
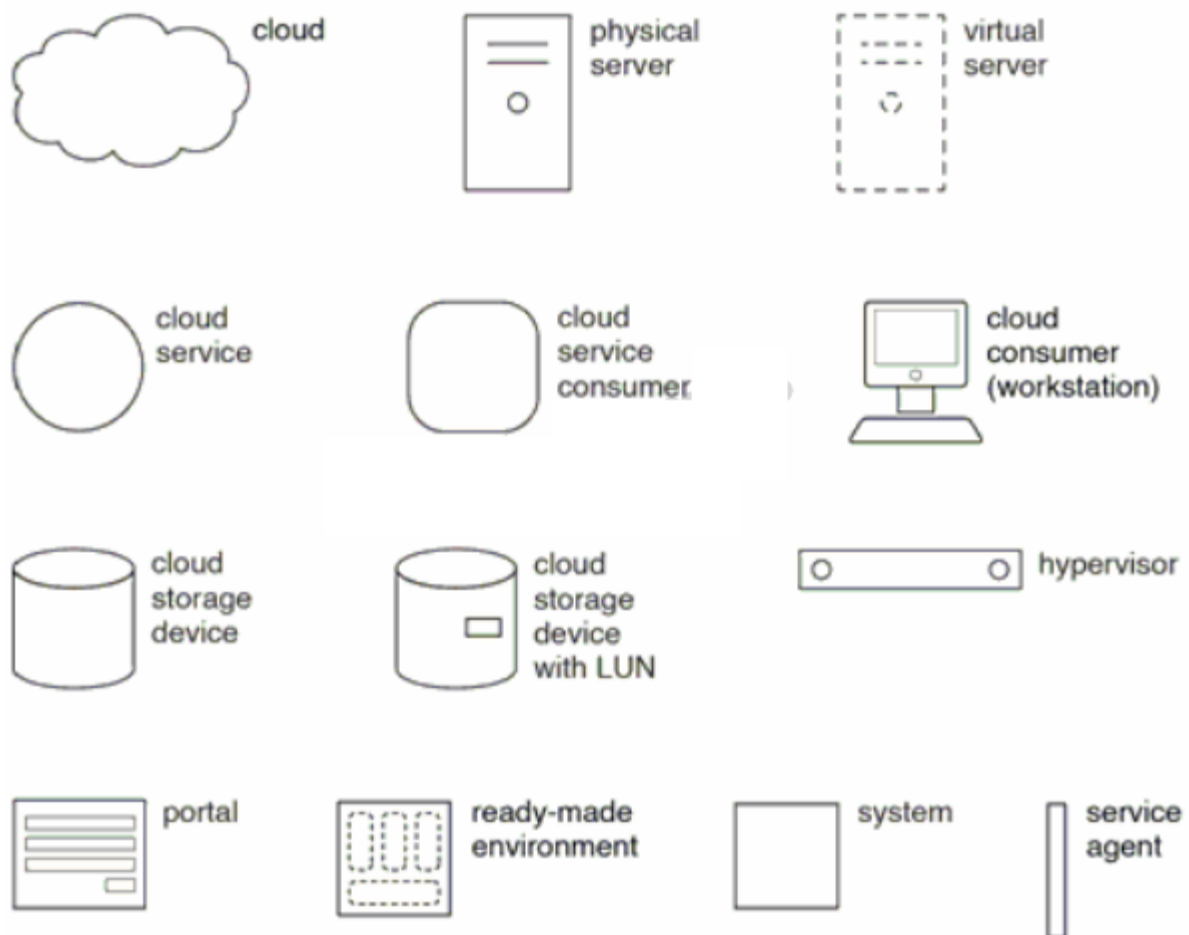

Question: 1

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).



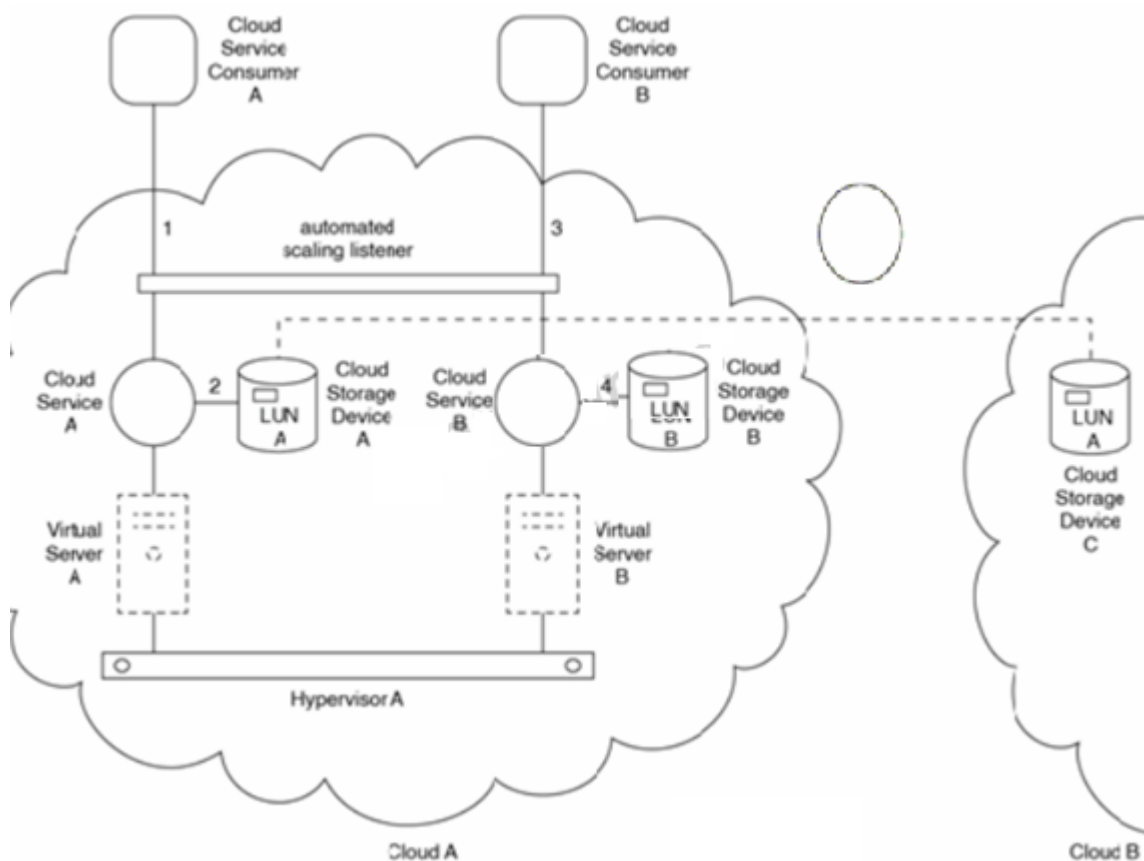


Cloud Service Consumer A belongs to Organization A, Organization A uses LUN A and LUN B on Cloud Storage Device A to store their important client account data. Cloud Storage Device A is a low-performance cloud storage device, which begins to cause performance issues as more data is added to LUNs A and 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?

- 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.

Question: 2

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 Hypervisor A, 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 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: 3

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

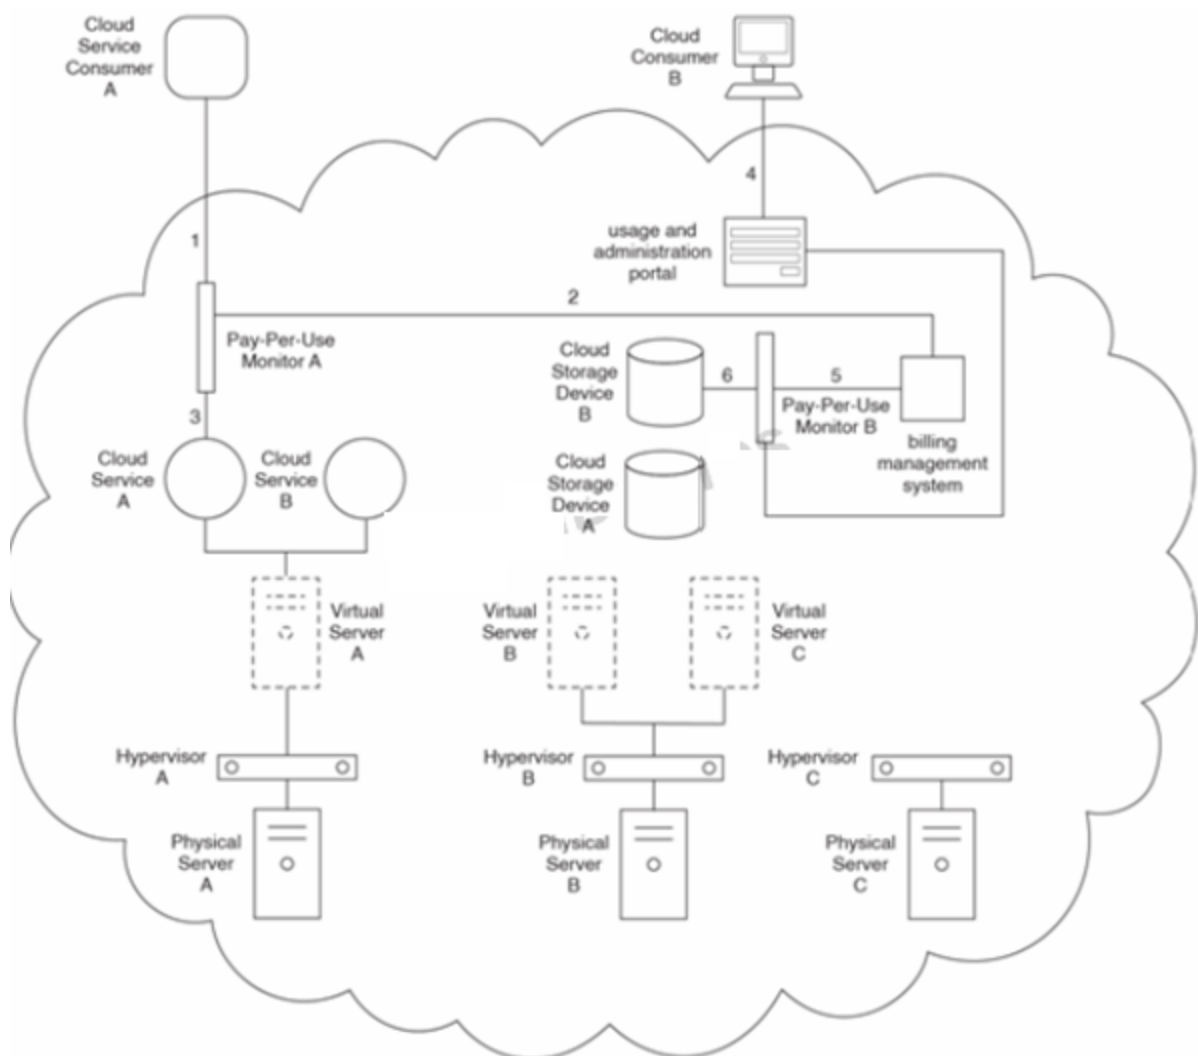
(2). Cloud Service A receives and processes the request

(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. 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: 4

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.



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?

- 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
