

AWS-Certified-Cloud-Practitioner Dumps

Amazon AWS Certified Cloud Practitioner

<https://www.certleader.com/AWS-Certified-Cloud-Practitioner-dumps.html>



NEW QUESTION 1

- (Topic 2)

A company wants to create a chatbot and integrate the chatbot with its current web application. Which AWS service will meet these requirements?

- A. AmazonKendra
- B. Amazon Lex
- C. AmazonTextract
- D. AmazonPolly

Answer: B

Explanation:

The AWS service that will meet the requirements of the company that wants to create a chatbot and integrate the chatbot with its current web application is Amazon Lex. Amazon Lex is a service that helps customers build conversational interfaces using voice and text. The company can use Amazon Lex to create a chatbot that can understand natural language and respond to user requests, using the same deep learning technologies that power Amazon Alexa. Amazon Lex also provides easy integration with other AWS services, such as Amazon Comprehend, Amazon Polly, and AWS Lambda, as well as popular platforms, such as Facebook Messenger, Slack, and Twilio. Amazon Lex helps customers create engaging and interactive chatbots for their web applications. Amazon Kendra, Amazon Textract, and Amazon Polly are not the best services to use for this purpose. Amazon Kendra is a service that helps customers provide accurate and natural answers to natural language queries using machine learning. Amazon Textract is a service that helps customers extract text and data from scanned documents using optical character recognition (OCR) and machine learning. Amazon Polly is a service that helps customers convert text into lifelike speech using deep learning. These services are more useful for different types of natural language processing and generation tasks, rather than creating and integrating chatbots.

NEW QUESTION 2

- (Topic 2)

A company wants to migrate its on-premises application to the AWS Cloud. The company is legally obligated to retain certain data in its on-premises data center. Which AWS service or feature will support this requirement?

- A. AWS Wavelength
- B. AWS Local Zones
- C. VMware Cloud on AWS
- D. AWS Outposts

Answer: D

Explanation:

AWS Outposts is a fully managed service that extends AWS infrastructure, AWS services, APIs, and tools to virtually any datacenter, co-location space, or on-premises facility for a truly consistent hybrid experience. AWS Outposts enables you to run AWS services in your on-premises data center, which can support the requirement of retaining certain data on-premises due to legal obligations.

NEW QUESTION 3

- (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 security. 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 cloud.

NEW QUESTION 4

- (Topic 2)

Which AWS service provides a highly accurate and easy-to-use enterprise search service that is powered by machine learning (ML)?

- A. Amazon Kendra
- B. Amazon SageMaker
- C. Amazon Augmented AI (Amazon A2I)
- D. Amazon Polly

Answer: A

Explanation:

Amazon Kendra is a service that provides a highly accurate and easy-to-use enterprise search service that is powered by machine learning. Kendra delivers powerful natural language search capabilities to your websites and applications so your end users can more easily find the information they need within the vast amount of content spread across your company. Amazon SageMaker is a service that provides a fully managed platform for data scientists and developers to quickly and easily build, train, and deploy machine learning models at any scale. Amazon Augmented AI (Amazon A2I) is a service that makes it easy to build the workflows required for human review of ML predictions. Amazon A2I brings human review to all developers, removing the undifferentiated heavy lifting associated with building human review systems or managing large numbers of human reviewers. Amazon Polly is a service that turns text into lifelike speech, allowing you to create applications that talk, and build entirely new categories of speech-enabled products. None of these services provide an enterprise search service that is powered by machine learning.

NEW QUESTION 5

- (Topic 2)

A company manages factory machines in real time. The company wants to use AWS technology to deploy its monitoring applications as close to the factory machines as possible.

Which AWS solution will meet these requirements with the LEAST latency?

- A. AWS Outposts
- B. Amazon EC2
- C. AWS App Runner
- D. AWS Batch

Answer: A

Explanation:

AWS Outposts is a fully managed service that extends AWS infrastructure, AWS services, APIs, and tools to virtually any datacenter, co-location space, or on-premises facility for a truly consistent hybrid experience. AWS Outposts enables you to run AWS services in your on-premises data center1.

NEW QUESTION 6

- (Topic 2)

A company is collecting user behavior patterns to identify how to meet goals for sustainability impact.

Which guidelines are best practices for the company to implement to meet these goals? (Select TWO.)

- A. Scale infrastructure with user load.
- B. Maximize the geographic distance between workloads and user locations.
- C. Eliminate creation and maintenance of unused assets.
- D. Scale resources with excess capacity and remove auto scaling.
- E. Scale infrastructure based on the number of users.

Answer: AC

Explanation:

To meet the goals for sustainability impact, the company should follow the best practices of scaling infrastructure with user load and eliminating creation and maintenance of unused assets. Scaling infrastructure with user load means adjusting the capacity of the infrastructure to match the demand of the users, which can reduce the energy consumption and carbon footprint of the system. Eliminating creation and maintenance of unused assets means avoiding the waste of resources and money on assets that are not needed or used, which can also improve the environmental and economic efficiency of the system3.

NEW QUESTION 7

- (Topic 2)

Which AWS service or tool helps companies measure the environmental impact of their AWS usage?

- A. AWS customer carbon footprint tool
- B. AWS Compute Optimizer
- C. Sustainability pillar
- D. OS-Climate (Open Source Climate Data Commons)

Answer: A

Explanation:

AWS customer carbon footprint tool is an AWS service or tool that helps companies measure the environmental impact of their AWS usage. It allows users to estimate the carbon emissions associated with their AWS resources and services, such as EC2, S3, and Lambda. It also provides recommendations and best practices to reduce the carbon footprint and improve the sustainability of their AWS workloads4. AWS Compute Optimizer is an AWS service that helps users optimize the performance and cost of their EC2 instances and Auto Scaling groups. It provides recommendations for optimal instance types, sizes, and configurations based on the workload characteristics and utilization metrics. It does not help users measure the environmental impact of their AWS usage. Sustainability pillar is a concept that refers to the ability of a system to operate in an environmentally friendly and socially responsible manner. It is not an AWS service or tool that helps users measure the environmental impact of their AWS usage. OS-Climate (Open Source Climate Data Commons) is an initiative that aims to provide open source data, tools, and platforms to accelerate climate action and innovation. It is not an AWS service or tool that helps users measure the environmental impact of their AWS usage.

NEW QUESTION 8

- (Topic 2)

A company wants to develop a shopping application that records customer orders. The application needs to use an AWS managed database service to store data.

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

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

Answer: A

Explanation:

A is correct because Amazon RDS is the AWS service that provides a managed relational database service that supports various database engines, such as MySQL, PostgreSQL, Oracle, and SQL Server. B is incorrect because Amazon Redshift is the AWS service that provides a managed data warehouse service that is optimized for analytical queries. C is incorrect because Amazon ElastiCache is the AWS service that provides a managed in-memory data store service that supports Redis and Memcached. D is incorrect because Amazon Neptune is the AWS service that provides a managed graph database service that supports property graph and RDF models.

NEW QUESTION 9

- (Topic 2)

Which design principle is included in the operational excellence pillar of the AWS Well- Architected Framework?

- A. Create annotated documentation.
- B. Anticipate failure.
- C. Ensure performance efficiency.
- D. Optimize costs.

Answer: A

Explanation:

Create annotated documentation is the design principle that is included in the operational excellence pillar of the AWS Well-Architected Framework. According to the AWS Well-Architected Framework whitepaper, creating annotated documentation means "documenting your workload so that the team understands the architecture, how to operate the workload, and how the workload delivers value to customers." Anticipate failure, ensure performance efficiency, and optimize costs are design principles that belong to other pillars of the AWS Well-Architected Framework, such as reliability, performance efficiency, and cost optimization.

NEW QUESTION 10

- (Topic 2)

A company is preparing to launch a redesigned website on AWS. Users from around the world will download digital handbooks from the website. Which AWS solution should the company use to provide these static files securely?

- A. Amazon Kinesis Data Streams
- B. Amazon CloudFront with Amazon S3
- C. Amazon EC2 instances with an Application Load Balancer
- D. Amazon Elastic File System (Amazon EFS)

Answer: B

Explanation:

Amazon CloudFront with Amazon S3 is a solution that allows you to provide static files securely to users from around the world. Amazon CloudFront is a fast content delivery network (CDN) service that securely delivers data, videos, applications, and APIs to customers globally with low latency, high transfer speeds, all within a developer-friendly environment. Amazon S3 is an object storage service that offers industry-leading scalability, data availability, security, and performance. You can use Amazon S3 to store and retrieve any amount of data from anywhere. You can also configure Amazon S3 to work with Amazon CloudFront to distribute your content to edge locations near your users for faster delivery and lower latency. Amazon Kinesis Data Streams is a service that enables you to build custom applications that process or analyze streaming data for specialized needs. This option is not relevant for providing static files securely. Amazon EC2 instances with an Application Load Balancer is a solution that allows you to distribute incoming traffic across multiple targets, such as EC2 instances, in multiple Availability Zones. This option is suitable for dynamic web applications, but not necessary for static files. Amazon Elastic File System (Amazon EFS) is a service that provides a simple, scalable, fully managed elastic NFS file system for use with AWS Cloud services and on- premises resources. This option is not relevant for providing static files securely.

NEW QUESTION 10

- (Topic 2)

Which AWS service or tool provides on-demand access to AWS security and compliance reports and AWS online agreements?

- A. AWS Artifact
- B. AWS Trusted Advisor
- C. Amazon Inspector
- D. AWS Billing console

Answer: A

Explanation:

AWS Artifact is the AWS service or tool that provides on-demand access to AWS security and compliance reports and AWS online agreements. AWS Trusted Advisor is a tool that provides real-time guidance to help users provision their resources following AWS best practices. Amazon Inspector is a service that helps users improve the security and compliance of their applications. AWS Billing console is a tool that helps users manage their AWS costs and usage. These concepts are explained in the AWS Cloud Practitioner Essentials course.

NEW QUESTION 14

- (Topic 2)

A company wants to push VPC Flow Logs to an Amazon S3 bucket. A company wants to optimize long-term compute costs of AWS Lambda functions and Amazon EC2 instances. Which AWS purchasing option should the company choose to meet these requirements?

- A. Dedicated Hosts
- B. Compute Savings Plans
- C. Reserved Instances
- D. Spot Instances

Answer: B

Explanation:

Compute Savings Plans are a flexible and cost-effective way to optimize long-term compute costs of AWS Lambda functions and Amazon EC2 instances. With Compute Savings Plans, customers can commit to a consistent amount of compute usage (measured in \$/hour) for a 1-year or 3-year term and receive a discount of up to 66% compared to On-Demand prices. Dedicated Hosts are physical servers with EC2 instance capacity fully dedicated to the customer's use. They are suitable for customers who have specific server-bound software licenses or compliance requirements. Reserved Instances are a pricing model that provides a significant discount (up to 75%) compared to On-Demand pricing and a capacity reservation for EC2 instances. They are available in 1-year or 3-year terms and different payment options. Spot Instances are spare EC2 instances that are available at up to 90% discount compared to On-Demand prices. They are suitable for customers who have flexible start and end times, can withstand interruptions, and can handle excess capacity.

NEW QUESTION 17

- (Topic 2)

A company wants to access a report about the estimated environmental impact of the company's AWS usage. Which AWS service or feature should the company use to meet this requirement?

- A. AWS Organizations
- B. IAM policy
- C. AWS Billing console
- D. Amazon Simple Notification Service (Amazon SNS)

Answer: C

Explanation:

The company should use the AWS Billing console to access a report about the estimated environmental impact of the company's AWS usage. The AWS Billing console provides customers with various tools and reports to manage and monitor their AWS costs and usage. One of the reports available in the AWS Billing console is the AWS Sustainability Dashboard, which shows the estimated carbon footprint and energy mix of the customer's AWS usage. The company can use this dashboard to measure and improve the sustainability of their cloud workloads. AWS Organizations, IAM policy, and Amazon Simple Notification Service (Amazon SNS) are not services or features that can provide a report about the estimated environmental impact of the company's AWS usage. AWS Organizations is a service that enables customers to centrally manage and govern their AWS accounts. IAM policy is a document that defines the permissions for an IAM identity (user, group, or role) or an AWS resource. Amazon SNS is a fully managed pub/sub messaging service that enables customers to send messages to subscribers or other AWS services.

NEW QUESTION 20

- (Topic 2)

A company wants to use Amazon EC2 instances for a stable production workload that will run for 1 year. Which instance purchasing option meets these requirements MOST cost-effectively?

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

Answer: B

Explanation:

B is correct because Reserved Instances are the instance purchasing option that offers the most cost-effective way to use Amazon EC2 instances for a stable production workload that will run for 1 year, as they provide significant discounts compared to On-Demand Instances in exchange for a commitment to use a specific amount of computing power for a period of time. A is incorrect because Dedicated Hosts are the instance purchasing option that allows customers to use physical servers that are fully dedicated to their use, which is more expensive and less flexible than Reserved Instances. C is incorrect because On-Demand Instances are the instance purchasing option that allows customers to pay for compute capacity by the hour or second with no long-term commitments, which is more suitable for short-term, variable, and unpredictable workloads. D is incorrect because Spot Instances are the instance purchasing option that allows customers to bid on spare Amazon EC2 computing capacity, which is more suitable for flexible, scalable, and fault-tolerant workloads that can tolerate interruptions.

NEW QUESTION 23

- (Topic 2)

Which AWS service offers a global content delivery network (CDN) that helps companies securely deliver websites, videos, applications, and APIs at high speeds with low latency?

- A. Amazon EC2
- B. Amazon CloudFront
- C. Amazon CloudWatch
- D. AWS CloudFormation

Answer: B

Explanation:

Amazon CloudFront is the AWS service that offers a global content delivery network (CDN) that helps companies securely deliver websites, videos, applications, and APIs at high speeds with low latency. Amazon CloudFront is a web service that speeds up distribution of static and dynamic web content, such as HTML, CSS, JavaScript, and image files, to users. Amazon CloudFront uses a global network of edge locations, located near users' geographic locations, to cache and serve content with high availability and performance. Amazon CloudFront also provides features such as AWS Shield for DDoS protection, AWS Certificate Manager for SSL/TLS encryption, AWS WAF for web application firewall, and AWS Lambda@Edge for customizing content delivery with serverless code. Amazon EC2, Amazon CloudWatch, and AWS CloudFormation are not services that offer a global CDN. Amazon EC2 is a service that provides scalable compute capacity in the cloud. Amazon CloudWatch is a service that provides monitoring and observability for AWS resources and applications. AWS CloudFormation is a service that provides a common language to model and provision AWS resources and their dependencies.

NEW QUESTION 27

- (Topic 2)

Which AWS service requires the customer to patch the guest operating system?

- A. AWS Lambda
- B. Amazon OpenSearch Service
- C. Amazon EC2
- D. Amazon ElastiCache

Answer: C

Explanation:

The AWS service that requires the customer to patch the guest operating system is Amazon EC2. 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. The customer is responsible for patching and updating the guest operating system and any applications that run on the EC2 instances, as part of the security in the cloud. AWS Lambda, Amazon

OpenSearch Service, and Amazon ElastiCache are not services that require the customer to patch the guest operating system. AWS Lambda is a serverless compute service that allows customers to run code without provisioning or managing servers. Amazon OpenSearch Service is a fully managed service that makes it easy to deploy, operate, and scale OpenSearch clusters in the AWS Cloud. Amazon ElastiCache is a fully managed service that provides in-memory data store and cache solutions, such as Redis and Memcached. These services are managed by AWS, and AWS is responsible for patching and updating the underlying infrastructure and software.

NEW QUESTION 32

- (Topic 1)

Which AWS service will help protect applications running on AWS from DDoS attacks?

- A. Amazon GuardDuty
- B. AWS WAF
- C. AWS Shield
- D. Amazon Inspector

Answer: C

Explanation:

AWS Shield is a managed Distributed Denial of Service (DDoS) protection service that safeguards applications running on AWS. AWS Shield 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.

NEW QUESTION 35

- (Topic 1)

Which AWS features will meet these requirements? (Select TWO.)

- A. Security groups
- B. Network ACLs
- C. S3 bucket policies
- D. IAM user policies
- E. S3 bucket versioning

Answer: CD

Explanation:

The correct answers are C and D because S3 bucket policies and IAM user policies are AWS features that will meet the requirements. S3 bucket policies are access policies that can be attached to Amazon S3 buckets to grant or deny permissions to the bucket and the objects it contains. S3 bucket policies can be used to control who has permission to read, write, or delete objects that the company stores in the S3 bucket. IAM user policies are access policies that can be attached to IAM users to grant or deny permissions to AWS resources and actions. IAM user policies can be used to control who has permission to read, write, or delete objects that the company stores in the S3 bucket. The other options are incorrect because they are not AWS features that will meet the requirements. Security groups and network ACLs are AWS features that act as firewalls to control inbound and outbound traffic to and from Amazon EC2 instances and subnets. Security groups and network ACLs do not control who has permission to read, write, or delete objects that the company stores in the S3 bucket. S3 bucket versioning is an AWS feature that enables users to keep multiple versions of the same object in the same bucket. S3 bucket versioning can be used to recover from accidental overwrites or deletions of objects, but it does not control who has permission to read, write, or delete objects that the company stores in the S3 bucket. Reference: Using Bucket Policies and User Policies, Security Groups for Your VPC, Network ACLs, [Using Versioning]

NEW QUESTION 38

- (Topic 1)

Which of the following are pillars of the AWS Well-Architected Framework? (Select TWO.)

- A. Availability
- B. Reliability
- C. Scalability
- D. Responsive design
- E. Operational excellence

Answer: BE

Explanation:

The correct answers to the questions are B and E because reliability and operational excellence are pillars of the AWS Well-Architected Framework. The AWS Well-Architected Framework is a set of best practices and guidelines for designing and operating reliable, secure, efficient, and cost-effective systems in the cloud. The AWS Well-Architected Framework consists of five pillars: operational excellence, security, reliability, performance efficiency, and cost optimization. Each pillar has a set of design principles that describe the characteristics of a well-architected system. Reliability is the pillar that focuses on the ability of a system to recover from failures and meet business and customer demand. Operational excellence is the pillar that focuses on the ability of a system to run and monitor processes that support business outcomes and continually improve. The other options are incorrect because they are not pillars of the AWS Well-Architected Framework. Availability, scalability, and responsive design are important aspects of cloud architecture, but they are not separate pillars in the framework. Availability and scalability are related to the reliability and performance efficiency pillars, while responsive design is related to the customer experience and user interface. Reference: AWS Well-Architected Framework

NEW QUESTION 42

- (Topic 1)

Which of the following is an AWS value proposition that describes a user's ability to scale infrastructure based on demand?

- A. Speed of innovation
- B. Resource elasticity
- C. Decoupled architecture
- D. Global deployment

Answer: B

Explanation:

Resource elasticity is an AWS value proposition that describes a user's ability to scale infrastructure based on demand. Resource elasticity means that the user can provision or deprovision resources quickly and easily, without any upfront commitment or long-term contract. Resource elasticity can help the user optimize the cost and performance of the application, as well as respond to changing business needs and customer expectations. Resource elasticity can be achieved by using services such as Amazon EC2, Amazon S3, Amazon RDS, Amazon DynamoDB, Amazon ECS, and AWS Lambda. [AWS Cloud Value Framework] AWS Certified Cloud Practitioner - aws.amazon.com

NEW QUESTION 44

- (Topic 1)

A company wants to centrally manage security policies and billing services within a multi-account AWS environment. Which AWS service should the company use to meet these requirements?

- A. AWS Identity and Access Management (IAM)
- B. AWS Organizations
- C. AWS Resource Access Manager (AWS RAM)
- D. AWS Config

Answer: B

Explanation:

AWS Organizations is a service that helps you centrally manage and govern your environment as you grow and scale your AWS resources. You can use AWS Organizations to create groups of accounts and apply policies to them. You can also use AWS Organizations to consolidate billing for multiple accounts. Therefore, the correct answer is B. You can learn more about AWS Organizations and its features from this page.

NEW QUESTION 46

- (Topic 1)

A company needs to store data across multiple Availability Zones in an AWS Region. The data will not be accessed regularly but must be immediately retrievable.

Which Amazon Elastic File System (Amazon EFS) storage class meets these requirements MOST cost effectively?

- A. EFS Standard
- B. EFS Standard-Infrequent Access(EFS Standard-IA)
- C. EFS One Zone
- D. EFS One Zone-Infrequent Access (EFS One Zone-IA)

Answer: B

Explanation:

EFS Standard-Infrequent Access (EFS Standard-IA) is the storage class that meets the requirements of storing data across multiple Availability Zones in an AWS Region, that will not be accessed regularly but must be immediately retrievable, most cost-effectively. EFS Standard-IA is designed for files that are accessed less frequently, but still require the same high performance, low latency, and high availability as EFS Standard. EFS Standard-IA has a lower storage cost than EFS Standard, but charges a small additional fee for each access. EFS One Zone and EFS One Zone-IA store data in a single Availability Zone, which reduces the availability and durability compared to EFS Standard and EFS Standard-IA.

NEW QUESTION 50

- (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 55

- (Topic 1)

Which AWS service or feature is used to send both text and email messages from distributed applications?

- A. Amazon Simple Notification Service (Amazon SNS)
- B. Amazon Simple Email Service (Amazon SES)
- C. Amazon CloudWatch alerts
- D. Amazon Simple Queue Service (Amazon SQS)

Answer: A

Explanation:

Amazon Simple Notification Service (Amazon SNS) is the AWS service or feature that is used to send both text and email messages from distributed applications. Amazon SNS is a fully managed pub/sub messaging service that enables the user to send messages to multiple subscribers or endpoints, such as email addresses, phone numbers, HTTP endpoints, AWS Lambda functions, and more. Amazon SNS can be used to send notifications, alerts, confirmations, and reminders from applications to users or other applications4.

NEW QUESTION 56

- (Topic 1)

An auditor needs to find out whether a specific AWS service is compliant with specific compliance frameworks. Which AWS service will provide this information?

- A. AWS Artifact
- B. AWS Trusted Advisor
- C. Amazon GuardDuty
- D. AWS Certificate Manager (ACM)

Answer: A

Explanation:

AWS Artifact is the service that will provide the information about whether a specific AWS service is compliant with specific compliance frameworks. AWS Artifact is a self-service portal that allows you to access, review, and download AWS security and compliance reports and agreements. You can use AWS Artifact to verify the compliance status of AWS services across various regions and compliance programs, such as ISO, PCI, SOC, FedRAMP, HIPAA, and more12

NEW QUESTION 61

- (Topic 1)

According to the AWS shared responsibility model, which of the following are AWS responsibilities? (Select TWO.)

- A. Network infrastructure and virtualization of infrastructure
- B. Security of application data
- C. Guest operating systems
- D. Physical security of hardware
- E. Credentials and policies

Answer: AD

Explanation:

The correct answers are A and D because network infrastructure and virtualization of infrastructure and physical security of hardware are AWS responsibilities 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 not AWS responsibilities according to the AWS shared responsibility model. Security of application data, guest operating systems, and credentials and policies are customer responsibilities according to the AWS shared responsibility model. Reference: [AWS Shared Responsibility Model]

NEW QUESTION 62

- (Topic 1)

Which AWS Well-Architected Framework concept represents a system's ability to remain functional when the system encounters operational problems?

- A. Consistency
- B. Elasticity
- C. Durability
- D. Latency

Answer: B

Explanation:

The AWS Well-Architected Framework is a set of best practices and guidelines for designing and operating systems in the cloud. The framework consists of five pillars: operational excellence, security, reliability, performance efficiency, and cost optimization. The concept of elasticity represents a system's ability to adapt to changes in demand by scaling resources up or down automatically. Therefore, the correct answer is B. You can learn more about the AWS Well-Architected Framework and its pillars from this page.

NEW QUESTION 64

- (Topic 1)

Which activity is a customer responsibility in the AWS Cloud according to the AWS shared responsibility model?

- A. Ensuring network connectivity from AWS to the internet
- B. Patching and fixing flaws within the AWS Cloud infrastructure
- C. Ensuring the physical security of cloud data centers
- D. Ensuring Amazon EBS volumes are backed up

Answer: D

Explanation:

The AWS shared responsibility model describes how AWS and the customer share responsibility for security and compliance of the AWS environment. AWS is responsible for the security of the cloud, which includes the physical security of AWS facilities, the infrastructure, hardware, software, and networking that run AWS services. The customer is responsible for security in the cloud, which includes the configuration of security groups, the encryption of customer data on AWS, the management of AWS Lambda infrastructure, and the management of network throughput of each AWS Region. One of the customer responsibilities is to ensure that Amazon EBS volumes are backed up.

NEW QUESTION 67

- (Topic 1)

Which AWS Support plan provides customers with access to an AWS technical account manager (TAM)?

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

Answer: D

Explanation:

The correct answer is D because AWS Enterprise Support is the support plan that provides customers with access to an AWS technical account manager (TAM). AWS Enterprise Support is the highest level of support plan offered by AWS, and it provides customers with the most comprehensive and personalized support experience. An AWS TAM is a dedicated technical resource who works closely with customers to understand their business and technical needs, provide proactive guidance, and coordinate support across AWS teams. The other options are incorrect because they are not support plans that provide customers with access to an AWS TAM. AWS Basic Support is the default and free support plan that provides customers with access to online documentation, forums, and account information. AWS Developer Support is the lowest level of paid support plan that provides customers with access to technical support during business hours, general guidance, and best practice recommendations. AWS Business Support is the intermediate level of paid support plan that provides customers with access to technical support 24/7, system health checks, architectural guidance, and case management. Reference: AWS Support Plans

NEW QUESTION 69

- (Topic 1)

Which AWS benefit is demonstrated by on-demand technology services that enable companies to replace upfront fixed expenses with variable expenses?

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

Answer: C

Explanation:

Pay-as-you-go pricing is an AWS benefit that demonstrates the ability of users to replace upfront fixed expenses with variable expenses. With pay-as-you-go pricing, users only pay for the resources they consume, without any long-term contracts or commitments. This can lower the total cost of ownership and increase the return on investment. Pay-as-you-go pricing also provides flexibility and scalability, as users can adjust their resource usage according to their changing needs and demands. AWS Cloud Value Framework AWS Certified Cloud Practitioner - aws.amazon.com

NEW QUESTION 73

- (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 77

SIMULATION - (Topic 1)

A company runs thousands of simultaneous simulations using AWS Batch. Each simulation is stateless, is fault tolerant, and runs for up to 3 hours.

Which pricing model enables the company to optimize costs and meet these requirements?

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

Answer: B

Explanation:

The correct answer is B because Spot Instances enable the company to optimize costs and meet the requirements. Spot Instances are spare 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 applications that can run for any duration. The other options are incorrect because they do not enable the company to optimize costs and meet the requirements. Reserved Instances are 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 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 Instances are EC2 instances that run on hardware that is dedicated to a single customer. Dedicated Instances are suitable for workloads that require regulatory compliance or data isolation. Reference: [Amazon EC2 Instance Purchasing Options]

NEW QUESTION 81

- (Topic 1)

A company has two AWS accounts in an organization in AWS Organizations for consolidated billing. All of the company's AWS resources are hosted in one AWS Region.

Account A has purchased five Amazon EC2 Standard Reserved Instances (RIs) and has four EC2 instances running. Account B has not purchased any RIs and also has four EC2 instances running. Which statement is true regarding pricing for these eight instances?

- A. The eight instances will be charged as regular instances.
- B. Four instances will be charged as RIs, and four will be charged as regular instances.
- C. Five instances will be charged as RIs, and three will be charged as regular instances.
- D. The eight instances will be charged as RIs.

Answer: B

Explanation:

The statement that is true regarding pricing for these eight instances is: four instances will be charged as RIs, and four will be charged as regular instances. Amazon EC2 Reserved Instances (RIs) are a pricing model that allows users to reserve EC2 instances for a specific term and benefit from discounted hourly rates and capacity reservation. RIs are purchased for a specific AWS Region, and can be shared across multiple accounts in an organization in AWS Organizations for consolidated billing. However, RIs are applied on a first-come, first-served basis, and there is no guarantee that all instances in the organization will be charged at the RI rate. In this case, Account A has purchased five RIs and has four instances running, so all four instances will be charged at the RI rate. Account B has not purchased any RIs and also has four instances running, so all four instances will be charged at the regular rate. The remaining RI in Account A will not be applied to any instance in Account B, and will be wasted.

NEW QUESTION 82

- (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 85

- (Topic 1)

Which AWS service aggregates, organizes, and prioritizes security alerts and findings from multiple AWS services?

- A. Amazon Detective
- B. Amazon Inspector
- C. Amazon Macie
- D. AWS Security Hub

Answer: D

Explanation:

The correct answer is D because AWS Security Hub is a service that aggregates, organizes, and prioritizes security alerts and findings from multiple AWS services, such as Amazon GuardDuty, Amazon Inspector, Amazon Macie, AWS Firewall Manager, and AWS IAM Access Analyzer. The other options are incorrect because they are not services that aggregate security alerts and findings from multiple AWS services. Amazon Detective is a service that helps users analyze and visualize security data to investigate and remediate potential issues. Amazon Inspector is a service that helps users find security vulnerabilities and deviations from best practices in their Amazon EC2 instances. Amazon Macie is a service that helps users discover, classify, and protect sensitive data stored in Amazon S3. Reference: AWS Security Hub FAQs

NEW QUESTION 88

- (Topic 1)

A company plans to migrate to AWS and wants to create cost estimates for its AWS use cases. Which AWS service or tool can the company use to meet these requirements?

- A. AWS Pricing Calculator
- B. Amazon CloudWatch
- C. AWS Cost Explorer
- D. AWS Budgets

Answer: A

Explanation:

AWS Pricing Calculator is a web-based planning tool that customers can use to create estimates for their AWS use cases. They can use it to model their solutions before building them, explore the AWS service price points, and review the calculations behind their estimates. Therefore, the correct answer is A. You can learn more about AWS Pricing Calculator and how it works from this page.

NEW QUESTION 90

- (Topic 1)

Which AWS database service provides in-memory data storage?

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

Answer: B

Explanation:

The correct answer is B because Amazon ElastiCache is a service that provides in-memory data storage. Amazon ElastiCache is a fully managed, scalable, and high-performance service that supports two popular open-source in-memory engines: Redis and Memcached. Amazon ElastiCache allows users to store and retrieve data from fast, low-latency, and high-throughput in-memory systems. Users can use Amazon ElastiCache to improve the performance of their applications by caching frequently accessed data, reducing database load, and enabling real-time data processing. The other options are incorrect because they are not services that provide in-memory data storage. Amazon DynamoDB is a service that provides key-value and document data storage. Amazon RDS is a service that provides relational data storage. Amazon Timestream is a service that provides time series data storage. Reference: Amazon ElastiCache FAQs

NEW QUESTION 95

- (Topic 1)

Which of the following is a recommended design principle of the AWS Well-Architected Framework?

- A. Reduce downtime by making infrastructure changes infrequently and in large increments.
- B. Invest the time to configure infrastructure manually.
- C. Learn to improve from operational failures.
- D. Use monolithic application design for centralization.

Answer: C

Explanation:

The correct answer is C because learning to improve from operational failures is a recommended design principle of the AWS Well-Architected Framework. The AWS Well-Architected Framework is a set of best practices and guidelines for designing and operating reliable, secure, efficient, and cost-effective systems in the cloud. The AWS Well-Architected Framework consists of five pillars: operational excellence, security, reliability, performance efficiency, and cost optimization. Each pillar has a set of design principles that describe the characteristics of a well-architected system. Learning to improve from operational failures is a design principle of the operational excellence pillar, which focuses on running and monitoring systems to deliver business value and continually improve supporting processes and procedures. The other options are incorrect because they are not recommended design principles of the AWS Well-Architected Framework. Reducing downtime by making infrastructure changes infrequently and in large increments is not a design principle of the AWS Well-Architected Framework, but rather a source of risk and inefficiency. A well-architected system should implement changes frequently and in small increments to minimize the impact and scope of failures. Investing the time to configure infrastructure manually is not a design principle of the AWS Well-Architected Framework, but rather a source of human error and inconsistency. A well-architected system should automate manual tasks to improve the speed and accuracy of operations. Using monolithic application design for centralization is not a design principle of the AWS Well-Architected Framework, but rather a source of complexity and rigidity. A well-architected system should use loosely coupled and distributed components to enable scalability and resilience. Reference: [AWS Well-Architected Framework]

NEW QUESTION 97

- (Topic 1)

Which of the following are components of an AWS Site-to-Site VPN connection? (Select TWO.)

- A. AWS Storage Gateway
- B. Virtual private gateway
- C. NAT gateway
- D. Customer gateway
- E. Internet gateway

Answer: BD

Explanation:

The correct answers are B and D because a virtual private gateway and a customer gateway are components of an AWS Site-to-Site VPN connection. A virtual private gateway is the AWS side of the VPN connection that attaches to the customer's VPC. A customer gateway is the customer side of the VPN connection that resides in the customer's network. The other options are incorrect because they are not components of an AWS Site-to-Site VPN connection. AWS Storage Gateway is a service that connects on-premises software applications with cloud-based storage. NAT gateway is a service that enables instances in a private subnet to connect to the internet or other AWS services, but prevents the internet from initiating a connection with those instances. Internet gateway is a service that enables communication between instances in a VPC and the internet. Reference: [What is AWS Site-to-Site VPN?]

NEW QUESTION 100

- (Topic 1)

Which AWS service or tool helps to centrally manage billing and allow controlled access to resources across AWS accounts?

- A. AWS Identity and Access Management (IAM)
- B. AWS Organizations
- C. AWS Cost Explorer
- D. AWS Budgets

Answer: B

Explanation:

AWS Organizations helps to centrally manage billing and allow controlled access to resources across AWS accounts. AWS Organizations is a service that enables the user to consolidate multiple AWS accounts into an organization that can be managed as a single unit. AWS Organizations allows the user to create groups of accounts and apply policies to them, such as service control policies (SCPs) that specify the services and actions that users and roles can access in the accounts. AWS Organizations also enables the user to use consolidated billing, which combines the usage and charges from all the accounts in the organization into a single bill.

NEW QUESTION 104

- (Topic 1)

What are some advantages of using Amazon EC2 instances to host applications in the AWS Cloud instead of on premises? (Select TWO.)

- A. EC2 includes operating system patch management
- B. EC2 integrates with Amazon VPC
- C. AWS CloudTrail, and AWS Identity and Access Management (IAM)
- D. EC2 has a 100% service level agreement (SLA).
- E. EC2 has a flexible, pay-as-you-go pricing model.
- F. EC2 has automatic storage cost optimization.

Answer: BD

Explanation:

Some of the advantages of using Amazon EC2 instances to host applications in the AWS Cloud instead of on premises are:

? EC2 integrates with Amazon VPC, AWS CloudTrail, and AWS Identity and Access Management (IAM). Amazon VPC lets you provision a logically isolated section of the AWS Cloud where you can launch AWS resources in a virtual network that you define. AWS CloudTrail enables governance, compliance, operational auditing, and risk auditing of your AWS account. AWS IAM enables you to manage access to AWS services and resources securely. Therefore, the correct answer is B. You can learn more about Amazon EC2 and its integration with other AWS services from this page.

? EC2 has a flexible, pay-as-you-go pricing model. You only pay for the compute capacity you use, and you can scale up and down as needed. You can also choose from different pricing options, such as On-Demand, Savings Plans, Reserved Instances, and Spot Instances, to optimize your costs. Therefore, the correct answer is D. You can learn more about Amazon EC2 pricing from this page.

The other options are incorrect because:

? EC2 does not include operating system patch management. You are responsible for managing and maintaining your own operating systems on EC2 instances. You can use AWS Systems Manager to automate common maintenance tasks, such as applying patches, or use Amazon EC2 Image Builder to create and maintain secure images. Therefore, the incorrect answer is A.

? EC2 does not have a 100% service level agreement (SLA). The EC2 SLA guarantees 99.99% availability for each EC2 Region, not for each individual instance. Therefore, the incorrect answer is C.

? EC2 does not have automatic storage cost optimization. You are responsible for choosing the right storage option for your EC2 instances, such as Amazon Elastic Block Store (EBS) or Amazon Elastic File System (EFS), and monitoring and optimizing your storage costs. You can use AWS Cost Explorer or AWS Trusted Advisor to analyze and reduce your storage spending. Therefore, the incorrect answer is E.

NEW QUESTION 106

- (Topic 1)

Which AWS service is a key-value database that provides sub-millisecond latency on a large scale?

- A. Amazon DynamoDB
- B. Amazon Aurora
- C. Amazon DocumentDB (with MongoDB compatibility)
- D. Amazon Neptune

Answer: A

Explanation:

The correct answer is A because Amazon DynamoDB is a key-value database that provides sub-millisecond latency on a large scale. Amazon DynamoDB is a fully managed, serverless, and scalable NoSQL database service that supports both key-value and document data models. The other options are incorrect because they are not key-value databases. Amazon Aurora is a relational database that is compatible with MySQL and PostgreSQL. Amazon DocumentDB (with MongoDB compatibility) is a document database that is compatible with MongoDB. Amazon Neptune is a graph database that supports property graph and RDF models. Reference: Amazon DynamoDB FAQs

NEW QUESTION 110

- (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 115

- (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 119

- (Topic 1)

When designing AWS workloads to be operational even when there are component failures, what is an AWS best practice?

- A. Perform quarterly disaster recovery tests.
- B. Place the main component on the us-east-1 Region.
- C. Design for automatic failover to healthy resources.
- D. Design workloads to fit on a single Amazon EC2 instance.

Answer: C

Explanation:

Designing for automatic failover to healthy resources is an AWS best practice when designing AWS workloads to be operational even when there are component failures. This means that you should architect your system to handle the loss of one or more components without impacting the availability or performance of your application. You can use various AWS services and features to achieve this, such as Auto Scaling, Elastic Load Balancing, Amazon Route 53, Amazon CloudFormation, and AWS CloudFormation4.

NEW QUESTION 123

- (Topic 1)

Which AWS service or tool can be used to consolidate payments for a company with multiple AWS accounts?

- A. AWS Cost and Usage Report
- B. AWS Organizations
- C. Cost Explorer
- D. AWS Budgets

Answer: B

Explanation:

AWS Organizations is an account management service that enables you to consolidate multiple AWS accounts into an organization that you create and centrally manage. AWS Organizations includes consolidated billing and account management capabilities that enable you to better meet the budgetary, security, and compliance needs of your business1.

NEW QUESTION 124

- (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 126

- (Topic 1)

A company is running applications on Amazon EC2 instances in the same AWS account for several different projects. The company wants to track the infrastructure costs for each of the projects separately. The company must conduct this tracking with the least possible impact to the existing infrastructure and with no additional cost.

What should the company do to meet these requirements?

- A. Use a different EC2 instance type for each project.
- B. Publish project-specific custom Amazon CloudWatch metrics for each application.
- C. Deploy EC2 instances for each project in a separate AWS account.
- D. Use cost allocation tags with values that are specific to each project.

Answer: D

Explanation:

The correct answer is D because cost allocation tags are a way to track the infrastructure costs for each of the projects separately. Cost allocation tags are key-value pairs that can be attached to AWS resources, such as EC2 instances, and used to categorize and group them for billing purposes. The other options are incorrect because they do not meet the requirements of the question. Use a different EC2 instance type for each project does not help to track the costs for each

project, and may impact the performance and compatibility of the applications. Publish project-specific custom Amazon CloudWatch metrics for each application does not help to track the costs for each project, and may incur additional charges for using CloudWatch. Deploy EC2 instances for each project in a separate AWS account does help to track the costs for each project, but it impacts the existing infrastructure and incurs additional charges for using multiple accounts. Reference: Using Cost Allocation Tags

NEW QUESTION 127

- (Topic 1)

Which duties are the responsibility of a company that is using AWS Lambda? (Select TWO.)

- A. Security inside of code
- B. Selection of CPU resources
- C. Patching of operating system
- D. Writing and updating of code
- E. Security of underlying infrastructure

Answer: AD

Explanation:

The duties that are the responsibility of a company that is using AWS Lambda are security inside of code and writing and updating of code. AWS Lambda is a serverless compute service that allows you to run code without provisioning or managing servers, scaling, or patching. AWS Lambda takes care of the security of the underlying infrastructure, such as the operating system, the network, and the firewall. However, the company is still responsible for the security of the code itself, such as encrypting sensitive data, validating input, and handling errors. The company is also responsible for writing and updating the code that defines the Lambda function, and choosing the runtime environment, such as Node.js, Python, or Java. AWS Lambda does not require the selection of CPU resources, as it automatically allocates them based on the memory configuration³⁴

NEW QUESTION 132

- (Topic 1)

A user wants to identify any security group that is allowing unrestricted incoming SSH traffic. Which AWS service can be used to accomplish this goal?

- A. Amazon Cognito
- B. AWS Shield
- C. Amazon Macie
- D. AWS Trusted Advisor

Answer: D

Explanation:

The correct answer to the question is D because AWS Trusted Advisor is an AWS service that can be used to accomplish the goal of identifying any security group that is allowing unrestricted incoming SSH traffic. AWS Trusted Advisor is a service that provides customers with recommendations that help them follow AWS best practices. Trusted Advisor evaluates the customer's AWS environment and identifies ways to optimize their AWS infrastructure, improve security and performance, reduce costs, and monitor service quotas. One of the checks that Trusted Advisor performs is the Security Groups - Specific Ports Unrestricted check, which flags security groups that allow unrestricted access to specific ports, such as port 22 for SSH. Customers can use this check to review and modify their security group rules to restrict SSH access to only authorized sources. Reference: Security Groups - Specific Ports Unrestricted

NEW QUESTION 135

- (Topic 1)

Which of the following is an advantage that users experience when they move on-premises workloads to the AWS Cloud?

- A. Elimination of expenses for running and maintaining data centers
- B. Price discounts that are identical to discounts from hardware providers
- C. Distribution of all operational controls to AWS
- D. Elimination of operational expenses

Answer: A

Explanation:

The advantage that users experience when they move on-premises workloads to the AWS Cloud is: elimination of expenses for running and maintaining data centers. By moving on-premises workloads to the AWS Cloud, users can reduce or eliminate the costs associated with owning and operating physical servers, storage, network equipment, and facilities. These costs include hardware purchase, maintenance, repair, power, cooling, security, and staff. Users can also benefit from the pay-as-you-go pricing model of AWS, which allows them to pay only for the resources they use, and scale up or down as needed.

NEW QUESTION 137

- (Topic 1)

Which AWS service provides the ability to host a NoSQL database in the AWS Cloud?

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

Answer: B

Explanation:

Amazon DynamoDB is a fully managed NoSQL database service that provides fast and predictable performance with seamless scalability. It supports both key-value and document data models, and allows you to create tables that can store and retrieve any amount of data, and serve any level of request traffic. You can also use DynamoDB Streams to capture data modification events in DynamoDB tables.

NEW QUESTION 139

- (Topic 1)

A company deploys its application to multiple AWS Regions and configures automatic failover between those Regions. Which cloud concept does this architecture represent?

- A. Security
- B. Reliability
- C. Scalability
- D. Cost optimization

Answer: B

Explanation:

Reliability is the cloud concept that this architecture represents. Reliability is the ability of a system to recover from infrastructure or service disruptions, dynamically acquire computing resources to meet demand, and mitigate disruptions such as misconfigurations or transient network issues. Deploying an application to multiple AWS Regions and configuring automatic failover between those Regions enhances the reliability of the application by reducing the impact of regional failures and increasing the availability of the application⁴

NEW QUESTION 143

- (Topic 1)

Which of the following are benefits that a company receives when it moves an on-premises production workload to AWS? (Select TWO.)

- A. AWS trains the company's staff on the use of all the AWS services.
- B. AWS manages all security in the cloud.
- C. AWS offers free support from technical account managers (TAMs).
- D. AWS offers high availability.
- E. AWS provides economies of scale.

Answer: DE

Explanation:

The correct answers are D and E because AWS offers high availability and AWS provides economies of scale are benefits that a company receives when it moves an on-premises production workload to AWS. High availability means that AWS has a global infrastructure that allows customers to deploy their applications and data across multiple regions and availability zones. This increases the fault tolerance and resilience of their applications and reduces the impact of failures. Economies of scale means that AWS can achieve lower variable costs than customers can get on their own. This allows customers to pay only for the resources they use and scale up or down as needed. The other options are incorrect because they are not benefits that a company receives when it moves an on-premises production workload to AWS. AWS trains the company's staff on the use of all the AWS services is not a benefit that a company receives when it moves an on-premises production workload to AWS. AWS does provide various learning resources and training courses for customers, but it does not train the company's staff on the use of all the AWS services. AWS manages all security in the cloud is not a benefit that a company receives when it moves an on-premises production workload to AWS. AWS is responsible for the security of the cloud, but the customer is responsible for the security in the cloud. AWS offers free support from technical account managers (TAMs) is not a benefit that a company receives when it moves an on-premises production workload to AWS. AWS does offer support from TAMs, but only for customers who have the AWS Enterprise Support plan, which is not free. Reference: What is Cloud Computing?, [AWS Shared Responsibility Model], [AWS Support Plans]

NEW QUESTION 146

- (Topic 3)

Which task can only an AWS account root user perform?

- A. Changing the AWS Support plan
- B. Deleting AWS resources
- C. Creating an Amazon EC2 instance key pair
- D. Configuring AWS WAF

Answer: A

Explanation:

The AWS account root user is the email address that you use to sign up for AWS. The root user has complete access to all AWS services and resources in the account. The root user can perform tasks that only the root user can do, such as changing the AWS Support plan, closing the account, and restoring IAM user permissions³⁴

NEW QUESTION 149

- (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 S3⁵. 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 151

- (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 155

- (Topic 3)

Which AWS service or tool helps users visualize, understand, and manage spending and usage over time?

- A. AWS Organizations
- B. AWS Pricing Calculator
- C. AWS Cost Explorer
- D. AWS Service Catalog

Answer: C

Explanation:

AWS Cost Explorer is the AWS service or tool that helps users visualize, understand, and manage spending and usage over time. AWS Cost Explorer is a web-based interface that allows users to access interactive graphs and tables that display their AWS costs and usage data. Users can create custom reports that analyze cost and usage data by various dimensions, such as service, region, account, tag, and more. Users can also view historical data for up to the last 12 months, forecast future costs for up to the next 12 months, and get recommendations for cost optimization. AWS Cost Explorer also provides preconfigured views that show common cost and usage scenarios, such as monthly spend by service, daily spend by linked account, and Reserved Instance utilization. Users can use AWS Cost Explorer to monitor their AWS spending and usage trends, identify cost drivers and anomalies, and optimize their resource allocation and budget planning. References: Cloud Cost Analysis - AWS Cost Explorer - AWS, Analyzing your costs with AWS Cost Explorer

NEW QUESTION 159

- (Topic 3)

A company wants to migrate a database from an on-premises environment to Amazon RDS.

After the migration is complete, which management task will the company still be responsible for?

- A. Hardware lifecycle management
- B. Application optimization
- C. Server maintenance
- D. Power, network, and cooling provisioning

Answer: B

Explanation:

Amazon RDS is a managed database service that handles most of the common database administration tasks, such as hardware provisioning, server maintenance, backup and recovery, patching, scaling, and replication. However, Amazon RDS does not optimize the application that interacts with the database. The company is still responsible for tuning the performance, security, and availability of the application according to its business requirements and best practices. References:

? What is Amazon Relational Database Service (Amazon RDS)?

? Perform common DBA tasks for Amazon RDS DB instances

NEW QUESTION 161

- (Topic 3)

Which of the following is an AWS Well-Architected Framework design principle for operational excellence in the AWS Cloud?

- A. Go global in minutes
- B. Make frequent, small, reversible changes
- C. Implement a strong foundation of identity and access management
- D. Stop spending money on hardware infrastructure for data center operations

Answer: B

Explanation:

Making frequent, small, reversible changes is one of the design principles for operational excellence in the AWS Cloud, as defined by the AWS Well-Architected Framework. This principle means that you should design your workloads to allow for rapid and safe changes, such as deploying updates, rolling back failures, and experimenting with new features. By making small and reversible changes, you can reduce the risk of errors, minimize the impact of failures, and increase the speed of recovery. References: 2: AWS Documentation - AWS Well-Architected Framework - Operational Excellence Pillar

NEW QUESTION 162

- (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 164

- (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 166

- (Topic 3)

Which AWS Cloud deployment model uses AWS Outposts as part of the application deployment infrastructure?

- A. On-premises
- B. Serverless
- C. Cloud-native
- D. Hybrid

Answer: D

Explanation:

AWS Outposts is a fully managed service that extends AWS infrastructure, services, APIs, and tools to customer premises. By providing local access to AWS managed infrastructure, AWS Outposts enables customers to build and run applications on premises using the same programming interfaces as in AWS Regions, while using local compute and storage resources for lower latency and local data processing needs. An Outpost is a pool of AWS compute and storage capacity deployed at a customer site. AWS operates, monitors, and manages this capacity as part of an AWS Region. You can create subnets on your Outpost and specify them when you create AWS resources such as EC2 instances, EBS volumes, ECS clusters, and RDS instances. Instances in Outpost subnets communicate with other instances in the AWS Region using private IP addresses, all within the same VPC. Outposts solutions allow you to extend and run native AWS services on premises, and is available in a variety of form factors, from 1U and 2U Outposts servers to 42U Outposts racks, and multiple rack deployments. With AWS Outposts, you can run some AWS services locally and connect to a broad range of services available in the local AWS Region². AWS Outposts is a hybrid cloud deployment model that uses AWS Outposts as part of the application deployment infrastructure. Hybrid cloud is a cloud computing environment that uses a mix of on-premises, private cloud, and public cloud services with orchestration between the platforms. Hybrid cloud provides businesses with greater flexibility, more deployment options, and optimized costs. By using AWS Outposts, customers can benefit from the fully managed infrastructure, services, APIs, and tools of AWS on premises, while still having access to the full range of AWS services available in the Region for a truly consistent hybrid experience³. References: On-Premises Private Cloud - AWS Outposts Family - AWS, What is AWS Outposts? - AWS Outposts

NEW QUESTION 168

- (Topic 2)

A company wants to migrate to the AWS Cloud. The company needs the ability to acquire resources when the resources are necessary. The company also needs the ability to release those resources when the resources are no longer necessary. Which architecture concept of the AWS Cloud meets these requirements?

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

Answer: A

Explanation:

The architecture concept of the AWS Cloud that meets the requirements of the company that wants to migrate to the AWS Cloud and needs the ability to acquire and release resources as needed is elasticity. Elasticity means that AWS customers can quickly and easily provision and scale up or down AWS resources as their demand changes, without any upfront costs or long-term commitments. AWS provides various tools and services that enable customers to achieve elasticity, such as Amazon EC2 Auto Scaling, Amazon CloudWatch, and AWS CloudFormation. Elasticity helps customers optimize their performance, availability, and cost efficiency. Availability, reliability, and durability are other architecture concepts of the AWS Cloud, but they are not directly related to the ability to

acquire and release resources as needed. Availability means that AWS customers can access their AWS resources and applications whenever and wherever they need them. Reliability means that AWS customers can depend on their AWS resources and applications to function correctly and consistently. Durability means that AWS customers can preserve their data and objects for long periods of time without loss or corruption¹²

NEW QUESTION 170

- (Topic 2)

A company wants to migrate its application to AWS. The company wants to replace upfront expenses with variable payment that is based on usage. What should the company do to meet these requirements?

- A. Use pay-as-you-go pricing.
- B. Purchase Reserved Instances.
- C. Pay less by using more.
- D. Rightsize instances.

Answer: A

Explanation:

Pay-as-you-go pricing is one of the main benefits of AWS. With pay-as-you-go pricing, you pay only for what you use, when you use it. There are no long-term contracts, termination fees, or complex licensing. You replace upfront expenses with lower variable costs and pay only for the resources you consume.

NEW QUESTION 172

- (Topic 2)

Which AWS service can a company use to securely store and encrypt passwords for a database?

- A. AWS Shield
- B. AWS Secrets Manager
- C. AWS Identity and Access Management (IAM)
- D. Amazon Cognito

Answer: B

Explanation:

AWS Secrets Manager is an AWS service that can be used to securely store and encrypt passwords for a database. It allows users to manage secrets, such as database credentials, API keys, and tokens, in a centralized and secure way. It also provides features such as automatic rotation, fine-grained access control, and auditing. AWS Shield is an AWS service that provides protection against Distributed Denial of Service (DDoS) attacks for AWS resources and services. It does not store or encrypt passwords for a database. AWS Identity and Access Management (IAM) is an AWS service that allows users to manage access to AWS resources and services. It can be used to create users, groups, roles, and policies that control who can do what in AWS. It does not store or encrypt passwords for a database. Amazon Cognito is an AWS service that provides user identity and data synchronization for web and mobile applications. It can be used to authenticate and authorize users, manage user profiles, and sync user data across devices. It does not store or encrypt passwords for a database.

NEW QUESTION 176

- (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 181

- (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 184

- (Topic 2)

Which AWS service is used to temporarily provide federated security credentials to a

- A. Amazon GuardDuty
- B. AWS Simple Token Service (AWS STS)
- C. AWS Secrets Manager
- D. AWS Certificate Manager

Answer: B

Explanation:

The AWS service that is used to temporarily provide federated security credentials to a user is AWS Security Token Service (AWS STS). AWS STS is a service that enables customers to request temporary, limited-privilege credentials for AWS Identity and Access Management (IAM) users or for users that they authenticate (federated users). The company can use AWS STS to grant federated users access to AWS resources without creating permanent IAM users or sharing long-term credentials. AWS STS helps customers manage and secure access to their AWS resources for federated users. Amazon GuardDuty, AWS Secrets Manager, and AWS Certificate Manager are not the best services to use for this purpose. Amazon GuardDuty is a threat detection service that monitors for malicious activity and unauthorized behavior across the AWS accounts and resources. AWS Secrets Manager is a service that helps customers manage and rotate secrets, such as database credentials, API keys, and passwords. AWS Certificate Manager is a service that helps customers provision, manage, and deploy public and private Secure Sockets Layer/Transport Layer Security (SSL/TLS) certificates for use with AWS services and internal connected resources. These services are more useful for different types of security and compliance tasks, rather than providing temporary federated security credentials to a user.

NEW QUESTION 187

- (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 troubleshooting.

NEW QUESTION 189

- (Topic 2)

A company is setting up AWS Identity and Access Management (IAM) on an AWS account. Which recommendation complies with IAM security best practices?

- A. Use the account root user access keys for administrative tasks.
- B. Grant broad permissions so that all company employees can access the resources they need.
- C. Turn on multi-factor authentication (MFA) for added security during the login process.
- D. Avoid rotating credentials to prevent issues in production applications.

Answer: C

Explanation:

C is correct because turning on multi-factor authentication (MFA) for added security during the login process is one of the IAM security best practices recommended by AWS. MFA adds an extra layer of protection on top of the user name and password, making it harder for attackers to access the AWS account. A is incorrect because using the account root user access keys for administrative tasks is not a good practice, as the root user has full access to all the resources in the AWS account and can cause irreparable damage if compromised. AWS recommends creating individual IAM users with the least privilege principle and using roles for applications that run on Amazon EC2 instances. B is incorrect because granting broad permissions so that all company employees can access the resources they need is not a good practice, as it increases the risk of unauthorized or accidental actions on the AWS resources. AWS recommends granting only the permissions that are required to perform a task and using groups to assign permissions to IAM users. D is incorrect because avoiding rotating credentials to prevent issues in production applications is not a good practice, as it increases the risk of credential leakage or compromise. AWS recommends rotating credentials regularly and using temporary security credentials from AWS STS when possible.

NEW QUESTION 190

- (Topic 2)

A company has an environment that includes Amazon EC2 instances, Amazon Lightsail, and on-premises servers. The company wants to automate the security updates for its operating systems and applications.

Which solution will meet these requirements with the LEAST operational effort?

- A. Use AWS Shield to identify and manage security events.
- B. Connect to each server by using a remote desktop connectio
- C. Run an update script.
- D. Use the AWS Systems Manager Patch Manager capability.
- E. Schedule Amazon GuardDuty to run on a nightly basis.

Answer: C

Explanation:

AWS Systems Manager Patch Manager is a capability that allows users to automate the security updates for their operating systems and applications. It enables users to scan their instances for missing patches, define patch baselines, schedule patching windows, and monitor patch compliance. It supports Amazon EC2 instances, Amazon Lightsail instances, and on-premises servers. AWS Shield is a service that provides protection against Distributed Denial of Service (DDoS) attacks for AWS resources and services. It does not automate the security updates for operating systems and applications. Connecting to each server by using a remote desktop connection and running an update script is a manual and time-consuming solution that requires a lot of operational effort. It is not a recommended best practice for automating the security updates for operating systems and applications. Amazon GuardDuty is a service that provides intelligent threat detection and continuous monitoring for AWS accounts and resources. It does not automate the security updates for operating systems and applications.

NEW QUESTION 195

- (Topic 2)

Which credential allows programmatic access to AWS resources for use from the AWS CLI or the AWS API?

- A. User name and password
- B. Access keys
- C. SSH public keys
- D. AWS Key Management Service (AWS KMS) keys

Answer: B

Explanation:

Access keys are long-term credentials that consist of an access key ID and a secret access key. You use access keys to sign programmatic requests that you make to AWS using the AWS CLI or AWS API¹. User name and password are credentials that you use to sign in to the AWS Management Console or the AWS Management Console mobile app². SSH public keys are credentials that you use to authenticate with EC2 instances that are launched from certain Linux AMIs³. AWS Key Management Service (AWS KMS) keys are customer master keys (CMKs) that you use to encrypt and decrypt your data and to control access to your data across AWS services and in your applications⁴.

NEW QUESTION 197

- (Topic 2)

A company needs to centralize its operational data. The company also needs to automate tasks across all of its Amazon EC2 instances.

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

- A. AWS Trusted Advisor
- B. AWS Systems Manager
- C. AWS CodeDeploy
- D. AWS Elastic Beanstalk

Answer: B

Explanation:

AWS Systems Manager is a service that enables users to centralize and automate the management of their AWS resources. It provides a unified user interface to view operational data, such as inventory, patch compliance, and performance metrics. It also allows users to automate common and repetitive tasks, such as patching, backup, and configuration management, across all of their Amazon EC2 instances¹. AWS Trusted Advisor is a service that provides best practices and recommendations to optimize the performance, security, and cost of AWS resources². AWS CodeDeploy is a service that automates the deployment of code and applications to Amazon EC2 instances or other compute services³. AWS Elastic Beanstalk is a service that simplifies the deployment and management of web applications using popular platforms, such as Java, PHP, and Node.js⁴.

NEW QUESTION 201

- (Topic 2)

A company is running an application on AWS. The company wants to identify and prevent the accidental

Which AWS service or feature will meet these requirements?

- A. Amazon GuardDuty
- B. Network ACL
- C. AWS WAF
- D. AWS Network Firewall

Answer: A

Explanation:

Amazon GuardDuty is a threat detection service that continuously monitors for malicious activity and unauthorized behavior to protect your AWS accounts, workloads, and data stored in Amazon S3. With the cloud, the collection and aggregation of account and network activities is simplified, but it can be time consuming for security teams to continuously analyze event log data for potential threats. With GuardDuty, you can automate anomaly detection and get actionable findings to help you protect your AWS resources⁴.

NEW QUESTION 202

- (Topic 2)

A company needs to design a solution for the efficient use of compute resources for an enterprise workload. The company needs to make informed decisions as its technology needs evolve.

Which pillar of the AWS Well-Architected Framework do these requirements represent?

- A. Operational excellence
- B. Performance efficiency
- C. Cost optimization
- D. Reliability

Answer: B

Explanation:

Performance efficiency is the pillar of the AWS Well-Architected Framework that represents the requirements of designing a solution for the efficient use of compute resources for an enterprise workload and making informed decisions as the technology needs evolve. It focuses on using the right resources and services for the workload, monitoring performance, and continuously improving the efficiency of the solution. Operational excellence is the pillar of the AWS Well-Architected Framework that represents the ability to run and monitor systems to deliver business value and to continually improve supporting processes and procedures. Cost optimization is the pillar of the AWS Well-Architected Framework that represents the ability to run systems to deliver business value at the lowest price point. Reliability is the pillar of the AWS Well-Architected Framework that represents the ability of a system to recover from infrastructure or service disruptions, dynamically acquire computing resources to meet demand, and mitigate disruptions such as misconfigurations or transient network issues.

NEW QUESTION 205

- (Topic 2)

A user is moving a workload from a local data center to an architecture that is distributed between the local data center and the AWS Cloud.

Which type of migration is this?

- A. On-premises to cloud native
- B. Hybrid to cloud native
- C. On-premises to hybrid
- D. Cloud native to hybrid

Answer: C

Explanation:

C is correct because moving a workload from a local data center to an architecture that is distributed between the local data center and the AWS Cloud is an example of an on-premises to hybrid migration. A hybrid cloud is a cloud computing environment that uses a mix of on-premises, private cloud, and public cloud services with orchestration between the platforms. A is incorrect because on-premises to cloud native migration is the process of moving a workload from a local data center to an architecture that is fully hosted and managed on the AWS Cloud. B is incorrect because hybrid to cloud native migration is the process of moving a workload from an architecture that is distributed between the local data center and the AWS Cloud to an architecture that is fully hosted and managed on the AWS Cloud. D is incorrect because cloud native to hybrid migration is the process of moving a workload from an architecture that is fully hosted and managed on the AWS Cloud to an architecture that is distributed between the local data center and the AWS Cloud.

NEW QUESTION 207

- (Topic 2)

A company wants its workload to perform consistently and correctly. Which benefit of AWS Cloud computing does this goal represent?

- A. Security
- B. Elasticity
- C. Pay-as-you-go pricing
- D. Reliability

Answer: D

Explanation:

Reliability is the benefit of AWS Cloud computing that ensures the workload performs consistently and correctly. According to the AWS Cloud Practitioner Essentials course, reliability means "the ability of a system to recover from infrastructure or service disruptions, dynamically acquire computing resources to meet demand, and mitigate disruptions such as misconfigurations or transient network issues."¹ Elasticity, security, and pay-as-you-go pricing are also benefits of AWS Cloud computing, but they do not directly relate to the goal of consistent and correct performance.

NEW QUESTION 210

- (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 213

- (Topic 2)

A large company has multiple departments. Each department has its own AWS account. Each department has purchased Amazon EC2 Reserved Instances. Some departments do not use all the Reserved Instances that they purchased, and other departments need more Reserved Instances than they purchased. The company needs to manage the AWS accounts for all the departments so that the departments can share the Reserved Instances. Which AWS service or tool should the company use to meet these requirements?

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

Answer: D

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 apply service control policies (SCPs) across multiple AWS accounts to restrict what services and actions users and roles can access. You can also use AWS Organizations to enable features such as consolidated billing, AWS Config rules and conformance packs, and AWS CloudFormation StackSets across multiple accounts³. One of the benefits of using AWS Organizations is that you can share your Reserved Instances (RIs) with all of the accounts in your organization. This enables you to take advantage of the billing benefits of RIs without having to specify which account will use them⁴. AWS Systems Manager is a service that gives you visibility and control of your infrastructure on AWS. Cost Explorer is a tool that enables you to visualize, understand, and manage your AWS costs and usage over time. 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 can help you manage the AWS accounts for all the departments so that the departments can share the Reserved Instances.

NEW QUESTION 217

- (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 219

- (Topic 2)

A company is planning a migration to the AWS Cloud and wants to examine the costs that are associated with different workloads. Which AWS tool will meet these requirements?

- A. AWS Budgets
- B. AWS Cost Explorer
- C. AWS Pricing Calculator
- D. AWS Cost and Usage Report

Answer: C

Explanation:

The AWS tool that will meet the requirements of the company that is planning a migration to the AWS Cloud and wants to examine the costs that are associated with different workloads is AWS Pricing Calculator. AWS Pricing Calculator is a tool that helps customers estimate the cost of using AWS services based on their requirements and preferences. The company can use AWS Pricing Calculator to compare the costs of different AWS services and configurations, such as Amazon EC2, Amazon S3, Amazon RDS, and more. AWS Pricing Calculator also provides detailed breakdowns of the cost components, such as compute, storage, network, and data transfer. AWS Pricing Calculator helps customers plan and optimize their cloud budget and migration strategy. AWS Budgets, AWS Cost Explorer, and AWS Cost and Usage Report are not the best tools to use for this purpose. AWS Budgets is a tool that helps customers monitor and manage their AWS spending and usage against predefined budget limits and thresholds. AWS Cost Explorer is a tool that helps customers analyze and visualize their AWS spending and usage trends over time. AWS Cost and Usage Report is a tool that helps customers access comprehensive and granular information about their AWS costs and usage in a CSV or Parquet file. These tools are more useful for tracking and optimizing the existing AWS costs and usage, rather than estimating the costs of different workloads³⁴.

NEW QUESTION 224

- (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 interruptions. 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 225

- (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 227

- (Topic 2)

A new AWS user who has little cloud experience wants to build an application by using AWS services. The user wants to learn how to implement specific AWS services from other customer examples. The user also wants to ask questions to AWS experts.

Which AWS service or resource will meet these requirements?

- A. AWS Online Tech Talks
- B. AWS documentation
- C. AWS Marketplace
- D. AWS Health Dashboard

Answer: A

Explanation:

AWS Online Tech Talks are online presentations that cover a broad range of topics at varying technical levels and provide a live Q&A session with AWS experts. They are a great resource for new AWS users who want to learn how to implement specific AWS services from other customer examples and ask questions to AWS experts. AWS documentation, AWS Marketplace, and AWS Health Dashboard do not offer the same level of interactivity and guidance as AWS Online Tech Talks. Source: AWS Online Tech Talks

NEW QUESTION 231

- (Topic 2)

Which options are common stakeholders for the AWS Cloud Adoption Framework (AWS CAF) platform perspective? (Select TWO.)

- A. Chief financial officers (CFOs)
- B. IT architects
- C. Chief information officers (CIOs)
- D. Chief data officers (CDOs)
- E. Engineers

Answer: BE

Explanation:

The common stakeholders for the AWS Cloud Adoption Framework (AWS CAF) platform perspective are IT architects and engineers. The AWS CAF is a guidance that helps organizations design and travel an accelerated path to successful cloud adoption. The AWS CAF organizes the cloud adoption process into six areas of focus, called perspectives, which are business, people, governance, platform, security, and operations. Each perspective is divided into capabilities, which are further divided into skills and responsibilities. The platform perspective focuses on the provisioning and management of the cloud infrastructure and services that support the business applications. The platform perspective capabilities are design, implementation, and optimization. The stakeholders for the platform perspective are the IT architects and engineers who are responsible for designing, implementing, and optimizing the cloud platform. Chief financial officers (CFOs), chief information officers (CIOs), and chief data officers (CDOs) are not the common stakeholders for the AWS CAF platform perspective. CFOs are the common stakeholders for the AWS CAF business perspective, which focuses on the value realization of the cloud adoption. CIOs are the common stakeholders for the AWS CAF governance perspective, which focuses on the alignment of the IT strategy and processes with the business strategy and goals. CDOs are the common stakeholders for the AWS CAF security perspective, which focuses on the protection of the information assets and systems in the cloud.

NEW QUESTION 234

- (Topic 2)

A company wants to improve its security and audit posture by limiting Amazon EC2 inbound access. According to the AWS shared responsibility model, which task is the responsibility of the customer?

- A. Protect the global infrastructure that runs all of the services offered in the AWS Cloud.
- B. Configure logical access controls for resources, and protect account credentials.
- C. Configure the security used by managed services.
- D. Patch and back up Amazon Aurora.

Answer: B

Explanation:

According to the AWS shared responsibility model, the customer is responsible for configuring logical access controls for resources, and protecting account credentials. This includes managing IAM user permissions, security group rules, network ACLs, encryption keys, and other aspects of access management¹. AWS is responsible for protecting the global infrastructure that runs all of the services offered in the AWS Cloud, such as the hardware, software, networking, and facilities. AWS is also responsible for configuring the security used by managed services, such as Amazon RDS, Amazon DynamoDB, and Amazon Aurora².

NEW QUESTION 236

- (Topic 2)

A company wants guidance to optimize the cost and performance of its current AWS environment. Which AWS service or tool should the company use to identify areas for optimization?

- A. Amazon QuickSight
- B. AWS Trusted Advisor
- C. AWS Organizations
- D. AWS Budgets

Answer: B

Explanation:

AWS Trusted Advisor is the AWS service or tool that the company should use to identify areas for optimization. According to the AWS Trusted Advisor User Guide, "AWS Trusted Advisor is an online tool that provides you real time guidance to help you provision your resources following AWS best practices. AWS Trusted Advisor checks help optimize your AWS infrastructure, increase security and performance, reduce your overall costs, and monitor service limits." Amazon QuickSight, AWS Organizations, and AWS Budgets are not designed to provide optimization recommendations for the current AWS environment.

NEW QUESTION 241

- (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 243

- (Topic 2)

Which group shares responsibility with AWS for security and compliance of AWS accounts and resources?

- A. Third-party vendors
- B. Customers
- C. Reseller partners
- D. Internet providers

Answer: B

Explanation:

Customers share responsibility with AWS for security and compliance of AWS accounts and resources. This is part of the AWS shared responsibility model, which 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 physical and environmental controls of the AWS global infrastructure, such as power, cooling, fire suppression, and physical access. The customer is responsible for the security in the cloud, which includes the configuration and management of the AWS resources and applications, such as identity and access management, encryption, firewall, and backup.

For more information, see AWS Shared Responsibility Model and AWS Cloud Security.

NEW QUESTION 244

- (Topic 2)

Which AWS service is designed to help users orchestrate a workflow process for a set of AWS Lambda functions?

- A. Amazon DynamoDB
- B. AWS CodePipeline
- C. AWS Batch
- D. AWS Step Functions

Answer: D

Explanation:

The AWS service that is designed to help users orchestrate a workflow process for a set of AWS Lambda functions is AWS Step Functions. AWS Step Functions is a service that helps users coordinate multiple AWS services into serverless workflows that can be triggered by events, such as messages, API calls, or schedules. AWS Step Functions allows users to create and visualize complex workflows that can include branching, parallel execution, error handling, retries, and timeouts. AWS Step Functions can integrate with AWS Lambda to orchestrate a sequence of Lambda functions that perform different tasks or logic. Amazon DynamoDB, AWS CodePipeline, and AWS Batch are not the best services to use for orchestrating a workflow process for a set of AWS Lambda functions. Amazon DynamoDB is a fully managed NoSQL database service that provides fast and consistent performance, scalability, and flexibility. AWS CodePipeline is a fully managed continuous delivery service that helps users automate the release process of their applications. AWS Batch is a fully managed service that helps users run batch computing workloads on the AWS Cloud.

NEW QUESTION 246

- (Topic 2)

A company has an application workload that is stateless by design and can sustain occasional downtime. The application performs massively parallel computations.

Which Amazon EC2 pricing model should the company choose for its application to reduce cost?

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

Answer: B

Explanation:

Amazon EC2 Spot Instances let you take advantage of unused EC2 capacity in the AWS cloud. Spot Instances are available at up to a 90% discount compared to On-Demand prices. You can use Spot Instances for various stateless, fault-tolerant, or flexible applications such as big data, containerized workloads, CI/CD, web servers, high-performance computing (HPC), and other test & development workloads. Spot Instances are well-suited for massively parallel computations, as they can provide large amounts of compute capacity at a low cost, and can be interrupted with a two-minute notice.

NEW QUESTION 251

- (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 252

- (Topic 2)

How should the company deploy the application to meet these requirements?

- A. In a single Availability Zone
- B. On AWS Direct Connect
- C. On Reserved Instances
- D. In multiple Availability Zones

Answer: D

Explanation:

Deploying the application in multiple Availability Zones is the best way to ensure high availability for the application. Availability Zones are isolated locations within an AWS Region that are engineered to be fault-tolerant from failures in other Availability Zones. By deploying the application in multiple Availability Zones, the company can reduce the impact of outages and increase the resilience of the application. Deploying the application in a single Availability Zone, on AWS Direct Connect, or on Reserved Instances does not provide the same level of high availability as deploying the application in multiple Availability Zones. Source: Availability Zones

NEW QUESTION 253

- (Topic 2)

A company moves a workload to AWS to run on Amazon EC2 instances. The company needs to run the workload in the most cost-effective way. What can the company do to meet this requirement?

- A. Use AWS Key Management Service (AWS KMS).
- B. Use multiple AWS accounts and consolidated billing.
- C. Use AWS CloudFormation to deploy the infrastructure.
- D. Rightsized all the EC2 instances that are used in the deployment.

Answer: D

Explanation:

Rightsizing all the EC2 instances that are used in the deployment is the best way to run the workload in the most cost-effective way. Rightsizing means choosing the optimal instance type and size for the workload based on the performance and capacity requirements. Rightsizing helps to avoid over-provisioning or under-provisioning of the EC2 instances, which can result in wasted resources or poor performance. Rightsizing also helps to take advantage of the different pricing models and features that AWS offers, such as On-Demand, Reserved, and Spot Instances, and Auto Scaling. For more information, see [Rightsizing Your Instances and \[Cost Optimization with AWS\]](#).

NEW QUESTION 257

- (Topic 2)

A company has a single Amazon EC2 instance. The company wants to adopt a highly available architecture. What can the company do to meet this requirement?

- A. Scale vertically to a larger EC2 instance size.
- B. Scale horizontally across multiple Availability Zones.
- C. Purchase an EC2 Dedicated Instance.
- D. Change the EC2 instance family to a compute optimized instance.

Answer: B

Explanation:

Scaling horizontally across multiple Availability Zones is a way to adopt a highly available architecture, as it increases the fault tolerance and resilience of the application. Scaling vertically to a larger EC2 instance size is a way to improve the performance of the application, but it does not improve the availability. Purchasing an EC2 Dedicated Instance is a way to isolate the instance from other AWS customers, but it does not improve the availability. Changing the EC2 instance family to a compute optimized instance is a way to optimize the instance type for the workload, but it does not improve the availability. These concepts are explained in the [AWS Well-Architected Framework](#)².

NEW QUESTION 258

- (Topic 2)

What is a characteristic of Convertible Reserved Instances (RIs)?

- A. Users can exchange Convertible RIs for other Convertible RIs from a different instance family.
- B. Users can exchange Convertible RIs for other Convertible RIs in different AWS Regions.
- C. Users can sell and buy Convertible RIs on the AWS Marketplace.
- D. Users can shorten the term of their Convertible RIs by merging them with other Convertible RIs.

Answer: A

Explanation:

Convertible Reserved Instances (RIs) are a type of Reserved Instance that allow you to change the attributes of the RI as long as the exchange results in the creation of Reserved Instances of equal or greater value. You can exchange Convertible RIs for other Convertible RIs from a different instance family, size, platform, tenancy, or scope (Region or Availability Zone)³.

NEW QUESTION 260

- (Topic 2)

A company that is planning to migrate to the AWS Cloud is based in an isolated area that has limited internet connectivity. The company needs to perform local data processing on premises. The company needs a solution that can operate without a stable internet connection. Which AWS service will meet these requirements?

- A. Amazon S3
- B. AWS Snowball Edge
- C. AWS StorageGateway
- D. AWS Backup

Answer: B

Explanation:

AWS Snowball Edge is a service that provides a physical device that can store up to 100 TB of data and perform local data processing on premises. It enables users to transfer data to and from the AWS Cloud in areas with limited or no internet connectivity. It also supports AWS Greengrass, which allows users to run AWS Lambda functions and other AWS services locally without a stable internet connection. Amazon S3 is a storage service that provides scalable, durable, and secure object storage. It requires a stable internet connection to transfer data to and from the AWS Cloud. AWS Storage Gateway is a service that provides a hybrid storage solution that connects on-premises applications to AWS Cloud storage services, such as Amazon S3, Amazon S3 Glacier, and Amazon EBS. It requires a stable internet connection to synchronize data between the on-premises and cloud storage. AWS Backup is a service that provides a centralized and automated solution to back up data across AWS services and on-premises resources. It requires a stable internet connection to transfer data to and from the AWS Cloud.

NEW QUESTION 263

- (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 267

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your AWS-Certified-Cloud-Practitioner Exam with Our Prep Materials Via below:

<https://www.certleader.com/AWS-Certified-Cloud-Practitioner-dumps.html>