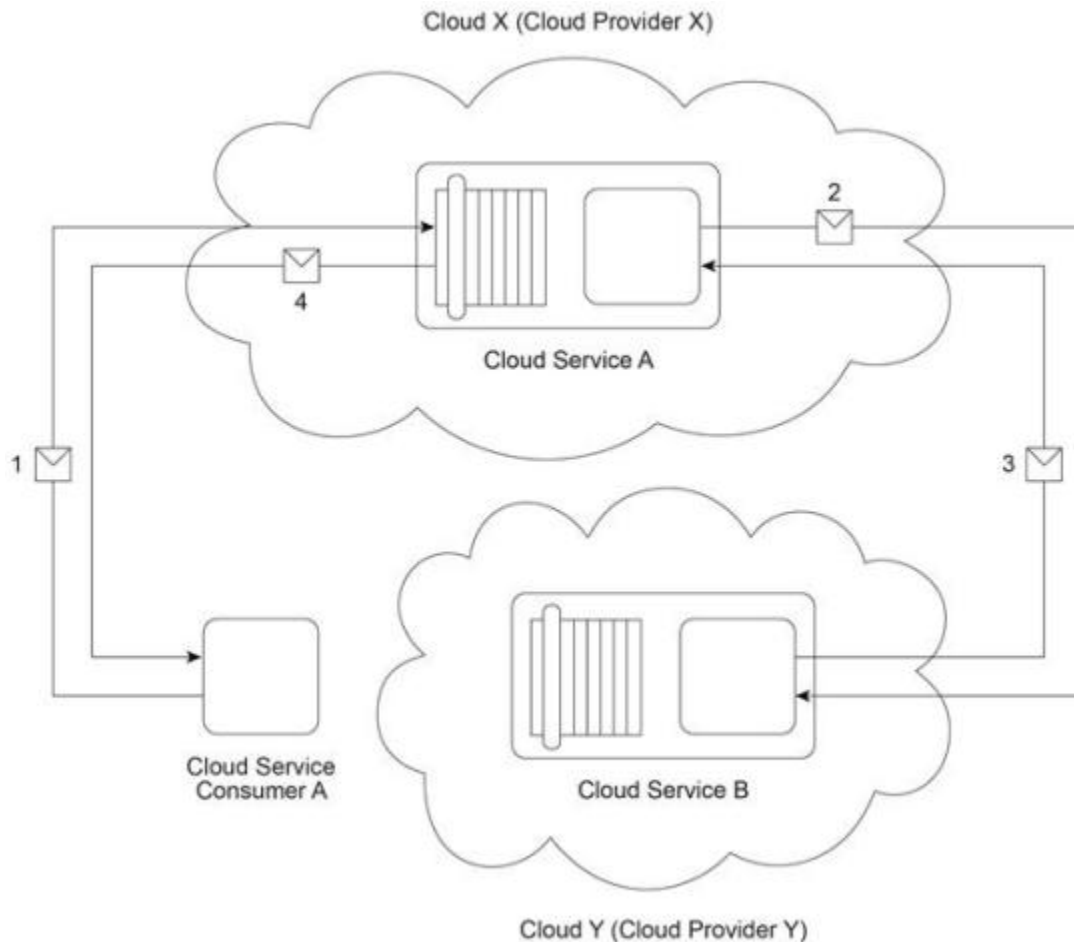


---

**Question: 1**

---

Cloud Service Consumer A invokes Cloud Service A from Cloud X (owned by Cloud Provider X) (1). To fulfill the request from Cloud Service Consumer A, Cloud Service A needs to invoke Cloud Service B that resides on Cloud Y (owned by Cloud Provider Y) (2). After completing its processing, Cloud Service B sends a response to Cloud Service A (3). Cloud Service A verifies the response and then finally sends its response to Cloud Service Consumer A (4).



The guaranteed availability of the Cloud Service A implementation is 95% and the guaranteed availability of the Cloud Service B implementation is 95%. Which of the following statements accurately describes the actual availability that Cloud Service Consumer A can receive based on the described scenario?

- A. Because Cloud Service Consumer A's response message is processed by two separate cloud services, the combined availability increases as follows:  
 $1 - (1 - 0.95) \times (1 - 0.95) = 0.9975$  or 99.75%
- B. Because Cloud Service A acts as both a cloud service and cloud service consumer in order to process Cloud Service Consumer B's request message, Cloud Service A forms a dependency on Cloud Service B. As a result, the combined availability decreases, as follows:  
 $0.95 \times 0.95 = 0.9025$  or 90.25%
- C. Cloud Service Consumer A benefits from redundant cloud service implementations, thereby increasing the guaranteed availability as follows:

$$1 - (1 - (0.95 - 0.1)) \times (1 - (0.95 - 0.1)) = 0.9775 \text{ or } 97.75\%$$

D. As a result of the dependency formed by Cloud Service A on Cloud Service B, the combined availability decreases significantly as follows:

$$(0.95 \times 0.95) - 0.1 = 0.8025 \text{ or } 80.25\%$$

---

**Answer: B**

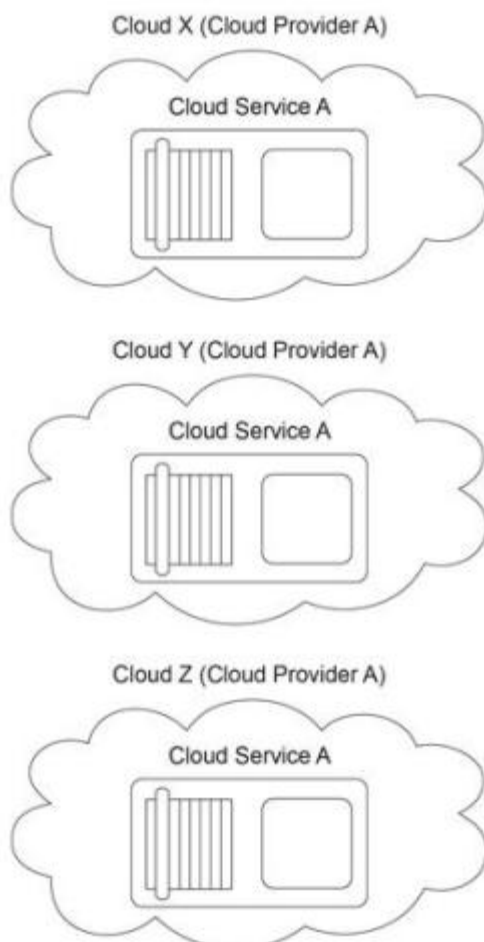
---

---

**Question: 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**

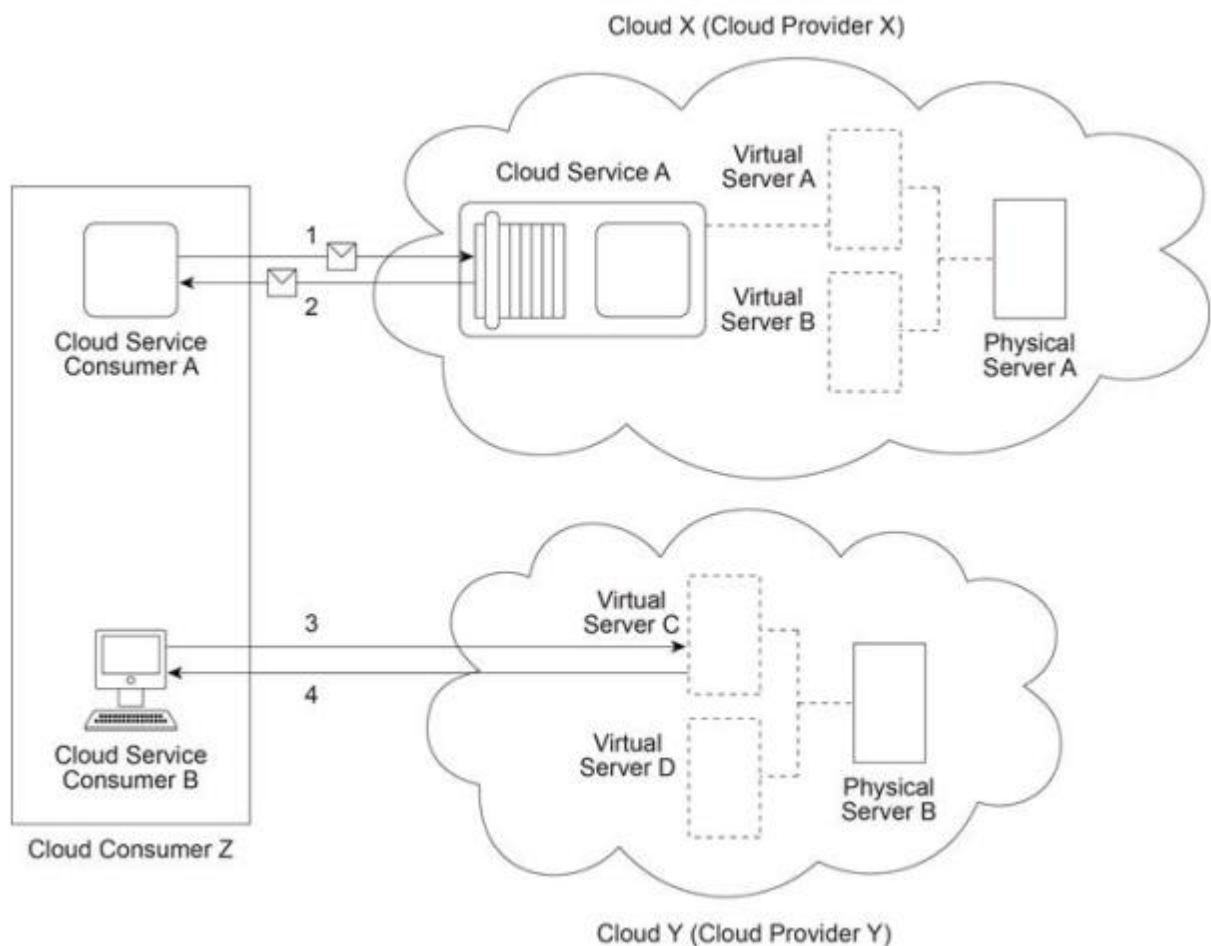
---

---

**Question: 3**

---

Cloud Provider X has deployed a virtualization environment in Cloud X comprised of Physical Server A hosting Virtual Servers A and B. Cloud Provider X implements Cloud Service A on Virtual Server A and makes it available to Cloud Service Consumer A, which interacts with Cloud Service A by sending and receiving messages (1, 2). Cloud Provider Y has deployed a virtualization environment comprised of Physical Server B hosting Virtual Servers C and D. Virtual Server C is made available to Cloud Service Consumer B, which interacts with Virtual Server C (3,4) in order to prepare for the deployment of a new cloud service that will be used internally by Cloud Provider Y to process data obtained from Cloud Service A .



Cloud Consumer Z and Cloud Provider X belong to the same organization. Cloud Provider Y is a third-party organization. Which of the following statements provides a valid scenario that accurately describes the involvement of cloud deployment models, cloud delivery models, roles and/or boundaries? (Note that the correct answer represents one of multiple valid scenarios that can exist.)

A. Cloud X is based on the private cloud deployment model. Cloud Service A is based on the SaaS delivery model. Cloud Y is based on the private cloud deployment model. Virtual Server C is being offered as part of the IaaS delivery model. A cloud resource administrator working for Cloud Consumer Z uses Cloud Service Consumer B to access Virtual Server C. Cloud Consumer Z is the cloud service owner of Cloud Service A. Cloud Consumer Z's organizational boundary encompasses Cloud Service Consumers A and B. Cloud Consumer Z's trust boundary encompasses Cloud Service Consumers A and B, Cloud Service A and Virtual Server C.

B. Cloud X is based on the private cloud deployment model. Cloud Service A is based on the SaaS delivery model. Cloud Y is based on the community cloud deployment model. Virtual Server C is being offered as part of the IaaS delivery model. A cloud resource administrator working for Cloud Consumer Z uses Cloud Service Consumer A to access Cloud Service A. Cloud Consumer Z's organizational and trust boundaries encompass Cloud Service Consumers A and B, Cloud Service A and Virtual Server C.

C. Cloud X is based on the private cloud deployment model. Cloud Service A is based on the SaaS delivery model. Cloud Y is based on the public cloud deployment model. Virtual Server C is being offered as part of the IaaS delivery model. A cloud resource administrator working for Cloud

Consumer Z uses Cloud Service Consumer B to access Virtual Server C . Cloud Consumer Z is the cloud service owner of Cloud Service A . Cloud Consumer Z's organizational boundary encompasses Cloud Service Consumers A and B. Cloud Consumer T s trust boundary encompasses Cloud Service Consumers A and B, Cloud Service A and Virtual Server C .

D. Cloud X is based on the private cloud deployment model. Cloud Service A is based on the SaaS delivery model. Cloud Y is based on the public cloud deployment model. Virtual Server C is being offered as part of the IaaS delivery model. A cloud resource administrator working for Cloud Consumer Z uses Cloud Service Consumer B to access Virtual Server C . Cloud Consumer Z' s trust boundary encompasses Cloud Service Consumers A and B, Cloud Service A and Virtual Server C . The organization that owns Cloud Consumer Z is the cloud service owner of Cloud Service A .

---

**Answer: D**

---

---

**Question: 4**


---

A company is planning to build and launch a new SaaS product that will be available for use by the general public. It intends to build the service on-premise and then deploy it in a public cloud. The company has the following set of four requirements for the implementation of the new service:

1. The cloud service needs to exchange messages primarily by using HTTP methods and other features provided by HTTP.
2. The cloud service needs to store highly structured data with potentially complex relationships.
3. The cloud service needs to be deployed on a dedicated virtual server that can be administered with a high level of control by the cloud consumer's own cloud resource administrator.
4. The cloud service needs to be deployed with a minimal amount of integration testing.

For this project, the company has a very limited budget. The company is assessing the IT resources that are offered by Clouds X and Y within the constraints of its limited budget.

Cloud X can offer an IaaS environment with very few proprietary characteristics that includes a database that supports only no relational storage, as well as support for the deployment and usage of REST services.



Cloud Y can offer a PaaS environment with a pre-configured virtual server that includes native support for WSDL and SOAP, as well as a database that supports only relational storage. The implementation of a new service within Cloud Y will require compliance to a high level of proprietary characteristics. As previously listed, the company has identified four specific implementation requirements for its new cloud service. Which of the following statements correctly identifies how many of the four requirements Clouds X and Y can directly fulfill?

- A. Cloud X fulfills 0 out of 4 requirements. Cloud Y fulfills 4 out of 4 requirements.
- B. Cloud X fulfills 1 out of 4 requirements. Cloud Y fulfills 3 out of 4 requirements.
- C. Cloud X fulfills 2 out of 4 requirements. Cloud Y fulfills 2 out of 4 requirements.
- D. Cloud X fulfills 3 out of 4 requirements. Cloud Y fulfills 1 out of 4 requirements.

---

**Answer: D**

---