



Amazon-Web-Services

Exam Questions CLF-C02

AWS Certified Cloud Practitioner

About ExamBible

[Your Partner of IT Exam](#)

Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1

- (Topic 3)

A company wants durable storage for static content and infinitely scalable data storage infrastructure at the lowest cost. Which AWS service should the company choose?

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

Answer: B

Explanation:

Amazon S3 is a service that provides durable storage for static content and infinitely scalable data storage infrastructure at the lowest cost. Amazon S3 is an object storage service that allows you to store and retrieve any amount of data from anywhere on the internet. Amazon S3 offers industry-leading scalability, availability, and performance, as well as 99.999999999% (11 9s) of durability and multi-AZ resilience. Amazon S3 also provides various storage classes that offer different levels of performance and cost optimization, such as S3 Standard, S3 Intelligent-Tiering, S3 Standard-Infrequent Access (S3 Standard-IA), S3 One Zone-Infrequent Access (S3 One Zone-IA), and S3 Glacier⁴⁵⁶. Amazon S3 is ideal for storing static content, such as images, videos, documents, and web pages, as well as building data lakes, backup and archive solutions, big data analytics, and machine learning applications⁴⁵⁶. References: 4: Cloud Storage on AWS, 5: Object Storage - Amazon Simple Storage Service (S3) - AWS, 6: Amazon S3 Documentation

NEW QUESTION 2

- (Topic 3)

Which abilities are benefits of the AWS Cloud? (Select TWO.)

- A. Trade variable expenses for capital expenses.
- B. Deploy globally in minutes.
- C. Plan capacity in advance of deployments.
- D. Take advantage of economies of scale.
- E. Reduce dependencies on network connectivity.

Answer: AB

Explanation:

The AWS Cloud offers many benefits, such as:

? Trade variable expenses for capital expenses: You can pay only for the resources you use, instead of investing in fixed costs upfront. This reduces the risk and complexity of planning and managing your IT infrastructure⁴

? Deploy globally in minutes: You can leverage the global infrastructure of AWS to deploy your applications and data in multiple regions and availability zones. This enables you to reach your customers faster, improve performance, and increase reliability⁵

NEW QUESTION 3

- (Topic 3)

Which AWS service provides this functionality?

- A. AWS IAM Identity Center (AWS Single Sign-On)
- B. AWS Systems Manager
- C. AWS Config
- D. AWS Control Tower

Answer: D

Explanation:

AWS Control Tower is a service that provides an easy way to set up and govern a secure, multi-account AWS environment. It automates the creation of accounts, organizational units, policies, and best practices based on the AWS Well-Architected Framework. AWS IAM Identity Center (AWS Single Sign-On) is a service that enables users to centrally manage access to multiple AWS accounts and business applications using a single sign-on experience. AWS Systems Manager is a service that provides operational management for AWS resources and applications. AWS Config is a service that enables users to assess, audit, and evaluate the configurations of AWS resources.

NEW QUESTION 4

- (Topic 3)

A developer has been hired by a large company and needs AWS credentials. Which are security best practices that should be followed? (Select TWO.)

- A. Grant the developer access to only the AWS resources needed to perform the job.
- B. Share the AWS account root user credentials with the developer.
- C. Add the developer to the administrator's group in AWS IAM.
- D. Configure a password policy that ensures the developer's password cannot be changed.
- E. Ensure the account password policy requires a minimum length.

Answer: AE

Explanation:

The security best practices that should be followed are A and E.

* A. Grant the developer access to only the AWS resources needed to perform the job. This is an example of the principle of least privilege, which means giving the minimum permissions necessary to achieve a task. This reduces the risk of unauthorized access, data leakage, or accidental damage to AWS resources. You can use AWS Identity and Access Management (IAM) to create users, groups, roles, and policies that grant fine-grained access to AWS resources¹².

* E. Ensure the account password policy requires a minimum length. This is a basic security measure that helps prevent brute-force attacks or guessing of passwords. A longer password is harder to crack than a shorter one. You can use IAM to configure a password policy that enforces a minimum password length, as well as other requirements such as complexity, expiration, and history³⁴.

- * B. Share the AWS account root user credentials with the developer. This is a bad practice that should be avoided. The root user has full access to all AWS resources and services, and can perform sensitive actions such as changing billing information, closing the account, or deleting all resources. Sharing the root user credentials exposes your account to potential compromise or misuse. You should never share your root user credentials with anyone, and use them only for account administration tasks⁵.
- * C. Add the developer to the administrator's group in IAM. This is also a bad practice that should be avoided. The administrator's group has full access to all AWS resources and services, which is more than what a developer needs to perform their job. Adding the developer to the administrator's group violates the principle of least privilege and increases the risk of unauthorized access, data leakage, or accidental damage to AWS resources. You should create a custom group for the developer that grants only the necessary permissions for their role¹².
- * D. Configure a password policy that ensures the developer's password cannot be changed. This is another bad practice that should be avoided. Preventing the developer from changing their password reduces their ability to protect their credentials and comply with security policies. For example, if the developer's password is compromised, they cannot change it to prevent further unauthorized access. Or if the company requires periodic password rotation, they cannot update their password to meet this requirement. You should allow the developer to change their password as needed, and enforce a password policy that sets reasonable rules for password management³⁴.

NEW QUESTION 5

- (Topic 3)

Which AWS service is a cloud security posture management (CSPM) service that aggregates alerts from various AWS services and partner products in a standardized format?

- A. AWS Security Hub
- B. AWS Trusted Advisor
- C. Amazon EventBridge
- D. Amazon GuardDuty

Answer: A

Explanation:

AWS Security Hub is a cloud security posture management (CSPM) service that performs security best practice checks, aggregates alerts, and enables automated remediation. Security Hub collects findings from the security services enabled across your AWS accounts, such as intrusion detection findings from Amazon GuardDuty, vulnerability scans from Amazon Inspector, and sensitive data identification findings from Amazon Macie. Security Hub also collects findings from partner security products using a standardized AWS Security Finding Format, eliminating the need for time-consuming data parsing and normalization efforts. Customers can designate an administrator account that can access all findings across their accounts. References: AWS Security Hub Overview, AWS Security Hub FAQs

NEW QUESTION 6

- (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 7

- (Topic 3)

A company needs to engage third-party consultants to help maintain and support its AWS environment and the company's business needs.

Which AWS service or resource will meet these requirements?

- A. AWS Support
- B. AWS Organizations
- C. AWS Service Catalog
- D. AWS Partner Network (APN)

Answer: D

Explanation:

The AWS service or resource that will meet these requirements is D. AWS Partner Network (APN).

AWS Partner Network (APN) is a global community of consulting and technology partners that offer a wide range of services and solutions for AWS customers. APN partners can help customers design, architect, build, migrate, and manage their workloads and applications on AWS. APN partners have access to various resources, training, tools, and support to enhance their AWS expertise and deliver value to customers¹².

AWS Support is a service that provides technical assistance and guidance for AWS customers. AWS Support offers different plans with varying levels of response time, access channels, and features. AWS Support does not directly engage third-party consultants, but rather connects customers with AWS experts and resources³.

AWS Organizations is a service that allows customers to manage multiple AWS accounts within a single organization. AWS Organizations enables customers to create groups of accounts, apply policies, automate account creation, and consolidate billing. AWS Organizations does not directly engage third-party consultants, but rather helps customers simplify and optimize their AWS account management⁴.

AWS Service Catalog is a service that allows customers to create and manage catalogs of IT services that are approved for use on AWS. AWS Service Catalog enables customers to control the configuration, deployment, and governance of their IT services. AWS Service Catalog does not directly engage third-party

consultants, but rather helps customers standardize and streamline their IT service delivery⁵.

References:

1: AWS Partner Network (APN) - Amazon Web Services (AWS) 2: Find an APN Partner - Amazon Web Services (AWS) 3: AWS Support – Amazon Web Services 4: AWS Organizations – Amazon Web Services 5: AWS Service Catalog – Amazon Web Services

NEW QUESTION 8

- (Topic 3)

A company is planning to migrate to the AWS Cloud and wants to become more responsive to customer inquiries and feedback. The company wants to focus on organizational transformation.

A company wants to give its customers the ability to view specific data that is hosted in Amazon S3 buckets. The company wants to keep control over the full datasets that the company shares with the customers.

Which S3 feature will meet these requirements?

- A. S3 Storage Lens
- B. S3 Cross-Region Replication (CRR)
- C. S3 Versioning
- D. S3 Access Points

Answer: D

Explanation:

S3 Access Points are a feature of Amazon S3 that allows you to easily manage access to specific data that is hosted in S3 buckets. S3 Access Points are unique hostnames that customers can use to access data in S3 buckets. You can create multiple access points for a single bucket, each with its own name and permissions. You can use S3 Access Points to provide different levels of access to different groups of customers, such as read-only or write-only access. You can also use S3 Access Points to enforce encryption or logging requirements for specific data. S3 Access Points help you keep control over the full datasets that you share with your customers, while simplifying the access management and improving the performance and scalability of your applications.

NEW QUESTION 9

- (Topic 3)

A cloud practitioner needs to obtain AWS compliance reports before migrating an environment to the AWS Cloud How can these reports be generated?

- A. Contact the AWS Compliance team
- B. Download the reports from AWS Artifact
- C. Open a case with AWS Support
- D. Generate the reports with Amazon Made

Answer: B

Explanation:

AWS Artifact is a service that provides on-demand access to security and compliance reports from AWS and Independent Software Vendors (ISVs) who sell their products on AWS Marketplace. You can use AWS Artifact to download auditor-issued reports, certifications, accreditations, and other third-party attestations of AWS compliance with various standards and regulations, such as PCI-DSS, HIPAA, FedRAMP, GDPR, and more¹²³⁴. You can also use AWS Artifact to review, accept, and manage your agreements with AWS and apply them to current and future accounts within your organization². References: 1: Cloud Compliance - Amazon Web Services

(AWS), 2: Security Compliance Management - AWS Artifact - AWS, 3: AWS Compliance Contact Us - Amazon Web Services, 4: AWS SECURITY AND COMPLIANCE QUICK REFERENCE GUIDE

NEW QUESTION 10

- (Topic 3)

A company needs to migrate a PostgreSQL database from on-premises to Amazon RDS. Which AWS service or tool should the company use to meet this requirement?

- A. Cloud Adoption Readiness Tool
- B. AWS Migration Hub
- C. AWS Database Migration Service (AWS DMS)
- D. AWS Application Migration Service

Answer: C

Explanation:

AWS Database Migration Service (AWS DMS) is a managed and automated service that helps you migrate your databases from your on-premises or cloud environment to AWS, either as a one-time migration or as a continuous replication. AWS DMS supports migration between 20-plus database and analytics engines, such as PostgreSQL, Oracle, MySQL, SQL Server, MongoDB, Amazon Aurora, Amazon RDS, Amazon Redshift, and Amazon S3. AWS DMS also provides schema conversion and validation tools, as well as monitoring and security features. AWS DMS is a cost-effective and reliable solution for database migration, as you only pay for the compute resources and additional log storage used during the migration process, and you can minimize the downtime and data loss with

Multi-AZ and ongoing replication¹²

To migrate a PostgreSQL database from on-premises to Amazon RDS using AWS DMS, you need to perform the following steps:

? Create an AWS DMS replication instance in the same AWS Region as your target Amazon RDS PostgreSQL DB instance. The replication instance is a server that runs the AWS DMS replication software and connects to your source and target endpoints. You can choose the instance type, storage, and network settings based on your migration requirements³

? Create a source endpoint that points to your on-premises PostgreSQL database.

You need to provide the connection details, such as the server name, port, database name, user name, and password. You also need to specify the engine name as postgres and the SSL mode as required⁴

? Create a target endpoint that points to your Amazon RDS PostgreSQL DB instance. You need to provide the connection details, such as the server name, port, database name, user name, and password. You also need to specify the engine name as postgres and the SSL mode as verify-full.

? Create a migration task that defines the migration settings and options, such as the replication instance, the source and target endpoints, the migration type (full load, full load and change data capture, or change data capture only), the table mappings, the task settings, and the task monitoring role. You can also use the AWS Schema Conversion Tool (AWS SCT) to convert your source schema to the target schema and apply it to the target endpoint before or after creating the migration task.

? Start the migration task and monitor its progress and status using the AWS DMS console, the AWS CLI, or the AWS DMS API. You can also use AWS CloudFormation to automate the creation and execution of the migration task.

The other options are not suitable for migrating a PostgreSQL database from on-premises to Amazon RDS. Cloud Adoption Readiness Tool is a tool that helps you assess your readiness for cloud adoption based on six dimensions: business, people, process, platform, operations, and security. It does not perform any database migration tasks. AWS Migration Hub is a service that helps you track and manage the progress of your application migrations across multiple AWS and partner services, such as AWS DMS, AWS Application Migration Service, AWS Server Migration Service, and CloudEndure Migration. It does not perform any database migration tasks itself, but rather integrates with other migration services. AWS Application Migration Service is a service that helps you migrate your applications from your on-premises or cloud environment to AWS without making any changes to the applications, their architecture, or the migrated servers. It does not support database migration, but rather replicates your servers as Amazon Machine Images (AMIs) and launches them as EC2 instances on AWS.

References: AWS Database Migration Service, What is AWS Database Migration Service?, Working with an AWS DMS replication instance, Creating source and target endpoints for PostgreSQL, [Creating a target endpoint for Amazon RDS for PostgreSQL], [Creating a migration task for AWS DMS], [AWS Schema Conversion Tool], [Starting a migration task for AWS DMS], [AWS CloudFormation], [Cloud Adoption Readiness Tool], [AWS Migration Hub], [AWS Application Migration Service]

NEW QUESTION 10

- (Topic 3)

A company wants its AWS usage to be more sustainable. The company wants to track, measure, review, and forecast polluting emissions that result from its AWS applications.

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

- A. AWS Health Dashboard
- B. AWS customer carbon footprint tool
- C. AWS Support Center
- D. Amazon QuickSight

Answer: B

Explanation:

AWS customer carbon footprint tool is a tool that helps customers measure and manage their carbon emissions from their AWS usage. It provides data on the carbon intensity, energy consumption, and estimated emissions of AWS services across regions and time periods. It also enables customers to review and forecast their emissions, and compare them with industry benchmarks. AWS Health Dashboard is a service that provides personalized information about the health and performance of AWS services and resources. AWS Support Center is a service that provides access to AWS support resources, such as cases, forums, and documentation. Amazon QuickSight is a service that provides business intelligence and analytics for AWS data sources.

NEW QUESTION 12

- (Topic 3)

A company needs to deploy applications in the AWS Cloud as quickly as possible. The company also needs to minimize the complexity that is related to the management of AWS resources.

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

- A. AWS config
- B. AWS Elastic Beanstalk
- C. Amazon EC2
- D. Amazon Personalize

Answer: B

Explanation:

AWS Elastic Beanstalk is the AWS service that allows customers to deploy applications in the AWS Cloud as quickly as possible. AWS Elastic Beanstalk automatically handles the deployment, from capacity provisioning, load balancing, and auto-scaling to application health monitoring. Customers can upload their code and Elastic Beanstalk will take care of the rest¹. AWS Elastic Beanstalk also minimizes the complexity that is related to the management of AWS resources. Customers can retain full control of the underlying AWS resources powering their applications and adjust the settings to suit their needs¹. Customers can also use the AWS Management Console, the AWS Command Line Interface (AWS CLI), or APIs to manage their applications¹.

AWS Config is the AWS service that enables customers to assess, audit, and evaluate the configurations of their AWS resources. AWS Config continuously monitors and records the configuration changes of the resources and evaluates them against desired configurations or best practices². AWS Config does not help customers deploy applications in the AWS Cloud as quickly as possible or minimize the complexity that is related to the management of AWS resources.

Amazon EC2 is the AWS service that provides secure, resizable compute capacity in the cloud. Customers can launch virtual servers called instances and choose from various configurations of CPU, memory, storage, and networking resources³. Amazon EC2 does not automatically handle the deployment or management of AWS resources for customers. Customers have to manually provision, configure, monitor, and scale their instances and other related resources.

Amazon Personalize is the AWS service that enables customers to create personalized recommendations for their users based on their behavior and preferences. Amazon Personalize uses machine learning to analyze data and deliver real-time recommendations⁴. Amazon Personalize does not help customers deploy applications in the AWS Cloud as quickly as possible or minimize the complexity that is related to the management of AWS resources.

NEW QUESTION 17

- (Topic 3)

A software engineer wants to launch a virtual machine (VM) and MySQL database on AWS.

Which AWS service will meet these requirements with the LEAST operational effort?

- A. Amazon Elastic Container Service (Amazon ECS)
- B. AWS Elastic Beanstalk
- C. Amazon Lightsail
- D. Amazon EC2

Answer: B

Explanation:

AWS Elastic Beanstalk is a service that enables you to quickly deploy and manage applications in the AWS Cloud without worrying about the infrastructure that runs those applications. You simply upload your application, and Elastic Beanstalk automatically handles the details of capacity provisioning, load balancing, scaling, and application health monitoring. Elastic Beanstalk supports several platform configurations for Java, .NET, PHP, Node.js, Python, Ruby, Go, and Docker web applications that can run on familiar servers such as Apache, Nginx, Passenger, and IIS. You can also use Elastic Beanstalk to launch a virtual machine (VM)

and MySQL database on AWS with the least operational effort. Amazon Elastic Container Service (Amazon ECS) is a fully managed container orchestration service that enables you to easily run, scale, and secure Docker containerized applications on AWS. However, it requires more operational effort than Elastic Beanstalk, as you need to define your application architecture and the specifications of the containers that run it. 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. It is designed for developers who have little or no prior cloud experience and want to launch and manage applications on AWS with minimal complexity. However, it does not support MySQL databases, and it requires more operational effort than Elastic Beanstalk, as you need to configure your VM and database settings. Amazon EC2 is a web service that provides secure, resizable compute capacity in the cloud. It allows you to launch a virtual machine (VM) and MySQL database on AWS, but it requires the most operational effort, as you need to provision, monitor, and manage your EC2 instances and database.

NEW QUESTION 20

- (Topic 3)

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

- A. Patch the Amazon RDS operating system.
- B. Upgrade the firmware of the network infrastructure.
- C. Manage data encryption.
- D. Maintain physical access control in an AWS Region.
- E. Grant least privilege access to IAM users.

Answer: CE

Explanation:

According to the AWS shared responsibility model, the customer is responsible for security in the cloud, which includes the tasks of managing data encryption and granting least privilege access to IAM users. Data encryption is the process of transforming data into an unreadable format that can only be accessed with a key or a password. The customer must decide whether to encrypt their data at rest (when it is stored on AWS) or in transit (when it is moving between AWS and the customer or between AWS services). The customer must also choose the encryption method, algorithm, and key management solution that best suit their needs. AWS provides various services and features that support data encryption, such as AWS Key Management Service (AWS KMS), AWS Certificate Manager (ACM), and AWS Encryption SDK. IAM users are entities that represent the people or applications that interact with AWS resources and services.

The customer must grant the IAM users the minimum permissions that they need to perform their tasks, and avoid giving them unnecessary or excessive access. This is known as the principle of least privilege, and it helps reduce the risk of unauthorized or malicious actions. The customer can use IAM policies, roles, groups, and permissions boundaries to manage the access of IAM users.

NEW QUESTION 24

- (Topic 3)

Which AWS service or feature is associated with a subnet in a VPC and is used to control inbound and outbound traffic?

- A. Amazon Inspector
- B. Network ACLs
- C. AWS Shield
- D. VPC Flow Logs

Answer: B

Explanation:

Network ACLs (network access control lists) are an optional layer of security for your VPC that act as a firewall for controlling traffic in and out of one or more subnets. You can use network ACLs to allow or deny traffic based on protocol, port, or source and destination IP address. Network ACLs are stateless, meaning that they do not track the traffic that flows through them. Therefore, you must create rules for both inbound and outbound traffic.

NEW QUESTION 28

- (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 29

- (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 Snowmobile¹²³⁴. 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 34

- (Topic 3)

Which AWS service provides the ability to manage infrastructure as code?

- A. AWS CodePipeline
- B. AWS CodeDeploy
- C. AWS Direct Connect
- D. AWS CloudFormation

Answer: D

Explanation:

The AWS service that provides the ability to manage infrastructure as code is AWS CloudFormation. Infrastructure as code is a process of defining and provisioning AWS resources using code or templates, rather than manual actions or scripts. AWS CloudFormation allows you to create and update stacks of AWS resources based on predefined templates that describe the desired state and configuration of the resources. AWS CloudFormation automates and simplifies the deployment and management of AWS resources, and ensures consistency and repeatability across different environments and regions. AWS CloudFormation also supports rollback, change sets, drift detection, and nested stacks features that help you to monitor and control the changes to your infrastructure¹.

NEW QUESTION 39

- (Topic 3)

Which AWS services can be used to store files? (Select TWO.)

- A. Amazon S3
- B. AWS Lambda
- C. Amazon Elastic Block Store (Amazon EBS)
- D. Amazon SageMaker
- E. AWS Storage Gateway

Answer: AC

Explanation:

Amazon S3 and Amazon EBS are two AWS services that can be used to store files . Amazon S3 is an object storage service that offers high scalability, durability, availability, and performance. Amazon EBS is a block storage service that provides persistent and low-latency storage volumes for Amazon EC2 instances. AWS Lambda, Amazon SageMaker, and AWS Storage Gateway are other AWS services that have different purposes, such as serverless computing, machine learning, and hybrid cloud storage .

NEW QUESTION 41

- (Topic 3)

A company is running its application in the AWS Cloud. The company wants to periodically review its AWS account for cost optimization opportunities.

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

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

Answer: A

Explanation:

AWS Cost Explorer is an AWS service or tool that the company can use to periodically review its AWS account for cost optimization opportunities. AWS Cost Explorer is a tool that enables the company to visualize, understand, and manage their AWS costs and usage over time. The company can use AWS Cost Explorer to access interactive graphs and tables that show the breakdown of their costs and usage by service, region, account, tag, and more. The company can also use AWS Cost Explorer to forecast their future costs, identify trends and anomalies, and discover potential savings by using Reserved Instances or Savings Plans.

NEW QUESTION 46

- (Topic 3)

Elasticity in the AWS Cloud refers to which of the following? (Select TWO.)

- A. How quickly an Amazon EC2 instance can be restarted
- B. The ability to rightsized resources as demand shifts
- C. The maximum amount of RAM an Amazon EC2 instance can use
- D. The pay-as-you-go billing model
- E. How easily resources can be procured when they are needed

Answer: BE

Explanation:

Elasticity in the AWS Cloud refers to the ability to acquire resources as you need them and release resources when you no longer need them. In the cloud, you want to do this automatically¹. This means that you can rightsized resources as demand shifts, and you can easily procure resources when they are needed. Elasticity is not related to how quickly an Amazon EC2 instance can be restarted, the maximum amount of RAM an Amazon EC2 instance can use, or the pay-as-you-go billing model. These are aspects of scalability, performance, and cost, respectively².

For more information on elasticity, you can refer to the following sources:

? Elasticity - AWS Well-Architected Framework

? Elastic - Reactive Systems on AWS

? What is the difference between scalability and elasticity?

NEW QUESTION 49

- (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)⁵.

NEW QUESTION 50

- (Topic 3)

What is a customer responsibility when using AWS Lambda according to the AWS shared responsibility model?

- A. Managing the code within the Lambda function
- B. Confirming that the hardware is working in the data center
- C. Patching the operating system
- D. Shutting down Lambda functions when they are no longer in use

Answer: A

Explanation:

According to the AWS shared responsibility model, AWS is responsible for the security of the cloud, while customers are responsible for the security in the cloud. This means that AWS is responsible for the physical servers, networking, and operating system that run Lambda functions, while customers are responsible for the security of their code and AWS IAM to the Lambda service and within their function¹. Customers need to manage the code within the Lambda function, such as writing, testing, debugging, deploying, and updating the code, as well as ensuring that the code does not contain any vulnerabilities or malicious code that could compromise the security or performance of the function²³. References: ²: AWS Lambda - Amazon Web Services (AWS), ³: AWS Lambda Documentation, ¹: Amazon CLF-C02: What is customer responsibility under AWS ... - PUPUWEB

NEW QUESTION 55

- (Topic 3)

What is the purpose of having an internet gateway within a VPC?

- A. To create a VPN connection to the VPC
- B. To allow communication between the VPC and the internet
- C. To impose bandwidth constraints on internet traffic
- D. To load balance traffic from the internet across Amazon EC2 instances

Answer: B

Explanation:

An internet gateway is a service that allows for internet traffic to enter into a VPC. Otherwise, a VPC is completely segmented off and then the only way to get to it is potentially through a VPN connection rather than through internet connection. An internet gateway is a logical connection between an AWS VPC and the internet. It supports IPv4 and IPv6 traffic. It does not cause availability risks or bandwidth constraints on your network traffic¹. An internet gateway enables resources in your public subnets (such as EC2 instances) to connect to the internet if the resource has a public IPv4 address or an IPv6 address. Similarly, resources on the internet can initiate a connection to resources in your subnet using the public IPv4 address or IPv6 address². An internet gateway also provides a target in your VPC route tables for internet-routable traffic. For communication using IPv4, the internet gateway also performs network address translation (NAT). For communication using IPv6, NAT is not needed because IPv6 addresses are public². To enable access to or from the internet for instances in a subnet in a VPC using an internet gateway, you must create an internet gateway and attach it to your VPC, add a route to your subnet's route table that directs internet-bound traffic to the internet gateway, ensure that instances in your subnet have a public IPv4 address or an IPv6 address, and ensure that your network access control lists and security group rules allow the desired internet traffic to flow to and from your instance². References: Connect to the internet using an internet gateway, AWS Internet Gateway and VPC Routing

NEW QUESTION 56

- (Topic 3)

A company wants to query its server logs to gain insights about its customers' experiences. Which AWS service will store this data MOST cost-effectively?

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

Answer: D

Explanation:

Amazon S3 is an AWS service that provides scalable, durable, and cost-effective object storage in the cloud. Amazon S3 can store any amount and type of data,

such as server logs, and offers various storage classes with different performance and pricing characteristics. Amazon S3 is the most cost-effective option for storing server logs, as it offers low-cost storage classes, such as S3 Standard-Infrequent Access (S3 Standard-IA) and S3 Intelligent-Tiering, that are suitable for infrequently accessed or changing access patterns data. Amazon S3 also integrates with other AWS services, such as Amazon Athena and Amazon OpenSearch Service, that can query the server logs directly from S3 without requiring any additional data loading or transformation. References: Amazon S3, Amazon S3 Storage Classes, Querying Data in Amazon S3

NEW QUESTION 59

- (Topic 3)

Which capabilities are in the platform perspective of the AWS Cloud Adoption Framework (AWS CAF)? (Select TWO.)

- A. Performance and capacity management
- B. Data engineering
- C. Continuous integration and continuous delivery (CI/CD)
- D. Infrastructure protection
- E. Change and release management

Answer: BC

Explanation:

These are two of the seven capabilities that are in the platform perspective of the AWS Cloud Adoption Framework (AWS CAF). The platform perspective helps you build an enterprise-grade, scalable, hybrid cloud platform, modernize existing workloads, and implement new cloud-native solutions¹. The other five capabilities are:

? Platform architecture – Establish and maintain guidelines, principles, patterns, and guardrails for your cloud environment.

? Platform engineering – Build a compliant multi-account cloud environment with enhanced security features, and packaged, reusable cloud products.

? Platform operations – Manage and optimize your cloud environment with automation, monitoring, and incident response.

? Application development – Develop and deploy cloud-native applications using modern architectures and best practices.

? Application migration – Migrate your existing applications to the cloud using proven methodologies and tools.

Performance and capacity management, infrastructure protection, and change and release management are not capabilities of the platform perspective. They are part of the operations perspective, which helps you achieve operational excellence in the cloud². The operations perspective comprises six capabilities:

? Performance and capacity management – Monitor and optimize the performance and capacity of your cloud workloads.

? Infrastructure protection – Protect your cloud infrastructure from unauthorized access, malicious attacks, and data breaches.

? Change and release management – Manage changes and releases to your cloud workloads using automation and governance.

? Configuration management – Manage the configuration of your cloud resources and applications using automation and version control.

? Incident management – Respond to incidents affecting your cloud workloads using best practices and tools.

? Service continuity management – Ensure the availability and resilience of your cloud workloads using backup, recovery, and disaster recovery strategies.

NEW QUESTION 61

- (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 64

- (Topic 3)

A company is moving to the AWS Cloud to reduce operational overhead for its application infrastructure.

Which IT operation will the company still be responsible for after the migration to AWS?

- A. Security patching of AWS Elastic Beanstalk
- B. Backups of data that is stored in Amazon Aurora
- C. Termination of Amazon EC2 instances that are managed by AWS Auto Scaling
- D. Configuration of IAM access controls

Answer: D

Explanation:

AWS Elastic Beanstalk, Amazon Aurora, and AWS Auto Scaling are managed services that reduce the operational overhead for the customers. AWS is responsible for security patching, backups, and termination of these services. However, the customers are still responsible for configuring IAM access controls to manage the permissions and policies for their AWS resources. This is part of the AWS shared responsibility model, which defines the security and compliance responsibilities of AWS and the customers. You can learn more about the AWS shared responsibility model from this whitepaper or this digital course.

NEW QUESTION 65

- (Topic 3)

Which cloud concept is demonstrated by using AWS Compute Optimizer?

- A. Security validation
- B. Rightsizing
- C. Elasticity
- D. Global reach

Answer: B

Explanation:

Rightsizing is the cloud concept that is demonstrated by using AWS Compute Optimizer. Rightsizing is the process of adjusting the type and size of your cloud resources to match the optimal performance and cost for your workloads. AWS Compute Optimizer is a service that analyzes the configuration and utilization metrics of your AWS resources, such as Amazon EC2 instances, Amazon EBS volumes, AWS Lambda functions, and Amazon ECS services on AWS Fargate. It reports whether your resources are optimal, and generates optimization recommendations to reduce the cost and improve the performance of your workloads. AWS Compute Optimizer uses machine learning to analyze your historical utilization data and compare it with the most cost-effective AWS alternatives. You can use the recommendations to evaluate the trade-offs between cost and performance, and decide when to move or resize your resources to achieve the best results. References: Workload Rightsizing - AWS Compute Optimizer - AWS, What is AWS Compute Optimizer? - AWS Compute Optimizer

NEW QUESTION 70

- (Topic 3)

An auditor is preparing for an annual security audit. The auditor requests certification details for a company's AWS hosted resources across multiple Availability Zones in the us-east-1 Region.

How should the company respond to the auditor's request?

- A. Open an AWS Support ticket to request that the AWS technical account manager (TAM) respond and help the auditor.
- B. Open an AWS Support ticket to request that the auditor receive approval to conduct an onsite assessment of the AWS data centers in which the company operates.
- C. Explain to the auditor that AWS does not need to be audited because the company's application is hosted in multiple Availability Zones.
- D. Use AWS Artifact to download the applicable report for AWS security control
- E. Provide the report to the auditor.

Answer: D

Explanation:

AWS Artifact is your go-to, central resource for compliance-related information that matters to you. It provides on-demand access to AWS' security and compliance reports and select online agreements. Reports available in AWS Artifact include our Service Organization Control (SOC) reports, Payment Card Industry (PCI) reports, and certifications from accreditation bodies across geographies and compliance verticals that validate the implementation and operating effectiveness of AWS security controls. Agreements available in AWS Artifact include the Business Associate Addendum (BAA) and the Nondisclosure Agreement (NDA). You can use AWS Artifact to download the applicable report for AWS security controls and provide it to the auditor.

NEW QUESTION 73

- (Topic 3)

A company is running an Amazon EC2 instance in a VPC.

An ecommerce company is using Amazon EC2 Auto Scaling groups to manage a fleet of web servers running on Amazon EC2.

This architecture follows which AWS Well-Architected Framework best practice?

- A. Secure the workload
- B. Decouple infrastructure components
- C. Design for failure
- D. Think parallel

Answer: C

Explanation:

Design for failure is one of the best practices of the AWS Well-Architected Framework. It means that the architecture should be resilient and fault-tolerant, and able to handle failures without impacting the availability and performance of the applications. By using Amazon EC2 Auto Scaling groups, the ecommerce company can design for failure by automatically scaling the number of EC2 instances up or down based on demand or health status. Amazon EC2 Auto Scaling groups can also distribute the EC2 instances across multiple Availability Zones, which are isolated locations within an AWS Region that have independent power, cooling, and network connectivity. This way, the company can ensure that their web servers can handle traffic spikes, recover from failures, and provide a consistent user experience

NEW QUESTION 76

- (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 80

- (Topic 3)

Which of the following is a benefit of using an AWS managed service?

- A. Reduced operational overhead for a company's IT staff
- B. Increased fixed costs that can be predicted by a finance team
- C. Removal of the need to have a backup strategy
- D. Removal of the need to follow compliance standards

Answer: A

Explanation:

This is a benefit of using an AWS managed service, such as Amazon S3, Amazon DynamoDB, or AWS Lambda. AWS managed services are fully managed by AWS, which means that AWS handles the provisioning, scaling, patching, backup, and recovery of the underlying infrastructure and software. This reduces the operational overhead for the company's IT staff, who can focus on their core business logic and innovation. You can learn more about the AWS managed services from this webpage or this digital course.

NEW QUESTION 83

- (Topic 3)

A company plans to migrate to the AWS Cloud. The company wants to use the AWS Cloud Adoption Framework (AWS CAF) to define and track business outcomes as part of its cloud transformation journey.

Which AWS CAF governance perspective capability will meet these requirements?

- A. Benefits management
- B. Risk management
- C. Application portfolio management
- D. Cloud financial management

Answer: A

Explanation:

The correct answer is A. Benefits management.

Benefits management is the AWS CAF governance perspective capability that helps you define and track business outcomes as part of your cloud transformation journey. Benefits management helps you align your cloud initiatives with your business objectives, measure the value and impact of your cloud investments, and communicate the benefits of cloud adoption to your stakeholders¹².

Risk management is the AWS CAF governance perspective capability that helps you identify and mitigate the potential risks associated with cloud adoption, such as security, compliance, legal, and operational risks¹².

Application portfolio management is the AWS CAF governance perspective capability that helps you assess and optimize your existing application portfolio for cloud migration or modernization. Application portfolio management helps you categorize your applications based on their business value and technical fit, prioritize them for cloud adoption, and select the best migration or modernization strategy for each application¹².

Cloud financial management is the AWS CAF governance perspective capability that helps you manage and optimize the costs and value of your cloud resources. Cloud financial management helps you plan and budget for cloud adoption, track and allocate cloud costs, implement cost optimization strategies, and report on cloud financial performance¹². References:

1: AWS Cloud Adoption Framework: Governance Perspective 2: All you need to know about AWS Cloud Adoption Framework — Governance Perspective

NEW QUESTION 84

- (Topic 3)

A company wants to use guidelines from the AWS Well-Architected Framework to limit human error and facilitate consistent responses to events.

Which of the following is a Well-Architected design principle that will meet these requirements?

- A. Use AWS CodeDeploy.
- B. Perform operations as code.
- C. Migrate workloads to a Dedicated Host.
- D. Use AWS Compute Optimizer.

Answer: B

Explanation:

This is a design principle of the operational excellence pillar of the AWS Well-Architected Framework. Performing operations as code means using scripts, templates, or automation tools to perform routine tasks, such as provisioning, configuration, deployment, and monitoring. This reduces human error, increases consistency, and enables faster recovery from failures. You can learn more about the operational excellence pillar from this whitepaper or this digital course.

NEW QUESTION 85

- (Topic 3)

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

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

Answer: D

Explanation:

Amazon ElastiCache is a fully managed in-memory data store and cache service that delivers sub-millisecond response times to applications. You can use ElastiCache as a primary data store for your applications, or as a cache to improve the performance of your existing databases. ElastiCache supports two popular open-source in- memory engines: Redis and Memcached⁵.

NEW QUESTION 86

- (Topic 3)

A company is launching a mobile app. The company wants customers to be able to use the app without upgrading their mobile devices.

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

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

Answer: C

Explanation:

Cost optimization is one of the five pillars of the AWS Well-Architected Framework. It focuses on avoiding unnecessary costs, understanding and controlling where money is being spent, selecting the most appropriate and right number of resource types, analyzing spend over time, and scaling to meet business needs without overspending.

NEW QUESTION 87

- (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 define⁶⁷.

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

breached⁸⁹. 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 91

- (Topic 3)

A company is assessing its AWS Business Support plan to determine if the plan still meets the company's needs. The company is considering switching to AWS Enterprise Support.

Which additional benefit will the company receive with AWS Enterprise Support?

- A. A full set of AWS Trusted Advisor checks
- B. Phone, email, and chat access to cloud support engineers 24 hours a day, 7 days a week
- C. A designated technical account manager (TAM) to assist in monitoring and optimization
- D. A consultative review and architecture guidance for the company's applications

Answer: C

Explanation:

AWS Enterprise Support provides customers with a designated technical account manager (TAM) who is a single point of contact for all technical and operational issues. The TAM provides consultative architectural and operational guidance delivered in the context of the customer's applications and use-cases to help them achieve the greatest value from AWS. The TAM also helps customers with proactive services, such as strategic business reviews, security improvement programs, guided Well-Architected reviews, cost optimization workshops, and more¹.

A full set of AWS Trusted Advisor checks is not an additional benefit of AWS Enterprise Support, as it is also included in the AWS Business Support plan². AWS Trusted Advisor is a tool that provides best practice recommendations for cost optimization, performance, security, fault tolerance, and service limits.

Phone, email, and chat access to cloud support engineers 24 hours a day, 7 days a week is not an additional benefit of AWS Enterprise Support, as it is also included in the AWS Business Support plan². Cloud support engineers can help customers with technical issues, such as troubleshooting, configuration, usage, and service features.

A consultative review and architecture guidance for the company's applications is not an additional benefit of AWS Enterprise Support, as it is also included in the AWS Business Support plan². Customers can request a consultative review from a solutions architect who will provide best practices and recommendations based on the customer's use-cases and goals.

NEW QUESTION 94

- (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 98

- (Topic 3)

Which AWS services are supported by Savings Plans? (Select TWO.)

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

Answer: AC

Explanation:

The AWS services that are supported by Savings Plans are:

? Amazon EC2: Amazon EC2 is a service that provides scalable computing capacity in the AWS cloud. You can use Amazon EC2 to launch virtual servers, configure security and networking, and manage storage. Amazon EC2 is eligible for both Compute Savings Plans and EC2 Instance Savings Plans¹².

? Amazon SageMaker: Amazon SageMaker is a service that helps you build and deploy machine learning models. You can use Amazon SageMaker to access Jupyter notebooks, use common machine learning algorithms, train and tune models, and deploy them to a hosted environment. Amazon SageMaker is eligible for SageMaker Savings Plans¹³.

The other options are not supported by Savings Plans. Amazon RDS, Amazon Redshift, and Amazon DynamoDB are database services that are eligible for Reserved Instances, but not Savings Plans⁴.

NEW QUESTION 103

- (Topic 3)

A company is building an application in the AWS Cloud. The company wants to use temporary credentials for the application to access other AWS resources.

Which AWS service will meet these requirements?

- A. AWS Key Management Service (Aws KMS)
- B. AWS CloudHSM
- C. Amazon Cognito
- D. AWS Security Token Service (Aws STS)

Answer: D

Explanation:

AWS Security Token Service (AWS STS) is a service that provides temporary security credentials to users or applications that need to access AWS resources. The temporary credentials have a limited lifetime and can be configured to last from a few minutes to several hours. The credentials are not stored with the user or application, but are generated dynamically and provided on request. The credentials work almost identically to long-term access key credentials, but have the advantage of not requiring distribution, rotation, or revocation¹.

AWS Key Management Service (AWS KMS) is a service that provides encryption and decryption services for data and keys. It does not provide temporary security credentials². AWS CloudHSM is a service that provides hardware security modules (HSMs) for cryptographic operations and key management. It does not provide temporary security credentials³.

Amazon Cognito is a service that provides user authentication and authorization for web and mobile applications. It can also provide temporary security credentials for authenticated users, but not for applications⁴.

NEW QUESTION 106

- (Topic 3)

Which option is the default pricing model for Amazon EC2 instances?

- A. On-Demand Instances
- B. Savings Plans
- C. Spot Instances
- D. Reserved Instances

Answer: A

Explanation:

On-Demand Instances are the default pricing model for Amazon EC2 instances. They allow users to pay for compute capacity by the second, with no long-term commitments or upfront payments. They are suitable for applications with short-term, irregular, or unpredictable workloads that cannot be interrupted³. Savings Plans are a pricing model that offer significant savings on Amazon EC2 and AWS Fargate usage, in exchange for a commitment to a consistent amount of usage (measured in \$/hour) for a 1- year or 3-year term. Spot Instances are a pricing model that offer spare Amazon EC2 compute capacity at up to 90% discount compared to On-Demand prices, but they can be interrupted by AWS with a two-minute notice when the demand exceeds the supply. Reserved Instances are a pricing model that offer up to 75% discount compared to On- Demand prices, in exchange for a commitment to use a specific instance type and size in a specific region for a 1-year or 3-year term.

NEW QUESTION 110

- (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 111

- (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 116

- (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 120

- (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 121

- (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 122

- (Topic 3)

A company wants an automated process to continuously scan its Amazon EC2 instances for software vulnerabilities. Which AWS service will meet these requirements?

- A. Amazon GuardDuty
- B. Amazon Inspector
- C. Amazon Detective
- D. Amazon Cognito

Answer: B

Explanation:

Amazon Inspector is the AWS service that can be used to perform vulnerability scans on AWS EC2 instances for software vulnerabilities automatically in a periodic fashion. Amazon Inspector automatically discovers EC2 instances and scans them for software vulnerabilities and unintended network exposure. Amazon Inspector uses AWS Systems Manager (SSM) and the SSM Agent to collect information about the software application inventory of the EC2 instances. This data is then scanned by Amazon Inspector for software vulnerabilities¹². Amazon Inspector also integrates with other AWS services, such as Amazon EventBridge and AWS Security Hub, to automate discovery, expedite vulnerability routing, and shorten mean time to remediate (MTTR) vulnerabilities².

NEW QUESTION 123

- (Topic 3)

Which AWS services make use of global edge locations'? (Select TWO.)

- A. AWS Fargate
- B. Amazon CloudFront
- C. AWS Global Accelerator
- D. AWS Wavelength
- E. Amazon VPC

Answer: BC

Explanation:

Amazon CloudFront and AWS Global Accelerator are two AWS services that make use of global edge locations. Edge locations are AWS sites that are deployed worldwide in major cities and places with a high population. Edge locations are used to cache data and reduce latency for end-user access¹. Amazon CloudFront is a content delivery network (CDN) service that securely delivers data, videos, applications, and APIs to customers globally with low latency and high transfer speeds. Amazon CloudFront uses a global network of over 200 edge locations and 13 regional edge caches to cache your content closer to your viewers, improving performance and reducing costs²³. AWS Global Accelerator is a networking service that improves the availability and performance of your applications with local or global users. AWS Global Accelerator uses the AWS global network to route user traffic to the optimal endpoint based on health, performance, and policies. AWS Global Accelerator uses over 100 edge locations to bring your application endpoints closer to your users, reducing network hops and improving user experience⁴⁵. References: 1: AWS for the Edge - Amazon Web Services (AWS), 2: Content Delivery Network (CDN) - Amazon CloudFront - AWS, 3: Amazon CloudFront Documentation, 4: AWS Global Accelerator - Amazon Web Services, 5: AWS Global Accelerator Documentation

NEW QUESTION 125

- (Topic 3)

A company wants to establish a private network connection between AWS and its corporate network. Which AWS service or feature will meet this requirement?

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

Answer: C

Explanation:

AWS Direct Connect is a cloud service solution that makes it easy to establish a dedicated network connection from your premises to AWS. Using AWS Direct Connect, you can establish private connectivity between AWS and your datacenter, office, or colocation environment, which in many cases can reduce your network costs, increase bandwidth throughput, and provide a more consistent network experience than internet-based connections¹². References: 1: Dedicated Network Connection - AWS Direct Connect - AWS, 2: What is AWS Direct Connect? - AWS Direct Connect

NEW QUESTION 126

- (Topic 3)

Which of the following is a fully managed graph database service on AWS?

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

Answer: D

Explanation:

Amazon Neptune is a fully managed graph database service on AWS. A graph database is a type of database that stores and queries data as a network of nodes and edges, representing entities and relationships. Graph databases are useful for applications that deal with highly connected data, such as social networks, recommendation engines, fraud detection, and knowledge graphs⁴⁵. Amazon Neptune is a fast, reliable, and scalable graph database service that supports two popular graph models: property graphs and RDF. Amazon Neptune also supports two open standards for querying graphs: Apache TinkerPop Gremlin and SPARQL. Amazon Neptune handles the heavy lifting of managing the database, such as provisioning, patching, backup, recovery, encryption, and replication⁴⁵⁶. References: 4: Managed Graph Database - Amazon Neptune - AWS, 5: Amazon Neptune – A Fully Managed Graph Database Service, 6: Working with AWS Neptune. Neptune is a fully-managed graph ... - Medium

NEW QUESTION 131

- (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 135

- (Topic 3)

Which perspective in the AWS Cloud Adoption Framework (AWS CAF) includes a capability for well-designed data and analytics architecture?

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

Answer: D

Explanation:

The correct answer is D. Platform.

The Platform perspective in the AWS Cloud Adoption Framework (AWS CAF) includes a capability for well-designed data and analytics architecture. This capability helps you design, implement, and optimize your data and analytics solutions on AWS, using services such as Amazon S3, Amazon Redshift, Amazon EMR, Amazon Kinesis, Amazon Athena, and Amazon QuickSight. A well-designed data and analytics architecture enables you to collect, store, process, analyze, and visualize data from various sources, and derive insights that can drive your business decisions¹².

The Security perspective does not include a capability for data and analytics architecture, but it does include a capability for data protection, which helps you secure your data at rest and in transit using encryption, key management, access control, and auditing¹³.

The Governance perspective does not include a capability for data and analytics architecture, but it does include a capability for data governance, which helps you manage the quality, availability, usability, integrity, and security of your data assets¹⁴.

The Operations perspective does not include a capability for data and analytics architecture, but it does include a capability for data operations, which helps you monitor, troubleshoot, and optimize the performance and availability of your data pipelines and workloads¹.

References:

1: Foundational capabilities - An Overview of the AWS Cloud Adoption Framework 2: [AWS Cloud Adoption Framework: Platform Perspective] 3: [AWS Cloud Adoption Framework: Security Perspective] 4: [AWS Cloud Adoption Framework: Governance Perspective] : [AWS Cloud Adoption Framework: Operations Perspective]

NEW QUESTION 140

- (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 145

- (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 150

- (Topic 3)

A company is using Amazon DynamoDB.

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

- A. Patch the operating system
- B. Provision hosts
- C. Manage database access permissions.
- D. Secure the operating system

Answer: C

Explanation:

According to the AWS shared responsibility model, AWS is responsible for the security of the cloud, while customers are responsible for the security in the cloud. This means that AWS is responsible for the physical servers, networking, and operating system that run DynamoDB, while customers are responsible for the security of their data and access to the database. Customers need to manage database access permissions, such as creating and managing AWS Identity and Access Management (IAM) policies and roles, and using encryption and key management options to protect their data¹²³. References: 1: Shared Responsibility Model - Amazon Web Services (AWS), 2: Security in Amazon DynamoDB - Amazon DynamoDB, 3: AWS Shared Responsibility Model - Introduction to DevOps ...

NEW QUESTION 153

- (Topic 3)

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

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

Answer: C

Explanation:

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

NEW QUESTION 157

- (Topic 3)

An ecommerce company wants to provide relevant product recommendations to its customers. The recommendations will include products that are frequently purchased with other products that the customer already purchased. The recommendations also will include products of a specific color and products from the

customer's favorite brand.

Which AWS service or feature should the company use to meet these requirements with the LEAST development effort?

- A. Amazon Comprehend
- B. Amazon Forecast
- C. Amazon Personalize
- D. Amazon SageMaker Studio

Answer: C

Explanation:

Amazon Personalize is a service that provides real-time personalized recommendations based on the user's behavior, preferences, and context. It can also incorporate metadata such as product color and brand to generate more relevant recommendations. Amazon Comprehend is a natural language processing (NLP) service that can analyze text for entities, sentiments, topics, and more. Amazon Forecast is a service that provides accurate time-series forecasting based on machine learning. Amazon SageMaker Studio is a web-based integrated development environment (IDE) for machine learning.

NEW QUESTION 160

- (Topic 3)

A user has a stateful workload that will run on Amazon EC2 for the next 3 years. What is the MOST cost-effective pricing model for this workload?

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

Answer: B

Explanation:

Reserved Instances are a pricing model that offers significant discounts on Amazon EC2 usage compared to On-Demand Instances. Reserved Instances are suitable for stateful workloads that have predictable and consistent usage patterns for a long-term period. By committing to a one-year or three-year term, customers can reduce their total cost of ownership and optimize their cloud spend. Reserved Instances also provide capacity reservation, ensuring that customers have access to the EC2 instances they need when they need them. References: AWS Pricing Calculator, Amazon EC2 Pricing, [AWS Cloud Practitioner Essentials: Module 3 - Compute in the Cloud]

NEW QUESTION 161

- (Topic 3)

Which AWS service converts text to lifelike voices?

- A. Amazon Transcribe
- B. Amazon Rekognition
- C. Amazon Polly
- D. Amazon Textract

Answer: C

Explanation:

Amazon Polly is a service that turns text into lifelike speech, allowing you to create applications that talk, and build entirely new categories of speech-enabled products. Polly's Text-to-Speech (TTS) service uses advanced deep learning technologies to synthesize natural sounding human speech¹. Amazon Polly supports dozens of languages and a wide range of natural-sounding voices. You can customize and control the speech output by using lexicons and SSML tags. You can also store and redistribute the speech output in standard audio formats like MP3 and OGG2.

Amazon Transcribe is a service that converts speech to text, enabling you to create text transcripts from audio or video files. It can recognize multiple speakers, different languages, accents, dialects, and background noises. It can also add punctuation and formatting to the transcripts. Amazon Transcribe is useful for applications such as subtitling, captioning, transcription, and voice search.

Amazon Rekognition is a service that provides image and video analysis using computer vision and deep learning. It can detect objects, faces, text, scenes, activities, and emotions in images and videos. It can also perform face recognition, face comparison, face search, celebrity recognition, and facial analysis.

Amazon Rekognition is useful for applications such as security, social media, e-commerce, and media and entertainment.

Amazon Textract is a service that extracts text and data from scanned documents using optical character recognition (OCR) and machine learning. It can identify the contents of fields in forms and tables, as well as the relationships between them. It can also preserve the layout and structure of the original document.

Amazon Textract is useful for applications such as data entry, document management, compliance, and analytics. References:

? Text to Speech Software – Amazon Polly – Amazon Web Services

? What is Text to Speech – Amazon Web Services (AWS)

? AWS Amazon Polly - Text to Speech Converter - CodeCanyon

? Amazon's Text-To-Speech AI Service Sounds More Natural And ... - Forbes

? Working with AWS Amazon Polly Text-to-Speech (TTS) Service

? [Automatic Speech Recognition - Amazon Transcribe - AWS]

? [Amazon Rekognition – Video and Image - AWS]

? [Extract Text & Data - OCR - Amazon Textract - AWS]

NEW QUESTION 162

- (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 167

- (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 169

- (Topic 3)

A company needs to set a maximum spending limit on AWS services each month. The company also needs to set up alerts for when the company reaches its spending limit.

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

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

Answer: D

Explanation:

AWS Budgets is a service that helps you plan your service usage, service costs, and instance reservations, and track how close your plan is to your budgeted amount. You can set custom budgets that alert you when you exceed (or are forecasted to exceed) your budgeted thresholds. You can also use AWS Budgets to set a maximum spending limit on AWS services each month and set up alerts for when you reach your spending limit. Cost Explorer is a service that enables you to visualize, understand, and manage your AWS costs and usage over time. You can use Cost Explorer to view charts and graphs that show how your costs are trending, identify areas that need further inquiry, and see the impact of your cost management actions. However, Cost Explorer does not allow you to set a maximum spending limit or alerts for your AWS services. AWS Trusted Advisor is a service 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 cost optimization opportunities, such as unused or underutilized resources, but it does not allow you to set a maximum spending limit or alerts for your AWS services. Service Quotas is a service that enables you to view and manage your quotas, also referred to as limits, from a central location. Quotas, also referred to as limits, are the maximum number of resources that you can create in your AWS account. However, Service Quotas does not allow you to set a maximum spending limit or alerts for your AWS services.

NEW QUESTION 172

- (Topic 3)

Which AWS service provides protection against DDoS attacks for applications that run in the AWS Cloud?

- A. Amazon VPC
- B. AWS Shield
- C. AWS Audit Manager
- D. AWS Config

Answer: B

Explanation:

AWS Shield is an AWS service that provides protection against distributed denial of service (DDoS) attacks for applications that run in the AWS Cloud. DDoS attacks are attempts to make an online service unavailable by overwhelming it with traffic from multiple sources. AWS Shield provides two tiers of protection: AWS Shield Standard and AWS Shield Advanced. AWS Shield Standard is automatically enabled for all AWS customers at no additional charge. It provides protection against common and frequently occurring network and transport layer DDoS attacks. AWS Shield Advanced is an optional paid service that provides additional protection against larger and more sophisticated DDoS attacks. AWS Shield Advanced also provides access to 24/7 DDoS response team, cost protection, and enhanced detection and mitigation capabilities

NEW QUESTION 174

- (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 178

- (Topic 3)

A company uses AWS for its web application. The company wants to minimize latency and perform compute operations for the application as close to end users as possible.

Which AWS service or infrastructure component will provide this functionality?

- A. AWS Regions
- B. Availability Zones
- C. Edge locations
- D. AWS Direct Connect

Answer: C

Explanation:

Edge locations are sites that Amazon CloudFront uses to cache copies of your content for faster delivery to users at any location. You can use Amazon CloudFront to deliver your entire website, including dynamic, static, streaming, and interactive content using a global network of edge locations. Requests for your content are automatically routed to the nearest edge location, so content is delivered with the best possible performance³. Edge locations can also host AWS Lambda functions to perform compute operations for your web application as close to end users as possible⁴.

NEW QUESTION 180

- (Topic 3)

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

Which AWS service meets these requirements?

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

Answer: D

Explanation:

The AWS service that meets the requirements of providing a graph database service that is scalable and highly available is Amazon Neptune. Amazon Neptune is a fast, reliable, and fully managed graph database service that supports property graph and RDF graph models. Amazon Neptune is designed to store billions of relationships and query the graph with milliseconds latency. Amazon Neptune also offers high availability and durability by replicating six copies of the data across three Availability Zones and continuously backing up the data to Amazon S3⁵. Amazon Aurora, Amazon Redshift, and Amazon DynamoDB are other AWS services that provide relational or non- relational database solutions, but they do not support graph database models.

NEW QUESTION 185

- (Topic 3)

A company wants to set up a high-speed connection between its data center and its applications that run on AWS. The company must not transfer data over the internet.

Which action should the company take to meet these requirements?

- A. Transfer data to AWS by using AWS Snowball.
- B. Transfer data to AWS by using AWS Storage Gateway.
- C. Set up a VPN connection between the data center and an AWS Region.
- D. Set up an AWS Direct Connect connection between the company network and AWS.

Answer: D

Explanation:

AWS Direct Connect is a cloud service solution that makes it easy to establish a dedicated network connection from a customer's premises to AWS. AWS Direct Connect does not involve the public internet, and therefore can reduce network costs, increase bandwidth throughput, and provide a more consistent network experience than internet-based connections. AWS Snowball is a petabyte-scale data transport service that uses secure devices to transfer large amounts of data into and out of the AWS Cloud. AWS Storage Gateway is a hybrid cloud storage service that gives customers on-premises access to virtually unlimited cloud storage. A VPN connection enables customers to establish a secure and private connection between their network and AWS.

NEW QUESTION 186

- (Topic 3)

Which database engines does Amazon Aurora support? (Select TWO.)

- A. Oracle
- B. Microsoft SQL Server
- C. MySQL
- D. PostgreSQL
- E. MongoDB

Answer: CD

Explanation:

Amazon Aurora is a relational database service that is compatible with MySQL and PostgreSQL engines. It delivers up to five times the performance of MySQL and up to three times the performance of PostgreSQL. It also provides high availability, scalability, security, and durability¹

NEW QUESTION 187

- (Topic 3)

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

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

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

Answer: B

Explanation:

Amazon RDS is a managed database service that handles most of the common database administration tasks, such as hardware provisioning, server maintenance, backup and recovery, patching, scaling, and replication. However, Amazon RDS does not optimize the application that interacts with the database. The company is still responsible for tuning the performance, security, and availability of the application according to its business requirements and best practices¹². References:

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

? Perform common DBA tasks for Amazon RDS DB instances

NEW QUESTION 190

- (Topic 3)

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

For how much time will the customer be billed?

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

Answer: C

Explanation:

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

NEW QUESTION 192

- (Topic 2)

A company wants to migrate to the AWS Cloud. The company needs the ability to acquire resources when the resources are necessary.

The company also needs the ability to release those resources when the resources are no longer necessary.

Which architecture concept of the AWS Cloud meets these requirements?

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

Answer: A

Explanation:

The architecture concept of the AWS Cloud that meets the requirements of the company that wants to migrate to the AWS Cloud and needs the ability to acquire and release resources as needed is elasticity. Elasticity means that AWS customers can quickly and easily provision and scale up or down AWS resources as their demand changes, without any upfront costs or long-term commitments. AWS provides various tools and services that enable customers to achieve elasticity, such as Amazon EC2 Auto Scaling, Amazon CloudWatch, and AWS CloudFormation. Elasticity helps customers optimize their performance, availability, and cost efficiency. Availability, reliability, and durability are other architecture concepts of the AWS Cloud, but they are not directly related to the ability to acquire and release resources as needed. Availability means that AWS customers can access their AWS resources and applications whenever and wherever they need them. Reliability means that AWS customers can depend on their AWS resources and applications to function correctly and consistently. Durability means that AWS customers can preserve their data and objects for long periods of time without loss or corruption¹²

NEW QUESTION 196

- (Topic 2)

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

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

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

Answer: B

Explanation:

Amazon EC2 On-Demand Instances are instances that you pay for by the second, with no long-term commitments or upfront payments⁴. This option is suitable for applications that have unpredictable or intermittent workloads, such as the one described in the question. Amazon EC2 Standard Reserved Instances are instances that you purchase for a one-year or three-year term, and pay a lower hourly rate compared to On-Demand Instances. This option is suitable for applications that have steady state or predictable usage. AWS Wavelength is a service that enables developers to build applications that deliver ultra-low latency

to mobile devices and users by deploying AWS compute and storage at the edge of the 5G network. This option is not relevant for the application described in the question. Application Load Balancer is a type of load balancer that operates at the application layer and distributes traffic based on the content of the request. This option is not a service or feature to host the application, but rather to balance the traffic among multiple instances.

NEW QUESTION 197

- (Topic 2)

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 199

- (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 204

- (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 207

- (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 209

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 211

- (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 214

- (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 218

- (Topic 2)

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

Which type of migration is this?

- A. On-premises to cloud native

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

Answer: C

Explanation:

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

NEW QUESTION 222

- (Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 223

- (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 225

- (Topic 2)

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

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

Answer: D

Explanation:

Amazon ElastiCache is a service that offers fully managed in-memory data store and cache services that deliver sub-millisecond response times to applications. You can use Amazon ElastiCache to improve the performance of your applications by retrieving data from fast, managed, in-memory data stores, instead of relying entirely on slower disk-based databases. Amazon Aurora is a relational database service that combines the performance and availability of high-end commercial

databases with the simplicity and cost-effectiveness of open source databases. Amazon RDS is a service that makes it easy to set up, operate, and scale a relational database in the cloud. Amazon DynamoDB is a key-value and document database that delivers single-digit millisecond performance at any scale. None of these services are in-memory data store services.

NEW QUESTION 228

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 230

- (Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 234

- (Topic 2)

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

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

Answer: BE

Explanation:

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

NEW QUESTION 236

- (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 241

.....

Relate Links

100% Pass Your CLF-C02 Exam with ExamBible Prep Materials

<https://www.exambible.com/CLF-C02-exam/>

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>