

Amazon-Web-Services

Exam Questions CLF-C02

AWS Certified Cloud Practitioner



NEW QUESTION 1

- (Topic 3)

A company wants to make an upfront commitment for continued use of its production Amazon EC2 instances in exchange for a reduced overall cost. Which pricing options meet these requirements with the LOWEST cost? (Select TWO.)

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

Answer: CD

Explanation:

Reserved Instances (RIs) are a pricing model that allows you to reserve EC2 instances for a specified period of time (one or three years) and receive a significant discount compared to On-Demand pricing. RIs are suitable for workloads that have predictable usage patterns and require a long-term commitment. You can choose between three payment options: All Upfront, Partial Upfront, or No Upfront. The more you pay upfront, the greater the discount¹.

Savings Plans are a flexible pricing model that can help you reduce your EC2 costs by up to 72% compared to On-Demand pricing, in exchange for a commitment to a consistent amount of usage (measured in \$/hour) for a one or three year term. Savings Plans apply to usage across EC2, AWS Lambda, and AWS Fargate. You can choose between two types of Savings Plans: Compute Savings Plans and EC2 Instance Savings Plans. Compute Savings Plans offer the most flexibility and apply to any instance family, size, OS, tenancy, or region. EC2 Instance Savings Plans offer the highest discount and apply to a specific instance family within a region².

Spot Instances are a pricing model that allows you to bid for unused EC2 capacity in the AWS cloud and are available at a discount of up to 90% compared to On-Demand pricing. Spot Instances are suitable for fault-tolerant or stateless workloads that can run on heterogeneous hardware and have flexible start and end times. However, Spot Instances are not guaranteed and can be interrupted by AWS at any time if the demand for capacity increases or your bid price is lower than the current Spot price³.

On-Demand Instances are a pricing model that allows you to pay for compute capacity by the hour or second with no long-term commitments. On-Demand Instances are suitable for short-term, spiky, or unpredictable workloads that cannot be interrupted, or for applications that are being developed or tested on EC2 for the first time. However, On-Demand Instances are the most expensive option among the four pricing models⁴.

Dedicated Hosts are physical EC2 servers fully dedicated for your use. Dedicated Hosts can help you reduce costs by allowing you to use your existing server-bound software licenses, such as Windows Server, SQL Server, and SUSE Linux Enterprise Server. Dedicated Hosts can be purchased On-Demand or as part of Savings Plans. Dedicated Hosts are suitable for workloads that need to run on dedicated physical servers or have strict licensing requirements. However, Dedicated Hosts are not the lowest cost option among the four pricing models.

NEW QUESTION 2

- (Topic 3)

What is the LEAST expensive AWS Support plan that provides the full set of AWS Trusted Advisor best practice checks for cost optimization?

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

Answer: B

Explanation:

AWS Business Support is the least expensive AWS Support plan that provides the full set of AWS Trusted Advisor best practice checks for cost optimization. AWS Trusted Advisor is a service that provides best practices and recommendations for cost optimization, performance, security, and fault tolerance. AWS Business Support also provides other benefits, such as 24/7 technical support, unlimited cases, and faster response times. AWS Enterprise Support is the most expensive AWS Support plan that provides the same benefits as AWS Business Support, plus additional benefits, such as a technical account manager and enterprise concierge support. AWS Developer Support and AWS Basic Support are cheaper AWS Support plans that provide only a limited set of AWS Trusted Advisor best practice checks for cost optimization .

NEW QUESTION 3

- (Topic 3)

A company encourages its teams to test failure scenarios regularly and to validate their understanding of the impact of potential failures. Which pillar of the AWS Well-Architected Framework does this philosophy represent?

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

Answer: A

Explanation:

This is the pillar of the AWS Well-Architected Framework that represents the philosophy of testing failure scenarios regularly and validating the understanding of the impact of potential failures. The operational excellence pillar covers the best practices for designing, running, monitoring, and improving systems in the AWS Cloud. Testing failure scenarios is one of the ways to improve the system's resilience, reliability, and recovery. You can learn more about the operational excellence pillar from this whitepaper or this digital course.

NEW QUESTION 4

- (Topic 3)

A company wants to ensure that all of its Amazon EC2 instances have compliant operating system patches. Which AWS service will meet these requirements?

- A. AWS Compute Optimizer
- B. AWS Elastic Beanstalk

- C. AWS AppSync
- D. AWS Systems Manager

Answer: D

Explanation:

AWS Systems Manager gives you visibility and control of your infrastructure on AWS. Systems Manager provides a unified user interface so you can view operational data from multiple AWS services and allows you to automate operational tasks across your AWS resources. You can use Systems Manager to apply OS patches, create system images, configure Windows and Linux operating systems, and execute PowerShell commands⁵. Systems Manager can help you ensure that all of your Amazon EC2 instances have compliant operating system patches by using the Patch Manager feature.

NEW QUESTION 5

- (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 6

- (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 7

- (Topic 3)

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

- A. Management of the ongoing security of user data
- B. Advisory solutions for AWS adoption
- C. Technical support 24 hours a day, 7 days a week
- D. Monitoring of monthly billing costs in AWS accounts

Answer: B

Explanation:

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

NEW QUESTION 8

- (Topic 3)

A company wants to minimize network latency between its Amazon EC2 instances. The EC2 instances do not need to be highly available. Which solution meets these requirements?

- A. Use EC2 instances in a single Availability Zone.
- B. Use Amazon CloudFront as the database for the EC2 instances.
- C. Use EC2 instances in the same edge location and the same Availability Zone.
- D. Use EC2 instances in the same edge location and the same AWS Region.

Answer: A

Explanation:

Using EC2 instances in a single Availability Zone is a solution that meets the requirements of minimizing network latency between the EC2 instances and not needing high availability. An Availability Zone is a physically isolated location within an AWS Region that has its own power, cooling, and network connectivity. EC2 instances within the same Availability Zone can communicate with each other using low-latency private IP addresses. However, EC2 instances in a single Availability Zone are not highly available, because they are vulnerable to failures or disruptions that affect the Availability Zone

NEW QUESTION 9

- (Topic 3)

Which AWS service or feature allows a user to establish a dedicated network connection between a company's on-premises data center and the AWS Cloud?

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

Answer: A

Explanation:

AWS Direct Connect is an AWS service that allows users to establish a dedicated network connection between their on-premises data center and the AWS Cloud. This connection bypasses the public internet and provides more predictable network performance, reduced bandwidth costs, and increased security. Users can choose from different port speeds and connection types, and use AWS Direct Connect to access AWS services in any AWS Region globally. Users can also use AWS Direct Connect in conjunction with AWS VPN to create a hybrid network architecture that combines the benefits of both private and public connectivity. References: AWS Direct Connect, [AWS Cloud Practitioner Essentials: Module 3 - Compute in the Cloud]

NEW QUESTION 10

- (Topic 3)

A company processes personally identifiable information (PII) and must keep data in the country where it was generated. The company wants to use Amazon EC2 instances for these workloads.

Which AWS service will meet these requirements?

- A. AWS Outposts
- B. AWS Storage Gateway
- C. AWS DataSync
- D. AWS OpsWorks

Answer: A

Explanation:

AWS Outposts is an AWS service that extends AWS infrastructure, services, APIs, and tools to virtually any datacenter, co-location space, or on-premises facility. AWS Outposts enables you to run Amazon EC2 instances and other AWS services locally, while maintaining a consistent and seamless connection to the AWS Cloud. AWS Outposts is ideal for workloads that require low latency, local data processing, or data residency. By using AWS Outposts, the company can process personally identifiable information (PII) and keep data in the country where it was generated, while leveraging the benefits of AWS

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 14

- (Topic 3)

A company has a MySQL database running on a single Amazon EC2 instance. The company now requires higher availability in the event of an outage.

Which set of tasks would meet this requirement?

- A. Add an Application Load Balancer in front of the EC2 instance.
- B. Configure EC2 Auto Recovery to move the instance to another Availability Zone.
- C. Migrate to Amazon RDS and enable Multi-AZ.
- D. Enable termination protection for the EC2 instance to avoid outages.

Answer: C

Explanation:

The set of tasks that would meet the requirement of having higher availability for a MySQL database running on a single Amazon EC2 instance is to migrate to Amazon RDS and enable Multi-AZ. Amazon RDS is a fully managed relational database service that supports MySQL and other popular database engines. By enabling Multi-AZ, users can have a primary database in one Availability Zone and a synchronous standby replica in another Availability Zone. In case of a planned or unplanned outage of the primary database, Amazon RDS automatically fails over to the standby replica with minimal disruption³. Adding an Application Load Balancer in front of the EC2 instance, configuring EC2 Auto Recovery to move the instance to another Availability Zone, or enabling termination protection for the EC2 instance would not provide higher availability for the database, as they do not address the single point of failure or data replication issues.

NEW QUESTION 17

- (Topic 3)

Which cloud computing advantage is a company applying when it uses AWS Regions to increase application availability to users in different countries?

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

Answer: D

Explanation:

Global reach is a cloud computing advantage that a company can apply when it uses AWS Regions to increase application availability to users in different countries. Global reach refers to the ability to deploy applications and services in multiple geographic locations around the world, and to serve customers with low latency and high performance. AWS has the largest and most reliable global infrastructure of any cloud provider, with 25 Regions and 81 Availability Zones across the Americas, Europe, Asia Pacific, Africa, and the Middle East¹²³. By using AWS Regions, a company can choose the best location for its application based on customer proximity, compliance requirements, and disaster recovery strategies²³. References: 1: AWS Global Infrastructure - Amazon Web Services (AWS), 2: Regions and Availability Zones - Amazon Elastic Compute Cloud, 3: AWS Infrastructure: Regions and Availability Zones Explained

NEW QUESTION 20

- (Topic 3)

A company wants to integrate its online shopping website with social media login credentials.

Which AWS service can the company use to make this integration?

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

Answer: C

Explanation:

Amazon Cognito is a service that enables you to add user sign-up and sign-in features to your web and mobile applications. Amazon Cognito also supports social and enterprise identity federation, which means you can allow your users to sign in with their existing credentials from identity providers such as Google, Facebook, Apple, and Amazon. Amazon Cognito integrates with OpenID Connect (OIDC) and Security Assertion Markup Language (SAML) 2.0 protocols to facilitate the authentication and authorization process. Amazon Cognito also provides advanced security features, such as adaptive authentication, user verification, and multi-factor authentication (MFA). References: Amazon Cognito, What is Amazon Cognito?

NEW QUESTION 24

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 28

- (Topic 3)

Which company needs to apply security rules to a subnet for Amazon EC2 instances. Which AWS service or feature provides this functionality?

- A. Network ACLs
- B. Security groups
- C. AWS Certificate Manager (ACM)

D. AWS Config

Answer: A

Explanation:

Network ACLs (network access control lists) are an AWS service or feature that provides the functionality of applying security rules to a subnet for EC2 instances. A subnet is a logical partition of an IP network within a VPC (virtual private cloud). A VPC is a logically isolated section of the AWS Cloud where the company can launch AWS resources in a virtual network that they define. A network ACL is a virtual firewall that controls the inbound and outbound traffic for one or more subnets. The company 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, the company must create rules for both inbound and outbound traffic⁴

NEW QUESTION 33

- (Topic 3)

A company needs a fully managed file server that natively supports Microsoft workloads and file systems. The file server must also support the SMB protocol. Which AWS service should the company use to meet these requirements?

- A. Amazon Elastic File System (Amazon EFS)
- B. Amazon FSx for Lustre
- C. Amazon FSx for Windows File Server
- D. Amazon Elastic Block Store (Amazon EBS)

Answer: C

Explanation:

Amazon FSx for Windows File Server is a fully managed file server that supports Microsoft workloads and file systems, including the SMB protocol. It provides features such as user quotas, end-user file restore, and Microsoft Active Directory integration. Amazon EFS is a fully managed file system that supports the NFS protocol, not SMB. Amazon FSx for Lustre is a fully managed file system that supports high-performance computing workloads, not Microsoft workloads. Amazon EBS is a block storage service that does not provide a file system or SMB support. References: Amazon FSx for Windows File Server, Amazon FSx for Lustre, Amazon EFS, Amazon EBS

NEW QUESTION 36

- (Topic 3)

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

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

Answer: B

Explanation:

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

NEW QUESTION 39

- (Topic 3)

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

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

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

Answer: D

Explanation:

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

NEW QUESTION 40

- (Topic 3)

A company wants to migrate its high-performance computing (HPC) application to Amazon EC2 instances. The application has multiple components. The application must have fault tolerance and must have the ability to fail over automatically.

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

- A. Multiple AWS Regions
- B. Multiple edge locations
- C. Multiple Availability Zones
- D. Regional edge caches

Answer: C

Explanation:

Using EC2 instances in multiple Availability Zones is an AWS infrastructure solution that meets the requirements of migrating a high performance computing (HPC) application to AWS with fault tolerance and failover capabilities, and with the least latency between components. An Availability Zone is a physically isolated location within an AWS Region that has its own power, cooling, and network connectivity. EC2 instances within the same Region can communicate with each other using low-latency private IP addresses. By using EC2 instances in multiple Availability Zones, the company can achieve fault tolerance and failover for their HPC application, because they can distribute the workload and data across different locations that are independent of each other. If one Availability Zone becomes unavailable or impaired, the company can redirect the traffic and data to another Availability Zone without affecting the performance and availability of the application⁵

NEW QUESTION 43

- (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 44

- (Topic 3)

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

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

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

Answer: B

Explanation:

Amazon Athena is a serverless, interactive analytics service that allows users to run SQL queries on data stored in Amazon S3. It is ideal for occasional queries on large datasets, as it does not require any server provisioning, configuration, or management. Users only pay for the queries they run, based on the amount of data scanned. Amazon Athena supports various data formats, such as CSV, JSON, Parquet, ORC, and Avro, and integrates with AWS Glue Data Catalog to create and manage schemas. Amazon Athena also supports querying data from other sources, such as on- premises or other cloud systems, using data connectors¹.

Amazon Redshift is a fully managed data warehouse service that allows users to run complex analytical queries on petabyte-scale data. However, it requires users to provision and maintain clusters of nodes, and pay for the storage and compute capacity they use. Amazon Redshift is more suitable for frequent and consistent queries on structured or semi-structured data².

Amazon Kinesis is a platform for streaming data on AWS, enabling users to collect, process, and analyze real-time data. It is not designed for querying data stored in Amazon S3. Amazon Kinesis consists of four services: Kinesis Data Streams, Kinesis Data Firehose, Kinesis Data Analytics, and Kinesis Video Streams³.

Amazon RDS is a relational database service that provides six database engines: Amazon Aurora, PostgreSQL, MySQL, MariaDB, Oracle Database, and SQL Server. It simplifies database administration tasks such as backup, patching, scaling, and replication. However, it is not optimized for querying data stored in Amazon S3. Amazon RDS is more suitable for transactional workloads that require high performance and availability⁴.

References:

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

? [Amazon Redshift – Data Warehouse Solution - AWS]

? [Amazon Kinesis - Streaming Data Platform - AWS]

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

NEW QUESTION 45

- (Topic 3)

A developer wants to deploy an application quickly on AWS without manually creating the required resources. Which AWS service will meet these requirements?

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

Answer: B

Explanation:

AWS Elastic Beanstalk is a service that allows you to deploy and manage applications on AWS without manually creating and configuring the required resources, such as EC2 instances, load balancers, security groups, databases, and more. AWS Elastic Beanstalk automatically handles the provisioning, scaling, load balancing, health monitoring, and updating of your application, while giving you full control over the underlying AWS resources if needed. AWS Elastic Beanstalk supports a variety of platforms and languages, such as Java, .NET, PHP, Node.js, Python, Ruby, Go, and Docker. You can use the AWS Management Console, the AWS CLI, the AWS SDKs, or the AWS Elastic Beanstalk API to create and manage your applications. You can also use AWS CodeStar, AWS CodeCommit, AWS CodeBuild, AWS CodeDeploy, and AWS CodePipeline to integrate AWS Elastic Beanstalk with your development and deployment workflows¹²

NEW QUESTION 49

- (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 51

- (Topic 3)

A company needs to implement identity management for a fleet of mobile apps that are running in the AWS Cloud.

Which AWS service will meet this requirement?

- A. Amazon Cognito
- B. AWS Security Hub
- C. AWS Shield
- D. AWS WAF

Answer: A

Explanation:

Amazon Cognito is a service that provides identity management for mobile and web applications, allowing users to sign up, sign in, and access AWS resources with different identity providers. AWS Security Hub is a service that provides a comprehensive view of the security posture of AWS accounts and resources. AWS Shield is a service that provides protection against distributed denial of service (DDoS) attacks. AWS WAF is a web application firewall that helps protect web applications from common web exploits.

NEW QUESTION 53

- (Topic 3)

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

Which AWS service would support this workload?

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

Answer: A

Explanation:

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

NEW QUESTION 57

- (Topic 3)

Which AWS service provides a single location to track the progress of application migrations?

- A. AWS Application Discovery Service
- B. AWS Application Migration Service
- C. AWS Service Catalog
- D. AWS Migration Hub

Answer: D

Explanation:

AWS Migration Hub is a service that provides a single location to track the progress of application migrations across multiple AWS and partner solutions. It allows you to choose the AWS and partner migration tools that best fit your needs, while providing visibility into the status of migrations across your portfolio of applications¹. AWS Migration Hub supports migration status updates from the following tools: AWS Application Migration Service, AWS Database Migration Service, CloudEndure Migration, Server Migration Service, and Migrate for Compute Engine¹.

The other options are not correct for the following reasons:

? AWS Application Discovery Service is a service that helps you plan your migration projects by automatically identifying servers, applications, and dependencies in your on-premises data centers². It does not track the progress of application migrations, but rather provides information to help you plan and scope your migrations.

? AWS Application Migration Service is a service that helps you migrate and modernize applications from any source infrastructure to AWS with minimal downtime and disruption³. It is one of the migration tools that can send status updates to AWS Migration Hub, but it is not the service that provides a single location to track the progress of application migrations.

? AWS Service Catalog is a service that allows you to create and manage catalogs of IT services that are approved for use on AWS⁴. It does not track the

progress of application migrations, but rather helps you manage the provisioning and governance of your IT services.

References:

- ? 1: What Is AWS Migration Hub? - AWS Migration Hub
- ? 2: What Is AWS Application Discovery Service? - AWS Application Discovery Service
- ? 3: App Migration Tool - AWS Application Migration Service - AWS
- ? 4: What Is AWS Service Catalog? - AWS Service Catalog

NEW QUESTION 61

- (Topic 3)

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

Which AWS service or tool will meet these requirements?

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

Answer: C

Explanation:

Amazon QuickSight Q is a natural language query feature that lets you ask questions about your data using everyday language and get answers in seconds. You can type questions such as “What are the total sales by region?” or “How did marketing campaign A perform?” and get answers in the form of relevant visualizations, such as charts or tables. You can also use Q to drill down into details, filter data, or perform calculations. Q uses machine learning to understand your data and your intent, and provides suggestions and feedback to help you refine your questions.

NEW QUESTION 66

- (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 70

- (Topic 3)

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

Which AWS service will meet the requirement?

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

Answer: B

Explanation:

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

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

- ? 1: What Is Amazon Macie? - Amazon Macie
- ? 2: What Is Amazon Inspector? - Amazon Inspector
- ? 3: What Is IAM? - AWS Identity and Access Management
- ? 4: What Is Amazon CloudWatch? - Amazon CloudWatch

NEW QUESTION 72

- (Topic 3)

A company wants to verify if multi-factor authentication (MFA) is enabled for all users within its AWS accounts.

Which AWS service or resource will meet this requirement?

- A. AWS Cost and Usage Report
- B. IAM credential reports
- C. AWS Artifact
- D. Amazon CloudFront reports

Answer: B

Explanation:

The AWS service or resource that will meet the requirement of verifying if multi-factor authentication (MFA) is enabled for all users within its AWS accounts is IAM credential reports. IAM credential reports are downloadable reports that list all the users in an AWS account and the status of their various credentials, including passwords, access keys, and MFA devices. Users can use IAM credential reports to audit the security status of their AWS accounts and identify any issues or risks⁴. AWS Cost and Usage Report, AWS Artifact, and Amazon CloudFront reports are other AWS services or resources that provide different types of information, such as billing, compliance, and content delivery, but they do not show the MFA status of the users.

NEW QUESTION 77

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 79

- (Topic 3)

A company must be able to develop, test, and launch an application in the AWS Cloud quickly. Which advantage of cloud computing will meet these requirements?

- A. Stop guessing capacity
- B. Trade fixed expense for variable expense
- C. Achieve economies of scale
- D. Increase speed and agility

Answer: D

Explanation:

One of the benefits of cloud computing is that it enables customers to increase speed and agility in developing, testing, and launching applications. Cloud computing provides on-demand access to a variety of IT resources, such as compute, storage, networking, databases, and analytics, without requiring upfront investments or long-term commitments. Customers can provision and release resources in minutes, scale up and down as needed, and experiment with new technologies and features. This allows customers to accelerate their innovation cycles, deliver faster time-to-market, and respond to changing customer needs and demands

NEW QUESTION 80

- (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 84

- (Topic 3)

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

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

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

E. AWS Service Catalog

Answer: AD

Explanation:

AWS Personal Health Dashboard and AWS Service Health Dashboard are two AWS services that can help the company to verify that underlying AWS services and general AWS infrastructure are operating normally. AWS Personal Health Dashboard provides a personalized view into the performance and availability of the AWS services you are using, as well as alerts that are automatically triggered by changes in the health of those services. In addition to event-based alerts, Personal Health Dashboard provides proactive notifications of scheduled activities, such as any changes to the infrastructure powering your resources, enabling you to better plan for events that may affect you. These notifications can be delivered to you via email or mobile for quick visibility, and can always be viewed from within the AWS Management Console. When you get an alert, it includes detailed information and guidance, enabling you to take immediate action to address AWS events impacting your resources³. AWS Service Health Dashboard provides a general status of AWS services, and the Service health view displays the current and historical status of all AWS services. This page shows reported service events for services across AWS Regions. You don't need to sign in or have an AWS account to access the AWS Service Health Dashboard – Service health page. You can also subscribe to RSS feeds for specific services or regions to receive notifications about service events⁴. References: Getting started with your AWS Health Dashboard – Your account health, Introducing AWS Personal Health Dashboard

NEW QUESTION 89

- (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 93

- (Topic 3)

Which of the following are general AWS Cloud design principles described in the AWS Well-Architected Framework?

- A. Consolidate key components into monolithic architectures.
- B. Test systems at production scale.
- C. Provision more capacity than a workload is expected to need.
- D. Drive architecture design based on data collected about the workload behavior and requirements.
- E. Make AWS Cloud architectural decisions static, one-time events.

Answer: BD

Explanation:

These are two of the general AWS Cloud design principles described in the AWS Well-Architected Framework. Testing systems at production scale means using tools such as AWS CloudFormation, AWS CodeDeploy, and AWS X-Ray to simulate real-world scenarios and measure the performance, scalability, and availability of the system. Driving architecture design based on data means using tools such as Amazon CloudWatch, AWS CloudTrail, and AWS Config to collect and analyze metrics, logs, and events about the system and use the insights to optimize the system's design and operation. You can learn more about the AWS Well-Architected Framework from this whitepaper or [this digital course].

NEW QUESTION 95

- (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 100

- (Topic 3)

A company is building an application that needs to deliver images and videos globally with minimal latency. Which approach can the company use to accomplish this in a cost effective manner?

- A. Deliver the content through Amazon CloudFront.
- B. Store the content on Amazon S3 and enable S3 cross-region replication.
- C. Implement a VPN across multiple AWS Regions.

D. Deliver the content through AWS PrivateLink.

Answer: A

Explanation:

Amazon CloudFront is a fast content delivery network (CDN) service that securely delivers data, videos, applications, and APIs to customers globally with low latency, high transfer speeds, all within a developer-friendly environment. It works seamlessly with services including AWS Shield for DDoS mitigation, Amazon S3, Elastic Load Balancing or Amazon EC2 as origins for your applications, and Lambda@Edge to run custom code closer to customers' users and to customize the user experience. By using CloudFront, you can cache your content at the edge locations that are closest to your end users, reducing the network latency and improving the performance of your application. CloudFront also offers a pay-as-you-go pricing model, so you only pay for the data transfer and requests that you use.

NEW QUESTION 105

- (Topic 3)

A company wants to use the AWS Cloud to deploy an application globally.

Which architecture deployment model should the company use to meet this requirement?

- A. Multi-Region
- B. Single-Region
- C. Multi-AZ
- D. Single-AZ

Answer: A

Explanation:

The architecture deployment model that the company should use to meet this requirement is A. Multi-Region.

A multi-region deployment model is a cloud computing architecture that distributes an application and its data across multiple geographic regions. A multi-region deployment model enables a company to achieve global reach, high availability, disaster recovery, and performance optimization. By deploying an application in multiple regions, a company can serve customers from the nearest region, reduce latency, increase redundancy, and comply with data sovereignty regulations¹².

A single-region deployment model is a cloud computing architecture that runs an application and its data within a single geographic region. A single-region deployment model is simpler and cheaper than a multi-region deployment model, but it has limited scalability, availability, and performance. A single-region deployment model may not be suitable for a company that wants to deploy an application globally, as it may face challenges such as network latency, regional outages, or regulatory compliance¹².

A multi-AZ (Availability Zone) deployment model is a cloud computing architecture that distributes an application and its data across multiple isolated locations within a single region. An Availability Zone is a physically separate location within an AWS Region that has independent power, cooling, and networking. A multi-AZ deployment model enhances the availability and durability of an application by providing redundancy and fault tolerance within a region³⁴.

A single-AZ deployment model is a cloud computing architecture that runs an application and its data within a single Availability Zone. A single-AZ deployment model is the simplest and most cost-effective option, but it has no redundancy or fault tolerance. A single-AZ deployment model may not be suitable for a company that wants to deploy an application globally, as it may face challenges such as network latency, regional outages, or regulatory compliance³⁴.

References:

1: AWS Cloud Computing - W3Schools 2: Understand the Different Cloud Computing Deployment Models Unit - Trailhead 3: Regions and Availability Zones - Amazon Elastic Compute Cloud 4: AWS Reference Architecture Diagrams

NEW QUESTION 106

- (Topic 3)

A company's IT team is managing MySQL database server clusters. The IT team has to patch the database and take backup snapshots of the data in the clusters. The company wants to move this workload to AWS so that these tasks will be completed automatically.

What should the company do to meet these requirements?

- A. Deploy MySQL database server clusters on Amazon EC2 instances.
- B. Use Amazon RDS with a MySQL database.
- C. Use an AWS CloudFormation template to deploy MySQL database servers on Amazon EC2 instances.
- D. Migrate all the MySQL database data to Amazon S3.

Answer: B

Explanation:

Amazon RDS is a service that makes it easy to set up, operate, and scale a relational database in the cloud. Amazon RDS supports MySQL as one of the database engines. By using Amazon RDS with a MySQL database, the company can offload the tasks of patching the database and taking backup snapshots to AWS. Amazon RDS automatically patches the database software and operating system of the database instances. Amazon RDS also automatically backs up the database and retains the backups for a user-defined retention period. The company can also restore the database to any point in time within the retention period. Deploying MySQL database server clusters on Amazon EC2 instances, using an AWS CloudFormation template to deploy MySQL database servers on Amazon EC2 instances, or migrating all the MySQL database data to Amazon S3 are not the best options to meet the requirements. These options would not automate the tasks of patching the database and taking backup snapshots, and would require more operational overhead from the company³

NEW QUESTION 110

- (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 115

- (Topic 3)

A user needs a relational database but does not have the resources to manage the hardware, resiliency, and replication. Which AWS service option meets the user's requirements'?

- A. Run MySQL on Amazon Elastic Container Service (Amazon ECS)
- B. Run MySQL on Amazon EC2
- C. Choose Amazon RDS for MySQL
- D. Choose Amazon ElastiCache for Redis

Answer: C

Explanation:

Amazon RDS for MySQL is a fully managed, open-source cloud database service that allows you to easily operate and scale your relational database of choice, including MySQL. With Amazon RDS for MySQL, you don't have to worry about the hardware, resiliency, and replication of your database, as Amazon RDS handles these tasks for you. Amazon RDS for MySQL also provides features such as automated backups, multi-AZ deployments, read replicas, encryption, monitoring, and more. Amazon RDS for MySQL is compatible with the MySQL Community Edition versions 5.7 and 8.0, which means that you can use the same code, applications, and tools that you already use with MySQL4567. References: 4: Hosted MySQL - Amazon RDS for MySQL - AWS, 5: Amazon RDS for MySQL - Amazon Relational Database Service, 6: Amazon RDS for MySQL —, 7: Managed SQL Database - Amazon Relational Database Service (RDS) - AWS

NEW QUESTION 118

- (Topic 3)

Which AWS service is always provided at no charge?

- A. Amazon S3
- B. AWS Identity and Access Management (IAM)
- C. Elastic Load Balancers
- D. AWS WAF

Answer: B

Explanation:

AWS Identity and Access Management (IAM) is a web service that helps you securely control access to AWS resources. You can use IAM to create and manage AWS users and groups, and use permissions to allow and deny their access to AWS resources. IAM is always provided at no charge12. References: 1: AWS Identity and Access Management (IAM) - Amazon Web Services (AWS), 2: Which aws service is always provided at no charge? - Brainly.in

NEW QUESTION 119

- (Topic 3)

A company wants a time-series database service that makes it easier to store and analyze trillions of events each day. Which AWS service will meet this requirement?

- A. Amazon Neptune
- B. Amazon Timestream
- C. Amazon Forecast
- D. Amazon DocumentDB (with MongoDB compatibility)

Answer: B

Explanation:

Amazon Timestream is a fast, scalable, and serverless time-series database service for IoT and other operational applications that makes it easy to store and analyze trillions of events per day up to 1,000 times faster and at as little as 1/10th the cost of relational databases1. Amazon Timestream saves you time and cost in managing the lifecycle of time series data, and its purpose-built query engine lets you access and analyze recent and historical data together with a single query1. Amazon Timestream has built-in time series analytics functions, helping you identify trends and patterns in near real time1. The other options are not suitable for storing and analyzing trillions of events per day. Amazon Neptune is a graph database service that supports highly connected data sets. Amazon Forecast is a machine learning service that generates accurate forecasts based on historical data. Amazon DocumentDB (with MongoDB compatibility) is a document database service that supports MongoDB workloads.

References:

? 1: Time Series Database – Amazon Timestream – Amazon Web Services

NEW QUESTION 120

- (Topic 3)

A company wants to migrate its on-premises relational databases to the AWS Cloud. The company wants to use infrastructure as close to its current geographical location as possible.

Which AWS service or resource should the company use to select its Amazon RDS deployment area?

- A. Amazon Connect
- B. AWS Wavelength
- C. AWS Regions
- D. AWS Direct Connect

Answer: C

Explanation:

AWS Regions are the AWS service or resource that the company should use to select its Amazon RDS deployment area. AWS Regions are separate geographic areas where AWS clusters its data centers. Each AWS Region consists of multiple, isolated, and physically separate Availability Zones within a geographic area. Each AWS Region is designed to be isolated from the other AWS Regions to achieve the highest possible fault tolerance and stability. AWS provides a more extensive global footprint than any other cloud provider, and to support its global footprint and ensure customers are served across the world, AWS opens new

Regions rapidly. AWS maintains multiple geographic Regions, including Regions in North America, South America, Europe, China, Asia Pacific, South Africa, and the Middle East. Amazon RDS is available in several AWS Regions worldwide. To create or work with an Amazon RDS DB instance in a specific AWS Region, you must use the corresponding regional service endpoint. You can choose the AWS Region that meets your latency or legal requirements. You can also use multiple AWS Regions to design a disaster recovery solution or to distribute your read workload. References: Global Infrastructure Regions & AZs - [aws.amazon.com](https://aws.amazon.com/global-infrastructure/), Regions, Availability Zones, and Local Zones - Amazon Relational Database Service

NEW QUESTION 122

- (Topic 3)

Which of the following is a software development framework that a company can use to define cloud resources as code and provision the resources through AWS CloudFormation?

- A. AWS CLI
- B. AWS Developer Center
- C. AWS Cloud Development Kit (AWS CDK)
- D. AWS CodeStar

Answer: C

Explanation:

AWS Cloud Development Kit (AWS CDK) is a software development framework that allows you to define cloud resources as code using familiar programming languages, such as TypeScript, Python, Java, .NET, and Go (in Developer Preview). You can use AWS CDK to model your application resources using high-level constructs that provide sensible defaults and best practices, or use low-level constructs that provide full access to the underlying AWS CloudFormation resources. AWS CDK synthesizes your code into AWS CloudFormation templates that you can deploy using the AWS CDK CLI or the AWS Management Console. AWS CDK also integrates with other AWS services, such as AWS CodeCommit, AWS CodeBuild, AWS CodePipeline, AWS Lambda, Amazon EC2, Amazon S3, and more, to help you automate your development and deployment processes. AWS CDK is an open-source framework that you can extend and contribute to. References: Cloud Development Framework - AWS Cloud Development Kit - [AWS](https://aws.amazon.com/cdk/), AWS Cloud Development Kit Documentation, AWS Cloud Development Kit - Wikipedia, AWS CDK Intro Workshop | AWS CDK Workshop

NEW QUESTION 123

- (Topic 3)

Which of the following is a pillar of the AWS Well-Architected Framework?

- A. Redundancy
- B. Operational excellence
- C. Availability
- D. Multi-Region

Answer: B

Explanation:

The AWS Well-Architected Framework helps cloud architects build secure, high-performing, resilient, and efficient infrastructure for their applications and workloads. Based on five pillars — operational excellence, security, reliability, performance efficiency, and cost optimization — the Framework provides a consistent approach for customers and partners to evaluate architectures, and implement designs that can scale over time. Operational excellence is one of the pillars of the Framework, and it focuses on running and monitoring systems to deliver business value, and continually improving processes and procedures.

NEW QUESTION 127

- (Topic 3)

A company wants a customized assessment of its current on-premises environment. The company wants to understand its projected running costs in the AWS Cloud.

Which AWS service or tool will meet these requirements?

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

Answer: D

Explanation:

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

NEW QUESTION 129

- (Topic 3)

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

- A. Observability
- B. Incident and problem management
- C. Incident response
- D. Infrastructure protection
- E. Availability and continuity

Answer: CD

Explanation:

The AWS Cloud Adoption Framework (AWS CAF) security perspective helps users achieve the confidentiality, integrity, and availability of their data and cloud workloads. It comprises nine capabilities that are grouped into three categories: preventive, detective, and responsive. Incident response and infrastructure protection are two of the capabilities in the responsive and preventive categories, respectively. Incident response helps users prepare for and respond to security incidents in a timely and effective manner, using tools and processes that leverage AWS features and services. Infrastructure protection helps users implement security controls and mechanisms to protect their cloud resources, such as network, compute, storage, and database, from unauthorized access or malicious attacks. References: Security perspective: compliance and assurance, AWS Cloud Adoption Framework

NEW QUESTION 133

- (Topic 3)

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

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

Answer: C

Explanation:

A network access control list (network ACL) is a feature that acts as a firewall for controlling traffic in and out of one or more subnets in a virtual private cloud (VPC). Network ACLs can be configured with rules that allow or deny traffic based on the source and destination IP addresses, ports, and protocols¹. AWS Security Hub is a service that provides a comprehensive view of the security posture of AWS accounts and resources². Security groups are features that act as firewalls for controlling traffic at the instance level³. AWS WAF is a web application firewall that helps protect web applications from common web exploits⁴.

NEW QUESTION 138

- (Topic 3)

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

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

Answer: B

Explanation:

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

NEW QUESTION 139

- (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 144

- (Topic 3)

A company has deployed an Amazon EC2 instance.

Which option is an AWS responsibility under the AWS shared responsibility model?

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

Answer: C

Explanation:

According to the AWS shared responsibility model, AWS is responsible for protecting the infrastructure that runs all of the services offered in the AWS Cloud, such as data centers, hardware, software, networking, and facilities¹. This includes the configuration of infrastructure devices, such as routers, switches, firewalls, and load balancers². Customers are responsible for managing their data, applications, operating systems, security groups, and other aspects of their AWS environment¹. Therefore, options A, B, and D are customer responsibilities, not AWS responsibilities. References: 1: AWS Well-Architected Framework - Elasticity;

2: Reactive Systems on AWS - Elastic

NEW QUESTION 146

- (Topic 3)

Which AWS service could an administrator use to provide desktop environments for several employees?

- A. AWS Organizations
- B. AWS Fargate
- C. AWS WAF
- D. AWS Workspaces

Answer: D

Explanation:

AWS Workspaces is a service that provides fully managed, secure, and reliable virtual desktops for your employees. You can access your personal Windows environment on various devices, such as Android, iOS, Fire, Mac, PC, Chromebook, and Linux. You can choose from different bundles of CPU, memory, storage, and software options to suit your needs. You can also integrate AWS Workspaces with your existing Active Directory, VPN, and security policies. AWS Workspaces helps you reduce the cost and complexity of managing your desktop infrastructure, while enhancing the productivity and security of your remote workers⁴⁵⁶. References: 4: Amazon WorkSpaces Client Download, 5: VDI Desktops - Amazon WorkSpaces Family - AWS, 6: Amazon WorkSpaces

NEW QUESTION 148

- (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 150

- (Topic 3)

A company wants to store data with high availability, encrypt the data at rest, and have direct access to the data over the internet.

Which AWS service will meet these requirements MOST cost-effectively?

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

Answer: C

Explanation:

Amazon Elastic File System (Amazon EFS) provides a simple, scalable, fully managed elastic NFS file system for use with AWS Cloud services and on-premises resources. It is built to scale on demand to petabytes without disrupting applications, growing and shrinking automatically as you add and remove files, eliminating the need to provision and manage capacity to accommodate growth. Amazon EFS offers two storage classes: the Standard storage class, and the Infrequent Access storage class (EFS IA).

EFS IA provides price/performance that is cost-optimized for files not accessed every day. Amazon EFS encrypts data at rest and in transit, and supports direct access over the internet⁴.

NEW QUESTION 151

- (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 156

- (Topic 3)

Which service enables customers to audit API calls in their AWS accounts'?

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

Answer: A

Explanation:

AWS CloudTrail is a service that provides a record of actions taken by a user, role, or an AWS service in your AWS account. CloudTrail captures all API calls for AWS services as events, including calls from the AWS Management Console, AWS SDKs, command line tools, and higher-level AWS services. You can use CloudTrail to monitor, audit, and troubleshoot your AWS account activity³⁴. AWS Trusted Advisor is a service that provides best practices recommendations for cost optimization, performance, security, and fault tolerance in your AWS account⁵. 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 practices⁶. AWS X-Ray is a service that helps you analyze and debug your applications by collecting data about the requests that your application serves, and providing tools to view, filter, and gain insights into that data⁷. References: Logging AWS Audit Manager API calls with CloudTrail, Logging AWS Account Management API calls using AWS CloudTrail, Review API calls in your AWS account using CloudTrail, Monitor the usage of AWS API calls using Amazon CloudWatch, Which service enables customers to audit API calls in their AWS ...

NEW QUESTION 159

- (Topic 3)

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

Which AWS pricing model or offering will meet these requirements?

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

Answer: B

Explanation:

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

NEW QUESTION 160

- (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 163

- (Topic 3)

Which AWS service requires the customer to be fully responsible for applying operating system patches?

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

Answer: D

Explanation:

Amazon EC2 is the AWS service that requires the customer to be fully responsible for applying operating system patches. Amazon EC2 is a 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 resources¹. Customers have full control and access to their instances, which means they are also responsible for managing and maintaining them, including applying operating system patches². Customers can use AWS Systems Manager Patch Manager, a feature of AWS Systems Manager, to automate the process of patching their EC2 instances with both security-related updates and other types of updates³.

NEW QUESTION 165

- (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 169

- (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 173

- (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 175

- (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 180

- (Topic 3)

Which Amazon S3 storage class is the MOST cost-effective for long-term storage?

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

Answer: A

Explanation:

Amazon S3 Glacier Deep Archive is the lowest-cost storage class in the cloud. It is designed for long-term data archiving that is rarely accessed. It offers a retrieval time of 12 hours and a durability of 99.999999999% (11 9's). It is ideal for data that must be retained for 7 years or longer to meet regulatory compliance requirements.

NEW QUESTION 184

- (Topic 3)

A company needs a bridge between technology and business to help evolve to a culture of continuous growth and learning.

Which perspective in the AWS Cloud Adoption Framework (AWS CAF) serves as this bridge?

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

Answer: A

Explanation:

The People perspective in the AWS Cloud Adoption Framework (AWS CAF) serves as a bridge between technology and business, accelerating the cloud journey to help organizations more rapidly evolve to a culture of continuous growth, learning, and where change becomes business-as-normal, with focus on culture, organizational structure, leadership, and workforce¹. References: People Perspective - AWS Cloud Adoption Framework

NEW QUESTION 186

- (Topic 3)

Which AWS services can limit manual errors by consistently provisioning AWS resources in multiple environments?

- A. AWS Config
- B. AWS CodeStar
- C. AWS CloudFormation
- D. AWS Cloud Development Kit (AWS CDK)
- E. AWS CodeBuild

Answer: CD

Explanation:

AWS CloudFormation and AWS Cloud Development Kit (AWS CDK) are AWS services that can limit manual errors by consistently provisioning AWS resources in multiple environments. AWS CloudFormation is a service that enables you to model and provision AWS resources using templates. You can use AWS CloudFormation to define the AWS resources and their dependencies that you need for your applications, and to automate the creation and update of those resources across multiple environments, such as development, testing, and production. AWS CloudFormation helps you ensure that your AWS resources are configured consistently and correctly, and that you can easily replicate or modify them as needed. AWS Cloud Development Kit (AWS CDK) is a service that enables you to use familiar programming languages, such as Python, TypeScript, Java, and C#, to define and provision AWS resources. You can use AWS CDK to write code that synthesizes into AWS CloudFormation templates, and to leverage the existing libraries and tools of your preferred language. AWS CDK helps you reduce the complexity and errors of writing and maintaining AWS CloudFormation templates, and to apply the best practices and standards of software development to your AWS infrastructure.

NEW QUESTION 189

- (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 194

- (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 198

- (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 202

- (Topic 3)

Which AWS services or features can a company use to connect the network of its on- premises data center to AWS? (Select TWO.)

- A. AWS VPN
- B. AWS Directory Service
- C. AWS Data Pipeline
- D. AWS Direct Connect
- E. AWS CloudHSM

Answer: AD

Explanation:

AWS VPN and AWS Direct Connect are two services that enable customers to connect their on-premises data center network to the AWS Cloud. AWS VPN establishes a secure and encrypted connection over the public internet, while AWS Direct Connect establishes a dedicated and private connection through a partner network. You can learn more about AWS VPN from [this webpage] or [this digital course]. You can learn more about AWS Direct Connect from [this webpage] or [this digital course].

NEW QUESTION 205

- (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 207

- (Topic 3)

A company wants to build a new web application by using AWS services. The application must meet the on-demand load for periods of heavy activity.

Which AWS services or resources provide the necessary workload adjustments to meet these requirements? (Select TWO.)

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

Answer: BD

Explanation:

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

NEW QUESTION 210

- (Topic 3)

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

Which AWS service will meet this requirement?

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

Answer: B

Explanation:

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

NEW QUESTION 213

- (Topic 3)

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

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

- A. Security
- B. Performance efficiency
- C. Cost optimization

D. Operational excellence

Answer: D

Explanation:

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

NEW QUESTION 215

- (Topic 3)

Which VPC component provides a layer of security at the subnet level?

- A. Security groups
- B. Network ACLs
- C. NAT gateways
- D. Route tables

Answer: B

Explanation:

Network ACLs are a feature that provide a layer of security at the subnet level by acting as a firewall to control traffic in and out of one or more subnets. Network ACLs can be configured with rules that allow or deny traffic based on the source and destination IP addresses, ports, and protocols⁵. Security groups are a feature that provide a layer of security at the instance level by acting as a firewall to control traffic to and from one or more instances. Security groups can be configured with rules that allow or deny traffic based on the source and destination IP addresses, ports, protocols, and security groups. NAT gateways are a feature that enable instances in a private subnet to connect to the internet or other AWS services, but prevent the internet from initiating a connection with those instances. Route tables are a feature that determine where network traffic from a subnet or gateway is directed.

NEW QUESTION 219

- (Topic 3)

Which AWS service or feature gives users the ability to capture information about network traffic in a VPC?

- A. VPC Flow Logs
- B. Amazon Inspector
- C. VPC route tables
- D. AWS CloudTrail

Answer: A

Explanation:

VPC Flow Logs is a feature that enables you to capture information about the IP traffic going to and from network interfaces in your VPC. Flow log data can be published to Amazon CloudWatch Logs, Amazon S3, or Amazon Kinesis Data Firehose. You can use VPC Flow Logs to diagnose network issues, monitor traffic patterns, detect security anomalies, and comply with auditing requirements³⁴. References: Logging IP traffic using VPC Flow Logs - Amazon Virtual Private Cloud, New – VPC Traffic Mirroring – Capture & Inspect Network Traffic | AWS News Blog

NEW QUESTION 223

- (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 227

- (Topic 3)

Which AWS Cloud benefit describes the ability to acquire resources as they are needed and release resources when they are no longer needed?

- A. Economies of scale
- B. Elasticity
- C. Agility
- D. Security

Answer: B

Explanation:

The AWS Cloud benefit that describes the ability to acquire resources as they are needed and release resources when they are no longer needed is elasticity. Elasticity means that users can quickly add and remove resources to match the demand of their applications, and only pay for what they use. Elasticity enables users to handle unpredictable workloads, reduce costs, and improve performance¹. Economies of scale, agility, and security are other benefits of the AWS Cloud, but they do not describe the specific ability of acquiring and releasing resources on demand.

NEW QUESTION 229

- (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 232

- (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 234

- (Topic 2)

Which controls are the responsibility of both AWS and AWS customers, according to the AWS shared responsibility model? (Select TWO.)

- A. Physical and environmental controls
- B. Patch management
- C. Configuration management
- D. Account structures
- E. Choice of the AWS Region where data is stored

Answer: BC

Explanation:

Patch management and configuration management are controls that are the responsibility of both AWS and AWS customers, according to the AWS shared responsibility model. Patch management is the process of applying updates to software and applications to fix vulnerabilities, bugs, or performance issues. Configuration management is the process of defining and maintaining the settings and parameters of systems and applications to ensure their consistency and reliability. AWS is responsible for patching and configuring the software and services that it manages, such as the AWS global infrastructure, the hypervisor, and the AWS managed services. The customer is responsible for patching and configuring the software and services that they manage, such as the guest operating system, the applications, and the AWS customer-managed services. Physical and environmental controls are the responsibility of AWS, according to the AWS shared responsibility model. Physical and environmental controls are the measures that protect the physical security and availability of the AWS global infrastructure, such as power, cooling, fire suppression, and access control. AWS is responsible for maintaining these controls and ensuring the resilience and reliability of the AWS Cloud. Account structures are the responsibility of the customer, according to the AWS shared responsibility model. Account structures are the ways that customers organize and manage their AWS accounts and resources, such as using AWS Organizations, IAM users and roles, resource tagging, and billing preferences. The customer is responsible for creating and configuring these structures and ensuring the security and governance of their AWS environment. Choice of the AWS Region where data is stored is the responsibility of the customer, according to the AWS shared responsibility model. AWS Regions are geographic areas that consist of multiple isolated Availability Zones. Customers can choose which AWS Region to store their data and run their applications, depending on their latency, compliance, and cost requirements. The customer is responsible for selecting the appropriate AWS Region and ensuring the data sovereignty and regulatory compliance of their data.

NEW QUESTION 237

- (Topic 2)

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

Which AWS feature or purchasing option will meet these requirements?

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

Answer: B

Explanation:

Consolidated billing is an AWS feature that allows users to combine the usage and costs of multiple AWS accounts into a single bill. This enables users to obtain billing discounts that are based on the company's use of AWS services, such as volume pricing tiers, Reserved Instance discounts, and Savings Plans discounts⁵. Resource tagging is an AWS feature that allows users to assign metadata to AWS resources, such as EC2 instances, S3 buckets, and Lambda

functions. This enables users to organize, track, and manage their AWS resources, such as filtering, grouping, and reporting. Pay-as-you-go pricing is an AWS pricing model that allows users to pay only for the resources and services they use, without any upfront or long-term commitments. This enables users to lower their costs by scaling up or down as needed, and avoiding over-provisioning or under-utilization. Spot Instances are spare EC2 instances that are available at up to 90% discount compared to On-Demand prices. They are suitable for workloads that can tolerate interruptions, such as batch processing, data analysis, and testing. Spot Instances are allocated based on the current supply and demand, and can be reclaimed by AWS with a two-minute notice when the demand exceeds the supply.

NEW QUESTION 240

- (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 242

- (Topic 2)

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

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

Answer: AB

Explanation:

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

NEW QUESTION 245

- (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 249

- (Topic 2)

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

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

Answer: C

Explanation:

C is correct because turning on multi-factor authentication (MFA) for added security during the login process is one of the IAM security best practices

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

NEW QUESTION 251

- (Topic 2)

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 254

- (Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 255

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 257

- (Topic 2)

Which AWS service can defend against DDoS attacks?

- A. AWS Firewall Manager
- B. AWS Shield Standard
- C. AWS WAF
- D. Amazon Inspector

Answer: B

Explanation:

AWS Shield Standard is a service that provides protection against Distributed Denial of Service (DDoS) attacks for all AWS customers at no additional charge. It

automatically detects and mitigates the most common and frequently occurring network and transport layer DDoS attacks that target AWS resources, such as Amazon EC2 instances, Elastic Load Balancers, Amazon CloudFront distributions, and Amazon Route 53 hosted zones. AWS Firewall Manager is a service that allows users to centrally configure and manage firewall rules across their AWS accounts and resources, such as AWS WAF web ACLs, AWS Shield Advanced protections, and Amazon VPC security groups. AWS WAF is a web application firewall that helps protect web applications from common web exploits, such as SQL injection, cross-site scripting, and bot attacks. Amazon Inspector is an automated security assessment service that helps improve the security and compliance of applications deployed on AWS. It analyzes the behavior of the applications and checks for vulnerabilities, exposures, and deviations from best practices.

NEW QUESTION 259

- (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 260

- (Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 262

- (Topic 2)

A company must store call recordings for 6 years. The storage system should be highly durable and cost-effective.

Which AWS service meets these requirements?

- A. AWS Snowball
- B. Amazon S3
- C. AWS Storage Gateway
- D. Amazon Kinesis

Answer: B

Explanation:

Amazon S3 is a service that provides highly durable and cost-effective object storage for a variety of use cases, including backup and archive, big data analytics, disaster recovery, and cloud applications. Amazon S3 offers 99.999999999% (11 9's) of durability, meaning that data is designed to withstand the loss of two facilities concurrently. Amazon S3 also offers several storage classes with different price and performance characteristics, such as S3 Glacier and S3 Glacier Deep Archive, which are ideal for long-term archival of data that is rarely accessed. AWS Snowball, AWS Storage Gateway, and Amazon Kinesis are not designed to provide the same level of durability and cost-effectiveness as Amazon S3 for storing call recordings for 6 years. Source: Amazon S3

NEW QUESTION 265

- (Topic 2)

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

Which AWS service will provide this information?

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

Answer: A

Explanation:

AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. AWS Config continuously monitors and records your AWS resource configurations and allows you to automate the evaluation of recorded configurations against desired configurations. AWS Config can help you determine when an EBS volume was removed from an EC2 instance by providing a timeline of configuration changes and compliance status. AWS Trusted Advisor, Amazon Timestream, and Amazon QuickSight do not provide the same level of configuration tracking and auditing as AWS Config. Source: AWS Config

NEW QUESTION 266

- (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 271

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

CLF-C02 Practice Exam Features:

- * CLF-C02 Questions and Answers Updated Frequently
- * CLF-C02 Practice Questions Verified by Expert Senior Certified Staff
- * CLF-C02 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * CLF-C02 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The CLF-C02 Practice Test Here](#)