

DUMPSBOSS.

Cloud Technology Lab

SOA C90.03

Version Demo

Total Demo Questions: 5

Total Premium Questions: 21

Buy Premium PDF

<https://dumpsboss.co>


support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

QUESTION NO: 1

Cloud Provider Y owns Cloud Y, which provides a set of cloud services, virtual servers, and one physical server. Cloud Services A and B are hosted on Virtual Server A, which is hosted by Physical Server A. Physical Server A also hosts Virtual Server B, which hosts redundant implementations of Cloud Services A and B for load balancing purposes. Cloud Service Consumer A is accessing Cloud Service A located on Virtual Server A. Cloud Service Consumer B is accessing Cloud Service B located on Virtual Server B. Cloud Service Consumers A and B are both owned by Cloud Consumer X.

Cloud Consumer X has recently complained that Cloud Service A is becoming increasingly unreliable. Cloud Provider Y launches an investigation that reveals that Virtual Server A has been crashing sporadically due to successful malicious intermediary attacks where malicious service agents have been inserting harmful data into messages sent by Cloud Service Consumer A to Cloud Service A.



Which of the following statements describes a solution that can mitigate the malicious intermediary attacks that have been occurring on Virtual Server A - and - can further help prevent the same types of attacks from occurring on Virtual Server B?

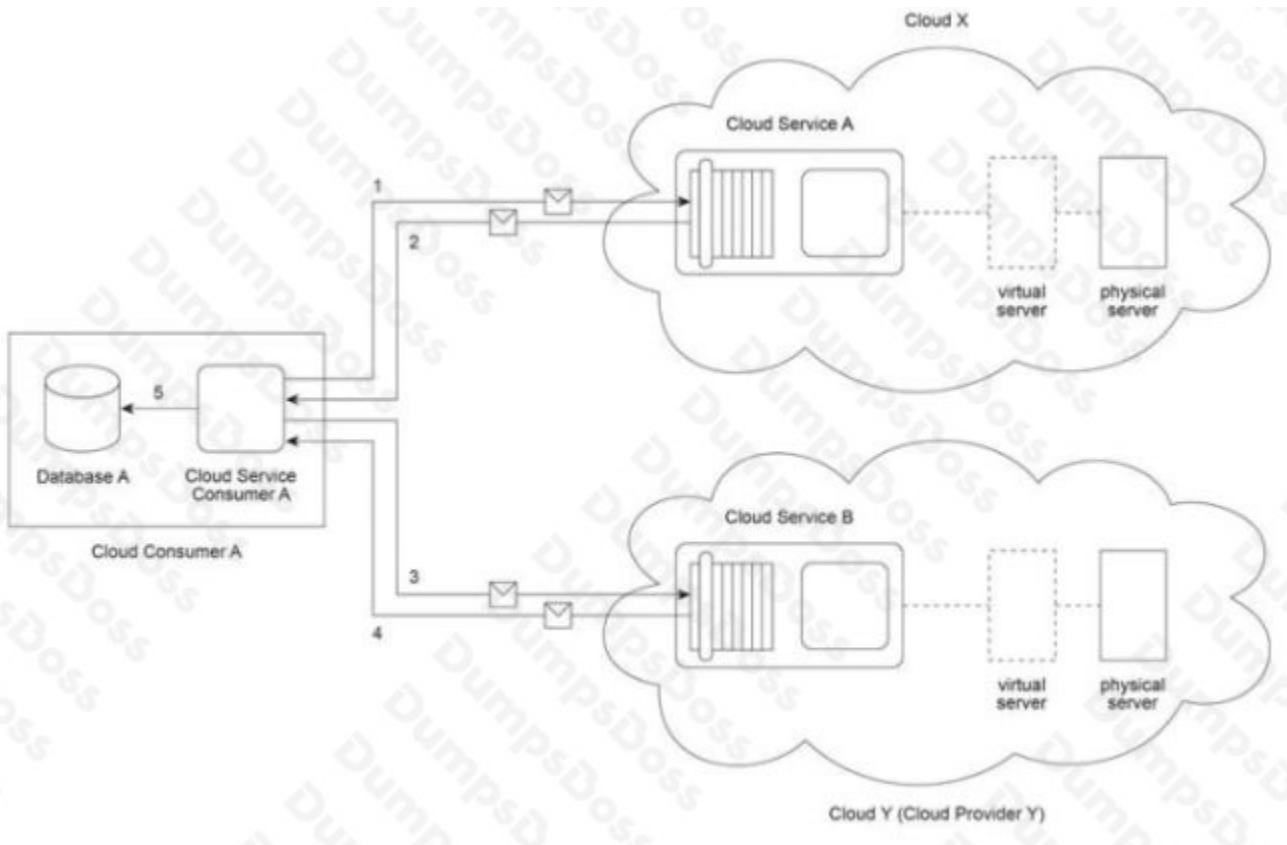
- A.** Cloud Provider Y can implement the cloud-based security group mechanism that establishes a logical trust boundary that encompasses Virtual Servers A and B.
- B.** Cloud Provider Y can implement the encryption and digital signatures mechanisms to protect the contents of messages from being augmented by malicious service agents.
- C.** Cloud Provider Y can implement the failover system mechanism by requiring that redundant copies of all service agents be deployed. This will reduce the chances of any one service agent performing malicious actions at runtime.

D. Cloud Provider Y can implement the identity and access management mechanism for Virtual Servers A and B by physically deploying this mechanism on Physical Server A and then further implementing the resource replication mechanism to propagate the identity and access management system across Virtual Servers A and B.

ANSWER: B

QUESTION NO: 2

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?

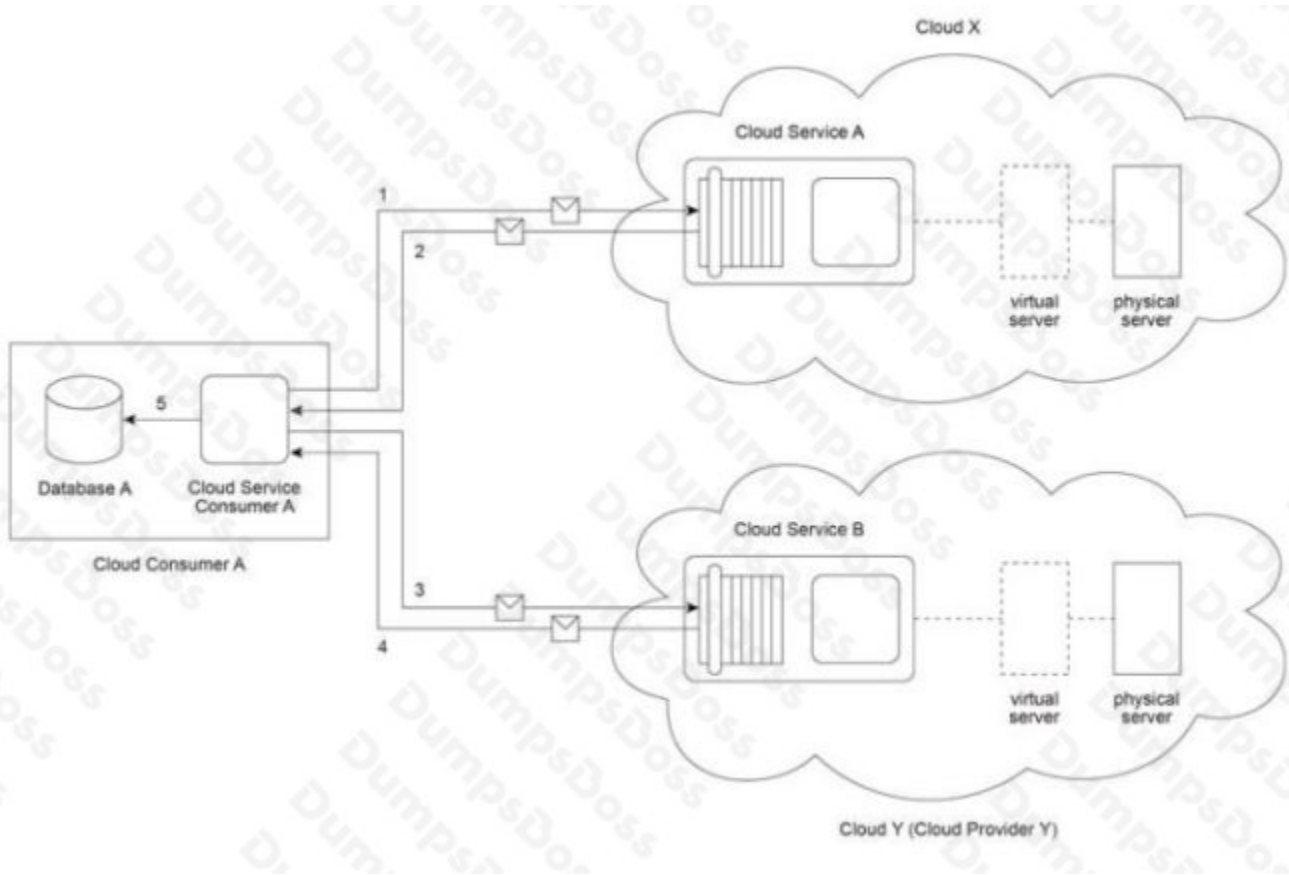
- A.** Moving Cloud Service A to Cloud Y will require that Database A also be moved to Cloud Y due to the need for Cloud Service A and Database A to share a common virtual server within the same organizational boundary, as required by the cloud-based security group. The move of Database A will increase the integration testing effort and, as a result, will also increase the overall integration costs.
- B.** Once Cloud Service A is deployed in Cloud Y, it may form dependencies upon proprietary parts of Cloud Y that may limit its mobility should it be decided to move it outside of Cloud Y in the future. This can incur further locked-in costs that need to be accounted for.
- C.** By moving Cloud Service A to Cloud Y, the SaaS delivery model will be established for Cloud Service A, thereby allowing the service implementation to build upon existing infrastructure from underlying PaaS and IaaS delivery models that would have been required for Cloud Service B to be implemented in Cloud Y.
- D.** Public Cloud Y charges for the use of its IT resources. Moving Cloud Service A to Cloud Y can therefore result in new on-going costs. Although Cloud Service A may be able to share some of the existing IT resources used by Cloud Service B, it will likely incur new on-going costs that need to be budgeted for.

ANSWER: B D

QUESTION NO: 3

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.



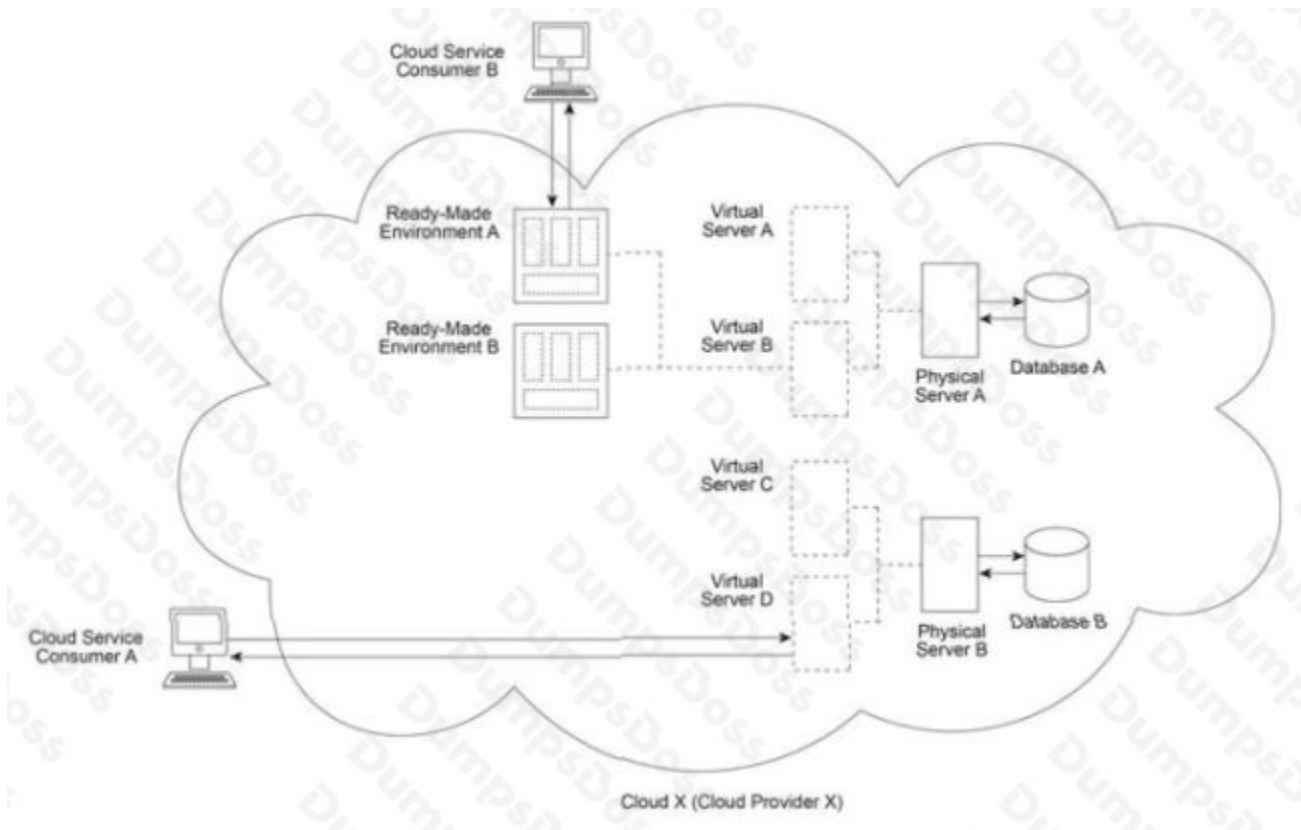
Which of the following statements describe valid financial considerations that can be taken into account for assessing the feasibility of this move?

- 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: 4

Cloud Provider X (which owns Cloud X) deploys two physical servers (Physical Servers A and B) and two databases (Databases A and B). Virtual Servers A and B are hosted by Physical Server A and Ready-Made Environments A and B are hosted by Virtual Server B. Virtual Servers C and D are hosted by Physical Server B. Cloud Service Consumer A regularly accesses Virtual Server D in order to test and deploy a new cloud service that was developed on-premise by the cloud consumer organization operating Cloud Service Consumer A. Cloud Service Consumer B (operated by a different cloud consumer organization) has been regularly accessing Ready-Made Environment A in order to develop and deploy a different new cloud service.



Cloud X is a private cloud that, to-date, has been set up within the cloud provider company to provision IT resources free of charge to internal cloud consumers, via PaaS and IaaS delivery models. The cloud consumers that have been operating Cloud Service Consumers A and B represent different IT departments within the company that have been working separately on the development of new cloud services. Cloud Service Consumer A has been accessing Virtual Server D to make necessary configurations and administration settings for the upcoming deployment of a new cloud service that was previously developed outside of Cloud X. Cloud Service Consumer B has been accessing Ready-Made Environment A to develop and now deploy a different new cloud service.

Cloud Provider X (which is represented by a separate IT department dedicated to governing and administering Cloud X) determines that it will need to introduce three specific enhancements to Cloud X in order to accommodate both upcoming cloud services. First, it needs to add a way to charge cloud consumers for the usage of cloud services. Secondly, it needs to add a way for new cloud services to be automatically scaled. Finally, it needs to add a way for a cloud consumer to be automatically notified when a cloud service encounters runtime loads that exceed its allocated usage threshold.

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

A. The pay-for-use monitor mechanism can be installed to address the requirement for cloud service usage charges to be tracked and recorded. The automated scaling listener mechanism can be implemented to address the requirement for cloud services to be automatically scaled. The multi-device broker mechanism can be implemented to address the requirement for notifications to be issued when cloud service loads exceed thresholds.

- B.** The pay-for-use monitor mechanism can be installed to address the requirement for cloud service usage to be tracked and recorded. The automated scaling listener mechanism can be implemented to address both the requirement for cloud services to be automatically scaled and for notifications to be issued when cloud service loads exceed thresholds.
- C.** The pay-for-use monitor mechanism can be implemented to address the requirement for cloud service usage charges to be tracked and recorded, for cloud services to be automatically scaled, and for notifications to be issued when cloud service loads exceed thresholds.
- D.** The state management database mechanism together with the virtual server mechanism and the automated scaling listener mechanism can be implemented to address the requirement for cloud service usage charges to be tracked and recorded, for cloud services to be automatically scaled, and for notifications to be issued when cloud service loads exceed thresholds.


ANSWER: B

QUESTION NO: 5

A cloud consumer is interested in leasing cloud-based virtual servers. It compares the virtual servers offered by Cloud Provider X and Cloud Provider Y. Cloud X (owned by Cloud Provider X) and Cloud Y (owned by Cloud Provider Y) both provide shared physical servers that host multiple virtual servers for other cloud consumers.

The virtual servers on Cloud X are accessed directly, whereas the virtual servers on Cloud Y are accessed via an automated scaling listener. On Cloud X, virtual servers are pre-configured to support a specific amount of concurrent cloud service consumers. When this threshold is exceeded, cloud service consumer requests are rejected. Due to the use of the automated scaling listener, virtual servers on Cloud Y can provide a greater level of elasticity.

The hourly cost to the cloud consumer to use a virtual server on Cloud X is half that of the cost to use a virtual server on Cloud Y. Within a one month period, Cloud Provider X bases its hourly charge on the maximum number of virtual servers used. Within a one month period, Cloud Provider Y bases its hourly charges on actual virtual server usage. Cloud Provider Y charges \$20 for each hour that a cloud consumer uses a virtual server.



The cloud consumer predicts its monthly usage requirements to be as follows:

Number of virtual servers	Usage (hours)
2	70
4	20
6	10

The cloud consumer is required choose the cloud provider with the lowest on-going cost based on its predicted usage. Which of the following statements accurately calculates the on-going usage costs of Cloud Providers X and Y and correctly states the cloud provider that the cloud consumer must choose?

- A.** The total usage duration is $(10 + 20 + 70)$ hours = 100 hours.
 The actual usage is $(10 \times 6) + (20 \times 4) + (70 \times 2)$ hours = 280 hours.
 The cost of using virtual servers from Cloud Provider X is $100 \times 6 \times \$10 = \$6,000$.
 The cost of using virtual servers from Cloud Provider Y is $280 \times \$20 = \$5,600$.
 The cloud consumer must therefore choose Cloud Provider Y.

B. The total usage duration is $(10 + 20 + 70)$ hours = 100 hours.
The actual usage is $((10 \times 6) + (20 \times 4) + (70 \times 2)) \times ((2 + 4 + 6) / 3)$ hours = 1,120 hours
The cost of using virtual servers from Cloud Provider X is $10 \times 6 \times \$10 = \$6,000$.
The cost of using virtual servers from Cloud Provider Y is $1,120 \times \$20 = \$22,400$.
The cloud consumer must therefore choose Cloud Provider X.

C. The total usage duration is $(10 + 20 + 70)$ hours = 100 hours.
The actual usage is $((10 \times 6) + (20 \times 4) + (70 \times 2)) \times ((2 + 4 + 6) / 3)$ hours = 1,120 hours
The cost of using virtual servers from Cloud Provider X is $6 \times 100 \times \$10 = \$6,000$.
The cost of using virtual servers from Cloud Provider Y is $1,120 \times \$20 = \$22,400$.
The cloud consumer must therefore choose Cloud Provider Y.

D. The total usage duration is $(10 + 20 + 70) \times 12$ hours = 1,200 hours.
The actual usage is $(10 \times 6) + (20 \times 4) + (70 \times 2)$ hours = 280 server hours.
The cost of using virtual servers from Cloud Provider X is $12 \times 100 \times 5 \times \$10 = \$60,000$.
The cost of using virtual servers from Cloud Provider Y is $280 \times \$20 = \$5,600$.
The cloud consumer must therefore choose Cloud Provider Y.

ANSWER: A