

Amazon

Exam Questions AWS-Certified-Cloud-Practitioner

Amazon AWS Certified Cloud Practitioner



NEW QUESTION 1

- (Topic 2)

Which task is the responsibility of AWS when using AWS services?

- A. Management of IAM user permissions
- B. Creation of security group rules for outbound access
- C. Maintenance of physical and environmental controls
- D. Application of Amazon EC2 operating system patches

Answer: C

Explanation:

AWS is responsible for maintaining the physical and environmental controls of the AWS Cloud, such as power, cooling, fire suppression, and physical security¹. The customer is responsible for managing the IAM user permissions, creating security group rules for outbound access, applying Amazon EC2 operating system patches, and other aspects of security in the cloud¹.

NEW QUESTION 2

- (Topic 2)

A company wants an in-memory data store that is compatible with open source in the cloud.

Which AWS service should the company use?

- A. Amazon DynamoDB
- B. Amazon ElastiCache
- C. Amazon Elastic Block Store (Amazon EBS)
- D. Amazon Redshift

Answer: B

Explanation:

Amazon ElastiCache is a fully managed in-memory data store service that is compatible with open source engines such as Redis and Memcached¹. It provides fast and scalable performance for applications that require high throughput and low latency¹. Amazon DynamoDB is a fully managed NoSQL database service that provides consistent and single-digit millisecond latency at any scale². Amazon EBS is a block storage service that provides persistent and durable storage volumes for Amazon EC2 instances³. Amazon Redshift is a fully managed data warehouse service that allows users to run complex analytic queries using SQL⁴.

NEW QUESTION 3

- (Topic 2)

Which AWS Cloud design principle does a company follow by using AWS CloudTrail?

- A. Recover automatically.
- B. Perform operations as code.
- C. Measure efficiency.
- D. Ensure traceability.

Answer: D

Explanation:

The company follows the AWS Cloud design principle of ensuring traceability by using AWS CloudTrail. AWS CloudTrail is a service that records the API calls and events made by or on behalf of the AWS account. The company can use AWS CloudTrail to monitor, audit, and analyze the activity and changes in their AWS resources and applications. AWS CloudTrail helps the company to achieve compliance, security, governance, and operational efficiency. Recovering automatically, performing operations as code, and measuring efficiency are other AWS Cloud design principles, but they are not directly related to using AWS CloudTrail. Recovering automatically means that the company can design their cloud workloads to handle failures gracefully and resume normal operations without manual intervention. Performing operations as code means that the company can automate the creation, configuration, and management of their cloud resources using scripts or templates. Measuring efficiency means that the company can monitor and optimize the performance and utilization of their cloud resources and applications³⁴.

NEW QUESTION 4

- (Topic 1)

Which AWS service or tool does AWS Control Tower use to create resources?

- A. AWS CloudFormation
- B. AWS Trusted Advisor
- C. AWS Directory Service
- D. AWS Cost Explorer

Answer: A

Explanation:

AWS Control Tower uses AWS CloudFormation to create resources in your landing zone. AWS CloudFormation is a service that helps you model and set up your AWS resources using templates. AWS Control Tower supports creating AWS::ControlTower::EnabledControl resources in AWS CloudFormation. Therefore, the correct answer is A. You can learn more about AWS Control Tower and AWS CloudFormation from this page.

NEW QUESTION 5

- (Topic 1)

A company needs to run its existing custom, nonproduction workloads in the AWS Cloud quickly and cost-effectively.

The workloads can recover from interruptions easily. Which pricing model should the company use?

- A. Reserved Instances

- B. On-Demand Instances
- C. Spot Instances
- D. Dedicated Hosts

Answer: C

Explanation:

The correct answer is C because Spot Instances are the pricing model that enables the company to run its existing custom, nonproduction workloads in the AWS Cloud quickly and cost-effectively. Spot Instances are spare Amazon EC2 instances that are available at up to 90% discount compared to On-Demand prices. Spot Instances are suitable for stateless, fault-tolerant, and flexible workloads that can recover from interruptions easily. The other options are incorrect because they are not the pricing model that enables the company to run its existing custom, nonproduction workloads in the AWS Cloud quickly and cost-effectively. Reserved Instances are Amazon EC2 instances that are reserved for a specific period of time (one or three years) in exchange for a lower hourly rate. Reserved Instances are suitable for steady-state or predictable workloads that run for a long duration. On-Demand Instances are Amazon EC2 instances that are launched and billed at a fixed hourly rate. On-Demand Instances are suitable for short-term, irregular, or unpredictable workloads that cannot be interrupted. Dedicated Hosts are physical servers that are dedicated to a single customer. Dedicated Hosts are suitable for workloads that require regulatory compliance or data isolation. Reference: Amazon EC2 Instance Purchasing Options

NEW QUESTION 6

- (Topic 1)

Which of the following are design principles for reliability in the AWS Cloud? (Select TWO.)

- A. Build architectures with tightly coupled resources.
- B. Use AWS Trusted Advisor to meet security best practices.
- C. Use automation to recover immediately from failure.
- D. Rightsize Amazon EC2 instances to ensure optimal performance.
- E. Simulate failures to test recovery processes.

Answer: CE

Explanation:

The design principles for reliability in the AWS Cloud are:

- ? Test recovery procedures. The best way to ensure that systems can recover from failures is to regularly test them using simulated scenarios. This can help identify gaps and improve the recovery process.
- ? Automatically recover from failure. By using automation, systems can detect and correct failures without human intervention. This can reduce the impact and duration of failures and improve the availability of the system.
- ? Scale horizontally to increase aggregate system availability. By adding more redundant resources to the system, the impact of individual resource failures can be reduced. This can also improve the performance and scalability of the system.
- ? Stop guessing capacity. By using monitoring and automation, systems can adjust the capacity based on the demand and performance metrics. This can prevent failures due to insufficient or excessive capacity and optimize the cost and efficiency of the system.
- ? Manage change in automation. By using automation, changes to the system can be applied in a consistent and controlled manner. This can reduce the risk of human errors and configuration drifts that can cause failures. AWS Well- Architected Framework

NEW QUESTION 7

- (Topic 1)

Which of the following is available to a company that has an AWS Business Support plan?

- A. AWS Support concierge
- B. AWS DDoS Response Team (DRT)
- C. AWS technical account manager (TAM)
- D. AWS Health API

Answer: D

Explanation:

AWS Health API is available to a company that has an AWS Business Support plan. The AWS Health API provides programmatic access to the AWS Health information that is presented in the AWS Personal Health Dashboard. The AWS Health API can help users get timely and personalized information about events that can affect the availability and performance of their AWS resources, such as scheduled maintenance, network issues, or service disruptions. The AWS Health API can also integrate with other AWS services, such as Amazon CloudWatch Events and AWS Lambda, to enable automated actions and notifications. AWS Health API OverviewAWS Support Plans

NEW QUESTION 8

- (Topic 1)

Which AWS Support plan assigns an AWS concierge agent to a company's account?

- A. AWS Basic Support
- B. AWS Developer Support
- C. AWS Business Support
- D. AWS Enterprise Support

Answer: D

Explanation:

AWS Enterprise Support is the AWS Support plan that assigns an AWS concierge agent to a company's account. AWS Enterprise Support is the highest level of support that AWS offers, and it provides the most comprehensive and personalized assistance. An AWS concierge agent is a dedicated technical account manager who acts as a single point of contact for the company and helps to optimize the AWS environment, resolve issues, and access AWS experts. For more information, see [AWS Support Plans] and [AWS Concierge Support].

NEW QUESTION 9

- (Topic 1)

A company is migrating a relational database server to the AWS Cloud. The company wants to minimize administrative overhead of database maintenance tasks. Which AWS service will meet these requirements?

- A. Amazon DynamoDB
- B. Amazon EC2
- C. Amazon Redshift
- D. Amazon RDS

Answer: D

Explanation:

Amazon RDS is the AWS service that will meet the requirements of migrating a relational database server to the AWS Cloud and minimizing administrative overhead of database maintenance tasks. Amazon RDS is a fully managed relational database service that handles routine database tasks, such as provisioning, patching, backup, recovery, failure detection, and repair. Amazon RDS supports several database engines, such as MySQL, PostgreSQL, Oracle, SQL Server, and Amazon Aurora5.

NEW QUESTION 10

- (Topic 1)

A cloud engineer needs to download AWS security and compliance documents for an upcoming audit. Which AWS service can provide the documents?

- A. AWS Trusted Advisor
- B. AWS Artifact
- C. AWS Well-Architected Tool
- D. AWS Systems Manager

Answer: B

Explanation:

AWS Artifact is the AWS service that can provide security and compliance documents for an upcoming audit. AWS Artifact is a self-service portal that allows users to access and download AWS compliance reports and agreements. These documents provide evidence of AWS's compliance with global, regional, and industry-specific security standards and regulations

NEW QUESTION 10

- (Topic 1)

A company is hosting a web application in a Docker container on Amazon EC2. AWS is responsible for which of the following tasks?

- A. Scaling the web application and services developed with Docker
- B. Provisioning or scheduling containers to run on clusters and maintain their availability
- C. Performing hardware maintenance in the AWS facilities that run the AWS Cloud
- D. Managing the guest operating system, including updates and security patches

Answer: C

Explanation:

AWS is responsible for performing hardware maintenance in the AWS facilities that run the AWS Cloud. This is part of the shared responsibility model, where AWS is responsible for the security of the cloud, and the customer is responsible for security in the cloud. AWS is also responsible for the global infrastructure that runs all of the services offered in the AWS Cloud, including the hardware, software, networking, and facilities that run AWS Cloud services3. The customer is responsible for the guest operating system, including updates and security patches, as well as the web application and services developed with Docker4.

NEW QUESTION 12

- (Topic 1)

In which of the following AWS services should database credentials be stored for maximum security?

- A. AWS Identity and Access Management (IAM)
- B. AWS Secrets Manager
- C. Amazon S3
- D. AWS Key Management Service (AWS KMS)

Answer: B

Explanation:

AWS Secrets Manager is the AWS service where database credentials should be stored for maximum security. AWS Secrets Manager helps to protect the secrets, such as database credentials, passwords, API keys, and tokens, that are used to access applications, services, and resources. AWS Secrets Manager enables secure storage, encryption, rotation, and retrieval of the secrets. AWS Secrets Manager also integrates with other AWS services, such as AWS Identity and Access Management (IAM), AWS Key Management Service (AWS KMS), and AWS Lambda. For more information, see [What is AWS Secrets Manager?] and [Getting Started with AWS Secrets Manager].

NEW QUESTION 15

- (Topic 1)

Which AWS service or tool provides users with the ability to monitor AWS service quotas?

- A. AWS CloudTrail
- B. AWS Cost and Usage Reports
- C. AWS Trusted Advisor
- D. AWS Budgets

Answer: C

Explanation:

The correct answer is C because AWS Trusted Advisor is an AWS service or tool that provides users with the ability to monitor AWS service quotas. AWS Trusted Advisor is an online tool that provides users with real-time guidance to help them provision their resources following AWS best practices. One of the categories of checks that AWS Trusted Advisor performs is service limits, which monitors the usage of each AWS service and alerts users when they are close to reaching the default limit. The other options are incorrect because they are not AWS services or tools that provide users with the ability to monitor AWS service quotas. AWS CloudTrail is a service that enables users to track user activity and API usage across their AWS account. AWS Cost and Usage Reports is a tool that enables users to access comprehensive information about their AWS costs and usage. AWS Budgets is a tool that enables users to plan their service usage, costs, and reservations. Reference: [AWS Trusted Advisor FAQs]

NEW QUESTION 18

- (Topic 1)

A company hosts an application on an Amazon EC2 instance. The EC2 instance needs to access several AWS resources, including Amazon S3 and Amazon DynamoDB.

What is the MOST operationally efficient solution to delegate permissions?

- A. Create an IAM role with the required permission
- B. Attach the role to the EC2 instance.
- C. Create an IAM user and use its access key and secret access key in the application.
- D. Create an IAM user and use its access key and secret access key to create a CLI profile in the EC2 instance.
- E. Create an IAM role with the required permission
- F. Attach the role to the administrative IAM user.

Answer: A

Explanation:

Creating an IAM role with the required permissions and attaching the role to the EC2 instance is the most operationally efficient solution to delegate permissions. An IAM role is an entity that defines a set of permissions for making AWS service requests. An IAM role can be assumed by an EC2 instance to access other AWS resources, such as Amazon S3 and Amazon DynamoDB, without having to store any credentials on the instance. This solution is more secure and scalable than using IAM users and their access keys. For more information, see [IAM Roles for Amazon EC2] and [Using an IAM Role to Grant Permissions to Applications Running on Amazon EC2 Instances].

NEW QUESTION 23

- (Topic 1)

A company has a social media platform in which users upload and share photos with other users. The company wants to identify and remove inappropriate photos. The company has no machine learning (ML) scientists and must build this detection capability with no ML expertise.

Which AWS service should the company use to build this capability?

- A. Amazon SageMaker
- B. Amazon Textract
- C. Amazon Rekognition
- D. Amazon Comprehend

Answer: C

Explanation:

Amazon Rekognition is the AWS service that the company should use to build the capability of identifying and removing inappropriate photos. Amazon Rekognition is a service that uses deep learning technology to analyze images and videos for various purposes, such as face detection, object recognition, text extraction, and content moderation. Amazon Rekognition can help users detect unsafe or inappropriate content in images and videos, such as nudity, violence, or drugs, and provide confidence scores for each label. Amazon Rekognition does not require any machine learning expertise, and users can easily integrate it with other AWS services

NEW QUESTION 24

- (Topic 1)

Which best practice for cost governance does this example show?

- A. Resource controls
- B. Cost allocation
- C. Architecture optimization
- D. Tagging enforcement

Answer: C

Explanation:

Architecture optimization is the best practice for cost governance that this example shows. Architecture optimization is the process of designing and implementing AWS solutions that are efficient, scalable, and cost-effective. By using specific AWS services to improve efficiency and reduce cost, the company is following the architecture optimization best practice. Some of the techniques for architecture optimization include using the right size and type of resources, leveraging elasticity and scalability, choosing the most suitable storage class, and using serverless and managed services.

NEW QUESTION 26

- (Topic 1)

What can a user accomplish using AWS CloudTrail?

- A. Generate an IAM user credentials report.
- B. Record API calls made to AWS services.
- C. Assess the compliance of AWS resource configurations with policies and guidelines.
- D. Ensure that Amazon EC2 instances are patched with the latest security update
- E. A company uses Amazon Workspaces.

Answer: B

Explanation:

AWS CloudTrail is an AWS service that enables users to accomplish the task of recording API calls made to AWS services. AWS CloudTrail is a service that tracks user activity and API usage across the AWS account. AWS CloudTrail records the details of every API call made to AWS services, such as the identity of the caller, the time of the call, the source IP address of the caller, the parameters and responses of the call, and more. Users can use AWS CloudTrail to audit, monitor, and troubleshoot their AWS resources and actions. The other options are incorrect because they are not tasks that users can accomplish using AWS CloudTrail. Generating an IAM user credentials report is a task that users can accomplish using IAM, which is an AWS service that enables users to manage access and permissions to AWS resources and services. Assessing the compliance of AWS resource configurations with policies and guidelines is a task that users can accomplish using AWS Config, which is an AWS service that enables users to assess, audit, and evaluate the configurations of their AWS resources. Ensuring that Amazon EC2 instances are patched with the latest security updates is a task that users can accomplish using AWS Systems Manager, which is an AWS service that enables users to automate operational tasks, manage configuration and compliance, and monitor system health and performance. Reference: AWS CloudTrail FAQs

NEW QUESTION 29

- (Topic 1)

What are the characteristics of Availability Zones? (Select TWO.)

- A. All Availability Zones in an AWS Region are interconnected with high-bandwidth, low- latency networking
- B. Availability Zones are physically separated by a minimum of distance of 150 km (100 miles).
- C. All traffic between Availability Zones is encrypted.
- D. Availability Zones within an AWS Region share redundant power, networking, and connectivity.
- E. Every Availability Zone contains a single data center.

Answer: AD

Explanation:

Availability Zones are physically separate locations within an AWS Region that are engineered to be isolated from failures. Each Availability Zone has independent power, cooling, and physical security, and is connected to other Availability Zones in the same Region by a low-latency network. Therefore, the correct answers are A and D. You can learn more about Availability Zones and their characteristics from this page.

NEW QUESTION 34

- (Topic 1)

Which AWS services or features can control VPC traffic? (Select TWO.)

- A. Security groups
- B. AWS Direct Connect
- C. Amazon GuardDuty
- D. Network ACLs
- E. Amazon Connect

Answer: AD

Explanation:

The AWS services or features that can control VPC traffic are security groups and network ACLs. Security groups are stateful firewalls that control the inbound and outbound traffic at the instance level. You can assign one or more security groups to each instance in a VPC, and specify the rules that allow or deny traffic based on the protocol, port, and source or destination. Network ACLs are stateless firewalls that control the inbound and outbound traffic at the subnet level. You can associate one network ACL with each subnet in a VPC, and specify the rules that allow or deny traffic based on the protocol, port, and source or destination. AWS Direct Connect, Amazon GuardDuty, and Amazon Connect are not services or features that can control VPC traffic. AWS Direct Connect is a service that establishes a dedicated network connection between your premises and AWS. Amazon GuardDuty is a service that monitors your AWS account and workloads for malicious or unauthorized activity. Amazon Connect is a service that provides a cloud-based contact center solution.

NEW QUESTION 39

- (Topic 1)

Which option is an advantage of AWS Cloud computing that minimizes variable costs?

- A. High availability
- B. Economies of scale
- C. Global reach
- D. Agility

Answer: B

Explanation:

Economies of scale is the advantage of AWS Cloud computing that minimizes variable costs. Economies of scale refers to the reduction in the cost per unit as the output increases. AWS Cloud computing leverages economies of scale by providing a large pool of shared resources that can be accessed on demand and paid for as needed. AWS Cloud computing also passes the cost savings to the customers by offering lower prices and discounts. For more information, see Economies of Scale and AWS Pricing.

NEW QUESTION 42

- (Topic 1)

An ecommerce company has migrated its IT infrastructure from an on-premises data center to the AWS Cloud.

Which AWS service is used to track, record, and audit configuration changes made to AWS resources?

- A. AWS Shield
- B. AWS Config
- C. AWS IAM
- D. Amazon Inspector

Answer: B

Explanation:

AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. AWS Config continuously monitors and records your AWS resource configurations and allows you to automate the evaluation of recorded configurations against desired configurations. With AWS Config, you can review changes in configurations and relationships between AWS resources, dive into detailed resource configuration histories, and determine your overall compliance against the configurations specified in your internal guidelines³.

NEW QUESTION 44

- (Topic 1)

Which AWS service meets this requirement?

- A. AWS CloudFormation
- B. AWS Elastic Beanstalk
- C. AWS Cloud9
- D. AWS CloudShell

Answer: A

Explanation:

AWS CloudFormation is a service that gives developers and businesses an easy way to create a collection of related AWS and third-party resources, and provision and manage them in an orderly and predictable fashion. You can use AWS CloudFormation's sample templates or create your own templates to describe the AWS and third-party resources, and any associated dependencies or runtime parameters, required to run your application.

NEW QUESTION 48

- (Topic 1)

Which feature of the AWS Cloud gives users the ability to pay based on current needs rather than forecasted needs?

- A. AWS Budgets
- B. Pay-as-you-go pricing
- C. Volume discounts
- D. Savings Plans

Answer: B

Explanation:

Pay-as-you-go pricing is the feature of the AWS Cloud that gives users the ability to pay based on current needs rather than forecasted needs. Pay-as-you-go pricing means that users only pay for the AWS services and resources they use, without any upfront or long-term commitments. This allows users to scale up or down their usage depending on their changing business requirements, and avoid paying for idle or unused capacity. Pay-as-you-go pricing also enables users to benefit from the economies of scale and lower costs of AWS as they grow their business⁵

NEW QUESTION 53

- (Topic 1)

company wants to protect its AWS Cloud information, systems, and assets while performing risk assessment and mitigation tasks. Which pillar of the AWS Well-Architected Framework is supported by these goals?

- A. Reliability
- B. Security
- C. Operational excellence
- D. Performance efficiency

Answer: B

Explanation:

The pillar of the AWS Well-Architected Framework that is supported by the goals of protecting AWS Cloud information, systems, and assets while performing risk assessment and mitigation tasks is security. Security is the ability to protect information, systems, and assets while delivering business value through risk assessments and mitigation strategies. The security pillar covers topics such as identity and access management, data protection, infrastructure protection, detective controls, incident response, and compliance

NEW QUESTION 55

- (Topic 1)

A company has been storing monthly reports in an Amazon S3 bucket. The company exports the report data into comma-separated values (.csv) files. A developer wants to write a simple query that can read all of these files and generate a summary report.

Which AWS service or feature should the developer use to meet these requirements with the LEAST amount of operational overhead?

- A. Amazon S3 Select
- B. Amazon Athena
- C. Amazon Redshift
- D. Amazon EC2

Answer: B

Explanation:

Amazon Athena is the AWS service that the developer should use to write a simple query that can read all of the .csv files stored in an Amazon S3 bucket and generate a summary report. Amazon Athena is an interactive query service that allows users to analyze data in Amazon S3 using standard SQL. Amazon Athena does not require any server setup or management, and users only pay for the queries they run. Amazon Athena can handle various data formats, including .csv, and can integrate with other AWS services such as Amazon QuickSight for data visualization

NEW QUESTION 59

- (Topic 1)

A company is migrating an application that includes an Oracle database to AWS. The company cannot rewrite the application. To which AWS service could the company migrate the database?

- A. Amazon Athena
- B. Amazon DynamoDB®
- C. Amazon RDS
- D. Amazon DocumentDB (with MongoDB compatibility)

Answer: C

Explanation:

Amazon Relational Database Service (Amazon RDS) is a service that provides fully managed relational database engines. Amazon RDS supports several database engines, including Oracle, MySQL, PostgreSQL, MariaDB, SQL Server, and Amazon Aurora. Amazon RDS can be used to migrate an application that includes an Oracle database to AWS without rewriting the application, as long as the application is compatible with the Oracle version and edition supported by Amazon RDS. Amazon RDS can also provide benefits such as high availability, scalability, security, backup and restore, and performance optimization. [Amazon RDS Overview] AWS Certified Cloud Practitioner - aws.amazon.com

NEW QUESTION 63

- (Topic 1)

Which AWS service or feature captures information about the network traffic to and from an Amazon EC2 instance?

- A. VPC Reachability Analyzer
- B. Amazon Athena
- C. VPC Flow Logs
- D. AWS X-Ray

Answer: C

Explanation:

The correct answer is C because VPC Flow Logs is an AWS service or feature that captures information about the network traffic to and from an Amazon EC2 instance. VPC Flow Logs is a feature that enables customers to capture information about the IP traffic going to and from network interfaces in their VPC. VPC Flow Logs can help customers to monitor and troubleshoot connectivity issues, such as traffic not reaching an instance or traffic being rejected by a security group. The other options are incorrect because they are not AWS services or features that capture information about the network traffic to and from an Amazon EC2 instance. VPC Reachability Analyzer is an AWS service or feature that enables customers to perform connectivity testing between resources in their VPC and identify configuration issues that prevent connectivity. Amazon Athena is an AWS service that enables customers to query data stored in Amazon S3 using standard SQL. AWS X-Ray is an AWS service that enables customers to analyze and debug distributed applications, such as those built using a microservices architecture.

Reference: VPC Flow Logs

NEW QUESTION 67

- (Topic 1)

Which AWS service should a cloud practitioner use to receive real-time guidance for provisioning resources, based on AWS best practices related to security, cost optimization, and service limits?

- A. AWS Trusted Advisor
- B. AWS Config
- C. AWS Security Hub
- D. AWS Systems Manager

Answer: A

Explanation:

AWS Trusted Advisor is the AWS service that provides real-time guidance for provisioning resources, based on AWS best practices related to security, cost optimization, and service limits. AWS Trusted Advisor inspects the user's AWS environment and provides recommendations for improving performance, security, and reliability, reducing costs, and following best practices. AWS Trusted Advisor also alerts the user when they are approaching or exceeding their service limits, and helps them request limit increases³.

NEW QUESTION 69

- (Topic 1)

A company wants to manage access and permissions for its third-party software as a service (SaaS) applications. The company wants to use a portal where end users can access assigned AWS accounts and AWS Cloud applications. Which AWS service should the company use to meet these requirements?

- A. Amazon Cognito
- B. AWS IAM Identity Center (AWS Single Sign-On)
- C. AWS Identity and Access Management (IAM)
- D. AWS Directory Service for Microsoft Active Directory

Answer: B

Explanation:

AWS IAM Identity Center (AWS Single Sign-On) is the AWS service that the company should use to meet the requirements of managing access and permissions for its third-party SaaS applications. AWS Single Sign-On is a cloud-based service that makes it easy to centrally manage single sign-on (SSO) access to multiple AWS accounts and business applications. You can use AWS Single Sign-On to enable your users to sign in to a user portal with their existing corporate credentials and access all of their assigned accounts and applications from one place⁴.

NEW QUESTION 73

- (Topic 1)

A company is launching a new application in the AWS Cloud. The application will run on an Amazon EC2 instance. More EC2 instances will be needed when the workload increases.

Which AWS service or tool can the company use to launch the number of EC2 instances that will be needed to handle the workload?

- A. Elastic Load Balancing
- B. Amazon EC2 Auto Scaling
- C. AWS App2Container (A2C)
- D. AWS Systems Manager

Answer: B

Explanation:

Amazon EC2 Auto Scaling is the AWS service or tool that can help the company launch the number of EC2 instances that will be needed to handle the workload. Amazon EC2 Auto Scaling automatically adjusts the capacity of the EC2 instances based on the demand and the predefined scaling policies. Amazon EC2 Auto Scaling also helps to improve availability and reduce costs by scaling in and out as needed. For more information, see [What is Amazon EC2 Auto Scaling?](#) and [\[Getting Started with Amazon EC2 Auto Scaling\]](#).

NEW QUESTION 77

- (Topic 3)

A company needs a graph database service that is scalable and highly available.

Which AWS service meets these requirements?

- A. Amazon Aurora
- B. Amazon Redshift
- C. Amazon DynamoDB
- D. Amazon Neptune

Answer: D

Explanation:

The AWS service that meets the requirements of providing a graph database service that is scalable and highly available is Amazon Neptune. Amazon Neptune is a fast, reliable, and fully managed graph database service that supports property graph and RDF graph models. Amazon Neptune is designed to store billions of relationships and query the graph with milliseconds latency. Amazon Neptune also offers high availability and durability by replicating six copies of the data across three Availability Zones and continuously backing up the data to Amazon S3. Amazon Aurora, Amazon Redshift, and Amazon DynamoDB are other AWS services that provide relational or non-relational database solutions, but they do not support graph database models.

NEW QUESTION 79

- (Topic 3)

A company wants to grant users in one AWS account access to resources in another AWS account. The users do not currently have permission to access the resources.

Which AWS service will meet this requirement?

- A. IAM group
- B. IAM role
- C. IAM tag
- D. IAM Access Analyzer

Answer: B

Explanation:

IAM roles are a way to delegate access to resources in different AWS accounts. IAM roles allow users to assume a set of permissions for a limited time without having to create or share long-term credentials. IAM roles can be used to grant cross-account access by creating a trust relationship between the accounts and specifying the permissions that the role can perform. Users can then switch to the role and access the resources in the other account using temporary security credentials provided by the role. References: [Cross account resource access in IAM](#), [IAM tutorial: Delegate access across AWS accounts using IAM roles](#), [How to Enable Cross-Account Access to the AWS Management Console](#)

NEW QUESTION 82

- (Topic 3)

A customer runs an On-Demand Amazon Linux EC2 instance for 3 hours, 5 minutes, and 6 seconds.

For how much time will the customer be billed?

- A. 3 hours, 5 minutes
- B. 3 hours, 5 minutes, and 6 seconds
- C. 3 hours, 6 minutes
- D. 4 hours

Answer: C

Explanation:

Amazon EC2 usage is calculated by either the hour or the second based on the size of the instance, operating system, and the AWS Region where the instances are launched. Pricing is per instance-hour consumed for each instance, from the time an instance is launched until it's terminated or stopped. Each partial instance-hour consumed is billed per-second for Linux instances and as a full hour for all other instance types¹. Therefore, the customer will be billed for 3 hours and 6 minutes for running an On-Demand Amazon Linux EC2 instance for 3 hours, 5 minutes, and 6 seconds. References: [Understand Amazon EC2 instance-hours billing](#)

NEW QUESTION 87

- (Topic 3)

Which AWS service will allow a user to set custom cost and usage limits, and will alert when the thresholds are exceeded?

- A. AWS Organizations
- B. AWS Budgets
- C. Cost Explorer
- D. AWS Trusted Advisor

Answer: B

Explanation:

AWS Budgets allows you to set custom budgets that alert you when your costs or usage exceed (or are forecasted to exceed) your budgeted amount. You can also use AWS Budgets to set reservation utilization or coverage targets and receive alerts when your utilization drops below the threshold you define. AWS Budgets provides you with a comprehensive view of your cost and usage, as well as your reservation utilization and coverage1.

NEW QUESTION 92

- (Topic 2)

A company has an application that runs periodically in an on-premises environment. The application runs for a few hours most days, but runs for 8 hours a day for a week at the end of each month.

Which AWS service or feature should be used to host the application in the AWS Cloud?

- A. Amazon EC2 Standard Reserved Instances
- B. Amazon EC2 On-Demand Instances
- C. AWS Wavelength
- D. Application Load Balancer

Answer: B

Explanation:

Amazon EC2 On-Demand Instances are instances that you pay for by the second, with no long-term commitments or upfront payments4. This option is suitable for applications that have unpredictable or intermittent workloads, such as the one described in the question. Amazon EC2 Standard Reserved Instances are instances that you purchase for a one-year or three-year term, and pay a lower hourly rate compared to On-Demand Instances. This option is suitable for applications that have steady state or predictable usage. AWS Wavelength is a service that enables developers to build applications that deliver ultra-low latency to mobile devices and users by deploying AWS compute and storage at the edge of the 5G network. This option is not relevant for the application described in the question. Application Load Balancer is a type of load balancer that operates at the application layer and distributes traffic based on the content of the request. This option is not a service or feature to host the application, but rather to balance the traffic among multiple instances.

NEW QUESTION 96

- (Topic 2)

A company runs a database on Amazon Aurora in the us-east-1 Region. The company has a disaster recovery requirement that the database be available in another Region.

Which solution meets this requirement with minimal disruption to the database operations?

- A. Perform an Aurora Multi-AZ deployment.
- B. Deploy Aurora cross-Region read replicas.
- C. Create Amazon Elastic Block Store (Amazon EBS) volume snapshots for Aurora and copy them to another Region.
- D. Deploy Aurora Replicas.

Answer: B

Explanation:

The solution that meets the requirement of the company that runs a database on Amazon Aurora in the us-east-1 Region and has a disaster recovery requirement that the database be available in another Region with minimal disruption to the database operations is to deploy Aurora cross-Region read replicas. Aurora cross-Region read replicas are secondary Aurora clusters that are created in a different AWS Region from the primary Aurora cluster, and are kept in sync with the primary cluster using physical replication. The company can use Aurora cross-Region read replicas to improve the availability and durability of the database, as well as to reduce the recovery time objective (RTO) and recovery point objective (RPO) in case of a regional disaster. Performing an Aurora Multi-AZ deployment, creating Amazon EBS volume snapshots for Aurora and copying them to another Region, and deploying Aurora Replicas are not the best solutions for this requirement. An Aurora Multi-AZ deployment is a configuration that creates one or more Aurora Replicas within the same AWS Region as the primary Aurora cluster, and provides automatic failover in case of an Availability Zone outage. However, this does not provide cross-Region disaster recovery. Creating Amazon EBS volume snapshots for Aurora and copying them to another Region is a manual process that requires stopping the database, creating the snapshots, copying them to the target Region, and restoring them to a new Aurora cluster. This process can cause significant downtime and data loss. Deploying Aurora Replicas is a configuration that creates one or more secondary Aurora clusters within the same AWS Region as the primary Aurora cluster, and provides read scaling and high availability. However, this does not provide cross-Region disaster recovery.

NEW QUESTION 97

- (Topic 2)

Which design principles should a company apply to AWS Cloud workloads to maximize sustainability and minimize environmental impact? (Select TWO.)

- A. Maximize utilization of Amazon EC2 instances.
- B. Minimize utilization of Amazon EC2 instances.
- C. Minimize usage of managed services.
- D. Force frequent application reinstallations by users.
- E. Reduce the need for users to reinstall applications.

Answer: AE

Explanation:

To maximize sustainability and minimize environmental impact, a company should apply the following design principles to AWS Cloud workloads: maximize utilization of Amazon EC2 instances and reduce the need for users to reinstall applications. Maximizing utilization of Amazon EC2 instances means that the company can optimize the performance and efficiency of their compute resources, and avoid wasting energy and money on idle or underutilized instances. The company can use features such as Amazon EC2 Auto Scaling, Amazon EC2 Spot Instances, and AWS Compute Optimizer to automatically adjust the number and

type of instances based on demand, cost, and performance. Reducing the need for users to reinstall applications means that the company can minimize the amount of data and bandwidth required to deliver their applications to users, and avoid unnecessary downloads and updates that consume energy and resources. The company can use services such as Amazon CloudFront, AWS AppStream 2.0, and AWS Amplify to deliver their applications faster, more securely, and more efficiently to users across the globe. Minimizing utilization of Amazon EC2 instances, minimizing usage of managed services, and forcing frequent application reinstallations by users are not design principles that would maximize sustainability and minimize environmental impact. Minimizing utilization of Amazon EC2 instances would reduce the performance and efficiency of the compute resources, and potentially increase the costs and complexity of the cloud workloads. Minimizing usage of managed services would increase the operational overhead and responsibility of the company, and potentially expose them to more security and reliability risks. Forcing frequent application reinstallations by users would increase the amount of data and bandwidth required to deliver the applications to users, and potentially degrade the user experience and satisfaction.

NEW QUESTION 100

- (Topic 2)

Which AWS service is always available free of charge to users?

- A. Amazon Athena
 - B. AWS Identity and Access Management (IAM)
 - C. AWS Secrets Manager
 - D. Amazon ElastiCache
- A company has only basic knowledge of AWS technologies.

Answer: B

Explanation:

AWS Identity and Access Management (IAM) is a web service that helps you securely control access to AWS resources for your users. You use IAM to control who can use your AWS resources (authentication) and what resources they can use and in what ways (authorization). IAM is always available free of charge to users.

NEW QUESTION 102

- (Topic 2)

A company needs help managing multiple AWS linked accounts that are reported on a consolidated bill.

Which AWS Support plan includes an AWS concierge whom the company can ask for assistance?

- A. AWS Developer Support
- B. AWS Enterprise Support
- C. AWS Business Support
- D. AWS Basic Support

Answer: B

Explanation:

AWS Enterprise Support is the AWS Support plan that includes an AWS concierge whom the company can ask for assistance. According to the AWS Support Plans page, AWS Enterprise Support provides "a dedicated Technical Account Manager (TAM) who provides advocacy and guidance to help plan and build solutions using best practices, coordinate access to subject matter experts, and proactively keep your AWS environment operationally healthy."² AWS Business Support, AWS Developer Support, and AWS Basic Support do not include a TAM or a concierge service.

NEW QUESTION 106

- (Topic 2)

A company has multiple AWS accounts that include compute workloads that cannot be interrupted. The company wants to obtain billing discounts that are based on the company's use of AWS services.

Which AWS feature or purchasing option will meet these requirements?

- A. Resource tagging
- B. Consolidated billing
- C. Pay-as-you-go pricing
- D. Spot Instances

Answer: B

Explanation:

Consolidated billing is an AWS feature that allows users to combine the usage and costs of multiple AWS accounts into a single bill. This enables users to obtain billing discounts that are based on the company's use of AWS services, such as volume pricing tiers, Reserved Instance discounts, and Savings Plans discounts.⁵ Resource tagging is an AWS feature that allows users to assign metadata to AWS resources, such as EC2 instances, S3 buckets, and Lambda functions. This enables users to organize, track, and manage their AWS resources, such as filtering, grouping, and reporting. Pay-as-you-go pricing is an AWS pricing model that allows users to pay only for the resources and services they use, without any upfront or long-term commitments. This enables users to lower their costs by scaling up or down as needed, and avoiding over-provisioning or under-utilization. Spot Instances are spare EC2 instances that are available at up to 90% discount compared to On-Demand prices. They are suitable for workloads that can tolerate interruptions, such as batch processing, data analysis, and testing. Spot Instances are allocated based on the current supply and demand, and can be reclaimed by AWS with a two-minute notice when the demand exceeds the supply.

NEW QUESTION 107

- (Topic 2)

A company has a compliance requirement to record and evaluate configuration changes, as well as perform remediation actions on AWS resources.

Which AWS service should the company use?

- A. AWS Config
- B. AWS Secrets Manager
- C. AWS CloudTrail
- D. AWS Trusted Advisor

Answer: A

Explanation:

AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. AWS Config continuously monitors and records your AWS resource configurations and allows you to automate the evaluation of recorded configurations against desired configurations. With AWS Config, you can review changes in configurations and relationships between AWS resources, dive into detailed resource configuration histories, and determine your overall compliance against the configurations specified in your internal guidelines. This can help you simplify compliance auditing, security analysis, change management, and operational troubleshooting¹.

NEW QUESTION 108

- (Topic 2)

A company wants to use Amazon EC2 instances to run a stateless and restartable process after business hours. Which AWS service provides DNS resolution?

- A. Amazon CloudFront
- B. Amazon VPC
- C. Amazon Route 53
- D. AWS Direct Connect

Answer: C

Explanation:

Amazon Route 53 is the AWS service that provides DNS resolution. DNS (Domain Name System) is a service that translates domain names into IP addresses. Amazon Route 53 is a highly available and scalable cloud DNS service that offers domain name registration, DNS routing, and health checking. Amazon Route 53 can route the traffic to various AWS services, such as Amazon EC2, Amazon S3, and Amazon CloudFront. Amazon Route 53 can also integrate with other AWS services, such as AWS Certificate Manager, AWS Shield, and AWS WAF. For more information, see [What is Amazon Route 53?] and [Amazon Route 53 Features].

NEW QUESTION 112

- (Topic 2)

A company does not want to rely on elaborate forecasting to determine its usage of compute resources. Instead, the company wants to pay only for the resources that it uses. The company also needs the ability to increase or decrease its resource usage to meet business requirements. Which pillar of the AWS Well-Architected Framework aligns with these requirements?

- A. Operational excellence
- B. Security
- C. Reliability
- D. Cost optimization

Answer: D

Explanation:

Cost optimization is the pillar of the AWS Well-Architected Framework that aligns with the requirements of not relying on elaborate forecasting and paying only for the resources that are used. The cost optimization pillar focuses on the ability of a system to deliver business value at the lowest price point. Cost optimization involves using the right AWS services and resources for the workload, measuring and monitoring the cost and usage, and continuously improving the cost efficiency. Cost optimization also leverages the benefits of the AWS Cloud, such as pay-as-you-go pricing, elasticity, and scalability. For more information, see [Cost Optimization Pillar] and [Cost Optimization].

NEW QUESTION 113

- (Topic 2)

An application runs on multiple Amazon EC2 instances that access a shared file system simultaneously. Which AWS storage service should be used?

- A. Amazon EBS
- B. Amazon EFS
- C. Amazon S3
- D. AWS Artifact

Answer: B

Explanation:

Amazon Elastic File System (Amazon EFS) is the AWS storage service that should be used for an application that runs on multiple Amazon EC2 instances that access a shared file system simultaneously. Amazon EFS is a fully managed service that provides a scalable, elastic, and highly available file system for Linux-based workloads. Amazon EFS supports the Network File System version 4 (NFSv4) protocol and allows multiple EC2 instances to read and write data to the same file system concurrently. Amazon EFS also integrates with other AWS services, such as AWS Backup, AWS CloudFormation, and AWS CloudTrail. For more information, see [What is Amazon Elastic File System?] and [Amazon EFS Use Cases].

NEW QUESTION 118

- (Topic 2)

Which AWS service or tool provides recommendations to help users get rightsized Amazon EC2 instances based on historical workload usage data?

- A. AWS Pricing Calculator
- B. AWS Compute Optimizer
- C. AWS App Runner
- D. AWS Systems Manager

Answer: B

Explanation:

The AWS service or tool that provides recommendations to help users get rightsized Amazon EC2 instances based on historical workload usage data is AWS Compute Optimizer. AWS Compute Optimizer is a service that analyzes the configuration and performance of the AWS resources, such as Amazon EC2

instances, and provides recommendations for optimal resource types and sizes based on the workload patterns and metrics. AWS Compute Optimizer helps users improve the performance, availability, and cost efficiency of their AWS resources. AWS Pricing Calculator, AWS App Runner, and AWS Systems Manager are not the best services or tools to use for this purpose. AWS Pricing Calculator is a tool that helps users estimate the cost of using AWS services based on their requirements and preferences. AWS App Runner is a service that helps users easily and quickly deploy web applications and APIs without managing any infrastructure. AWS Systems Manager is a service that helps users automate and manage the configuration and operation of their AWS resources and applications³⁴

NEW QUESTION 121

- (Topic 2)

A company needs to host a highly available application in the AWS Cloud. The application runs infrequently for short periods of time. Which AWS service will meet these requirements with the LEAST amount of operational overhead?

- A. Amazon EC2
- B. AWS Fargate
- C. AWS Lambda
- D. Amazon Aurora

Answer: C

Explanation:

The AWS service that will meet the requirements of the company that needs to host a highly available application in the AWS Cloud that runs infrequently for short periods of time with the least amount of operational overhead is AWS Lambda. AWS Lambda is a serverless compute service that allows customers to run code without provisioning or managing servers. The company can use AWS Lambda to create and deploy their application as functions that are triggered by events, such as API calls, messages, or schedules. AWS Lambda automatically scales the compute resources based on the demand, and customers only pay for the compute time they consume. AWS Lambda also simplifies the management and maintenance of the application, as customers do not need to worry about the underlying infrastructure, security, or availability. Amazon EC2, AWS Fargate, and Amazon Aurora are not the best services to use for this purpose. Amazon EC2 is a service that provides scalable compute capacity in the cloud, and allows customers to launch and run virtual servers, called instances, with a variety of operating systems, configurations, and specifications. Amazon EC2 requires customers to provision and manage the instances, and pay for the instance hours they use, regardless of the application usage. AWS Fargate is a serverless compute engine for containers that allows customers to run containerized applications without managing servers or clusters. AWS Fargate requires customers to specify the amount of CPU and memory resources for each container, and pay for the resources they allocate, regardless of the application usage.

Amazon Aurora is a fully managed relational database service that provides high performance, availability, and compatibility. Amazon Aurora is not a compute service, and it is not suitable for hosting an application that runs infrequently for short periods of time¹²

NEW QUESTION 125

- (Topic 2)

What is an AWS responsibility under the AWS shared responsibility model?

- A. Configure the security group rules that determine which ports are open on an Amazon EC2 Linux instance.
 - B. Ensure the security of the internal network in the AWS data centers.
 - C. Patch the guest operating system with the latest security patches on Amazon EC2.
 - D. Turn on server-side encryption for Amazon S3 buckets.
- A company wants to deploy its critical application on AWS and maintain high availability.

Answer: B

Explanation:

Under the AWS shared responsibility model, AWS is responsible for ensuring the security of the internal network in the AWS data centers, as well as the physical security of the hardware and facilities that run AWS services. AWS customers are responsible for configuring the security group rules that determine which ports are open on an EC2 Linux instance, patching the guest operating system with the latest security patches on EC2, and turning on server-side encryption for S3 buckets. Source: AWS Shared Responsibility Model

NEW QUESTION 127

- (Topic 2)

A company needs to launch an Amazon EC2 instance.

Which of the following can the company use during the launch process to configure the root volume of the EC2 instance?

- A. Amazon EC2 Auto Scaling
- B. Amazon Data Lifecycle Manager (Amazon DLM)
- C. Amazon Machine Image (AMI)
- D. Amazon Elastic Block Store (Amazon EBS) volume

Answer: C

Explanation:

Amazon Machine Image (AMI) is the option that the company can use during the launch process to configure the root volume of the EC2 instance. An AMI is a template that contains the software configuration, such as the operating system, applications, and settings, required to launch an EC2 instance. An AMI also specifies the volume size and type of the root device for the instance. The company can choose an AMI provided by AWS, the AWS Marketplace, or the AWS community, or create a custom AMI. For more information, see [Amazon Machine Images (AMI)] and [Launching an Instance Using the Launch Instance Wizard].

NEW QUESTION 128

- (Topic 2)

Which service is an AWS in-memory data store service?

- A. Amazon Aurora
- B. Amazon RDS
- C. Amazon DynamoDB
- D. Amazon ElastiCache

Answer: D

Explanation:

Amazon ElastiCache is a service that offers fully managed in-memory data store and cache services that deliver sub-millisecond response times to applications. You can use Amazon ElastiCache to improve the performance of your applications by retrieving data from fast, managed, in-memory data stores, instead of relying entirely on slower disk-based databases. Amazon Aurora is a relational database service that combines the performance and availability of high-end commercial databases with the simplicity and cost-effectiveness of open source databases. Amazon RDS is a service that makes it easy to set up, operate, and scale a relational database in the cloud. Amazon DynamoDB is a key-value and document database that delivers single-digit millisecond performance at any scale. None of these services are in-memory data store services.

NEW QUESTION 133

- (Topic 2)

A company wants to move its data warehouse application to the AWS Cloud. The company wants to run and scale its analytics services without needing to provision and manage data warehouse clusters. Which AWS service will meet these requirements?

- A. Amazon Redshift provisioned data warehouse
- B. Amazon Redshift Serverless
- C. Amazon Athena
- D. Amazon S3

Answer: B

Explanation:

Amazon Redshift Serverless is the AWS service that will meet the requirements of the company that wants to move its data warehouse application to the AWS Cloud and run and scale its analytics services without needing to provision and manage data warehouse clusters. Amazon Redshift Serverless is a new feature of Amazon Redshift, which is a fully managed data warehouse service that allows customers to run complex queries and analytics on large volumes of structured and semi-structured data. Amazon Redshift Serverless automatically scales the compute and storage resources based on the workload demand, and customers only pay for the resources they consume. Amazon Redshift Serverless also simplifies the management and maintenance of the data warehouse, as customers do not need to worry about choosing the right cluster size, resizing the cluster, or distributing the data across the nodes. Amazon Redshift provisioned data warehouse, Amazon Athena, and Amazon S3 are not the best services to meet the requirements of the company. Amazon Redshift provisioned data warehouse requires customers to choose the number and type of nodes for their cluster, and manually resize the cluster if their workload changes. Amazon Athena is a serverless query service that allows customers to analyze data stored in Amazon S3 using standard SQL, but it is not a data warehouse service that can store and organize the data. Amazon S3 is a scalable object storage service that can store any amount and type of data, but it is not a data warehouse service that can run complex queries and analytics on the data.

NEW QUESTION 137

- (Topic 2)

A developer wants to use an Amazon S3 bucket to store application logs that contain sensitive data. Which AWS service or feature should the developer use to restrict read and write access to the S3 bucket?

- A. Security groups
- B. Amazon CloudWatch
- C. AWS CloudTrail
- D. ACLs

Answer: D

Explanation:

ACLs are an AWS service or feature that the developer can use to restrict read and write access to the S3 bucket. ACLs are access control lists that grant basic permissions to other AWS accounts or predefined groups. They can be used to grant read or write access to an S3 bucket or an object3. Security groups are virtual firewalls that control the inbound and outbound traffic for Amazon EC2 instances. They are not a service or feature that can be used to restrict access to an S3 bucket. Amazon CloudWatch is a service that provides monitoring and observability for AWS resources and applications. It can be used to collect and analyze metrics, logs, events, and alarms. It is not a service or feature that can be used to restrict access to an S3 bucket. AWS CloudTrail is a service that provides governance, compliance, and audit for AWS accounts and resources. It can be used to track and record the API calls and user activity in AWS. It is not a service or feature that can be used to restrict access to an S3 bucket.

NEW QUESTION 138

- (Topic 2)

A company has developed a distributed application that recovers gracefully from interruptions. The application periodically processes large volumes of data by using multiple Amazon EC2 instances. The application is sometimes idle for months. Which EC2 instance purchasing option is MOST cost-effective for this use case?

- A. Reserved Instances
- B. Spot Instances
- C. Dedicated Instances
- D. On-Demand Instances

Answer: B

Explanation:

Spot Instances are instances that use spare EC2 capacity that is available for up to 90% off the On-Demand price. Because Spot Instances can be interrupted by EC2 with two minutes of notification when EC2 needs the capacity back, you can use them for applications that have flexible start and end times, or that can withstand interruptions5. This option is most cost-effective for the use case described in the question. Reserved Instances are instances that you purchase for a one-year or three-year term, and pay a lower hourly rate compared to On-Demand Instances. This option is suitable for applications that have steady state or predictable usage. Dedicated Instances are instances that run on hardware that's dedicated to a single customer within an Amazon VPC. This option is suitable for applications that have stringent regulatory or compliance requirements. On-Demand Instances are instances that you pay for by the second, with no long-term commitments or upfront payments. This option is suitable for applications that have unpredictable or intermittent workloads.

NEW QUESTION 143

- (Topic 2)

Which AWS service or feature can be used to control inbound and outbound traffic on an Amazon EC2 instance?

- A. Internet gateways
- B. AWS Identity and Access Management (IAM)
- C. Network ACLs
- D. Security groups

Answer: D

Explanation:

D is correct because security groups are the AWS service or feature that can be used to control inbound and outbound traffic on an Amazon EC2 instance. Security groups act as a virtual firewall for the EC2 instance, allowing users to specify which protocols, ports, and source or destination IP addresses are allowed or denied. A is incorrect because internet gateways are the AWS service or feature that enable communication between instances in a VPC and the internet. They do not control the traffic on an EC2 instance. B is incorrect because AWS Identity and Access Management (IAM) is the AWS service or feature that enables users to manage access to AWS services and resources securely. It does not control the traffic on an EC2 instance. C is incorrect because network ACLs are the AWS service or feature that provide an optional layer of security for the VPC that acts as a firewall for controlling traffic in and out of one or more subnets. They do not control the traffic on an EC2 instance.

NEW QUESTION 148

- (Topic 2)

Which benefit of AWS Cloud computing provides lower latency between users and applications?

- A. Agility
- B. Economies of scale
- C. Global reach
- D. Pay-as-you-go pricing

Answer: C

Explanation:

Global reach is the benefit of AWS Cloud computing that provides lower latency between users and applications. Global reach means that AWS customers can deploy their applications and data in multiple regions around the world, and deliver them to users with high performance and availability. AWS has the largest global infrastructure of any cloud provider, with 25 geographic regions and 81 Availability Zones, as well as 216 Points of Presence in 84 cities across 42 countries. Customers can choose the optimal locations for their applications and data based on their business requirements, such as compliance, data sovereignty, and customer proximity. Agility, economies of scale, and pay-as-you-go pricing are other benefits of AWS Cloud computing, but they do not directly provide lower latency between users and applications. Agility means that AWS customers can quickly and easily provision and scale up or down AWS resources as needed, without upfront costs or long-term commitments. Economies of scale means that AWS customers can benefit from the lower costs and higher efficiency that AWS achieves by operating at a massive scale and passing the savings to the customers. Pay-as-you-go pricing means that AWS customers only pay for the AWS resources they use, without any upfront costs or long-term contracts.

NEW QUESTION 150

- (Topic 2)

A company is running an order processing system on Amazon EC2 instances. The company wants to migrate microservices-based application. Which combination of AWS services can the application use to meet these requirements? (Select TWO.)

- A. Amazon Simple Queue Service (Amazon SQS)
- B. AWS Lambda
- C. AWS Migration Hub
- D. AWS AppSync
- E. AWS Application Migration Service

Answer: AB

Explanation:

The combination of AWS services that the application can use to migrate to a microservices-based application are Amazon Simple Queue Service (Amazon SQS) and AWS Lambda. Amazon SQS is a fully managed message queuing service that enables customers to decouple and scale microservices, distributed systems, and serverless applications. The application can use Amazon SQS to send, store, and receive messages between the microservices, ensuring that each message is processed only once and in the right order. AWS Lambda is a serverless compute service that allows customers to run code without provisioning or managing servers. The application can use AWS Lambda to create and deploy microservices as functions that are triggered by events, such as messages from Amazon SQS. AWS Migration Hub, AWS AppSync, and AWS Application Migration Service are not the best services to use for migrating to a microservices-based application. AWS Migration Hub is a service that provides a single location to track the progress of application migrations across multiple AWS and partner solutions. AWS AppSync is a service that simplifies the development of GraphQL APIs for real-time and offline data synchronization. AWS Application Migration Service is a service that enables customers to migrate their on-premises applications to AWS without making any changes to the applications, servers, or databases.

NEW QUESTION 151

- (Topic 2)

A company has set up a VPC in its AWS account and has created a subnet in the VPC. The company wants to make the subnet public. Which AWS features should the company use to meet this requirement? (Select TWO.)

- A. Amazon VPC internet gateway
- B. Amazon VPC NAT gateway
- C. Amazon VPC route tables
- D. Amazon VPC network ACL
- E. Amazon EC2 security groups

Answer: AC

Explanation:

To make a subnet public, the company should use an Amazon VPC internet gateway and an Amazon VPC route table. An internet gateway is a horizontally scaled, redundant, and highly available VPC component that allows communication between your VPC and the internet. A route table contains a set of rules, called routes, that are used to determine where network traffic from your subnet or gateway is directed. To enable internet access for a subnet, you need to attach an internet gateway to your VPC and add a route to the internet gateway in the route table associated with the subnet.

NEW QUESTION 155

- (Topic 2)

A company needs to host a web server on Amazon EC2 instances for at least 1 year. The web server cannot tolerate interruption. Which EC2 instance purchasing option will meet these requirements MOST cost-effectively?

- A. On-Demand Instances
- B. Partial Upfront Reserved Instances
- C. Spot Instances
- D. No Upfront Reserved Instances

Answer: B

Explanation:

The most cost-effective EC2 instance purchasing option for the company that needs to host a web server on Amazon EC2 instances for at least 1 year and cannot tolerate interruption is Partial Upfront Reserved Instances. Reserved Instances are a pricing model that offer significant discounts compared to On-Demand Instances in exchange for a commitment to use a specific amount of compute capacity for a fixed period of time (1 or 3 years). Partial Upfront Reserved Instances require customers to pay a portion of the total cost upfront, and the remaining cost in monthly installments over the term. This option offers a lower effective hourly rate than No Upfront Reserved Instances, which require no upfront payment but have higher monthly payments. On-Demand Instances and Spot Instances are not the best options for the company. On-Demand Instances are a pricing model that offer the most flexibility and no long-term commitment, but have the highest hourly rate. Spot Instances are a pricing model that offer the lowest cost, but are subject to interruption based on supply and demand.

NEW QUESTION 158

- (Topic 2)

A manufacturing company has a critical application that runs at a remote site that has a slow internet connection. The company wants to migrate the workload to AWS. The application is sensitive to latency and interruptions in connectivity. The company wants a solution that can host this application with minimum latency. Which AWS service or feature should the company use to meet these requirements?

- A. Availability Zones
- B. AWS Local Zones
- C. AWS Wavelength
- D. AWS Outposts

Answer: D

Explanation:

AWS Outposts is a service that offers fully managed and configurable compute and storage racks built with AWS-designed hardware that allow you to run your workloads on premises and seamlessly connect to AWS services in the cloud. AWS Outposts is ideal for workloads that require low latency, local data processing, or local data storage. With AWS Outposts, you can use the same AWS APIs, tools, and infrastructure across on premises and the cloud to deliver a truly consistent hybrid experience. Availability Zones are isolated locations within each AWS Region that are engineered to be fault-tolerant and provide high availability. AWS Local Zones are extensions of AWS Regions that are placed closer to large population, industry, and IT centers where no AWS Region exists today. AWS Wavelength is a service that enables developers to build applications that deliver ultra-low latency to mobile devices and users by deploying AWS compute and storage at the edge of the 5G network. None of these services or features can help you host a critical application with minimum latency at a remote site that has a slow internet connection.

NEW QUESTION 160

- (Topic 2)

A company provides a software as a service (SaaS) application. The company has a new customer that is based in a different country. The new customer's data needs to be hosted in that country. Which AWS service or infrastructure component should the company use to meet this requirement?

- A. AWS Shield
- B. Amazon S3 Object Lock
- C. AWS Regions
- D. Placement groups

Answer: C

Explanation:

AWS Regions are geographic areas around the world where AWS has clusters of data centers. Each AWS Region consists of multiple, isolated, and physically separate AZ's within a geographic area. By hosting the customer's data in a specific AWS Region, the company can meet the requirement of hosting the data in the customer's country. AWS Shield is a service that provides always-on detection and automatic inline mitigations that minimize application downtime and latency, so there is no need to engage AWS Support to benefit from DDoS protection. Amazon S3 Object Lock is a feature that allows you to store objects using a write-once-read-many (WORM) model. You can use it to prevent an object from being deleted or overwritten for a fixed amount of time or indefinitely. Placement groups are logical grouping of instances within a single Availability Zone. Placement groups enable applications to participate in a low-latency, 10 Gbps network. None of these services or infrastructure components can help the company host the customer's data in a different country.

NEW QUESTION 162

- (Topic 2)

A company is running workloads for multiple departments within a single VPC. The company needs to be able to bill each department for its resource usage. Which action should the company take to accomplish this goal with the LEAST operational overhead?

- A. Add a department tag to each resource and configure cost allocation tags.
- B. Move each department resource to its own VPC.

- C. Move each department resource to its own AWS account.
- D. Use AWS Organizations to get a billing report for each department.

Answer: A

Explanation:

Adding a department tag to each resource and configuring cost allocation tags is an action that can help you accomplish the goal of billing each department for its resource usage with the least operational overhead. Tags are simple labels consisting of a key and an optional value that you can assign to AWS resources. You can use tags to organize your resources and track your AWS costs on a detailed level. Cost allocation tags enable you to track your AWS costs on a detailed level. After you activate cost allocation tags, AWS uses the cost allocation tags to organize your resource costs on your cost allocation report, to make it easier for you to categorize and track your AWS costs. Moving each department resource to its own VPC or its own AWS account is an action that can help you isolate and control the resources for each department, but it would incur more operational overhead than using tags. Using AWS Organizations to get a billing report for each department is an action that can help you consolidate billing and payment across multiple AWS accounts, but it would not help you bill each department for its resource usage within a single VPC.

NEW QUESTION 165

- (Topic 2)

Which perspective of the AWS Cloud Adoption Framework (AWS CAF) connects technology and business?

- A. Operations
- B. People
- C. Security
- D. Governance

Answer: D

Explanation:

The perspective of the AWS Cloud Adoption Framework (AWS CAF) that connects technology and business is governance. The governance perspective focuses on the alignment of the IT strategy and processes with the business strategy and goals, as well as the management of the IT budget, risk, and compliance. The governance perspective capabilities are portfolio management, business performance management, and IT governance. The governance perspective helps organizations ensure that their cloud adoption delivers the expected business value and outcomes, and that their cloud solutions are secure, reliable, and compliant. Operations, people, and security are other perspectives of the AWS CAF, but they do not directly connect technology and business. The operations perspective focuses on the management and monitoring of the cloud resources and applications, as well as the automation and optimization of the operational processes. The people perspective focuses on the development and empowerment of the human resources, as well as the transformation of the organizational culture and structure. The security perspective focuses on the protection of the information assets and systems in the cloud, as well as the implementation of the security policies and controls.

NEW QUESTION 170

- (Topic 2)

Which aspect of security is the customer's responsibility, according to the AWS shared responsibility model?

- A. Patch and configuration management
- B. Service and communications protection or zone security
- C. Physical and environmental controls
- D. Awareness and training

Answer: A

Explanation:

According to the AWS shared responsibility model, AWS is responsible for the security of the cloud, while the customer is responsible for the security in the cloud. This means that AWS provides the physical and environmental controls, the service and communications protection, and the awareness and training for its employees, while the customer provides the patch and configuration management, the identity and access management, the data encryption, and the firewall configuration for its resources.

NEW QUESTION 171

- (Topic 2)

A company plans to migrate its on-premises workload to AWS. Before the migration, the company needs to estimate its future AWS service costs. Which AWS service or tool should the company use to meet this requirement?

- A. AWS Trusted Advisor
- B. AWS Budgets
- C. AWS Pricing Calculator
- D. AWS Cost Explorer

Answer: C

Explanation:

AWS Pricing Calculator is the AWS service or tool that the company should use to estimate its future AWS service costs before the migration. AWS Pricing Calculator is a web-based tool that allows the company to create cost estimates for various AWS services and scenarios. AWS Pricing Calculator helps the company to compare the costs of running the workload on premises versus on AWS, and to optimize the costs by choosing the best options for the workload. AWS Pricing Calculator also provides a detailed breakdown of the cost components and a downloadable report. For more information, see [AWS Pricing Calculator] and [Getting Started with AWS Pricing Calculator].

NEW QUESTION 176

- (Topic 2)

Which statements explain the business value of migration to the AWS Cloud? (Select TWO.)

- A. The migration of enterprise applications to the AWS Cloud makes these applications automatically available on mobile devices.S
- B. AWS availability and security provide the ability to improve service level agreements (SLAs) while reducing risk and unplanned downtime.

- C. Companies that migrate to the AWS Cloud eliminate the need to plan for high availability and disaster recovery.
- D. Companies that migrate to the AWS Cloud reduce IT costs related to infrastructure, freeing budget for reinvestment in other areas.
- E. Applications are modernized because migration to the AWS Cloud requires companies to rearchitect and rewrite all enterprise applications.

Answer: BD

Explanation:

B and D are correct because AWS availability and security enable customers to improve their SLAs while reducing risk and unplanned downtime¹, and AWS reduces IT costs related to infrastructure, allowing customers to reinvest in other areas². A is incorrect because migrating to the AWS Cloud does not automatically make applications available on mobile devices, as it depends on the application design and compatibility. C is incorrect because companies that migrate to the AWS Cloud still need to plan for high availability and disaster recovery, as AWS is a shared responsibility model³. E is incorrect because migrating to the AWS Cloud does not require companies to rearchitect and rewrite all enterprise applications, as AWS offers different migration strategies depending on the application complexity and business objectives⁴.

NEW QUESTION 181

- (Topic 2)

A company wants to migrate its applications to the AWS Cloud. The company plans to identify and prioritize any business transformation opportunities and evaluate its AWS Cloud readiness. Which AWS service or tool should the company use to meet these requirements?

- A. AWS Cloud Adoption Framework (AWS CAF)
- B. AWS Managed Services (AMS)
- C. AWS Well-Architected Framework
- D. AWS Migration Hub

Answer: A

Explanation:

AWS Cloud Adoption Framework (AWS CAF) is a service or tool that helps users migrate their applications to the AWS Cloud. It provides guidance and best practices to identify and prioritize any business transformation opportunities and evaluate their AWS Cloud readiness. It also helps users align their business and technical perspectives, create an actionable roadmap, and measure their progress. AWS Managed Services (AMS) is a service that provides operational services for AWS infrastructure and applications. It helps users reduce their operational overhead and risk, and focus on their core business. It does not help users identify and prioritize any business transformation opportunities and evaluate their AWS Cloud readiness. AWS Well-Architected Framework is a tool that helps users design and implement secure, high-performing, resilient, and efficient solutions on AWS. It provides a set of questions and best practices across five pillars: operational excellence, security, reliability, performance efficiency, and cost optimization. It does not help users identify and prioritize any business transformation opportunities and evaluate their AWS Cloud readiness. AWS Migration Hub is a service that provides a single location to track and manage the migration of applications to AWS. It helps users discover their on-premises servers, group them into applications, and choose the right migration tools. It does not help users identify and prioritize any business transformation opportunities and evaluate their AWS Cloud readiness.

NEW QUESTION 185

- (Topic 2)

Which AWS service or tool offers consolidated billing?

- A. AWS Artifact
 - B. AWS Budgets
 - C. AWS Organizations
 - D. AWS Trusted Advisor
- A company wants to limit its employees' AWS access to a portfolio of predefined AWS resources.

Answer: C

Explanation:

AWS Organizations is a service that enables you to consolidate multiple AWS accounts into an organization that you create and centrally manage. With AWS Organizations, you can create a single payment method for all the AWS accounts in your organization through consolidated billing. Consolidated billing enables you to see a combined view of AWS charges incurred by all accounts in your organization, as well as get a detailed cost report for each individual AWS account associated with your organization. AWS Artifact is a service that provides on-demand access to AWS' security and compliance reports and select online agreements. AWS Budgets is a service that enables you to plan your service usage, service costs, and instance reservations. AWS Trusted Advisor is a service that provides real-time guidance to help you provision your resources following AWS best practices. None of these services or tools offer consolidated billing.

NEW QUESTION 187

- (Topic 2)

An ecommerce company wants to design a highly available application that will be hosted on multiple Amazon EC2 instances. How should the company deploy the EC2 instances to meet these requirements?

- A. Across multiple edge locations
- B. Across multiple VPCs
- C. Across multiple Availability Zones
- D. Across multiple AWS accounts

Answer: C

Explanation:

The company should deploy the EC2 instances across multiple Availability Zones to design a highly available application. Availability Zones are isolated locations within an AWS Region that are engineered to be fault-tolerant and operate independently of each other. By deploying the EC2 instances across multiple Availability Zones, the company can ensure that their application can withstand the failure of an entire Availability Zone and continue to operate with minimal disruption. Deploying the EC2 instances across multiple edge locations, VPCs, or AWS accounts will not provide the same level of availability and fault tolerance as Availability Zones. Edge locations are part of the Amazon CloudFront service, which is a content delivery network (CDN) that caches and serves web content to users. VPCs are virtual networks that isolate the AWS resources within an AWS Region. AWS accounts are the primary units of ownership and access control for AWS resources¹².

NEW QUESTION 188

- (Topic 2)

Which benefit of the AWS Cloud helps companies achieve lower usage costs because of the aggregate usage of all AWS users?

- A. No need to guess capacity
- B. Ability to go global in minutes
- C. Economies of scale
- D. Increased speed and agility

Answer: C

Explanation:

The benefit of the AWS Cloud that helps companies achieve lower usage costs because of the aggregate usage of all AWS users is economies of scale. Economies of scale means that AWS can achieve lower costs and higher efficiency by operating at a massive scale and passing the savings to the customers. AWS leverages the aggregate usage of all AWS users to negotiate better prices with hardware vendors, optimize power consumption, and improve operational processes. As a result, AWS can offer lower and more flexible pricing options to the customers, such as pay-as-you-go, reserved, and spot pricing models. No need to guess capacity, ability to go global in minutes, and increased speed and agility are other benefits of the AWS Cloud, but they are not directly related to the aggregate usage of all AWS users. No need to guess capacity means that AWS customers can avoid the risk of over-provisioning or under-provisioning resources, and scale up or down as needed. Ability to go global in minutes means that AWS customers can deploy their applications and data in multiple regions around the world, and deliver them to users with high performance and availability. Increased speed and agility means that AWS customers can quickly and easily provision and access AWS resources, and accelerate their innovation and time to market.

NEW QUESTION 193

- (Topic 2)

A company is using Amazon RDS.

A company is launching a critical business application in an AWS Region. How can the company increase resilience for this application?

- A. Deploy a copy of the application in another AWS account.
- B. Deploy the application by using multiple VPCs.
- C. Deploy the application by using multiple subnets.
- D. Deploy the application by using multiple Availability Zones.

Answer: D

Explanation:

Deploying the application by using multiple Availability Zones is the best way to increase resilience for the application. According to the Amazon RDS User Guide, "Amazon RDS provides high availability and failover support for DB instances using Multi-AZ deployments. In a Multi-AZ deployment, Amazon RDS automatically provisions and maintains a synchronous standby replica in a different Availability Zone. The primary DB instance is synchronously replicated across Availability Zones to a standby replica to provide data redundancy, eliminate I/O freezes, and minimize latency spikes during system backups."⁴ Deploying a copy of the application in another AWS account, using multiple VPCs, or using multiple subnets do not provide the same level of resilience as using multiple Availability Zones.

NEW QUESTION 195

- (Topic 2)

A company is building an application that will receive millions of database queries each second. The company needs the data store for the application to scale to meet these needs.

Which AWS service will meet this requirement?

- A. Amazon DynamoDB
- B. AWS Cloud9
- C. Amazon ElastiCache for Memcached
- D. Amazon Neptune

Answer: A

Explanation:

Amazon DynamoDB is the AWS service that will meet the requirement of building an application that will receive millions of database queries each second. Amazon DynamoDB is a fully managed NoSQL database service that provides fast and consistent performance, scalability, and durability. Amazon DynamoDB can handle any level of request traffic and automatically scale up or down the capacity based on the demand. Amazon DynamoDB also supports in-memory caching with Amazon DynamoDB Accelerator (DAX) to improve the response time and reduce the cost. For more information, see [What is Amazon DynamoDB?](#) and [Amazon DynamoDB Features](#).

NEW QUESTION 196

- (Topic 2)

Which AWS service is always free of charge for users?

- A. Amazon S3
- B. Amazon Aurora
- C. Amazon EC2
- D. AWS Identity and Access Management (IAM)

Answer: D

Explanation:

AWS Identity and Access Management (IAM) is a service that allows users to manage access to AWS resources and services. It enables users to create and manage users, groups, roles, and policies that control who can do what in AWS. IAM is always free of charge for users, as there is no additional cost for using IAM with any AWS service¹. Amazon S3 is a storage service that provides scalable, durable, and secure object storage. Amazon S3 has a free tier that offers 5 GB of storage, 20,000 GET requests, and 2,000 PUT requests per month for one year. However, users are charged for any additional usage beyond the free tier limits². Amazon Aurora is a relational database service that is compatible with MySQL and PostgreSQL. Amazon Aurora has a free tier that offers 750 hours of Aurora Single-AZ db.t2.small database usage and 20 GB of storage per month for one year. However, users are charged for any additional usage beyond the free tier limits³. Amazon EC2 is a compute service that provides resizable virtual servers. Amazon EC2 has a free tier that offers 750 hours of Linux and Windows t2.micro instances per month for one year. However, users are charged for any additional usage beyond the free tier limits⁴.

NEW QUESTION 201

- (Topic 1)

Which options does AWS make available for customers who want to learn about security in the cloud in an instructor-led setting? (Select TWO.)

- A. AWS Trusted Advisor
- B. AWS Online Tech Talks
- C. AWS Blog
- D. AWS Forums
- E. AWS Classroom Training

Answer: BE

Explanation:

The correct answers are B and E because AWS Online Tech Talks and AWS Classroom Training are options that AWS makes available for customers who want to learn about security in the cloud in an instructor-led setting. AWS Online Tech Talks are live, online presentations that cover a broad range of topics at varying technical levels. AWS Online Tech Talks are delivered by AWS experts and feature live Q&A sessions with the audience. AWS Classroom Training are in-person or virtual courses that are led by accredited AWS instructors. AWS Classroom Training offer hands-on labs, exercises, and best practices to help customers gain confidence and skills on AWS. The other options are incorrect because they are not options that AWS makes available for customers who want to learn about security in the cloud in an instructor-led setting. AWS Trusted Advisor is an AWS service that provides real-time guidance to help customers follow AWS best practices for security, performance, cost optimization, and fault tolerance. AWS Blog is an AWS resource that provides news, announcements, and insights from AWS experts and customers. AWS Forums are AWS resources that enable customers to interact with other AWS users and get feedback and support. Reference: AWS Online Tech Talks, AWS Classroom Training

NEW QUESTION 202

- (Topic 1)

How can an AWS user conduct security assessments of Amazon EC2 instances, NAT gateways, and Elastic Load Balancers in a way that is approved by AWS?

- A. Flood a target with requests.
- B. Use Amazon Inspector.
- C. Perform penetration testing.
- D. Use the AWS Service Health Dashboard.

Answer: B

Explanation:

Amazon Inspector is an automated security assessment service that helps improve the security and compliance of applications deployed on AWS. Amazon Inspector automatically assesses applications for exposure, vulnerabilities, and deviations from best practices. After performing an assessment, Amazon Inspector produces a detailed list of security findings prioritized by level of severity.

NEW QUESTION 204

- (Topic 1)

A company wants to migrate its on-premises data warehouse to AWS. The information in the data warehouse is used to populate analytics dashboards. Which AWS service should the company use for the data warehouse?

- A. Amazon ElastiCache
- B. Amazon Aurora
- C. Amazon RDS
- D. Amazon Redshift

Answer: D

Explanation:

The AWS service that the company should use for the data warehouse is Amazon Redshift. Amazon Redshift is a fully managed, petabyte-scale data warehouse service that is optimized for analytical queries. It can integrate with various data sources and business intelligence tools to provide fast and cost-effective insights. Amazon Redshift also offers high availability, scalability, security, and compliance features. [Amazon Redshift Overview]

NEW QUESTION 209

- (Topic 1)

Which AWS service gives users the ability to provision a dedicated and private network connection from their internal network to AWS?

- A. AWS CloudHSM
- B. AWS Direct Connect
- C. AWS VPN
- D. Amazon Connect

Answer: B

Explanation:

AWS Direct Connect gives users the ability to provision a dedicated and private network connection from their internal network to AWS. AWS Direct Connect links the user's internal network to an AWS Direct Connect location over a standard Ethernet fiber-optic cable. One end of the cable is connected to the user's router, the other to an AWS Direct Connect router. With this connection in place, the user can create virtual interfaces directly to the AWS cloud and Amazon Virtual Private Cloud (Amazon VPC), bypassing internet service providers in the network path.

NEW QUESTION 210

- (Topic 1)

Which AWS service provides highly durable object storage?

- A. Amazon S3

- B. Amazon Elastic File System (Amazon EFS)
- C. Amazon Elastic Block Store (Amazon EBS)
- D. Amazon FSx

Answer: A

Explanation:

Amazon S3 is the AWS service that provides highly durable object storage. Amazon S3 is designed to provide 99.999999999% durability of objects over a given year.

This means that you can store your data with high confidence that it will not be lost. Amazon S3 also provides high availability, scalability, security, and performance for your data. You can use Amazon S3 to store and retrieve any amount of data, at any time, from anywhere on the web.

NEW QUESTION 211

- (Topic 1)

A company needs to continuously monitor its environment to analyze network and account activity and identify potential security threats. Which AWS service should the company use to meet these requirements?

- A. AWS Artifact
- B. Amazon Macie
- C. AWS Identity and Access Management (IAM)
- D. Amazon GuardDuty

Answer: D

Explanation:

Amazon GuardDuty is a service that provides intelligent threat detection and continuous monitoring for the AWS environment. It analyzes network and account activity using machine learning and threat intelligence to identify potential security threats, such as unauthorized access, compromised credentials, malicious hosts, and reconnaissance activities. It also generates detailed and actionable findings that can be viewed on the AWS Management Console or sent to other AWS services, such as Amazon CloudWatch Events and AWS Lambda, for further analysis or remediation. Amazon GuardDuty Overview AWS Certified Cloud Practitioner - aws.amazon.com

NEW QUESTION 216

- (Topic 1)

An application is running on multiple Amazon EC2 instances. The company wants to make the application highly available by configuring a load balancer with requests forwarded to the EC2 instances based on URL paths.

Which AWS load balancer will meet these requirements and take the LEAST amount of effort to deploy?

- A. Network Load Balancer
- B. Application Load Balancer
- C. AWS OpsWorks Load Balancer
- D. Custom Load Balancer on Amazon EC2

Answer: B

Explanation:

The correct answer is B because Application Load Balancer is an AWS load balancer that will meet the requirements and take the least amount of effort to deploy. Application Load Balancer is a type of Elastic Load Balancing that operates at the application layer (layer 7) of the OSI model and routes requests to targets based on the content of the request. Application Load Balancer supports advanced features, such as path-based routing, host-based routing, and HTTP header-based routing. The other options are incorrect because they are not AWS load balancers that will meet the requirements and take the least amount of effort to deploy. Network Load Balancer is a type of Elastic Load Balancing that operates at the transport layer (layer 4) of the OSI model and routes requests to targets based on the destination IP address and port. Network Load Balancer does not support path-based routing. AWS OpsWorks Load Balancer is not an AWS load balancer, but rather a feature of AWS OpsWorks that enables users to attach an Elastic Load Balancing load balancer to a layer of their stack. Custom Load Balancer on Amazon EC2 is not an AWS load balancer, but rather a user-defined load balancer that runs on an Amazon EC2 instance. Custom Load Balancer on Amazon EC2 requires more effort to deploy and maintain than an AWS load balancer. Reference: Elastic Load Balancing

NEW QUESTION 218

- (Topic 1)

A company needs to test a new application that was written in Python. The code will activate when new images are stored in an Amazon S3 bucket. The application will put a watermark on each image and then will store the images in a different S3 bucket.

Which AWS service should the company use to conduct the test with the LEAST amount of operational overhead?

- A. Amazon EC2
- B. AWS CodeDeploy
- C. AWS Lambda
- D. Amazon Lightsail

Answer: C

Explanation:

AWS Lambda is a compute service that lets you run code without provisioning or managing servers. AWS Lambda executes your code only when needed and scales automatically, from a few requests per day to thousands per second. You pay only for the compute time you consume - there is no charge when your code is not running. With AWS Lambda, you can run code for virtually any type of application or backend service - all with zero administration. AWS Lambda runs your code on a high-availability compute infrastructure and performs all of the administration of the compute resources, including server and operating system maintenance, capacity provisioning and automatic scaling, code monitoring and logging

NEW QUESTION 222

- (Topic 1)

A company is developing an application that uses multiple AWS services. The application needs to use

temporary, limited-privilege credentials for authentication with other AWS APIs. Which AWS service or feature should the company use to meet these authentication requirements?

- A. Amazon API Gateway
- B. IAM users
- C. AWS Security Token Service (AWS STS)
- D. IAM instance profiles

Answer: C

Explanation:

AWS Security Token Service (AWS STS) is a service that enables applications to request temporary, limited-privilege credentials for authentication with other AWS APIs. AWS STS can be used to grant access to AWS resources to users who are federated (using IAM roles), switched (using IAM users), or cross-account (using IAM roles). AWS STS can also be used to assume a role within the same account or a different account. The credentials issued by AWS STS are short-term and have a limited scope, which can enhance the security and compliance of the application. AWS STS Overview AWS Certified Cloud Practitioner - aws.amazon.com

NEW QUESTION 227

- (Topic 1)

Which AWS solution gives companies the ability to use protocols such as NFS to store and retrieve objects in Amazon S3?

- A. Amazon FSx for Lustre
- B. AWS Storage Gateway volume gateway
- C. AWS Storage Gateway file gateway
- D. Amazon Elastic File System (Amazon EFS)

Answer: C

Explanation:

AWS Storage Gateway file gateway allows companies to use protocols such as NFS and SMB to store and retrieve objects in Amazon S3. File gateway provides a seamless integration between on-premises applications and Amazon S3, and enables low-latency access to data through local caching. File gateway also supports encryption, compression, and lifecycle management of the objects in Amazon S3. For more information, see What is AWS Storage Gateway? and File Gateway.

NEW QUESTION 229

- (Topic 1)

What does the Amazon S3 Intelligent-Tiering storage class offer?

- A. Payment flexibility by reserving storage capacity
- B. Long-term retention of data by copying the data to an encrypted Amazon Elastic Block Store (AmazonEBS) volume
- C. Automatic cost savings by moving objects between tiers based on access pattern changes
- D. Secure, durable, and lowest cost storage for data archival

Answer: C

Explanation:

The Amazon S3 Intelligent-Tiering storage class offers automatic cost savings by moving objects between tiers based on access pattern changes. This storage class is designed for data with unknown or changing access patterns. It has two access tiers: frequent access and infrequent access. Objects are stored in the frequent access tier by default, and are moved to the infrequent access tier after 30 consecutive days of no access. If an object in the infrequent access tier is accessed, it is moved back to the frequent access tier. There are no retrieval fees in S3 Intelligent-Tiering, and no additional tiering fees when objects are moved between access tiers within the S3 Intelligent-Tiering storage class1.

NEW QUESTION 234

- (Topic 1)

A company moves its infrastructure from on premises to the AWS Cloud. The company can now provision additional Amazon EC2 instances whenever the instances are required. With this ability, the company can launch new marketing campaigns in 3 days instead of 3 weeks. Which benefit of the AWS Cloud does this scenario demonstrate?

- A. Cost savings
- B. Improved operational resilience
- C. Increased business agility
- D. Enhanced security

Answer: C

Explanation:

Increased business agility is the benefit of the AWS Cloud that this scenario demonstrates. Business agility refers to the ability of a company to adapt to changing customer needs, market conditions, and competitive pressures. Moving to the AWS Cloud enables business agility by providing faster access to resources, lower upfront costs, and greater scalability and flexibility. By using the AWS Cloud, the company can launch new marketing campaigns in 3 days instead of 3 weeks, which shows that it can respond to customer feedback more quickly and efficiently. For more information, see Benefits of Cloud Computing and [Business Agility].

NEW QUESTION 239

- (Topic 1)

Which tasks are customer responsibilities according to the AWS shared responsibility model? (Select TWO.)

- A. Determine application dependencies with operating systems.
- B. Provide user access with AWS Identity and Access Management (IAM).
- C. Secure the data center in an Availability Zone.
- D. Patch the hypervisor.

E. Provide network availability in Availability Zones.

Answer: B

Explanation:

The correct answer to the question is B because providing user access with AWS Identity and Access Management (IAM) is a customer responsibility according to the AWS shared responsibility model. The AWS shared responsibility model is a framework that defines the division of responsibilities between AWS and the customer for security and compliance. AWS is responsible for the security of the cloud, which includes the global infrastructure, such as the regions, availability zones, and edge locations; the hardware, software, networking, and facilities that run the AWS services; and the virtualization layer that separates the customer instances and storage. The customer is responsible for the security in the cloud, which includes the customer data, the guest operating systems, the applications, the identity and access management, the firewall configuration, and the encryption. IAM is an AWS service that enables customers to manage access and permissions to AWS resources and services. Customers are responsible for creating and managing IAM users, groups, roles, and policies, and ensuring that they follow the principle of least privilege. Reference: AWS Shared Responsibility Model

NEW QUESTION 244

- (Topic 1)

Which tasks are the responsibility of AWS, according to the AWS shared responsibility model? (Select TWO.)

- A. Patch AWS network devices.
- B. Set user password rules.
- C. Provide physical security for compute resources.
- D. Configure security groups.
- E. Patch the operating system of an Amazon EC2 instance.

Answer: AC

Explanation:

The correct answers are A and C because patching AWS network devices and providing physical security for compute resources are tasks that are the responsibility of AWS, according to the AWS shared responsibility model. The AWS shared responsibility model is a framework that defines the division of responsibilities between AWS and the customer for security and compliance. AWS is responsible for the security of the cloud, which includes the global infrastructure, such as the regions, availability zones, and edge locations; the hardware, software, networking, and facilities that run the AWS services; and the virtualization layer that separates the customer instances and storage. The customer is responsible for the security in the cloud, which includes the customer data, the guest operating systems, the applications, the identity and access management, the firewall configuration, and the encryption. The other options are incorrect because they are tasks that are the responsibility of the customer, according to the AWS shared responsibility model. Setting user password rules, configuring security groups, and patching the operating system of an Amazon EC2 instance are all tasks that the customer has to perform to secure their AWS environment. Reference: AWS Shared Responsibility Model

NEW QUESTION 248

- (Topic 1)

Which AWS service will help a company identify the user who deleted an Amazon EC2 instance yesterday?

- A. Amazon CloudWatch
- B. AWS Trusted Advisor
- C. AWS CloudTrail
- D. Amazon Inspector

Answer: C

Explanation:

The correct answer is C because AWS CloudTrail is a service that will help a company identify the user who deleted an Amazon EC2 instance yesterday. AWS CloudTrail is a service that enables users to track user activity and API usage across their AWS account. AWS CloudTrail records the details of every API call made to AWS services, such as the identity of the caller, the time of the call, the source IP address of the caller, the parameters and responses of the call, and more. Users can use AWS CloudTrail to audit, monitor, and troubleshoot their AWS resources and actions. The other options are incorrect because they are not services that will help a company identify the user who deleted an Amazon EC2 instance yesterday. Amazon CloudWatch is a service that enables users to collect, analyze, and visualize metrics, logs, and events from their AWS resources and applications. AWS Trusted Advisor is a service that provides real-time guidance to help users follow AWS best practices for security, performance, cost optimization, and fault tolerance. Amazon Inspector is a service that helps users find security vulnerabilities and deviations from best practices in their Amazon EC2 instances. Reference: AWS CloudTrail FAQs

NEW QUESTION 249

- (Topic 1)

A developer needs to build an application for a retail company. The application must provide real-time product recommendations that are based on machine learning.

Which AWS service should the developer use to meet this requirement?

- A. AWS Health Dashboard
- B. Amazon Personalize
- C. Amazon Forecast
- D. Amazon Transcribe

Answer: B

Explanation:

Amazon Personalize is a fully managed machine learning service that customers can use to generate personalized recommendations for their users. It can also generate user segments based on the users' affinity for certain items or item metadata. Amazon Personalize uses the customers' data to train and deploy custom recommendation models that can be integrated into their applications. Therefore, the correct answer is B. You can learn more about Amazon Personalize and its use cases from this page.

NEW QUESTION 253

- (Topic 3)

A company wants to build a new web application by using AWS services. The application must meet the on-demand load for periods of heavy activity. Which AWS services or resources provide the necessary workload adjustments to meet these requirements? (Select TWO.)

- A. Amazon Machine Image (AMI)
- B. Amazon EC2 Auto Scaling
- C. Amazon EC2 instance
- D. AWS Lambda
- E. EC2 Image Builder

Answer: BD

Explanation:

Amazon EC2 Auto Scaling helps you ensure that you have the correct number of Amazon EC2 instances available to handle the load for your application. You create collections of EC2 instances, called Auto Scaling groups. You can specify the minimum number of instances in each Auto Scaling group, and Amazon EC2 Auto Scaling ensures that your group never goes below this size. You can specify the maximum number of instances in each Auto Scaling group, and Amazon EC2 Auto Scaling ensures that your group never goes above this size⁴. AWS Lambda lets you run code without provisioning or managing servers. You pay only for the compute time you consume. With Lambda, you can run code for virtually any type of application or backend service - all with zero administration. Just upload your code and Lambda takes care of everything required to run and scale your code with high availability. You can set up your code to automatically trigger from other AWS services or call it directly from any web or mobile app.

NEW QUESTION 255

- (Topic 3)

A company is migrating its workloads to the AWS Cloud. The company must retain full control of patch management for the guest operating systems that host its applications.

Which AWS service should the company use to meet these requirements?

- A. Amazon DynamoDB
- B. Amazon EC2
- C. AWS Lambda
- D. Amazon RDS

Answer: B

Explanation:

Amazon EC2 is the AWS service that the company should use to meet its requirements of retaining full control of patch management for the guest operating systems that host its applications. Amazon EC2 is a service that provides secure, resizable compute capacity in the cloud. Users can launch virtual servers, called instances, that run various operating systems, such as Linux, Windows, macOS, and more. Users have full administrative access to their instances and can install and configure any software, including patches and updates, on their instances. Users are responsible for managing the security and maintenance of their instances, including patching the guest operating system and applications. Users can also use AWS Systems Manager to automate and simplify the patching process for their EC2 instances. AWS Systems Manager is a service that helps users manage their AWS and on-premises resources at scale. Users can use AWS Systems Manager Patch Manager to scan their instances for missing patches, define patch baselines and maintenance windows, and apply patches automatically or manually across their instances. Users can also use AWS Systems Manager to monitor the patch compliance status and patching history of their instances. References: What is Amazon EC2?, AWS Systems Manager Patch Manager

NEW QUESTION 257

- (Topic 3)

A company wants a key-value NoSQL database that is fully managed and serverless. Which AWS service will meet these requirements?

- A. Amazon DynamoDB
- B. Amazon RDS
- C. Amazon Aurora
- D. Amazon Memory DB for Redis

Answer: A

Explanation:

Amazon DynamoDB is a key-value and document database that delivers single-digit millisecond performance at any scale. It is a fully managed, serverless database that does not require provisioning, patching, or backup. It offers built-in security, backup and restore, and in-memory caching³. Amazon RDS is a relational database service that makes it easy to set up, operate, and scale a relational database in the cloud. It provides cost-efficient and resizable capacity while automating time-consuming administration tasks such as hardware provisioning, database setup, patching, and backups. However, it is not a key-value NoSQL database, and it is not serverless, as it requires you to choose an instance type and size⁴. Amazon Aurora is a MySQL and PostgreSQL-compatible relational database built for the cloud, that combines the performance and availability of traditional enterprise databases with the simplicity and cost-effectiveness of open source databases. However, it is also not a key-value NoSQL database, and it is not serverless, as it requires you to choose an instance type and size. Amazon MemoryDB for Redis is a Redis-compatible, durable, in-memory database service that delivers ultra-fast performance and multi-AZ reliability for the most demanding applications. However, it is also not a key-value NoSQL database, and it is not serverless, as it requires you to choose a node type and size.

NEW QUESTION 261

- (Topic 3)

Which characteristic of the AWS Cloud helps users eliminate underutilized CPU capacity'?

- A. Agility
- B. Elasticity
- C. Reliability
- D. Durability

Answer: B

Explanation:

Elasticity is a characteristic of the AWS Cloud that helps users eliminate underutilized CPU capacity. Elasticity refers to the ability to dynamically provision and de-provision computing resources as per demand, ensuring that the application or service always has the required resources to operate efficiently. Elasticity helps

users optimize performance and costs, as they only pay for the resources they use and avoid wasting resources when the demand is low³⁴⁵. References: 3: Which characteristic of the AWS cloud helps users eliminate ... , 4: AWS Elastic Load Balancing and Application Load Balancer, 5: Which characteristic of the AWS Cloud helps users eliminate ...

NEW QUESTION 263

- (Topic 3)

What can a cloud practitioner use to retrieve AWS security and compliance documents and submit them as evidence to an auditor or regulator?

- A. AWS Certificate Manager
- B. AWS Systems Manager
- C. AWS Artifact
- D. Amazon Inspector

Answer: C

Explanation:

AWS Artifact is a service that provides on-demand access to AWS security and compliance documents, such as AWS ISO certifications, Payment Card Industry (PCI) reports, and Service Organization Control (SOC) reports. You can download these documents and submit them as evidence to your auditors or regulators to demonstrate the security and compliance of the AWS infrastructure and services that you use. AWS Artifact also allows you to review, accept, and manage AWS agreements, such as the Business Associate Addendum (BAA) for customers who are subject to the Health Insurance Portability and Accountability Act (HIPAA). References: AWS Artifact, What is AWS Artifact?

NEW QUESTION 268

- (Topic 3)

Which AWS service or feature can a company use to apply security rules to specific Amazon EC2 instances?

- A. Network ACLs
- B. Security groups
- C. AWS Trusted Advisor
- D. AWS WAF

Answer: B

Explanation:

Security groups are the AWS service or feature that can be used to apply security rules to specific Amazon EC2 instances. Security groups are virtual firewalls that control the inbound and outbound traffic for one or more instances. Customers can create security groups and add rules that reflect the role of the instance that is associated with the security group. For example, a web server instance needs security group rules that allow inbound HTTP and HTTPS access, while a database instance needs rules that allow access for the type of database¹². Security groups are stateful, meaning that the responses to allowed inbound traffic are also allowed, regardless of the outbound rules¹. Customers can assign multiple security groups to an instance, and the rules from each security group are effectively aggregated to create one set of rules¹.

Network ACLs are another AWS service or feature that can be used to control the traffic for a subnet. Network ACLs are stateless, meaning that they do not track the traffic that they allow. Therefore, customers must add rules for both inbound and outbound traffic³. Network ACLs are applied at the subnet level, not at the instance level.

AWS Trusted Advisor is an AWS service that provides best practice recommendations for security, performance, cost optimization, and fault tolerance. AWS Trusted Advisor does not apply security rules to specific Amazon EC2 instances, but it can help customers identify security gaps and improve their security posture⁴.

AWS WAF is an AWS service that helps protect web applications from common web exploits, such as SQL injection, cross-site scripting, and bot attacks. AWS WAF does not apply security rules to specific Amazon EC2 instances, but it can be integrated with other AWS services, such as Amazon CloudFront, Amazon API Gateway, and Application Load Balancer.

NEW QUESTION 271

- (Topic 3)

A company needs to block SQL injection attacks.

Which AWS service or feature can meet this requirement?

- A. AWS WAF
- B. AWS Shield
- C. Network ACLs
- D. Security groups

Answer: A

Explanation:

AWS WAF is a web application firewall that helps protect web applications from common web exploits, such as SQL injection attacks. It allows customers to create custom rules that block malicious requests. AWS Shield is a managed service that protects against distributed denial of service (DDoS) attacks, not SQL injection attacks. Network ACLs and security groups are network-level security features that filter traffic based on IP addresses and ports, not web requests or SQL queries. References: [AWS WAF], [AWS Shield], [Network ACLs], [Security groups]

NEW QUESTION 273

- (Topic 3)

Which maintenance task is the customer's responsibility, according to the AWS shared responsibility model?

- A. Physical connectivity among Availability Zones
- B. Network switch maintenance
- C. Hardware updates and firmware patches
- D. Amazon EC2 updates and security patches

Answer: D

Explanation:

According to the AWS shared responsibility model, customers are responsible for managing their data, applications, operating systems, security groups, and other aspects of their AWS environment. This includes installing updates and security patches of the guest operating system and any application software or utilities installed by the customer on the instances. AWS is responsible for protecting the infrastructure that runs all of the services offered in the AWS Cloud, such as data centers, hardware, software, networking, and facilities. This includes the physical connectivity among Availability Zones, the network switch maintenance, and the hardware updates and firmware patches. Therefore, option D is the correct answer, and options A, B, and C are AWS responsibilities, not customer responsibilities. References: : AWS Well-Architected Framework - Elasticity; : Reactive Systems on AWS - Elastic

NEW QUESTION 274

- (Topic 3)

Which scenarios represent the concept of elasticity on AWS? (Select TWO.)

- A. Scaling the number of Amazon EC2 instances based on traffic
- B. Resizing Amazon RDS instances as business needs change
- C. Automatically directing traffic to less-utilized Amazon EC2 instances
- D. Using AWS compliance documents to accelerate the compliance process
- E. Having the ability to create and govern environments using code

Answer: AB

Explanation:

These are two scenarios that represent the concept of elasticity on AWS. Elasticity means the ability to adjust the resources and capacity of the system in response to changes in demand or environment. Scaling the number of Amazon EC2 instances based on traffic means using services such as AWS Auto Scaling or Elastic Load Balancing to add or remove instances as the traffic increases or decreases. Resizing Amazon RDS instances as business needs change means using the Amazon RDS console or API to modify the instance type, storage type, or storage size of the database as the workload grows or shrinks. You can learn more about the concept of elasticity on AWS from [this webpage] or [this digital course].

NEW QUESTION 279

- (Topic 3)

A company is expecting a short-term spike in internet traffic for its application. During the traffic increase, the application cannot be interrupted. The company also needs to minimize cost and maximize flexibility.

A company needs to use a serverless interactive query service to analyze data in Amazon S3. The query service must support standard SQL.

Which AWS service will meet these requirements?

- A. Amazon Redshift
- B. AWS Glue
- C. Amazon Athena
- D. Amazon Kinesis Data Streams

Answer: C

Explanation:

Amazon Athena is a serverless interactive query service that makes it easy to analyze data in Amazon S3 using standard SQL. Athena is ideal for quick, ad-hoc querying but it can also handle complex analysis, including large joins, window functions, and arrays. Athena scales automatically—executing queries in parallel—so results are fast, even with large datasets and complex queries. Amazon Redshift is a fully managed, petabyte-scale data warehouse service that can run complex analytic queries against structured and semi-structured data using standard SQL. However, it is not a serverless service and requires provisioning and managing clusters of nodes. AWS Glue is a fully managed extract, transform, and load (ETL) service that makes it easy to prepare and load your data for analytics. However, it is not a query service and does not support standard SQL. Amazon Kinesis Data Streams is a service that enables you to build custom applications that process or analyze streaming data for specialized needs. However, it is not a query service and does not support standard SQL.

NEW QUESTION 280

- (Topic 3)

Which AWS service offers object storage?

- A. Amazon RDS
- B. Amazon Elastic File System (Amazon EFS)
- C. Amazon S3
- D. Amazon DynamoDB

Answer: C

Explanation:

Amazon S3 is the AWS service that offers object storage. Object storage is a technology that stores and manages data in an unstructured format called objects. Each object consists of the data, metadata, and a unique identifier. Object storage is ideal for storing large amounts of unstructured data, such as photos, videos, email, web pages, sensor data, and audio files¹. Amazon S3 provides industry-leading scalability, data availability, security, and performance for object storage². Amazon RDS is the AWS service that offers relational database storage. Relational database storage is a technology that stores and manages data in a structured format called tables. Each table consists of rows and columns that define the attributes and values of the data. Relational database storage is ideal for storing structured or semi-structured data, such as customer records, inventory, transactions, and analytics³.

Amazon Elastic File System (Amazon EFS) is the AWS service that offers file storage. File storage is a technology that stores and manages data in a hierarchical format called files and folders. Each file consists of the data and metadata, and each folder consists of files or subfolders. File storage is ideal for storing shared data that can be accessed by multiple users or applications, such as home directories, content repositories, media libraries, and configuration files⁴.

Amazon DynamoDB is the AWS service that offers NoSQL database storage. NoSQL database storage is a technology that stores and manages data in a flexible format called documents or key-value pairs. Each document or key-value pair consists of the data and metadata, and can have different attributes and values depending on the schema. NoSQL database storage is ideal for storing dynamic or unstructured data that requires high performance, scalability, and availability, such as web applications, social media, gaming, and IoT.

NEW QUESTION 285

- (Topic 3)

A company must archive Amazon S3 data that the company's business units no longer need to access.

Which S3 storage class will meet this requirement MOST cost-effectively?

- A. S3 Glacier Instant Retrieval
- B. S3 Glacier Flexible Retrieval
- C. S3 Glacier Deep Archive
- D. S3 One Zone-Infrequent Access (S3 One Zone-IA)

Answer: C

Explanation:

S3 Glacier Deep Archive is Amazon S3's lowest-cost storage class and supports long-term retention and digital preservation for data that may be accessed once or twice in a year. It is designed for customers — particularly those in highly-regulated industries, such as the Financial Services, Healthcare, and Public Sectors — that retain data sets for 7-10 years or longer to meet regulatory compliance requirements. Customers can store large amounts of data at a very low cost, and reliably access it with a wait time of 12 hours.

NEW QUESTION 286

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

AWS-Certified-Cloud-Practitioner Practice Exam Features:

- * AWS-Certified-Cloud-Practitioner Questions and Answers Updated Frequently
- * AWS-Certified-Cloud-Practitioner Practice Questions Verified by Expert Senior Certified Staff
- * AWS-Certified-Cloud-Practitioner Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * AWS-Certified-Cloud-Practitioner Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The AWS-Certified-Cloud-Practitioner Practice Test Here](#)