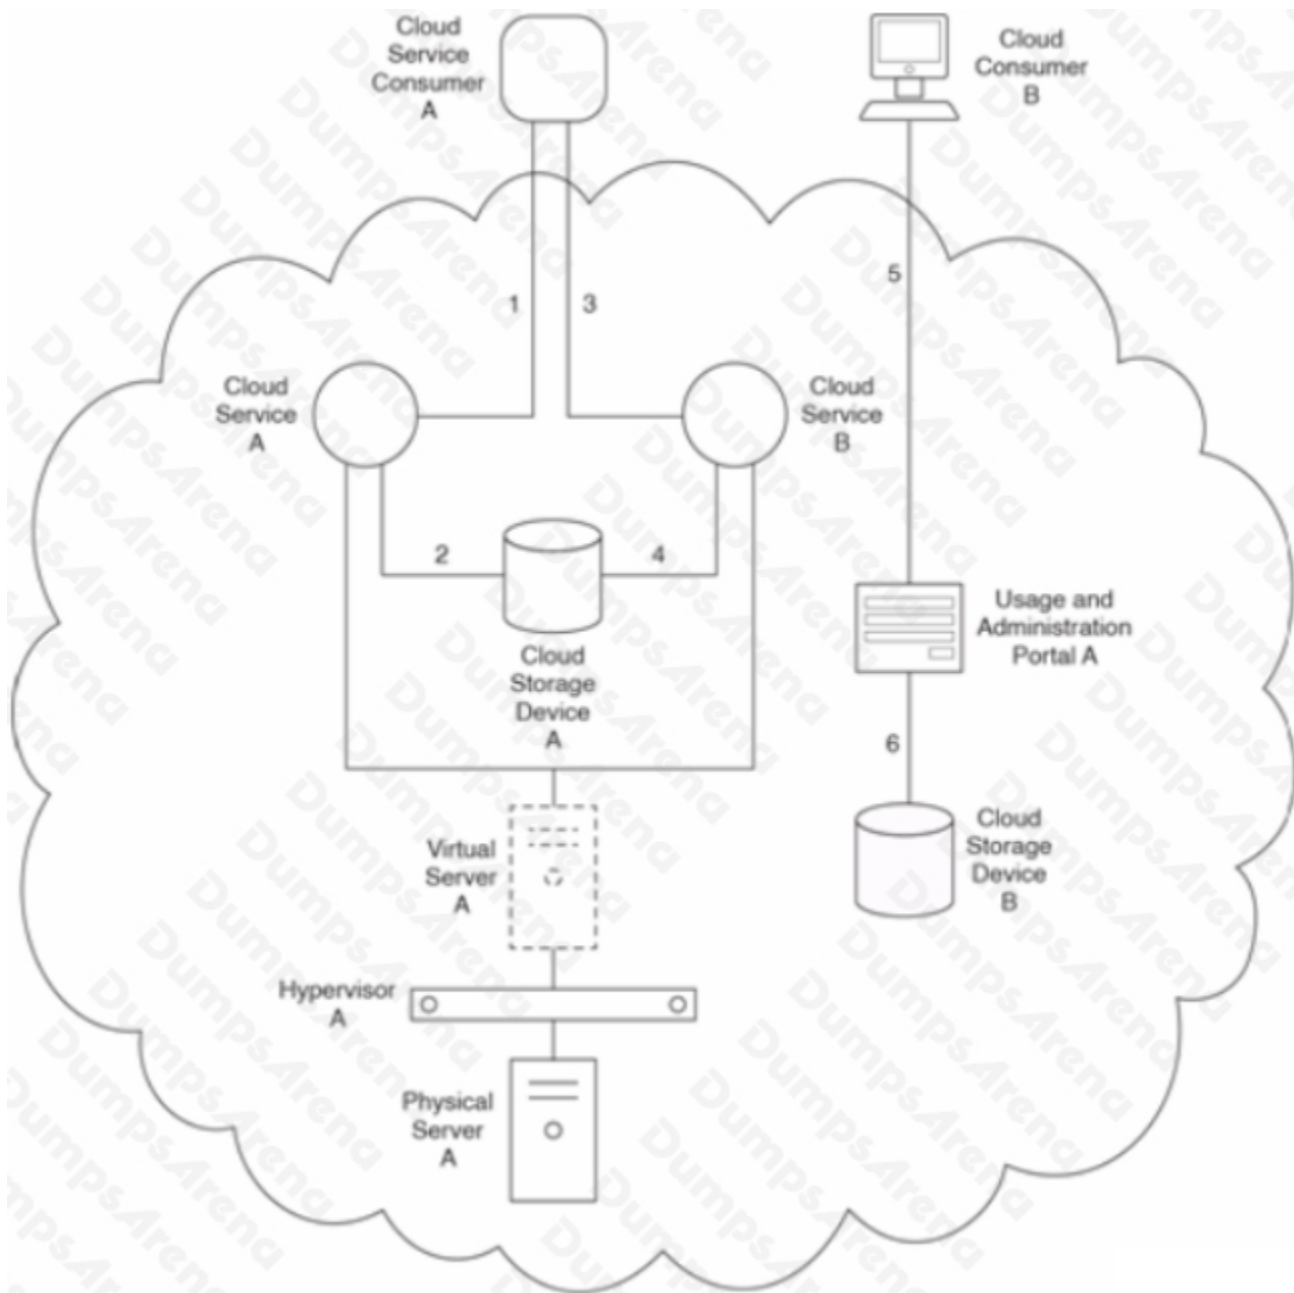
- A.** The Storage Maintenance Window pattern can be applied so that future outages of Cloud Storage Device A do not occur unexpectedly. The Resource Pooling and Resource Reservation patterns can be further applied to establish a resource pool on Cloud A that has resources reserved specifically for Cloud Service 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 NO: 5**

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 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?

- 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 Pay-as-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**