

Amazon-Web-Services

Exam Questions CLF-C01

AWS Certified Cloud Practitioner



NEW QUESTION 1

- (Topic 3)

A company wants to monitor for misconfigured security groups that are allowing unrestricted access to specific ports. Which AWS service will meet this requirement?

- A. AWS Trusted Advisor
- B. Amazon CloudWatch
- C. Amazon GuardDuty
- D. AWS Health Dashboard

Answer: A

Explanation:

AWS Trusted Advisor is an online tool that provides you real time guidance to help you provision your resources following AWS best practices, including security and performance. It can help you monitor for misconfigured security groups that are allowing unrestricted access to specific ports. Amazon CloudWatch is a service that monitors your AWS resources and the applications you run on AWS. Amazon GuardDuty is a threat detection service that continuously monitors for malicious activity and unauthorized behavior. AWS Health Dashboard provides relevant and timely information to help you manage events in progress, and provides proactive notification to help you plan for scheduled activities.

NEW QUESTION 2

- (Topic 3)

A company has all of its servers in the us-east-1 Region. The company is considering the deployment of additional servers different Region. Which AWS tool should the company use to find pricing information for other Regions?

- A. Cost Explorer
- B. AWS Budgets
- C. AWS Purchase Order Management
- D. AWS Pricing Calculator

Answer: D

Explanation:

AWS Pricing Calculator lets customers explore AWS services, and create an estimate for the cost of their use cases on AWS. AWS Pricing Calculator can also compare the costs of different AWS Regions and configurations. Cost Explorer is a tool that enables customers to visualize, understand, and manage their AWS costs and usage over time.

AWS Budgets gives customers the ability to set custom budgets that alert them when their costs or usage exceed (or are forecasted to exceed) their budgeted amount. AWS Purchase Order Management is a feature that allows customers to pay for their AWS invoices using purchase orders.

NEW QUESTION 3

- (Topic 3)

Which option is a customer responsibility under the AWS shared responsibility model?

- A. Maintenance of underlying hardware of Amazon EC2 instances
- B. Application data security
- C. Physical security of data centers
- D. Maintenance of VPC components

Answer: B

Explanation:

The option that is a customer responsibility under the AWS shared responsibility model is B. Application data security.

According to the AWS shared responsibility model, AWS is responsible for the security of the cloud, while the customer is responsible for the security in the cloud. This means that AWS manages the security of the underlying infrastructure, such as the hardware, software, networking, and facilities that run the AWS services, while the customer manages the security of their applications, data, and resources that they use on top of AWS¹². Application data security is one of the customer responsibilities under the AWS shared responsibility model. This means that the customer is responsible for protecting their application data from unauthorized access, modification, deletion, or leakage. The customer can use various AWS services and features to help with application data security, such as encryption, key management, access control, logging, and auditing¹². Maintenance of underlying hardware of Amazon EC2 instances is not a customer responsibility under the AWS shared responsibility model. This is part of the AWS responsibility to secure the cloud. AWS manages the physical servers that host the Amazon EC2 instances and ensures that they are updated, patched, and replaced as needed¹³.

Physical security of data centers is not a customer responsibility under the AWS shared responsibility model. This is also part of the AWS responsibility to secure the cloud. AWS operates and controls the facilities where the AWS services are hosted and ensures that they are protected from unauthorized access, environmental hazards, fire, and theft¹⁴. Maintenance of VPC components is not a customer responsibility under the AWS shared responsibility model. This is a shared responsibility between AWS and the customer. AWS provides the VPC service and ensures that it is secure and reliable, while the customer configures and manages their own VPCs and related components, such as subnets, route tables, security groups, network ACLs, gateways, and endpoints¹⁵.

References:

1: Shared Responsibility Model - Amazon Web Services (AWS) 2: AWS Cloud Computing - W3Schools 3: [Amazon EC2 FAQs - Amazon Web Services] 4: [AWS Security - Amazon Web Services] 5: [Amazon Virtual Private Cloud (VPC) - Amazon Web Services]

NEW QUESTION 4

- (Topic 3)

A company wants high levels of detection and near-real-time (NRT) mitigation against large and sophisticated distributed denial of service (DDoS) attacks on applications running on AWS.

Which AWS service should the company use?

- A. Amazon GuardDuty
- B. Amazon Inspector
- C. AWS Shield Advanced
- D. Amazon Macie

Answer: C

Explanation:

AWS Shield Advanced is a service that provides high levels of detection and near-real-time (NRT) mitigation against large and sophisticated distributed denial of service (DDoS) attacks on applications running on AWS. AWS Shield Advanced also provides you with 24x7 access to the AWS DDoS Response Team (DRT) and protection against DDoS attacks of any size or duration¹. Amazon GuardDuty is a service that provides threat detection for your AWS accounts and workloads, but it does not offer DDoS protection³. Amazon Inspector is a service that helps you improve the security and compliance of your applications deployed on AWS by automatically assessing them for vulnerabilities and deviations from best practices. Amazon Macie is a service that uses machine learning and pattern matching to discover and protect your sensitive data in AWS.

NEW QUESTION 5

- (Topic 3)

A company wants to migrate its on-premises workloads to the AWS Cloud. The company wants to separate workloads for chargeback to different departments. Which AWS services or features will meet these requirements? (Select TWO.)

- A. Placement groups
- B. Consolidated billing
- C. Edge locations
- D. AWS Config
- E. Multiple AWS accounts

Answer: BE

Explanation:

Consolidated billing is a feature of AWS Organizations that enables customers to consolidate billing and payment for multiple AWS accounts. With consolidated billing, customers can group multiple AWS accounts under one payer account, making it easier to manage billing and track costs across multiple accounts. Consolidated billing also offers benefits such as volume discounts, Reserved Instance discounts, and Savings Plans discounts. Consolidated billing is offered at no additional cost.

Multiple AWS accounts is a feature of AWS Organizations that enables customers to create and manage multiple AWS accounts from a central location. With multiple AWS accounts, customers can isolate workloads for different departments, projects, or environments, and apply granular access controls and policies to each account. Multiple AWS accounts also helps customers improve security, compliance, and governance of their AWS resources⁵⁶. References: 5:

Consolidated billing for AWS Organizations - AWS

Billing, 6: Understanding Consolidated Bills - AWS Billing, 7: AWS Consolidated Billing: Tutorial & Best Practices, 8: Simplifying Your Bills With Consolidated Billing on AWS - Aimably, 9: AWS Consolidated Billing - W3Schools

NEW QUESTION 6

- (Topic 3)

An ecommerce company has migrated its IT infrastructure from an on-premises data center to the AWS Cloud. Which cost is the company's direct responsibility?

- A. Cost of application software licenses
- B. Cost of the hardware infrastructure on AWS
- C. Cost of power for the AWS servers
- D. Cost of physical security for the AWS data center

Answer: A

Explanation:

The cost of application software licenses is the company's direct responsibility when it migrates its IT infrastructure from an on-premises data center to the AWS Cloud. Application software licenses are the agreements that grant users the right to use specific software products, such as operating systems, databases, or applications. Depending on the type and terms of the license, users may need to pay a fee to the software vendor or provider to use the software legally and access its features and updates. When users migrate their IT infrastructure to the AWS Cloud, they can choose to buy new licenses from AWS, bring their own licenses (BYOL), or use a combination of both. However, regardless of the option they choose, they are still responsible for complying with the license terms and paying the license fees to the software vendor or provider. AWS does not charge users for the application software licenses they bring or buy, but only for the AWS resources they use to run their applications. Therefore, the cost of application software licenses is the only cost among the options that is the company's direct responsibility. The other costs are either included in the AWS service fees or covered by AWS.

References: AWS License Manager Pricing, Software licensing: The blind spot in public cloud costs, Cost Optimization tips for SQL Server Licenses on AWS, Microsoft Licensing on AWS

NEW QUESTION 7

- (Topic 3)

Which AWS service or storage class provides low-cost, long-term data storage?

- A. Amazon S3 Glacier Deep Archive
- B. AWS Snowball
- C. Amazon MQ
- D. AWS Storage Gateway

Answer: A

Explanation:

Amazon S3 Glacier Deep Archive is a storage class within Amazon S3 that provides the lowest-cost, long-term data storage for data that is rarely accessed. AWS Snowball is a service that provides a physical device for transferring large amounts of data into and out of AWS. Amazon MQ is a service that provides managed message broker service for Apache ActiveMQ. AWS Storage Gateway is a service that provides hybrid cloud storage for on-premises applications.

NEW QUESTION 8

- (Topic 3)

Which of the following is a benefit that AWS Professional Services provides?

- A. Management of the ongoing security of user data

- B. Advisory solutions for AWS adoption
- C. Technical support 24 hours a day, 7 days a week
- D. Monitoring of monthly billing costs in AWS accounts

Answer: B

Explanation:

AWS Professional Services is a team of experts that help customers achieve their desired outcomes using the AWS Cloud. One of the benefits that AWS Professional Services provides is advisory solutions for AWS adoption, which include guidance on cloud strategy, architecture, migration, and innovation². Management of the ongoing security of user data, technical support 24 hours a day, 7 days a week, and monitoring of monthly billing costs in AWS accounts are not benefits that AWS Professional Services provides, as they are either the responsibility of the customer or the features of other AWS services or support plans³

NEW QUESTION 9

- (Topic 3)

Which AWS service or feature offers security for a VPC by acting as a firewall to control traffic in and out of subnets?

- A. AWS Security Hub
- B. Security groups
- C. Network ACL
- D. AWSWAF

Answer: C

Explanation:

A network access control list (network ACL) is a feature that acts as a firewall for controlling traffic in and out of one or more subnets in a virtual private cloud (VPC). AWS Security Hub is a service that provides a comprehensive view of the security posture of AWS accounts and resources. Security groups are features that act as firewalls for controlling traffic at the instance level. AWS WAF is a web application firewall that helps protect web applications from common web exploits.

NEW QUESTION 10

- (Topic 3)

A company wants to receive alerts to monitor its overall operating costs for its AWS public cloud infrastructure.

Which AWS offering will meet these requirements?

- A. Amazon EventBridge
- B. Compute Savings Plans
- C. AWS Budgets
- D. Migration Evaluator

Answer: C

Explanation:

AWS Budgets is a service that enables you to plan your service usage, service costs, and instance reservations. You can use AWS Budgets to create 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 monitor how close your usage and costs are to meeting your reservation purchases¹

NEW QUESTION 10

- (Topic 3)

A company is storing sensitive customer data in an Amazon S3 bucket. The company wants to protect the data from accidental deletion or overwriting.

Which S3 feature should the company use to meet these requirements?

- A. S3 Lifecycle rules
- B. S3 Versioning
- C. S3 bucket policies
- D. S3 server-side encryption

Answer: B

Explanation:

S3 Versioning is a feature that allows you to keep multiple versions of an object in the same bucket. You can use S3 Versioning to protect your data from accidental deletion or overwriting by enabling it on a bucket or a specific object. S3 Versioning also allows you to restore previous versions of an object if needed. S3 Lifecycle rules are used to automate the transition of objects between storage classes or to expire objects after a certain period of time. S3 bucket policies are used to control access to the objects in a bucket. S3 server-side encryption is used to encrypt the data at rest in S3. References: S3 Versioning, S3 Lifecycle rules, S3 bucket policies, S3 server-side encryption

NEW QUESTION 11

- (Topic 3)

Which AWS service gives users the ability to discover and protect sensitive data that is stored in Amazon S3 buckets?

- A. Amazon Macie
- B. Amazon Detective
- C. Amazon GuardDuty
- D. AWS IAM Access Analyzer

Answer: A

Explanation:

Amazon Macie is a data security and privacy service offered by AWS that uses machine learning and pattern matching to discover the sensitive data stored within Amazon S3. You can define your own custom type of sensitive data category that might be unique to your business or use case. Macie also provides you with dashboards and alerts that give you visibility into how your data is being accessed or moved. Macie helps you protect your data by enabling you to apply data protection techniques such as encryption, deletion, access control, and auditing. References: Strengthen the security of sensitive data stored in Amazon S3 by using additional AWS services, Security best practices for Amazon S3, Sensitive Data Protection on AWS, Sensitive Data Protection on Amazon Web Services

NEW QUESTION 13

- (Topic 3)

Which AWS service can provide a dedicated network connection with consistent low latency from on premises to the AWS Cloud?

- A. Amazon VPC
- B. Amazon Kinesis Data Streams
- C. AWS Direct Connect
- D. Amazon OpenSearch Service

Answer: C

Explanation:

AWS Direct Connect is a service that provides a dedicated network connection from on premises to the AWS Cloud. It can reduce network costs, increase bandwidth throughput, and provide a more consistent network experience than internet-based connections. It can also provide low latency for applications that require real-time data transfer. Amazon VPC is a service that provides a logically isolated section of the AWS Cloud where users can launch AWS resources in a virtual network that they define. Amazon Kinesis Data Streams is a service that provides a scalable and durable stream of data records for real-time data processing. Amazon OpenSearch Service is a service that provides a fully managed, scalable, and secure search and analytics solution that is compatible with Elasticsearch.

NEW QUESTION 18

- (Topic 3)

A team of researchers is going to collect data at remote locations around the world. Many locations do not have internet connectivity. The team needs to capture the data in the field, and transfer it to the AWS Cloud later. Which AWS service will support these requirements?

- A. AWS Outposts
- B. AWS Transfer Family
- C. AWS Snow Family
- D. AWS Migration Hub

Answer: C

Explanation:

AWS Snow Family is a group of devices that transport data in and out of AWS. AWS Snow Family devices are physical devices that can transfer up to exabytes of data. One exabyte is 1 000 000 000 000 megabytes. AWS Snow Family devices are designed for use in remote locations where internet connectivity is limited or unavailable. You can use these devices to collect and process data at the edge, and then ship them back to AWS for data upload. AWS Snow Family consists of three types of devices: AWS Snowcone, AWS Snowball, and AWS Snowmobile. References: 1: Edge Computing Devices, Secure Data Transfer - AWS Snow Family - AWS, 2: AWS Snow Family Documentation, 3: AWS Snow Family - W3Schools, 4: AWS Snow Family: Data Storage, Migration, and Computation

NEW QUESTION 19

- (Topic 3)

Which AWS service provides storage that can be mounted across multiple Amazon EC2 instances?

- A. Amazon Workspaces
- B. Amazon Elastic File System (Amazon EFS)
- C. AWS Database Migration Service (AWS DMS)
- D. AWS Snowball Edge

Answer: B

Explanation:

Amazon EFS is a fully managed service that provides scalable and elastic file storage for multiple Amazon EC2 instances. Amazon EFS supports the Network File System (NFS) protocol, which allows multiple EC2 instances to access the same file system concurrently. You can learn more about Amazon EFS from this webpage or this digital course.

NEW QUESTION 20

- (Topic 3)

A company wants to integrate natural language processing (NLP) into business intelligence (BI) dashboards. The company wants to ask questions and receive answers with relevant visualizations.

Which AWS service or tool will meet these requirements?

- A. Amazon Macie
- B. Amazon Rekognition
- C. Amazon QuickSight Q
- D. Amazon Lex

Answer: C

Explanation:

Amazon QuickSight Q is a natural language query feature that allows users to ask questions about their data and receive answers in the form of relevant visualizations. Amazon Macie is a data security and data privacy service that uses machine learning and pattern matching to discover and protect sensitive data in AWS. Amazon Rekognition is a computer vision service that can analyze images and videos for faces, objects, scenes, text, and more. Amazon Lex is a

service for building conversational interfaces using voice and text4.

NEW QUESTION 24

- (Topic 3)

A company has created an AWS Cost and Usage Report and wants to visualize the report. Which AWS service should the company use to ingest and display this information?

- A. Amazon QuickSight
- B. Amazon Pinpoint
- C. Amazon Neptune
- D. Amazon Kinesis

Answer: A

Explanation:

Amazon QuickSight is an AWS service that provides business intelligence and data visualization capabilities. Amazon QuickSight enables you to ingest, analyze, and display data from various sources, such as AWS Cost and Usage Reports, Amazon S3, Amazon Athena, Amazon Redshift, and Amazon RDS. You can use Amazon QuickSight to create interactive dashboards and charts that show insights and trends from your data. You can also share your dashboards and charts with other users or embed them into your applications.

NEW QUESTION 26

- (Topic 3)

A company has a physical tape library to store data backups. The tape library is running out of space. The company needs to extend the tape library's capacity to the AWS Cloud.

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

- A. Amazon Elastic File System (Amazon EFS)
- B. Amazon Elastic Block Store (Amazon EBS)
- C. Amazon S3
- D. AWS Storage Gateway

Answer: D

Explanation:

AWS Storage Gateway is a hybrid cloud storage service that provides on- premises access to virtually unlimited cloud storage. You can use AWS Storage Gateway to simplify storage management and reduce costs for key hybrid cloud storage use cases. One of these use cases is tape-based backup, which allows you to store data backups on virtual tapes in the AWS Cloud. You can use the Tape Gateway feature of AWS Storage Gateway to extend your existing physical tape library to the AWS Cloud. Tape Gateway provides a virtual tape infrastructure that scales seamlessly with your backup needs and eliminates the operational burden of provisioning, scaling, and maintaining a physical tape infrastructure¹²³. References: 1: Cloud Storage Appliances, Hybrid Device - AWS Storage Gateway - AWS, 2: AWS Storage Gateway Documentation, 3: AWS Storage Gateway Features | Amazon Web Services

NEW QUESTION 29

- (Topic 3)

Which AWS services or features give users the ability to create a network connection between two VPCs? (Select TWO.)

- A. VPC endpoints
- B. Amazon Route 53
- C. VPC peering
- D. AWS Direct Connect
- E. AWS Transit Gateway

Answer: CE

Explanation:

VPC peering and AWS Transit Gateway are two AWS services or features that give users the ability to create a network connection between two VPCs. VPC peering is a networking connection between two VPCs that enables you to route traffic between them privately. You can create a VPC peering connection between your own VPCs, with a VPC in another AWS account, or with a VPC in a different AWS Region. Traffic between peered VPCs never traverses the public internet. VPC peering does not support transitive peering relationships, which means that if VPC A is peered with VPC B, and VPC B is peered with VPC C, then VPC A and VPC C are not automatically peered⁷⁸⁹. AWS Transit Gateway is a networking service that acts as a regional router for your VPCs and on- premises networks. You can attach up to 5,000 VPCs and VPN connections to a single transit gateway and route traffic between them. AWS Transit Gateway simplifies the management and scalability of your network architecture, as you only need to create and manage a single connection from the central transit gateway to each connected network. AWS Transit Gateway supports transitive routing, which means that any network that is attached to the transit gateway can communicate with any other network that is attached to the same transit gateway . References: 7: VPC peering - Amazon Virtual Private Cloud, 8: Connect VPCs using VPC peering - Amazon Virtual Private Cloud, 9: Amazon VPC-to-Amazon VPC connectivity options - Amazon Virtual Private Cloud, : [AWS Transit Gateway - Amazon Web Services], : [Connect VPCs using AWS Transit Gateway - Amazon Virtual Private Cloud], : [AWS Transit Gateway: Simplify Your Network Architecture]

NEW QUESTION 32

- (Topic 3)

An IT engineer needs to access AWS services from an on-premises application. Which credentials or keys does the application need for authentication?

- A. AWS account user name and password
- B. IAM access key and secret
- C. Amazon EC2 key pairs
- D. AWS Key Management Service (AWS KMS) keys

Answer: B

Explanation:

IAM 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. If you need to access AWS services from an on-premises application, you can use IAM access keys to authenticate your requests. AWS account user name and password are used to sign in to the AWS Management Console. Amazon EC2 key pairs are used to connect to your EC2 instances using SSH. AWS Key Management Service (AWS KMS) keys are used to encrypt and decrypt your data using the AWS Encryption SDK or the AWS CLI.

NEW QUESTION 33

- (Topic 3)

Which AWS service or feature enables users to encrypt data at rest in Amazon S3?

- A. IAM policies
- B. Server-side encryption
- C. Amazon GuardDuty
- D. Client-side encryption

Answer: B

Explanation:

Server-side encryption is an encryption option that Amazon S3 provides to encrypt data at rest in Amazon S3. With server-side encryption, Amazon S3 encrypts an object before saving it to disk in its data centers and decrypts it when you download the objects. You have three server-side encryption options to choose from: SSE-S3, SSE-C, and SSE-KMS. SSE-S3 uses keys that are managed by Amazon S3. SSE-C allows you to manage your own encryption keys. SSE-KMS uses keys that are managed by AWS Key Management Service (AWS KMS).

NEW QUESTION 35

- (Topic 3)

A company's application has high customer usage during certain times of the day. The company wants to reduce the number of Amazon EC2 instances that run when application usage is low.

Which AWS service or instance purchasing option should the company use to meet this requirement?

- A. EC2 Instance Savings Plans
- B. Spot Instances
- C. Reserved Instances
- D. Amazon EC2 Auto Scaling

Answer: D

Explanation:

Amazon EC2 Auto Scaling is an AWS service that can help users reduce the number of Amazon EC2 instances that run when application usage is low. Amazon EC2 Auto Scaling allows users to create scaling policies that automatically adjust the number of EC2 instances based on the demand or a schedule. EC2 Instance Savings Plans, Spot Instances, and Reserved Instances are instance purchasing options that can help users save money on EC2 usage, but they do not automatically scale the number of instances according to the application usage.

NEW QUESTION 39

- (Topic 3)

A company has 5 TB of data stored in Amazon S3. The company plans to occasionally run queries on the data for analysis.

Which AWS service should the company use to run these queries in the MOST cost-effective manner?

- A. Amazon Redshift
- B. Amazon Athena
- C. Amazon Kinesis
- D. Amazon RDS

Answer: B

Explanation:

Amazon Athena is a serverless, interactive analytics service that allows users to run SQL queries on data stored in Amazon S3. It is ideal for occasional queries on large datasets, as it does not require any server provisioning, configuration, or management. Users only pay for the queries they run, based on the amount of data scanned. Amazon Athena supports various data formats, such as CSV, JSON, Parquet, ORC, and Avro, and integrates with AWS Glue Data Catalog to create and manage schemas. Amazon Athena also supports querying data from other sources, such as on-premises or other cloud systems, using data connectors.

Amazon Redshift is a fully managed data warehouse service that allows users to run complex analytical queries on petabyte-scale data. However, it requires users to provision and maintain clusters of nodes, and pay for the storage and compute capacity they use. Amazon Redshift is more suitable for frequent and consistent queries on structured or semi-structured data.

Amazon Kinesis is a platform for streaming data on AWS, enabling users to collect, process, and analyze real-time data. It is not designed for querying data stored in Amazon S3. Amazon Kinesis consists of four services: Kinesis Data Streams, Kinesis Data Firehose, Kinesis Data Analytics, and Kinesis Video Streams.

Amazon RDS is a relational database service that provides six database engines: Amazon Aurora, PostgreSQL, MySQL, MariaDB, Oracle Database, and SQL Server. It simplifies database administration tasks such as backup, patching, scaling, and replication. However, it is not optimized for querying data stored in Amazon S3. Amazon RDS is more suitable for transactional workloads that require high performance and availability.

References:

? Interactive SQL - Serverless Query Service - Amazon Athena - AWS

? [Amazon Redshift – Data Warehouse Solution - AWS]

? [Amazon Kinesis - Streaming Data Platform - AWS]

? [Amazon Relational Database Service (RDS) – AWS]

NEW QUESTION 42

- (Topic 3)

A development team wants to deploy multiple test environments for an application in a fast repeatable manner.

Which AWS service should the team use?

- A. Amazon EC2
- B. AWS CloudFormation

- C. Amazon QuickSight
- D. Amazon Elastic Container Service (Amazon ECS)

Answer: B

Explanation:

AWS CloudFormation is a service that allows you to model and provision your AWS resources using templates. You can define your infrastructure as code and automate the creation and update of your resources. AWS CloudFormation also supports nested stacks, change sets, and rollback features to help you manage complex and dynamic environments³⁴. References:

- ? AWS CloudFormation
- ? AWS Certified Cloud Practitioner Exam Guide

NEW QUESTION 44

- (Topic 3)

A company has teams that have different job roles and responsibilities. The company's employees often change teams. The company needs to manage permissions for the employees so that the permissions are appropriate for the job responsibilities.

Which IAM resource should the company use to meet this requirement with the LEAST operational overhead?

- A. IAM user groups
- B. IAM roles
- C. IAM instance profiles
- D. IAM policies for individual users

Answer: B

Explanation:

IAM roles are a way of granting temporary permissions to entities that need to access AWS resources, such as users, applications, or services. IAM roles allow customers to assign permissions to entities without having to create or manage IAM users or credentials for them. IAM roles can be assumed by different entities depending on the trust policy attached to the role. For example, IAM roles can be assumed by IAM users in the same or different AWS accounts, AWS services such as EC2 or Lambda, or external identities such as federated users or web identities. IAM roles can also be switched by IAM users to temporarily change their permissions. IAM roles are recommended for managing permissions for employees who often change teams, because they allow customers to define permissions based on job roles and responsibilities, and easily assign or revoke them as needed. IAM roles also reduce the operational overhead of creating, updating, or deleting IAM users or credentials for each employee or team change.

NEW QUESTION 48

- (Topic 3)

A company runs a MySQL database in its on-premises data center. The company wants to run a copy of this database in the AWS Cloud.

Which AWS service would support this workload?

- A. Amazon RDS
- B. Amazon Neptune
- C. Amazon ElastiCache for Redis
- D. Amazon Quantum Ledger Database (Amazon QLDB)

Answer: A

Explanation:

Amazon Relational Database Service (Amazon RDS) is a web service that makes it easier to set up, operate, and scale a relational database in the cloud. It provides cost-efficient and resizable capacity, while automating time-consuming administration tasks such as hardware provisioning, database setup, patching, and backups. Amazon RDS supports six popular database engines: Amazon Aurora, PostgreSQL, MySQL, MariaDB, Oracle Database, and SQL Server. Amazon RDS can support running a copy of a MySQL database in the AWS Cloud, as it offers compatibility, scalability, and availability features.

NEW QUESTION 52

- (Topic 3)

Which AWS service uses AWS Compute Optimizer to provide sizing recommendations based on workload metrics?

- A. Amazon EC2
- B. Amazon RDS
- C. Amazon Lightsail
- D. AWS Step Functions

Answer: A

Explanation:

Amazon EC2 is a web service that provides secure, resizable compute capacity in the cloud. It allows you to launch virtual servers, called instances, with different configurations of CPU, memory, storage, and networking resources. AWS Compute Optimizer analyzes the specifications and utilization metrics of your Amazon EC2 instances and generates recommendations for optimal instance types that can reduce costs and improve performance. You can view the recommendations on the AWS Compute Optimizer console or the Amazon EC2 console¹².

Amazon RDS, Amazon Lightsail, and AWS Step Functions are not supported by AWS Compute Optimizer. Amazon RDS is a managed relational database service that lets you set up, operate, and scale a relational database in the cloud. Amazon Lightsail is an easy- to-use cloud platform that offers everything you need to build an application or website, plus a cost-effective, monthly plan. AWS Step Functions lets you coordinate multiple AWS services into serverless workflows so you can build and update apps quickly³.

NEW QUESTION 53

- (Topic 3)

A company hosts a large amount of data in AWS. The company wants to identify if any of the data should be considered sensitive.

Which AWS service will meet the requirement?

- A. Amazon Inspector

- B. Amazon Macie
- C. AWS Identity and Access Management (IAM)
- D. Amazon CloudWatch

Answer: B

Explanation:

Amazon Macie is a fully managed service that uses machine learning and pattern matching to help you detect, classify, and better protect your sensitive data stored in the AWS Cloud¹. Macie can automatically discover and scan your Amazon S3 buckets for sensitive data such as personally identifiable information (PII), financial information, healthcare information, intellectual property, and credentials¹. Macie also provides you with a dashboard that shows the type, location, and volume of sensitive data in your AWS environment, as well as alerts and findings on potential security issues¹.

The other options are not suitable for identifying sensitive data in AWS. Amazon Inspector is a service that helps you find security vulnerabilities and deviations from best practices in your Amazon EC2 instances². AWS Identity and Access Management (IAM) is a service that helps you manage access to your AWS resources by creating users, groups, roles, and policies³. Amazon CloudWatch is a service that helps you monitor and troubleshoot your AWS resources and applications by collecting metrics, logs, events, and alarms⁴. References:

? 1: What Is Amazon Macie? - Amazon Macie

? 2: What Is Amazon Inspector? - Amazon Inspector

? 3: What Is IAM? - AWS Identity and Access Management

? 4: What Is Amazon CloudWatch? - Amazon CloudWatch

NEW QUESTION 57

- (Topic 3)

A company deployed an application on an Amazon EC2 instance. The application ran as expected for 6 months. In the past week, users have reported latency issues. A system administrator found that the CPU utilization was at 100% during business hours. The company wants a scalable solution to meet demand. Which AWS service or feature should the company use to handle the load for its application during periods of high demand?

- A. Auto Scaling groups
- B. AWS Global Accelerator
- C. Amazon Route 53
- D. An Elastic IP address

Answer: A

Explanation:

Auto Scaling groups are a feature that allows users to automatically scale the number of Amazon EC2 instances up or down based on demand or a predefined schedule. Auto Scaling groups can help improve the performance and availability of applications by adjusting the capacity in response to traffic fluctuations¹. AWS Global Accelerator is a service that improves the availability and performance of applications by routing traffic through AWS edge locations². Amazon Route 53 is a service that provides scalable and reliable domain name system (DNS) service³. An Elastic IP address is a static IPv4 address that can be associated with an Amazon EC2 instance⁴.

NEW QUESTION 58

- (Topic 3)

A company is hosting an application in the AWS Cloud. The company wants to verify that underlying AWS services and general AWS infrastructure are operating normally.

Which combination of AWS services can the company use to gather the required information? (Select TWO.)

- A. AWS Personal Health Dashboard
- B. AWS Systems Manager
- C. AWS Trusted Advisor
- D. AWS Service Health Dashboard
- E. AWS Service Catalog

Answer: AD

Explanation:

AWS Personal Health Dashboard and AWS Service Health Dashboard are two AWS services that can help the company to verify that underlying AWS services and general AWS infrastructure are operating normally. AWS Personal Health Dashboard provides a personalized view into the performance and availability of the AWS services you are using, as well as alerts that are automatically triggered by changes in the health of those services. In addition to event-based alerts, Personal Health Dashboard provides proactive notifications of scheduled activities, such as any changes to the infrastructure powering your resources, enabling you to better plan for events that may affect you. These notifications can be delivered to you via email or mobile for quick visibility, and can always be viewed from within the AWS Management Console. When you get an alert, it includes detailed information and guidance, enabling you to take immediate action to address AWS events impacting your resources³. AWS Service Health Dashboard provides a general status of AWS services, and the Service health view displays the current and historical status of all AWS services. This page shows reported service events for services across AWS Regions. You don't need to sign in or have an AWS account to access the AWS Service Health Dashboard – Service health page. You can also subscribe to RSS feeds for specific services or regions to receive notifications about service

events⁴. References: Getting started with your AWS Health Dashboard – Your account health, Introducing AWS Personal Health Dashboard

NEW QUESTION 59

- (Topic 3)

A company wants to allow users to authenticate and authorize multiple AWS accounts by using a single set of credentials.

Which AWS service or resource will meet this requirement?

- A. AWS Organizations
- B. IAM user
- C. AWS IAM Identity Center (AWS Single Sign-On)
- D. AWS Control Tower

Answer: C

Explanation:

AWS IAM Identity Center (AWS Single Sign-On) is a cloud-based service that makes it easy to centrally manage single sign-on (SSO) access to multiple AWS accounts and business applications. You can use AWS SSO to enable your users to sign in to the AWS Management Console or the AWS Command Line Interface (AWS CLI) with their existing corporate credentials². You can also manage SSO access and user permissions across all your AWS accounts in AWS Organizations³. References: AWS Single Sign-On - AWS Documentation, AWS Organizations - AWS Documentation

NEW QUESTION 63

- (Topic 3)

What does the concept of agility mean in AWS Cloud computing? (Select TWO.)

- A. The speed at which AWS resources are implemented
- B. The speed at which AWS creates new AWS Regions
- C. The ability to experiment quickly
- D. The elimination of wasted capacity
- E. The low cost of entry into cloud computing

Answer: AC

Explanation:

Agility in AWS Cloud computing means the ability to rapidly provision and deprovision AWS resources as needed, and the ability to experiment quickly with new ideas and solutions. Agility helps businesses to respond to changing customer demands, market opportunities, and competitive threats, and to innovate faster and cheaper. Agility also reduces the risk of failure, as businesses can test and validate their assumptions before committing to large-scale deployments. Some of the benefits of agility in AWS Cloud computing are:

? The speed at which AWS resources are implemented: AWS provides a variety of services and tools that allow you to create, configure, and launch AWS resources in minutes, using the AWS Management Console, the AWS Command Line Interface (AWS CLI), the AWS Software Development Kits (AWS SDKs), or the AWS CloudFormation templates. You can also use the AWS Cloud Development Kit (AWS CDK) to define your AWS resources as code using familiar programming languages, and synthesize them into AWS CloudFormation templates. You can also use the AWS Service Catalog to create and manage standardized portfolios of AWS resources that meet your organizational policies and best practices. AWS also offers on-demand, pay-as-you-go pricing models, so you only pay for the resources you use, and you can scale them up or down as your needs change¹²³⁴⁵

? The ability to experiment quickly: AWS enables you to experiment quickly with new ideas and solutions, without having to invest in upfront capital or long-term commitments. You can use AWS to create and test multiple prototypes, hypotheses, and minimum viable products (MVPs) in parallel, and measure their performance and feedback. You can also use AWS to leverage existing services and solutions, such as AWS Marketplace, AWS Solutions, and AWS Quick Starts, that can help you accelerate your innovation process. AWS also supports a culture of experimentation and learning, by providing tools and resources for continuous integration and delivery (CI/CD), testing, monitoring, and analytics.

References: Six advantages of cloud computing - Overview of Amazon Web Services, AWS Cloud Development Kit (AWS CDK), AWS Service Catalog, AWS Pricing, AWS CloudFormation, [Experimentation and Testing - AWS Well-Architected Framework], [AWS Marketplace], [AWS Solutions], [AWS Quick Starts], [AWS Developer Tools]

NEW QUESTION 64

- (Topic 3)

An ecommerce company wants to distribute traffic between the Amazon EC2 instances that host its website. Which AWS service or resource will meet these requirements?

- A. Application Load Balancer
- B. AWS WAF
- C. AWS CloudHSM
- D. AWS Direct Connect

Answer: A

Explanation:

This is the AWS service or resource that will meet the requirements of distributing traffic between the Amazon EC2 instances that host the website. Application Load Balancer is a type of Elastic Load Balancing that distributes incoming application traffic across multiple targets, such as Amazon EC2 instances, containers, IP addresses, and Lambda functions. Application Load Balancer operates at the application layer (layer 7) of the OSI model and supports advanced features such as path-based routing, host-based routing, health checks, and SSL termination. You can learn more about Application Load Balancer from [this webpage] or [this digital course].

NEW QUESTION 65

- (Topic 3)

Which AWS service or feature identifies whether an Amazon S3 bucket or an IAM role has been shared with an external entity?

- A. AWS Service Catalog
- B. AWS Systems Manager
- C. AWS IAM Access Analyzer
- D. AWS Organizations

Answer: C

Explanation:

AWS IAM Access Analyzer is a service that helps you identify the resources in your organization and accounts, such as Amazon S3 buckets or IAM roles, that are shared with an external entity. This lets you identify unintended access to your resources and data, which is a security risk. IAM Access Analyzer uses logic-based reasoning to analyze the resource-based policies in your AWS environment. For each instance of a resource shared outside of your account, IAM Access Analyzer generates a finding. Findings include information about the access and the external principal granted to it³⁴⁵. References: 3: Using AWS Identity and Access Management Access Analyzer, 4: IAM Access Analyzer - Amazon Web Services (AWS), 5: Welcome - IAM Access Analyzer

NEW QUESTION 67

- (Topic 3)

Which task does AWS perform automatically?

- A. Encrypt data that is stored in Amazon DynamoDB.

- B. Patch Amazon EC2 instances.
- C. Encrypt user network traffic.
- D. Create TLS certificates for users' websites.

Answer: B

Explanation:

AWS performs some tasks automatically to help you manage and secure your AWS resources. One of these tasks is patching Amazon EC2 instances. AWS provides two options for patching your EC2 instances: managed instances and patch baselines. Managed instances are a group of EC2 instances or on-premises servers that you can manage using AWS Systems Manager. Patch baselines define the patches that AWS Systems Manager applies to your instances. You can use AWS Systems Manager to automate the process of patching your instances based on a schedule or a maintenance window.

NEW QUESTION 69

- (Topic 3)

An ecommerce company wants to use Amazon EC2 Auto Scaling to add and remove EC2 instances based on CPU utilization. Which AWS service or feature can initiate an Amazon EC2 Auto Scaling action to achieve this goal?

- A. Amazon Simple Queue Service (Amazon SQS)
- B. Amazon Simple Notification Service (Amazon SNS)
- C. AWS Systems Manager
- D. Amazon CloudWatch alarm

Answer: D

Explanation:

Amazon CloudWatch alarm is an AWS service or feature that can initiate an Amazon EC2 Auto Scaling action based on CPU utilization. Amazon CloudWatch is a monitoring and observability service that collects and tracks metrics, logs, events, and alarms for your AWS resources and applications. Amazon CloudWatch alarms are actions that you can configure to send notifications or automatically make changes to the resources you are monitoring based on rules that you define.

Amazon EC2 Auto Scaling is a service that helps you maintain application availability and allows you to automatically add or remove EC2 instances according to definable conditions. You can create dynamic scaling policies that track a specific CloudWatch metric, such as CPU utilization, and define what action to take when the associated CloudWatch alarm is in ALARM. When the policy is in effect, Amazon EC2 Auto Scaling adjusts the group's desired capacity up or down when the threshold of an alarm is

breached. References: 6: Cloud Monitoring - Amazon CloudWatch - AWS, 7: Amazon

CloudWatch Documentation, 8: Dynamic scaling for Amazon EC2 Auto Scaling, 9: Amazon EC2 Auto Scaling Documentation

NEW QUESTION 71

- (Topic 3)

Which options are AWS Cloud Adoption Framework (AWS CAF) people perspective capabilities? (Select TWO.)

- A. Organizational alignment
- B. Portfolio management
- C. Organization design
- D. Risk management
- E. Modern application development

Answer: AC

Explanation:

The AWS Cloud Adoption Framework (AWS CAF) people perspective capabilities are the organizational skills and processes that enable effective cloud adoption. According to the AWS CAF people perspective whitepaper¹, there are seven capabilities in this perspective, two of which are:

? Organizational alignment: This capability helps you align your organizational structure, roles, and responsibilities to support your cloud transformation goals and objectives. It involves assessing your current and desired state of alignment, identifying gaps and misalignments, and designing and implementing changes to optimize your cloud performance¹.

? Organization design: This capability helps you design and evolve your organization to enable agility, innovation, and collaboration in the cloud. It involves defining your cloud operating model, identifying the skills and competencies needed for cloud roles, and creating career paths and development plans for your cloud workforce¹.

The other options are not capabilities in the AWS CAF people perspective. Portfolio management, risk management, and modern application development are capabilities in the AWS CAF business perspective, governance perspective, and platform perspective respectively².

References:

? 1: AWS Cloud Adoption Framework: People Perspective - AWS Cloud Adoption Framework: People Perspective

? 2: AWS Cloud Adoption Framework - AWS Cloud Adoption Framework

NEW QUESTION 72

- (Topic 3)

A company is using Amazon DynamoDB for its application database.

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

- A. Classify data.
- B. Configure access permissions.
- C. Manage encryption options.
- D. Provide public endpoints to store and retrieve data.
- E. Manage the infrastructure layer and the operating system.

Answer: DE

Explanation:

According to the AWS shared responsibility model, AWS is responsible for security of the cloud, while customers are responsible for security in the cloud. This means that AWS is responsible for protecting the infrastructure that runs AWS services, such as hardware, software, networking, and facilities. Customers are responsible for managing their data, classifying their assets, and using IAM tools to apply the appropriate permissions. For abstracted services, such as Amazon

DynamoDB, AWS operates the infrastructure layer, the operating system, and platforms, and provides customers with public endpoints to store and retrieve data. Customers are responsible for classifying their data, managing their encryption options, and configuring their access permissions. References: Shared Responsibility Model, Security and compliance in Amazon DynamoDB, [AWS Cloud Practitioner Essentials: Module 2 - Security in the Cloud]

NEW QUESTION 76

- (Topic 3)

A company wants its Amazon EC2 instances to share the same geographic area but use multiple independent underlying power sources. Which solution achieves this goal?

- A. Use EC2 instances in a single Availability Zone.
- B. Use EC2 instances in multiple AWS Regions.
- C. Use EC2 instances in multiple Availability Zones in the same AWS Region.
- D. Use EC2 instances in the same edge location and the same AWS Region.

Answer: C

Explanation:

The solution that achieves the goal of having Amazon EC2 instances share the same geographic area but use multiple independent underlying power sources is to use EC2 instances in multiple Availability Zones in the same AWS Region. An Availability Zone is a physically isolated location within an AWS Region that has its own power, cooling, and network connectivity. An AWS Region is a geographical area that consists of two or more Availability Zones. By using multiple Availability Zones, users can increase the fault tolerance and resilience of their applications, as well as reduce latency for end users³. Using EC2 instances in a single Availability Zone, multiple AWS Regions, or the same edge location and the same AWS Region would not meet the requirement of having multiple independent power sources.

NEW QUESTION 81

- (Topic 3)

A company wants a customized assessment of its current on-premises environment. The company wants to understand its projected running costs in the AWS Cloud. Which AWS service or tool will meet these requirements?

- A. AWS Trusted Advisor
- B. Amazon Inspector
- C. AWS Control Tower
- D. Migration Evaluator

Answer: D

Explanation:

Migration Evaluator is an AWS service that provides a customized assessment of your current on-premises environment and helps you build a data-driven business case for migration to AWS. Migration Evaluator collects and analyzes data from your on-premises servers, such as CPU, memory, disk, network, and utilization metrics, and compares them with the most cost-effective AWS alternatives. Migration Evaluator also helps you understand your existing software licenses and running costs, and provides recommendations for Bring Your Own License (BYOL) and License Included (LI) options in AWS. Migration Evaluator generates a detailed report that shows your projected running costs in the AWS Cloud, along with potential savings and benefits. You can use this report to support your decision-making and planning for cloud migration. References: Cloud Business Case & Migration Plan - Amazon Migration Evaluator - AWS, Getting started with Migration Evaluator

NEW QUESTION 86

- (Topic 3)

Which benefit does AWS offer exclusively to users who have an AWS Enterprise Support plan?

- A. Access to a technical project manager
 - B. Access to a technical account manager (TAM)
 - C. Access to a cloud support engineer
 - D. Access to a solutions architect
- A company wants to automatically set up and govern a multi-account AWS environment.

Answer: B

Explanation:

AWS Enterprise Support plan is the highest level of support that AWS offers to its customers. One of the exclusive benefits of this plan is the access to a technical account manager (TAM), who is a dedicated point of contact for guidance, advocacy, and support². A technical project manager, a cloud support engineer, and a solutions architect are not exclusive benefits of the AWS Enterprise Support plan, as they are also available to customers with lower-tier support plans or through other AWS services or programs^{3,4,5}.

NEW QUESTION 87

- (Topic 3)

A company has deployed an Amazon EC2 instance. Which option is an AWS responsibility under the AWS shared responsibility model?

- A. Managing and encrypting application data
- B. Installing updates and security patches of guest operating system
- C. Configuration of infrastructure devices
- D. Configuration of security groups on each instance

Answer: C

Explanation:

According to the AWS shared responsibility model, AWS is responsible for protecting the infrastructure that runs all of the services offered in the AWS Cloud, such as data centers, hardware, software, networking, and facilities¹. This includes the configuration of infrastructure devices, such as routers, switches, firewalls, and load balancers². Customers are responsible for managing their data, applications, operating systems, security groups, and other aspects of their AWS

environment1. Therefore, options A, B, and D are customer responsibilities, not AWS responsibilities. References: 1: AWS Well-Architected Framework - Elasticity; 2: Reactive Systems on AWS - Elastic

NEW QUESTION 92

- (Topic 3)

A company website is experiencing DDoS attacks.

Which AWS service can help protect the company website against these attacks?

- A. AWS Resource Access Manager
- B. AWS Amplify
- C. AWS Shield
- D. Amazon GuardDuty

Answer: C

Explanation:

AWS Shield is a managed DDoS protection service that safeguards applications running on AWS from distributed denial of service (DDoS) attacks. DDoS attacks are malicious attempts to disrupt the normal functioning of a website or application by overwhelming it with a large volume of traffic from multiple sources. AWS Shield provides two tiers of protection: Standard and Advanced. AWS Shield Standard is automatically enabled for all AWS customers at no additional cost. It protects your AWS resources, such as Amazon CloudFront, AWS Global Accelerator, and Amazon Route 53, from the most common and frequently occurring network and transport layer DDoS attacks. AWS Shield Advanced is an optional paid service that provides additional protection for your AWS resources and applications, such as Amazon Elastic Compute Cloud (Amazon EC2), Elastic Load Balancing (ELB), Amazon Simple Storage Service (Amazon S3), Amazon Relational Database Service (Amazon RDS), and AWS Elastic Beanstalk. AWS Shield Advanced offers enhanced detection and mitigation capabilities, 24/7 access to the AWS DDoS Response Team (DRT), real-time visibility and reporting, and cost protection against DDoS-related spikes in your AWS bill¹²

References: AWS Shield, What is a DDOS Attack & How to Protect Your Site Against One

NEW QUESTION 96

- (Topic 3)

AWS has the ability to achieve lower pay-as-you-go pricing by aggregating usage across hundreds of thousands of users.

This describes which advantage of the AWS Cloud?

- A. Launch globally in minutes
- B. Increase speed and agility
- C. High economies of scale
- D. No guessing about compute capacity

Answer: C

Explanation:

AWS has the ability to achieve lower pay-as-you-go pricing by aggregating usage across hundreds of thousands of users. This means that AWS can leverage its massive scale and purchasing power to reduce the costs of infrastructure, hardware, software, and operations. These savings are then passed on to the customers, who only pay for the resources they use. You can learn more about the AWS pricing model from [this webpage] or [this digital course].

NEW QUESTION 100

- (Topic 3)

Which mechanism allows developers to access AWS services from application code?

- A. AWS Software Development Kit
- B. AWS Management Console
- C. AWS CodePipeline
- D. AWS Config

Answer: A

Explanation:

AWS Software Development Kit (SDK) is a set of platform-specific building tools for developers. It allows developers to access AWS services from application code using familiar programming languages. It provides pre-built components and libraries that can be incorporated into applications, as well as tools to debug, monitor, and optimize performance². References: What is SDK? - SDK Explained - AWS

NEW QUESTION 103

- (Topic 3)

A company is moving an on-premises data center to the AWS Cloud. The company must migrate 50 petabytes of file storage data to AWS with the least possible operational overhead.

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

- A. AWS Snowmobile
- B. AWS Snowball Edge
- C. AWS Data Exchange
- D. AWS Database Migration Service (AWS DMS)

Answer: A

Explanation:

The AWS service that the company should use to meet these requirements is A. AWS Snowmobile.

AWS Snowmobile is a service that allows you to migrate large amounts of data to AWS using a 45-foot long ruggedized shipping container that can store up to 100 petabytes of data. AWS Snowmobile is designed for situations where you need to move massive amounts of data to the cloud in a fast, secure, and cost-effective way. AWS Snowmobile has the least possible operational overhead because it eliminates the need to buy, configure, or manage hundreds or thousands of storage devices¹².

AWS Snowball Edge is a service that allows you to migrate data to AWS using a physical device that can store up to 80 terabytes of data and has compute and

storage capabilities to run applications on the device. AWS Snowball Edge is suitable for situations where you have limited or intermittent network connectivity, or where bandwidth costs are high. However, AWS Snowball Edge has more operational overhead than AWS Snowmobile because you need to request multiple devices and transfer your data onto them using the client³.

AWS Data Exchange is a service that allows you to find, subscribe to, and use third-party data in the cloud. AWS Data Exchange is not a data migration service, but rather a data marketplace that enables data providers and data consumers to exchange data sets securely and efficiently⁴.

AWS Database Migration Service (AWS DMS) is a service that helps migrate databases to AWS. AWS DMS does not migrate file storage data, but rather supports various database platforms and engines as sources and targets⁵.

References:

1: AWS Snowmobile – Move Exabytes of Data to the Cloud in Weeks 2: AWS Snowmobile

- Amazon Web Services 3: Automated Software Vulnerability Management - Amazon Inspector - AWS 4: AWS Data Exchange - Find, subscribe to, and use third-party data in ... 5: AWS Database Migration Service – Amazon Web Services

NEW QUESTION 104

- (Topic 3)

A company wants to define a central data protection policy that works across AWS services for compute, storage, and database resources.

Which AWS service will meet this requirement?

- A. AWS Batch
- B. AWS Elastic Disaster Recovery
- C. AWS Backup
- D. Amazon FSx

Answer: C

Explanation:

The AWS service that will meet this requirement is C. AWS Backup.

AWS Backup is a service that allows you to define a central data protection policy that works across AWS services for compute, storage, and database resources. You can use AWS Backup to create backup plans that specify the frequency, retention, and lifecycle of your backups, and apply them to your AWS resources using tags or resource IDs. AWS Backup supports various AWS services, such as Amazon EC2, Amazon EBS, Amazon RDS, Amazon DynamoDB, Amazon EFS, Amazon FSx, and AWS Storage Gateway¹². AWS Batch is a service that allows you to run batch computing workloads on AWS. AWS Batch does not provide a central data protection policy, but rather enables you to optimize the allocation and utilization of your compute resources³.

AWS Elastic Disaster Recovery is a service that allows you to prepare for and recover from disasters using AWS. AWS Elastic Disaster Recovery does not provide a central data protection policy, but rather helps you minimize downtime and data loss by replicating your applications and data to AWS⁴.

Amazon FSx is a service that provides fully managed file storage for Windows and Linux applications. Amazon FSx does not provide a central data protection policy, but rather offers features such as encryption, snapshots, backups, and replication to protect your file systems⁵.

References:

1: AWS Backup – Centralized backup across AWS services 3: AWS Batch – Run Batch Computing Jobs on AWS 2: Data Protection Reference Architectures with AWS Backup 4: AWS Elastic Disaster Recovery – Prepare for and recover from disasters using AWS 5: Amazon FSx – Fully managed file storage for Windows and Linux applications

NEW QUESTION 108

- (Topic 3)

Which AWS service or feature can the company use to limit the access to AWS services for member accounts?

- A. AWS Identity and Access Management (IAM)
- B. Service control policies (SCPs)
- C. Organizational units (OUs)
- D. Access control lists (ACLs)

Answer: B

Explanation:

Service control policies (SCPs) are a type of organization policy that you can use to manage permissions in your organization. SCPs offer central control over the maximum available permissions for all accounts in your organization, allowing you to ensure your accounts stay within your organization's access control guidelines². SCPs are available only in an organization that has all features enabled².

NEW QUESTION 111

- (Topic 3)

A company wants to migrate its PostgreSQL database to AWS. The company does not use the database frequently.

Which AWS service or resource will meet these requirements with the LEAST management overhead?

- A. PostgreSQL on Amazon EC2
- B. Amazon RDS for PostgreSQL
- C. Amazon Aurora PostgreSQL-Compatible Edition
- D. Amazon Aurora Serverless

Answer: D

Explanation:

Amazon Aurora Serverless is an on-demand, auto-scaling configuration for Amazon Aurora PostgreSQL-Compatible Edition. It is a fully managed service that automatically scales up and down based on the application's actual needs. Amazon Aurora Serverless is suitable for applications that have infrequent, intermittent, or unpredictable database workloads, and that do not require the full power and range of options provided by provisioned Aurora clusters. Amazon Aurora Serverless eliminates the need to provision and manage database instances, and reduces the management overhead associated with database administration tasks such as scaling, patching, backup, and recovery. References: Amazon Aurora Serverless, Choosing between Aurora Serverless and provisioned Aurora DB clusters, [AWS Cloud Practitioner Essentials: Module 4 - Databases in the Cloud]

NEW QUESTION 113

- (Topic 3)

A company seeks cost savings in exchange for a commitment to use a specific amount of an AWS service or category of AWS services for 1 year or 3 years.

Which AWS pricing model or offering will meet these requirements?

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

Answer: B

Explanation:

Savings Plans are an AWS pricing model or offering that can meet the requirements of seeking cost savings in exchange for a commitment to use a specific amount of an AWS service or category of AWS services for 1 year or 3 years. Savings Plans are flexible plans that offer significant discounts on AWS compute usage, such as EC2, Lambda, and Fargate. The company can choose from two types of Savings Plans: Compute Savings Plans and EC2 Instance Savings Plans. Compute Savings Plans provide the most flexibility and apply to any eligible compute usage, regardless of instance family, size, region, operating system, or tenancy. EC2 Instance Savings Plans provide more savings and apply to a specific instance family within a region. The company can select the amount of compute usage per hour (e.g., \$10/hour) that they want to commit to for the duration of the plan (1 year or 3 years). The company will pay the discounted Savings Plan rate for the amount of usage that matches their commitment, and the regular on-demand rate for any usage beyond that

NEW QUESTION 116

- (Topic 3)

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

- A. High availability
- B. Performance efficiency
- C. Cost optimization
- D. Going global in minutes
- E. Continuous development

Answer: BC

Explanation:

The AWS Well-Architected Framework is a set of six pillars and lenses that help cloud architects design and run workloads in the cloud. The six pillars are: operational excellence, security, reliability, performance efficiency, cost optimization, and sustainability. Each pillar has a set of design principles and best practices that guide the architectural decisions. High availability is not a separate pillar, but a quality that can be achieved by applying the principles of the reliability pillar. Going global in minutes and continuous development are not pillars of the framework, but possible benefits of using AWS services and following the framework's recommendations. References: AWS Well-Architected - Build secure, efficient cloud applications, AWS Well-Architected Framework, The 6 Pillars of the AWS Well-Architected Framework

NEW QUESTION 118

- (Topic 3)

Which AWS service offers object storage?

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

Answer: C

Explanation:

Amazon S3 is the AWS service that offers object storage. Object storage is a technology that stores and manages data in an unstructured format called objects. Each object consists of the data, metadata, and a unique identifier. Object storage is ideal for storing large amounts of unstructured data, such as photos, videos, email, web pages, sensor data, and audio files¹. Amazon S3 provides industry-leading scalability, data availability, security, and performance for object storage². Amazon RDS is the AWS service that offers relational database storage. Relational database storage is a technology that stores and manages data in a structured format called tables. Each table consists of rows and columns that define the attributes and values of the data. Relational database storage is ideal for storing structured or semi-structured data, such as customer records, inventory, transactions, and analytics³. Amazon Elastic File System (Amazon EFS) is the AWS service that offers file storage. File storage is a technology that stores and manages data in a hierarchical format called files and folders. Each file consists of the data and metadata, and each folder consists of files or subfolders. File storage is ideal for storing shared data that can be accessed by multiple users or applications, such as home directories, content repositories, media libraries, and configuration files⁴. Amazon DynamoDB is the AWS service that offers NoSQL database storage. NoSQL database storage is a technology that stores and manages data in a flexible format called documents or key-value pairs. Each document or key-value pair consists of the data and metadata, and can have different attributes and values depending on the schema. NoSQL database storage is ideal for storing dynamic or unstructured data that requires high performance, scalability, and availability, such as web applications, social media, gaming, and IoT.

NEW QUESTION 119

- (Topic 3)

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

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

Answer: AB

Explanation:

These are two scenarios that represent the concept of elasticity on AWS. Elasticity means the ability to adjust the resources and capacity of the system in response to changes in demand or environment. Scaling the number of Amazon EC2 instances based on traffic means using services such as AWS Auto Scaling or Elastic Load Balancing to add or remove instances as the traffic increases or decreases. Resizing Amazon RDS instances as business needs change means

using the Amazon RDS console or API to modify the instance type, storage type, or storage size of the database as the workload grows or shrinks. You can learn more about the concept of elasticity on AWS from [this webpage] or [this digital course].

NEW QUESTION 124

- (Topic 3)

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

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

Answer: D

Explanation:

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

NEW QUESTION 126

- (Topic 3)

A company needs to block SQL injection attacks.

Which AWS service or feature can meet this requirement?

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

Answer: A

Explanation:

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

NEW QUESTION 130

- (Topic 3)

Which of the following is a fully managed MySQL-compatible database?

- A. Amazon S3
- B. Amazon DynamoDB
- C. Amazon Redshift
- D. Amazon Aurora

Answer: D

Explanation:

Amazon Aurora is a fully managed MySQL-compatible database that combines the performance and availability of traditional enterprise databases with the simplicity and cost-effectiveness of open-source databases. Amazon Aurora is part of the Amazon Relational Database Service (Amazon RDS) family, which means it inherits the benefits of a fully managed service, such as automated backups, patches, scaling, monitoring, and security. Amazon Aurora also offers up to five times the throughput of standard MySQL, as well as high availability, durability, and fault tolerance with up to 15 read replicas, cross-Region replication, and self-healing storage. Amazon Aurora is compatible with the latest versions of MySQL, as well as PostgreSQL, and supports various features and integrations that enhance its functionality and usability. References: Amazon Aurora, Amazon RDS, AWS — Amazon Aurora Overview

NEW QUESTION 131

- (Topic 3)

A company is planning to host its workloads on AWS.

Which AWS service requires the company to update and patch the guest operating system?

- A. Amazon DynamoDB
- B. Amazon S3
- C. Amazon EC2
- D. Amazon Aurora

Answer: C

Explanation:

Amazon EC2 is an AWS service that provides scalable, secure, and resizable compute capacity in the cloud. Amazon EC2 allows customers to launch and manage virtual servers, called instances, that run a variety of operating systems and applications. Customers have full control over the configuration and management of their instances, including the guest operating system. Therefore, customers are responsible for updating and patching the guest operating system on their EC2 instances, as well as any other software or utilities installed on the instances. AWS provides tools and services, such as AWS Systems Manager and AWS OpsWorks, to help customers automate and simplify the patching process. References: Shared Responsibility Model, Shared responsibility model, [Amazon EC2]

NEW QUESTION 133

- (Topic 3)

Which AWS Support plan is the minimum recommended tier for users who have production workloads on AWS?

- A. AWS Developer Support
- B. AWS Enterprise Support
- C. AWS Business Support
- D. AWS Enterprise On-Ramp Support

Answer: C

Explanation:

AWS Business Support is the minimum recommended tier for users who have production workloads on AWS. AWS Business Support provides 24x7 access to cloud support engineers via phone, chat, or email, as well as a guaranteed response time of less than one hour for urgent issues. AWS Business Support also includes access to AWS Trusted Advisor, a tool that provides real-time guidance to help you provision your resources following AWS best practices⁴.

NEW QUESTION 138

- (Topic 3)

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 140

- (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 142

- (Topic 3)

A company has designed its AWS Cloud infrastructure to run its workloads effectively. The company also has protocols in place to continuously improve supporting processes.

Which pillar of the AWS Well-Architected Framework does this scenario represent?

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

Answer: D

Explanation:

The scenario represents the operational excellence pillar of the AWS Well-Architected Framework, which focuses on running and monitoring systems to deliver business value and continually improve supporting processes and procedures¹. Security, performance efficiency, cost optimization, and reliability are the other four pillars of the framework¹.

NEW QUESTION 146

- (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 149

- (Topic 3)

A company deployed an Amazon EC2 instance last week. A developer realizes that the EC2 instance is no longer running. The developer reviews a list of provisioned EC2 instances, and the EC2 instance is no longer on the list.

What can the developer do to generate a recent history of the EC2 instance?

- A. Run Cost Explorer to identify the start time and end time of the EC2 instance.
- B. Use Amazon Inspector to find out when the EC2 instance was stopped.
- C. Perform a search in AWS CloudTrail to find all EC2 instance-related events.
- D. Use AWS Secrets Manager to display hidden termination logs of the EC2 instance.

Answer: C

Explanation:

AWS CloudTrail is a service that enables governance, compliance, operational auditing, and risk auditing of a customer's AWS account. AWS CloudTrail allows customers to track user activity and API usage across their AWS infrastructure. AWS CloudTrail can also provide a history of EC2 instance events, such as launch, stop, terminate, and reboot. Cost Explorer is a tool that enables customers to visualize, understand, and manage their AWS costs and usage over time. Amazon Inspector is an automated security assessment service that helps improve the security and compliance of applications deployed on AWS. AWS Secrets Manager helps customers protect secrets needed to access their applications, services, and IT resources.

NEW QUESTION 150

- (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 corruption.

NEW QUESTION 152

- (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 155

- (Topic 2)

Which AWS solution provides the ability for a company to run AWS services in the company's on-premises data center?

- A. AWS Direct Connect
- B. AWS Outposts
- C. AWS Systems Manager hybrid activations
- D. AWS Storage Gateway

Answer: B

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 158

- (Topic 2)

A retail company has recently migrated its website to AWS. The company wants to ensure that it is protected from SQL injection attacks. The website uses an Application Load Balancer to distribute traffic to multiple Amazon EC2 instances.

Which AWS service or feature can be used to create a custom rule that blocks SQL injection attacks?

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

Answer: B

Explanation:

AWS WAF is a web application firewall that helps protect your web applications or APIs against common web exploits that may affect availability, compromise security, or consume excessive resources. AWS WAF gives you control over how traffic reaches your applications by enabling you to create security rules that block common attack patterns, such as SQL injection or cross-site scripting, and rules that filter out specific traffic patterns you define². You can use AWS WAF to create a custom rule that blocks SQL injection attacks on your website.

NEW QUESTION 161

- (Topic 2)

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

Which AWS service should the company use?

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

Answer: A

Explanation:

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

NEW QUESTION 163

- (Topic 2)

Which AWS services allow users to monitor and retain records of account activities that include governance, compliance, and auditing?

(Select TWO.)

- A. Amazon CloudWatch
- B. AWS CloudTrail
- C. Amazon GuardDuty
- D. AWS Shield
- E. AWS WAF

Answer: AB

Explanation:

Amazon CloudWatch and AWS CloudTrail are the AWS services that allow users to monitor and retain records of account activities that include governance, compliance, and auditing. Amazon CloudWatch is a service that collects and tracks metrics, collects and monitors log files, and sets alarms. AWS CloudTrail is a service that enables governance, compliance, operational auditing, and risk auditing of your AWS account. Amazon GuardDuty, AWS Shield, and AWS WAF are AWS services that provide security and protection for AWS resources, but they do not monitor and retain records of account activities. These concepts are explained in the AWS Cloud Practitioner Essentials course³.

NEW QUESTION 167

- (Topic 2)

Which of the following is entirely the responsibility of AWS, according to the AWS shared responsibility model?

- A. Security awareness and training
- B. Development of an IAM password policy
- C. Patching of the guest operating system
- D. Physical and environmental controls

Answer: D

Explanation:

Physical and environmental controls are entirely the responsibility of AWS, according to the AWS shared responsibility model. The AWS shared responsibility model 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. For more information, see [AWS Shared Responsibility Model] and [AWS Cloud Security].

NEW QUESTION 172

- (Topic 2)

Which AWS service provides the SIMPLEST way for the company to establish a website on AWS?

- A. Amazon Elastic File System (Amazon EFS)
- B. AWS Elastic Beanstalk
- C. AWS Lambda
- D. Amazon Lightsail

Answer: D

Explanation:

Amazon Lightsail is an easy-to-use cloud platform that offers you everything needed to build an application or website, plus a cost-effective, monthly plan. Whether you're new to the cloud or looking to get on the cloud quickly with AWS infrastructure you trust, we've got you covered. Lightsail provides the simplest way for the company to establish a website on AWS.

NEW QUESTION 176

- (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 resources.

NEW QUESTION 180

- (Topic 2)

A company is reviewing the design of an application that will be migrated from on premises to a single Amazon EC2 instance. What should the company do to make the application highly available?

- A. Provision additional EC2 instances in other Availability Zones.
- B. Configure an Application Load Balancer (ALB). Assign the EC2 instance as the ALB's target.
- C. Use an Amazon Machine Image (AMI) to create the EC2 instance.
- D. Provision the application by using an EC2 Spot Instance.

Answer: A

Explanation:

Provisioning additional EC2 instances in other Availability Zones is a way to make the application highly available, as it reduces the impact of failures and increases fault tolerance. Configuring an Application Load Balancer and assigning the EC2 instance as the ALB's target is a way to distribute traffic among multiple instances, but it does not make the application highly available if there is only one instance. Using an Amazon Machine Image to create the EC2 instance is a way to launch a virtual server with a preconfigured operating system and software, but it does not make the application highly available by itself. Provisioning the application by using an EC2 Spot Instance is a way to use spare EC2 capacity at up to 90% off the On-Demand price, but it does not make the application highly available, as Spot Instances can be interrupted by EC2 with a two-minute notification.

NEW QUESTION 184

- (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 185

- (Topic 2)

A user discovered that an Amazon EC2 instance is missing an Amazon Elastic Block Store (Amazon EBS) data volume. The user wants to determine when the EBS volume was removed.

Which AWS service will provide this information?

- A. AWS Config
- B. AWS Trusted Advisor
- C. Amazon Timestream
- D. Amazon QuickSight

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. AWS Config can help you determine when an EBS volume was removed from an EC2 instance by providing a timeline of configuration changes and compliance status. AWS Trusted Advisor, Amazon Timestream, and Amazon QuickSight do not provide the same level of configuration tracking and auditing as AWS Config. Source: AWS Config

NEW QUESTION 190

- (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 191

- (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 194

- (Topic 2)

Which task can a company perform by using security groups in the AWS Cloud?

- A. Allow access to an Amazon EC2 instance through only a specific port.
- B. Deny access to malicious IP addresses at a subnet level.
- C. Protect data that is cached by Amazon CloudFront.
- D. Apply a stateless firewall to an Amazon EC2 instance.

Answer: A

Explanation:

Security groups are virtual firewalls that control the inbound and outbound traffic for Amazon EC2 instances. They can be used to allow access to an Amazon EC2 instance through only a specific port, such as port 22 for SSH or port 80 for HTTP. Security groups cannot deny access to malicious IP addresses at a subnet level, as they only allow or deny traffic based on the rules defined by the customer. To block malicious IP addresses, customers can use network ACLs, which are stateless firewalls that can be applied to subnets. Security groups cannot protect data that is cached by Amazon CloudFront, as they only apply to EC2 instances. To protect data that is cached by Amazon CloudFront, customers can use encryption, signed URLs, or signed cookies. Security groups are not stateless firewalls, as they track the state of the traffic and automatically allow the response traffic to flow back to the source. Stateless firewalls do not track the state of the traffic and require rules for both inbound and outbound traffic.

NEW QUESTION 195

- (Topic 2)

Which encryption types can be used to protect objects at rest in Amazon S3? (Select TWO.)

- A. Server-side encryption with AmazonS3 managed encryption keys (SSE-S3)
- B. Server-side encryption with AWS KMSmanaged keys (SSE-KMS)
- C. TLS
- D. SSL
- E. Transparent Data Encryption (TDE)

Answer: AB

Explanation:

Server-side encryption with Amazon S3 managed encryption keys (SSE-S3) and server-side encryption with AWS KMS managed keys (SSE-KMS) are the encryption types that can be used to protect objects at rest in Amazon S3. Server-side encryption means that Amazon S3 encrypts the objects before saving them on disks and decrypts them when they are downloaded. SSE-S3 uses one master key per bucket that is managed by Amazon S3. SSE-KMS uses a customer master key (CMK) that is stored in AWS Key Management Service (AWS KMS) and provides additional benefits, such as audit trails and key rotation. For more information, see Protecting Data Using Server-Side Encryption and Protecting Data Using Encryption.

NEW QUESTION 196

- (Topic 2)

Which AWS service or tool should a company use to forecast AWS spending?

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

Answer: D

Explanation:

Cost Explorer is an AWS service or tool that can be used to forecast AWS spending. It allows users to analyze their AWS costs and usage using interactive graphs and tables. It also provides features such as filtering, grouping, and forecasting to help users plan their future spending. Amazon DevPay is an AWS service that allows developers to sell applications that are built on AWS services. It handles the billing and metering for the customers of the applications and collects payments from them. It is not a tool for forecasting AWS spending. AWS Organizations is an AWS service that allows users to centrally manage and govern their AWS accounts. It provides features such as creating groups of accounts, applying policies, and automating account creation. It is not a tool for forecasting AWS spending. AWS Trusted Advisor is an AWS service that provides best practices and recommendations to optimize the performance, security, and cost of AWS resources. It can help users identify opportunities to reduce their AWS costs, but it is not a tool for forecasting AWS spending.

NEW QUESTION 201

- (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 206

- (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 207

- (Topic 2)

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

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

Answer: B

Explanation:

The most cost-effective EC2 instance purchasing option for the company that needs to host a web server on Amazon EC2 instances for at least 1 year and cannot tolerate interruption is Partial Upfront Reserved Instances. Reserved Instances are a pricing model that offer significant discounts compared to On-Demand Instances in exchange for a commitment to use a specific amount of compute capacity for a fixed period of time (1 or 3 years). Partial Upfront Reserved Instances require customers to pay a portion of the total cost upfront, and the remaining cost in monthly installments over the term. This option offers a lower effective hourly rate than No Upfront Reserved Instances, which require no upfront payment but have higher monthly payments. On-Demand Instances and Spot Instances are not the best options for the company. On-Demand Instances are a pricing model that offer the most flexibility and no long-term commitment, but have the highest hourly rate. Spot Instances are a pricing model that offer the lowest cost, but are subject to interruption based on supply and demand³⁴

NEW QUESTION 211

- (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 experience⁵. 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 212

- (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 213

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 218

- (Topic 2)

A company wants to create multiple isolated networks in the same AWS account. Which AWS service or component will provide this functionality?

- A. AWS Transit Gateway
- B. Internet gateway
- C. Amazon VPC
- D. Amazon EC2

Answer: C

Explanation:

Amazon Virtual Private Cloud (Amazon VPC) is the AWS service that allows customers to create multiple isolated networks in the same AWS account. A VPC is a logically isolated section of the AWS Cloud where customers can launch AWS resources in a virtual network that they define. Customers can create multiple VPCs within an AWS account, each with its own IP address range, subnets, route tables, security groups, network access control lists, gateways, and other components. AWS Transit Gateway, Internet gateway, and Amazon EC2 are not services or components that provide the functionality of creating multiple isolated networks in the same AWS account. AWS Transit Gateway is a service that enables customers to connect their Amazon VPCs and their on-premises networks to a single gateway. An Internet gateway is a component that enables communication between instances in a VPC and the Internet. Amazon EC2 is a service that provides scalable compute capacity in the cloud³⁴

NEW QUESTION 223

- (Topic 2)

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

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

Answer: A

Explanation:

According to the AWS shared responsibility model, AWS is responsible for the security of the cloud, while the customer is responsible for the security in the cloud. This means that AWS provides the physical and environmental controls, the service and communications protection, and the awareness and training for its employees, while the customer provides the patch and configuration management, the identity and access management, the data encryption, and the firewall configuration for its resources³.

NEW QUESTION 228

- (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 Framework².

NEW QUESTION 231

- (Topic 2)

A company is hosting a web application on Amazon EC2 instances. The company wants to implement custom conditions to filter and control inbound web traffic.

Which AWS service will meet these requirements?

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

Answer: B

Explanation:

The AWS service that will meet the requirements of the company that is hosting a web application on Amazon EC2 instances and wants to implement custom conditions to filter and control inbound web traffic is AWS WAF. AWS WAF is a web application firewall that helps protect web applications from common web exploits that could affect availability, compromise security, or consume excessive resources. The company can use AWS WAF to create custom rules that block malicious requests that match certain patterns, such as SQL injection or cross-site scripting. AWS WAF can be applied to web applications that are behind an Application Load Balancer, Amazon CloudFront, or Amazon API Gateway. Amazon GuardDuty, Amazon Macie, and AWS Shield 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. Amazon Macie is a data security and data privacy service that uses machine learning and pattern matching to discover, classify, and protect sensitive data stored in Amazon S3. AWS Shield is a managed distributed denial of service (DDoS) protection service that safeguards web applications running on AWS. These services are more useful for detecting and preventing different types of threats and attacks, rather than filtering and controlling inbound web traffic based on custom conditions.

NEW QUESTION 235

- (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 240

- (Topic 2)

A developer needs to maintain a development environment infrastructure and a production environment infrastructure in a repeatable fashion. Which AWS service should the developer use to meet these requirements?

- A. AWS Ground Station
- B. AWS Shield
- C. AWS IoT Device Defender
- D. AWS CloudFormation

Answer: D

Explanation:

AWS CloudFormation is a service that allows you to model and provision your AWS and third-party application resources in a repeatable and predictable way. You can use AWS CloudFormation to create, update, and delete a collection of resources as a single unit, called a stack. You can also use AWS CloudFormation to manage your development and production environments in a consistent and efficient manner.

NEW QUESTION 242

- (Topic 2)

A company is using Amazon RDS.

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

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

Answer: D

Explanation:

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

NEW QUESTION 245

- (Topic 2)

A company provides a web-based ecommerce service that runs in two Availability Zones within a single AWS Region. The web service distributes content that is

stored in the Amazon S3 Standard storage class. The company wants to improve the web service's performance globally. What should the company do to meet this requirement?

- A. Change the S3 storage class to S3 Intelligent-Tiering.
- B. Deploy an Amazon CloudFront distribution to cache web server content in edge locations.
- C. Use Amazon API Gateway for the web service.
- D. Migrate the website ecommerce servers to Amazon EC2 with enhanced networking.

Answer: B

Explanation:

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. CloudFront can cache web server content in edge locations, which are located closer to the end users, to improve the web service's performance globally.

NEW QUESTION 250

- (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 251

- (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 255

- (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 257

- (Topic 2)

A company needs Amazon EC2 instances for a workload that can tolerate interruptions.

Which EC2 instance purchasing option meets this requirement with the LARGEST discount compared to On-Demand prices?

- A. Spot Instances
- B. Convertible Reserved Instances

- C. Standard Reserved Instances
- D. Dedicated Hosts

Answer: A

Explanation:

Spot Instances are spare Amazon EC2 instances that are available at up to 90% discount compared to On-Demand prices. They are suitable for workloads that can tolerate interruptions, such as batch processing, data analysis, and testing. Spot Instances are allocated based on the current supply and demand, and can be reclaimed by AWS with a two-minute notice when the demand exceeds the supply. Convertible Reserved Instances are a type of Reserved Instances that provide a significant discount (up to 54%) compared to On-Demand prices and a capacity reservation for Amazon EC2 instances. They are available in 1-year or 3-year terms and allow users to change the instance family, size, operating system, or tenancy during the term. Standard Reserved Instances are another type of Reserved Instances that provide a larger discount (up to 75%) compared to On-Demand prices and a capacity reservation for Amazon EC2 instances. They are available in 1-year or 3-year terms and do not allow users to change the instance attributes during the term. Dedicated Hosts are physical servers with Amazon EC2 instance capacity fully dedicated to the user's use. They are suitable for users who have specific server-bound software licenses or compliance requirements.

NEW QUESTION 261

- (Topic 2)

A company wants to migrate its Microsoft SQL Server database management system from on premises to the AWS Cloud. Which AWS service should the company use to reduce management overhead for this environment?

- A. Amazon Elastic Container Service (Amazon ECS)
- B. Amazon SageMaker
- C. Amazon RDS
- D. Amazon Athena

Answer: C

Explanation:

Amazon Relational Database Service (Amazon RDS) is the AWS service that the company should use to migrate its Microsoft SQL Server database management system from on premises to the AWS Cloud. Amazon RDS is a fully managed service that provides a scalable, secure, and high-performance relational database platform. Amazon RDS supports several database engines, including Microsoft SQL Server. Amazon RDS reduces the management overhead for the database environment by taking care of tasks such as provisioning, patching, backup, recovery, and monitoring. For more information, see [What is Amazon Relational Database Service \(Amazon RDS\)?](#) and [Amazon RDS for SQL Server](#).

NEW QUESTION 263

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

CLF-C01 Practice Exam Features:

- * CLF-C01 Questions and Answers Updated Frequently
- * CLF-C01 Practice Questions Verified by Expert Senior Certified Staff
- * CLF-C01 Most Realistic Questions that Guarantee you a Pass on Your First Try
- * CLF-C01 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The CLF-C01 Practice Test Here](#)