DUMPSARENA

Cloud Technology Lab

SOA C90.03

Version Demo

Total Demo Questions: 5

Total Premium Questions: 21

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

QUESTION NO: 1

Cloud Service Consumer A accesses Cloud Service A (1) that resides in Cloud X. a private cloud owned by the same organization acting as Cloud Consumer A. Cloud Service A processes the message from Cloud Service Consumer A and then sends back a response with the requested data (2). Next, Cloud Service Consumer A sends a message containing some of this data to Cloud Service B (3), which resides in public Cloud Y that is owned by Cloud Provider Y. After processing the message. Cloud Service B sends back a response with additional data to Cloud Service Consumer A (4). Finally, Cloud Service Consumer A writes the data it collected from Cloud Services A and B to Database A (5).

Recently, Cloud Service Consumer A has been required to access Cloud Services A and B at a significantly higher rate, sometimes over 1,000 times within a given workday. This increased usage has not affected Cloud Service B' s performance. Cloud Service A, however, has been generating runtime exceptions, and responses to Cloud Service Consumer A have become increasingly slow and unreliable. It is determined that this decline in performance is due to infrastructure limitations within private Cloud X' s environment. Instead of investing in new infrastructure for Cloud X, it is decided to explore the feasibility of moving Cloud Service A to Cloud Y instead.

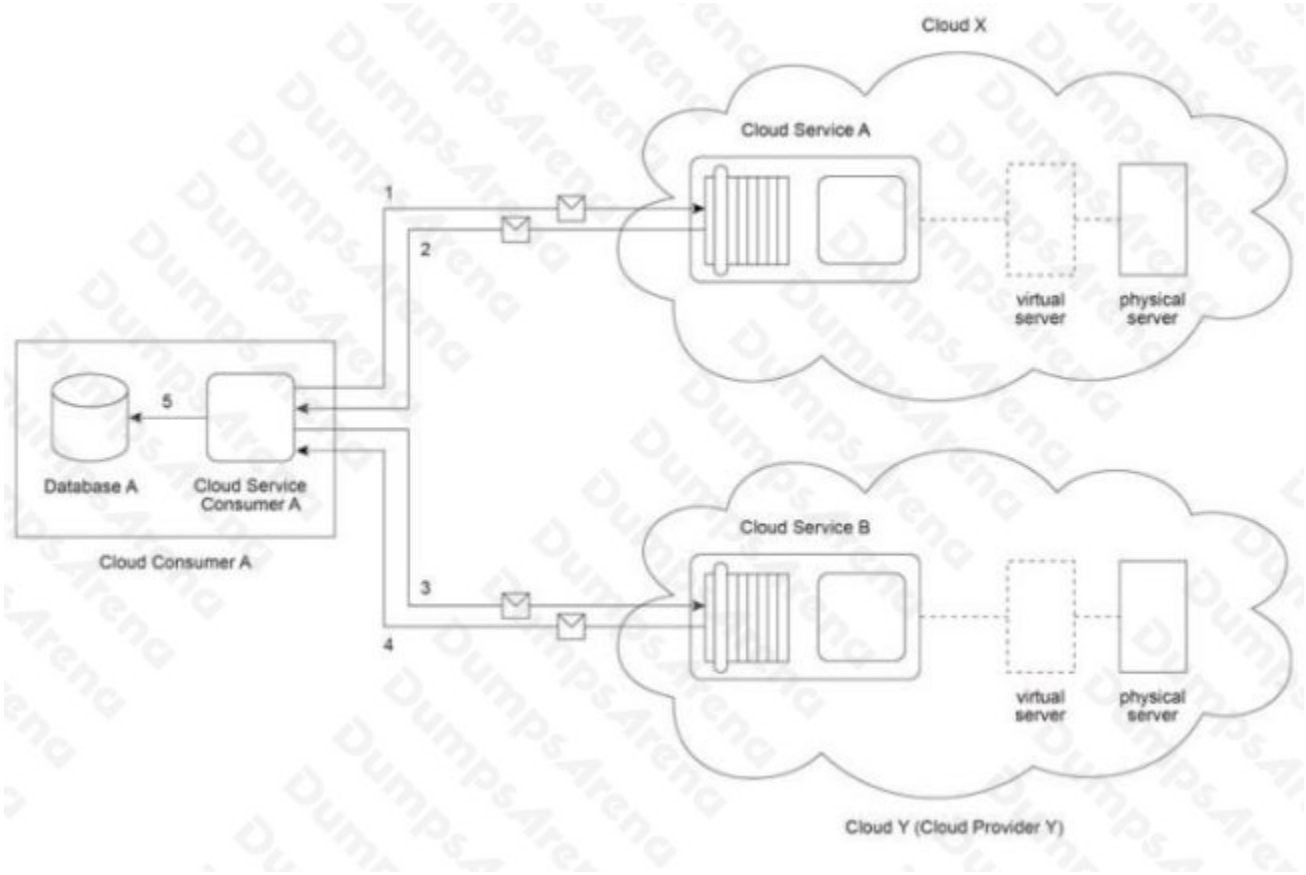

Which of the following statements describe valid financial considerations that can be taken into account for assessing the feasibility of this move?

A. Moving Cloud Service A to Cloud Y will result in the need for Cloud Service A to undergo integration testing to determine how well it can function within Cloud Y and what changes may need to be made to Cloud Service A in order for it to behave as expected. The integration testing and the subsequent changes required for Cloud Service A to function correctly within Cloud Y will incur integration costs that need to be budgeted for.

B. Because, in this scenario, the cost of capital is comprised of the up-front costs added to the on-going costs, the cost of capital required to move Cloud Service A to Cloud Y will be higher than upgrading Cloud X to accommodate Cloud Service A' s increased usage.

C. If the existing infrastructure that currently resides in private Cloud X was purchased specifically in support of Cloud Service A, then there may be a financial loss resulting from moving Cloud Service A out of Cloud X. This can be considered sunk costs that need to be evaluated.

D. By moving Cloud Service A to Cloud Y, there may be a decrease in operational governance control over the Cloud Service A implementation. This can increase locked-in costs because Cloud Consumer A may be forced to form dependencies upon proprietary tools used to configure and maintain the Cloud Service A implementation.

ANSWER: A C D

QUESTION NO: 2

The cloud service owner of Cloud Service A is evaluating Clouds X, Y and Z to determine which cloud environment can offer the greatest level of reliability. All three clouds are geographically dispersed across three separate time zones. As a result, each cloud experiences usage peaks at different times. Based on the metrics provided, the greater the usage of a cloud, the lower its reliability. When the cloud service owner complains to Cloud Provider A (the owner of all three clouds) that none of the clouds provide an adequate level of reliability, Cloud Provider A suggests a solution that increases resiliency.

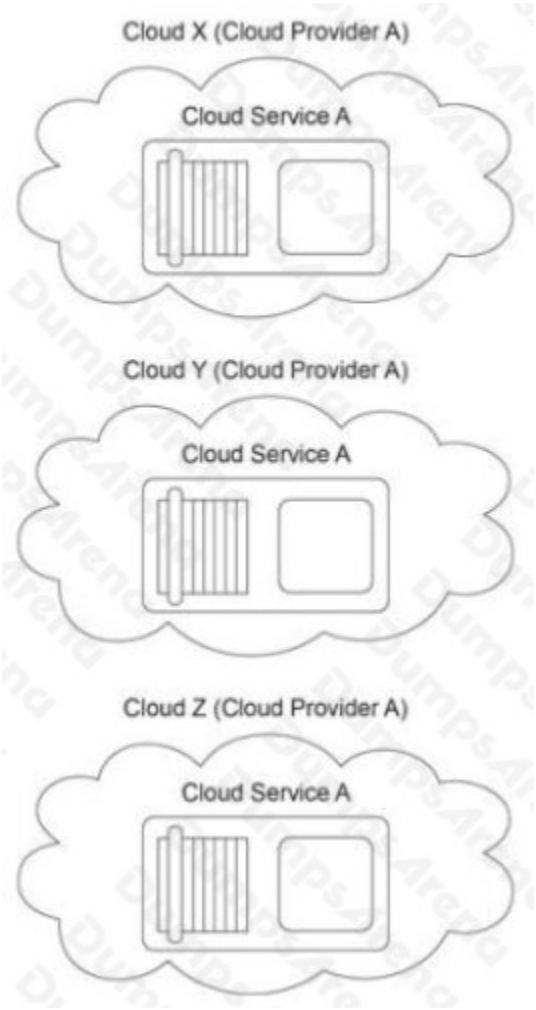

Which of the following statements accurately describes a solution that can be used to fulfill the resiliency requirements of Cloud Service A?

A. Redundant implementations of Cloud Service A are deployed in all three clouds. The failover system mechanism and a special type of automated scaling listener mechanism are implemented to establish a system whereby one redundant Cloud Service A implementation will automatically take over from another.

B. A cloud balancing solution is established, whereby an automated scaling listener mechanism is implemented on each cloud in such a way that every cloud can automatically scale out to another cloud. As a result, if reliability problems occur on any one cloud, the subsequent requests will be scaled out to another cloud in a manner that is transparent to cloud service consumers.

C. A failover system mechanism is implemented on Cloud X, which acts as the primary point of contact for cloud service consumers. Upon failure conditions occurring, the Cloud Service A implementation on Cloud X automatically hands over control of current and future message requests from cloud service consumers to Cloud Y. Cloud Y retains control of cloud service consumer communication until the next failure condition occurs, at which point it hands over control to Cloud Z. Finally, if a failure condition occurs in Cloud Z. control is handed back to Cloud X.

D. A cloud balancing solution is established, whereby a resource replication mechanism is implemented on each cloud. This allows Cloud Service A to be automatically replicated across cloud environments, thereby enabling each implementation of Cloud Service A to take the place of another, whenever failure conditions occur.

ANSWER: A

QUESTION NO: 3

The cloud service owner of Cloud Service A is evaluating Clouds X, Y and Z to determine which cloud environment can offer the greatest level of reliability. All three clouds are geographically dispersed across three separate time zones. As a result, each cloud experiences usage peaks at different times. Based on the metrics provided, the greater the usage of a cloud, the lower its reliability. When the cloud service owner complains to Cloud Provider A (the owner of all three clouds) that none of the clouds provide an adequate level of reliability, Cloud Provider A suggests a solution that increases resiliency.

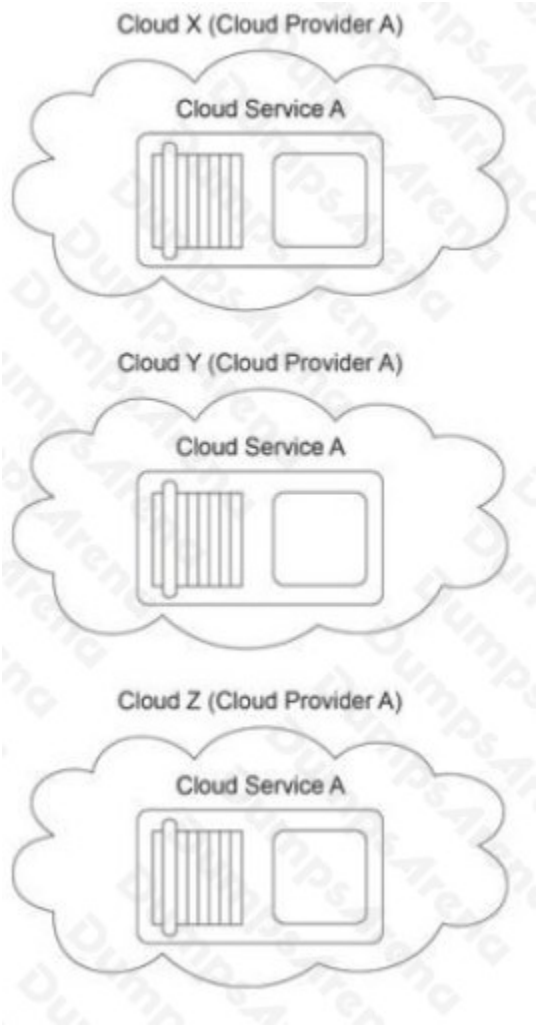

Which of the following statements accurately describe how a cloud balancing solution can be implemented to fulfill the resiliency requirements of Cloud Service A?

A. Cloud Service A is redundantly deployed in advance within Clouds X. Y and Z and is further supplemented with failover system mechanisms and specialized types of automated scaling listener mechanisms.

B. Using a PaaS environment, the cloud service owner can configure a primary Cloud Service A implementation on Cloud X so that when failure conditions occur, message requests from cloud service consumers are automatically routed to a redundant on-premise implementation of Cloud Service A. This requires the implementation of the failover system and resource replication mechanisms.

C. A type of automated scaling listener that can also be considered a cloud monitor is implemented in Clouds X, Y and Z to establish a system whereby each cloud can assume control of cloud service consumer message exchanges. This results in resilient cloud balancing, as opposed to on-demand cloud balancing.

D. Clouds X, Y and Z are equipped with failover system mechanisms and specialized types of automated scaling listener mechanisms in order to establish cross-cloud resiliency. Cloud Service A is configured so that Clouds X, Y and Z can dynamically generate redundant instances on-demand.

ANSWER: A D

QUESTION NO: 4

Cloud Service A is being made available on public Cloud X by Cloud Provider X via the SaaS delivery model. Cloud Service A is hosted by Physical Server A that also hosts cloud services being used by different cloud service consumers (and owned by different cloud service owners).

Cloud Service Consumers A and B access Cloud Service A on a regular basis and Physical Server A has been able to accommodate the resulting usage load After reviewing the implementation environment for Cloud Service A. Cloud Provider X determines it needs to enhance the ubiquitous access and multi-tenancy characteristics within this part of Cloud X.

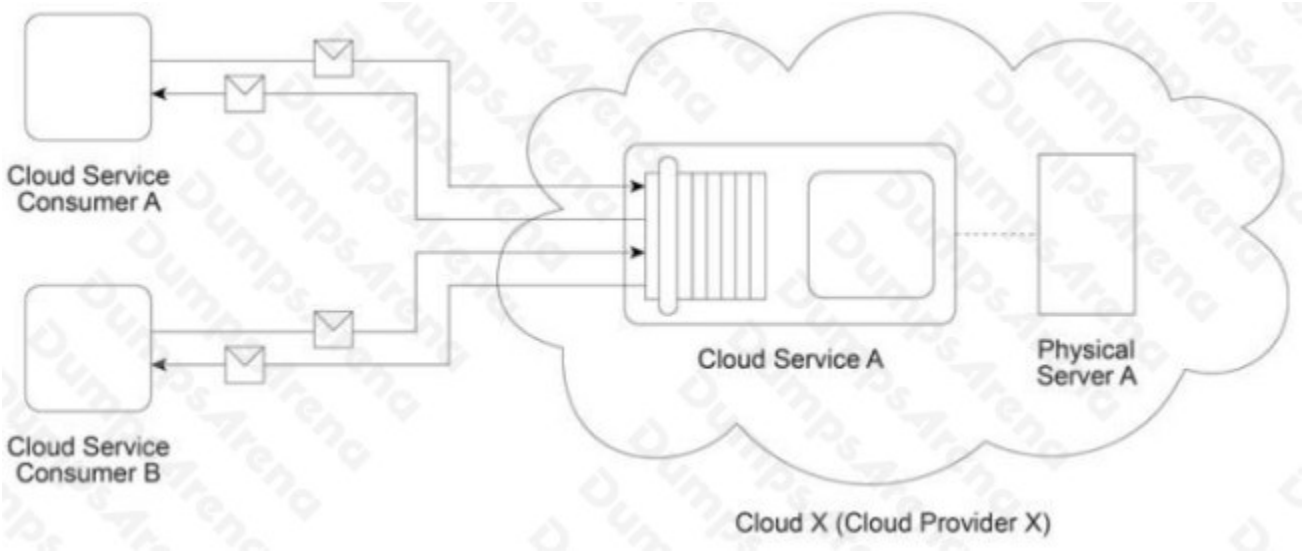

Which of the following statements accurately describes a solution that fulfills all of these requirements?

A. The multi-device broker mechanism can be implemented to address ubiquitous access requirements by broadening the range of cloud service consumers that can access Cloud Service The resilient computing mechanism can be implemented to establish the multitenancy capabilities of Cloud Service A so that it will be able to resiliency accommodate the additional cloud service consumer devices that will gain access to it.

B. The resilient computing mechanism can be implemented to address ubiquitous access requirements by making Cloud Service A available to a broader range of cloud service consumer devices, including desktops running Web browsers and various mobile devices. The resource replication mechanism can be implemented to enable multi-tenancy within the Cloud Service A implementation.

C. The resource replication mechanism can be introduced to address ubiquitous access requirements by broadening the range of cloud service consumers that can use Cloud Service A via standard service replication. The virtual server

mechanism can be implemented to enable multi-tenancy via each service replication resulting from the application of the resource replication mechanism.

D. The multi-device broker mechanism can be implemented to address ubiquitous access requirements by broadening the range of cloud service consumers that can access Cloud Service A. The virtual server mechanism can be implemented to establish the multitenancy capabilities of Cloud Service A.

ANSWER: D

QUESTION NO: 5

Organization A has been expanding and, as a result, is outgrowing the processing capacity of its on-premise Service A implementation. It is determined that this is due to usage thresholds of Service A and complex data processing limitations in Database A. The diagram depicts Organization A's current on-premise environment, where Service Consumers A. B and C attempt to access Service A at the same time. Service Consumer A successfully accesses Service A, which then successfully retrieves the requested data (1). Service Consumer B successfully accesses Service A, but due to the complex data structure, the request for the data times out and fails (2). Finally, Service Consumer C attempts to access Service A, but is rejected because Service A is unable to accept more concurrent requests.

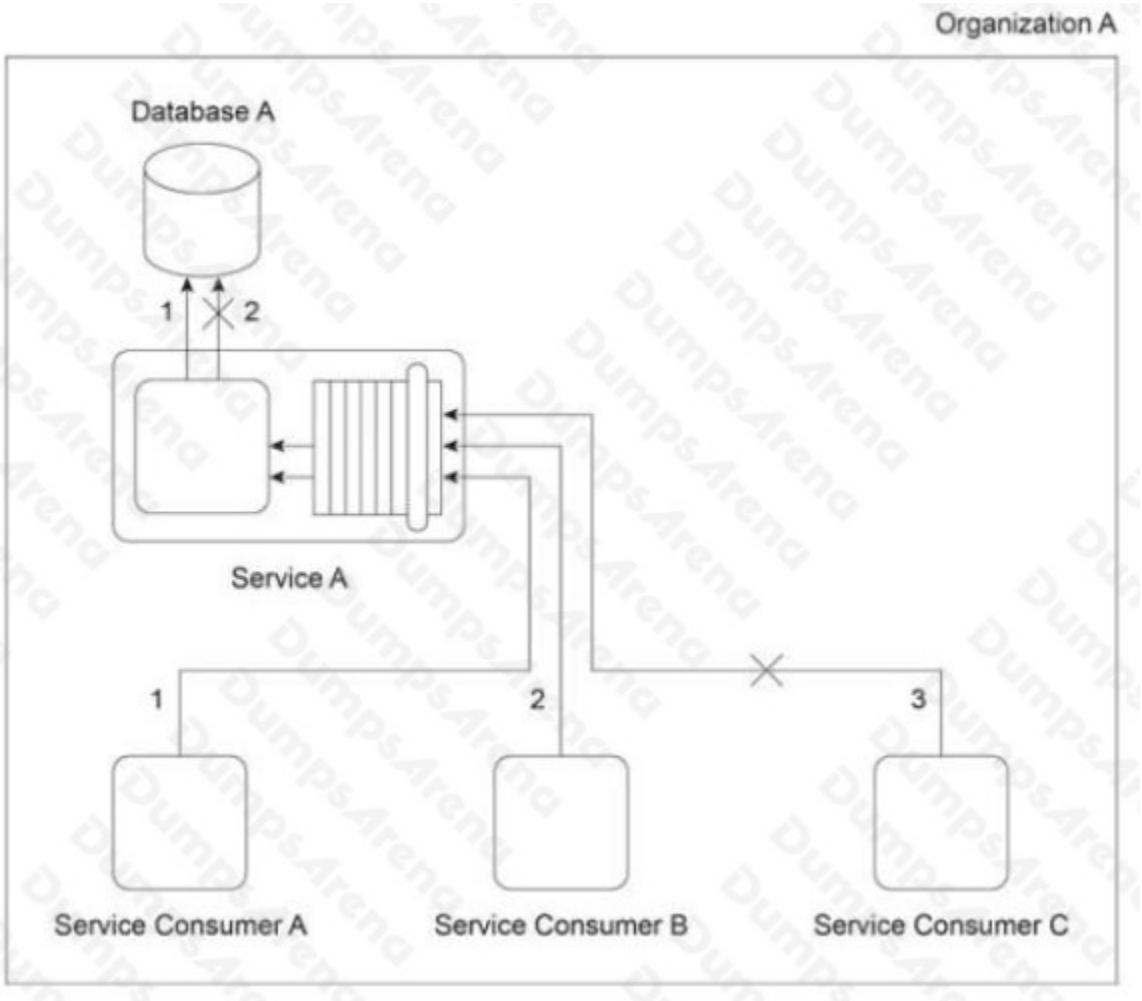

Organization A is required to continue using its on-premise Service A implementation, with the exception of Database A, which does not need to remain on-premise. Database A is dedicated to Service A and is comprised of relational data. Which of the following statements provides a solution that uses cloud-based IT resources to solve the performance limitations of Service A and Database A?

A. A failover system can be implemented in a hybrid architecture comprised of Organization A's existing on-premise environment and a public cloud environment. The failover system would span both environments so that when Service A is unable to process request messages from Service Consumers A, B or C, the failover system can automatically route messages to a redundant implementation of Service A residing in the public cloud. Similarly, when Database A is unable to process a data access request from Service the failover system can automatically route this request to a redundant implementation of Database A, also residing in the public cloud.

B. The state management database and resource replication mechanisms can be implemented to establish redundant implementations of Service A and Database A in both on-premise and cloud environments. Using resource replication, a cloud-based duplicate of Service A (Cloud Service A) will be established in a public cloud and will remain in synch with Service A via regular replication cycles. Using the resource replication mechanism together with the state management database mechanism allows for Database A to be dynamically replicated in an independent state management database that has redundant implementations in both on-premise and cloud environments. The state management database can be further optimized to support non-relational data to improve data access performance.

C. A cloud bursting solution can be implemented, whereby a redundant copy of Service A is implemented within a public cloud. This cloud-based, redundant implementation of Service A is referred to as Cloud Service A. A copy of Database A is also implemented within the cloud and both the on-premise and cloud-based copies of Database A are redesigned to be non¬relational in order to improve data access performance. Service A continues to act as a first point of contact for Service Consumers A, B and An automated scaling listener is deployed so that when Service A's thresholds are met, requests are automatically routed to Cloud Service A.

D. None of the above

ANSWER: C