PassQueen
a.k.a dumpsqueen.com

AWS Certified Cloud Practitioner

Amazon AWS CLF-C02

Version Demo

Total Demo Questions: 122

Total Premium Questions: 1220

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Cloud Concepts	192
Topic 2, Security and Compliance	356
Topic 3, Cloud Technology and Services	479
Topic 4, Billing, Pricing, and Support	192
Topic 5, Mix Questions	1
Total	1220

QUESTION NO: 1

Which combination of AWS services can be used to move a commercial relational database to an Amazon-managed open-source database? (Select TWO.)

A. AWS Database Migration Service (AWS DMS)

B. AWS software development kits (SDKs)

C. AWS Schema Conversion Tool

D. AWS Systems Manager

E. Amazon EMR

ANSWER: A C

Explanation:

AWS Database Migration Service (AWS DMS) and AWS Schema Conversion Tool are used together for heterogeneous database migrations, such as moving from a commercial database engine to an Amazon-managed open-source-compatible database. AWS Schema Conversion Tool assesses the source and target databases and converts database schema objects and application SQL code where conversion is possible. This preparation is important because different database engines can use different SQL dialects, data types, and object definitions. After the target schema has been prepared, AWS DMS migrates the source data to the target database. AWS DMS can also support ongoing replication by capturing and applying changes after the initial load, helping reduce downtime during a migration. For example, this combination can support a migration from a commercial relational database to Amazon Aurora PostgreSQL-Compatible Edition or Amazon RDS for PostgreSQL. AWS documents AWS DMS as a service for migrating databases to AWS, including heterogeneous migrations, and identifies AWS Schema Conversion Tool as the service that converts schemas and code between database engines. See the AWS DMS User Guide and the AWS Schema Conversion Tool User Guide.

QUESTION NO: 2

Which AWS service can run a managed PostgreSQL database that provides online transaction processing (OLTP)?

A. Amazon DynamoDB

B. Amazon Athena

C. Amazon RDS

D. Amazon EMR

ANSWER: C

Explanation:

Amazon RDS is correct because it is AWS's managed relational database service and supports PostgreSQL as a database engine. Amazon RDS for PostgreSQL is designed for relational workloads, including online transaction processing (OLTP), where applications perform frequent, concurrent, low-latency transactions such as inserting orders, updating account records, and querying application data. AWS manages much of the operational work involved in running the database, including hardware provisioning, database software installation, backups, patching, failure detection, and recovery. This lets customers use PostgreSQL features and standard PostgreSQL-compatible tools while reducing database administration overhead.

For production OLTP applications, Amazon RDS also offers capabilities such as Multi-AZ deployments for improved availability and durability, automated backups and point-in-time restore, read replicas for scaling read traffic, and monitoring through Amazon CloudWatch. PostgreSQL is a relational database engine, so it supports SQL queries, structured schemas, ACID transactions, and relationships between tables—core characteristics commonly needed by transactional application workloads. AWS identifies Amazon RDS for PostgreSQL as a fully managed PostgreSQL database service, and its RDS documentation describes the service's supported engines and operational management features. See Amazon RDS for PostgreSQL and Working with PostgreSQL on Amazon RDS.

QUESTION NO: 3

A company wants to set up cloud-based customer service centers.

Which AWS service or resource will meet this requirement?

A. Amazon Athena

B. Amazon Connect

C. AWS Direct Connect

D. AWS Enterprise Support

ANSWER: B

Explanation:

Amazon Connect is the appropriate service because it is AWS's fully managed cloud contact center solution. It enables an organization to create and operate customer service centers without provisioning traditional on-premises telephony or contact-center infrastructure. A company can create an Amazon Connect instance, claim or port phone numbers, configure interactive voice response and contact flows, establish queues and routing profiles, and provide agents with a browser-based workspace for handling customer interactions.

The service supports common customer-service channels and capabilities, including voice calls, chat, task management, agent routing, real-time and historical metrics, call recording, and integrations with AWS services and customer relationship management systems. These managed features allow the company to implement a scalable customer support operation while AWS manages the underlying contact-center platform. This directly meets the requirement for cloud-based customer service centers. AWS describes Amazon Connect as an easy-to-use cloud contact center that helps businesses deliver customer service at scale. See the Amazon Connect product page and the Amazon Connect Administrator Guide for service capabilities and setup concepts.

QUESTION NO: 4

A developer who has no AWS Cloud experience wants to use AWS technology to build a web application.

Which AWS service should the developer use to start building the application?

A. Amazon SageMaker

B. AWS Lambda

C. Amazon Lightsail

D. Amazon Elastic Container Service (Amazon ECS)

ANSWER: C

Explanation:

Amazon Lightsail is the best starting point because it is designed to make launching and managing common web applications straightforward for users with little or no prior cloud experience. Lightsail provides an easy-to-use interface and combines commonly needed resources, such as virtual private servers, storage, databases, networking, and static IP addresses, into simplified bundles with predictable monthly pricing. A developer can select a blueprint for a supported operating system or application stack, create an instance in a few steps, and begin deploying a website or web application without first needing to assemble and configure numerous individual AWS services.

This approach helps a new AWS user focus on building the application while still using AWS infrastructure. As requirements grow, Lightsail also supports connections to other AWS services, allowing the application to expand gradually. AWS describes Lightsail as a service for getting started quickly with cloud resources and notes that it is well suited to simple web applications and websites. See the Amazon Lightsail product page and the Amazon Lightsail User Guide.

QUESTION NO: 5

In which situations should a company create an 1AM user instead of an 1AM role? (Select TWO.)

A. When an application that runs on Amazon EC2 instances requires access to other AWS services

B. When the company creates AWS access credentials for individuals

C. When the company creates an application that runs on a mobile phone that makes requests to AWS

D. When the company needs to add users to 1AM groups

E. When users are authenticated in the corporate network and want to be able to use AWS without having to sign in a second time

ANSWER: B D

Explanation:

“When the company creates AWS access credentials for individuals” is an appropriate use of an IAM user because an IAM user represents a specific person or workload that requires long-term AWS credentials. An IAM user can be given a unique sign-in name and password for AWS Management Console access, and, where programmatic access is required, access keys. Permissions can then be assigned according to the principle of least privilege and managed for that individual identity.

“When the company needs to add users to 1AM groups” is also appropriate because IAM groups are collections of IAM users. Groups simplify permission administration by allowing policies to be attached to the group, with the resulting permissions applying to the member users. This is useful when several individual users need a shared set of permissions based on a job function or team. AWS documents that IAM users are identities with long-term credentials, while groups are specifically a way to organize IAM users and manage their permissions collectively. See the AWS IAM users documentation and the AWS IAM groups documentation.

QUESTION NO: 6

Which services can be used to deploy applications on AWS? (Select TWO.)

A. AWS Elastic Beanstalk

B. AWS Config

C. AWS OpsWorks
Q

D. AWS Application Discovery Service

E. Amazon Kinesis

ANSWER: A C

Explanation:

AWS Elastic Beanstalk is an application deployment service that lets developers upload application code while AWS provisions and operates the supporting environment. It can create resources such as Amazon EC2 instances, load balancers, Auto Scaling groups, and monitoring integrations, while allowing customers to adjust the underlying configuration as needed. It supports common web application platforms and container-based applications, making it a direct fit for deploying an application without manually building each component of the environment.

AWS OpsWorks
Q refers to AWS OpsWorks, which provides managed configuration-management capabilities and application deployment automation. In particular, AWS OpsWorks Stacks supports defining application layers, instances, and deployment behavior through lifecycle events and Chef-based automation. This enables teams to deploy application versions consistently across fleets of instances and manage the configuration required by those applications. Both services therefore address application

deployment, though they use different operational models: Elastic Beanstalk emphasizes managed application environments, while OpsWorks emphasizes configuration automation and controlled deployment workflows.

For additional details, see the AWS Elastic Beanstalk Developer Guide and the AWS OpsWorks Stacks User Guide.

QUESTION NO: 7

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: A B

Explanation:

AWS CloudTrail and Amazon CloudWatch together provide visibility and retained records that support operational oversight, governance, compliance, and auditing. AWS CloudTrail records activity in an AWS account, including actions taken through the AWS Management Console, AWS CLI, SDKs, and AWS services. CloudTrail event history provides a searchable record of management events, while trails can deliver ongoing event records to Amazon S3 for retention and analysis. This creates an audit trail of who performed an action, when it occurred, the source IP address, and the resources involved. Amazon CloudWatch monitors AWS resources and applications by collecting metrics, logs, and events. CloudWatch Logs can centrally collect log data and retain it for a configured period, enabling teams to monitor operational activity, investigate events, create alarms, and support ongoing compliance evidence collection. Used together, CloudTrail supplies the account-activity audit record and CloudWatch provides monitoring, log analysis, and alerting capabilities around operational and application events. AWS documents CloudTrail's governance, compliance, and operational-auditing use cases in the AWS CloudTrail User Guide. CloudWatch logging and retention capabilities are described in the Amazon CloudWatch Logs User Guide.

QUESTION NO: 8

Which of the following are benefits of using AWS Trusted Advisor? (Choose two.)

A. Providing high-performance container orchestration

B. Creating and rotating encryption keys

C. Detecting underutilized resources to save costs

D. Improving security by proactively monitoring the AWS environment

E. Implementing enforced tagging across AWS resources

ANSWER: C D

Explanation:

AWS Trusted Advisor analyzes an AWS account against AWS best-practice checks and provides recommendations in categories such as cost optimization and security. Detecting underutilized resources to save costs is a core benefit because Trusted Advisor identifies opportunities to reduce unnecessary spend, including resources that are idle or insufficiently used. These recommendations help customers review their environment and take informed action to optimize costs.

Improving security by proactively monitoring the AWS environment is also a benefit because Trusted Advisor security checks identify configurations that may increase risk, such as overly permissive access or missing protective settings. The service presents findings and recommended actions so customers can continually assess their AWS environment against established AWS guidance. Trusted Advisor therefore supports ongoing operational review of both cost efficiency and security posture, although customers remain responsible for evaluating and applying the recommendations. AWS documents the available check categories and recommendations in the AWS Trusted Advisor User Guide. For an overview of how Trusted Advisor helps optimize costs, improve performance, strengthen security, and increase fault tolerance, see AWS Trusted Advisor.

QUESTION NO: 9

Which pillar of the AWS Well-Architected Framework includes the AWS shared responsibility model?

A. Operational excellence

B. Performance efficiency

C. Reliability

D. Security

ANSWER: D

Explanation:

The correct answer is **Security**. The AWS shared responsibility model is a foundational security concept that defines how security and compliance responsibilities are divided between AWS and the customer. AWS is responsible for security *of* the cloud, including the facilities, physical hardware, networking, and foundational services that operate the AWS Cloud. Customers are responsible for security *in* the cloud, including configuration of AWS services, identity and access management, data protection, operating systems and applications where applicable, and network controls.

The Security pillar of the AWS Well-Architected Framework focuses on protecting data, systems, and assets while delivering business value through risk assessments and mitigation strategies. Its design principles include establishing a strong identity foundation, enabling traceability, protecting data in transit and at rest, keeping people away from data through automation, and preparing for security events. Understanding the shared responsibility model helps customers make appropriate security decisions based on the AWS services they use and their deployment choices. AWS identifies this model as a key concept within the Security pillar. See the AWS Well-Architected Security Pillar and the AWS Shared Responsibility Model.

QUESTION NO: 10

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: B

Explanation:

3 hours, 5 minutes, and 6 seconds is correct because Amazon EC2 On-Demand instances running Amazon Linux are billed in one-second increments, after a minimum charge of 60 seconds. Since this instance ran substantially longer than one minute, AWS measures and bills the actual elapsed runtime rather than rounding usage up to the next full minute or hour. The billable duration is therefore the complete period from launch until the instance enters the stopped or terminated state: 3 hours, 5 minutes, and 6 seconds. Per-second billing for eligible EC2 On-Demand instances gives customers more

precise cost alignment for workloads with variable durations, including short-lived compute jobs and development environments. AWS pricing documentation states that Linux-based EC2 instances are billed per second, with a 60-second minimum, unless a specific instance or software licensing arrangement uses different pricing rules. For this Amazon Linux workload, the stated runtime is the chargeable usage duration. See the Amazon EC2 On-Demand Pricing page and the Amazon EC2 instance lifecycle documentation for billing and instance-state context.

QUESTION NO: 11

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

A. Trade variable expenses for capital expenses.

B. Deploy globally in minutes.

C. Plan capacity in advance of deployments.

D. Take advantage of economies of scale.

E. Reduce dependencies on network connectivity.

ANSWER: B D

Explanation:

“Deploy globally in minutes” is a core AWS Cloud benefit because AWS provides a broad global infrastructure of Regions and Availability Zones. Customers can provision resources in multiple geographic locations quickly, placing applications nearer to users and supporting resilient architectures without building and operating physical data centers in each location. AWS describes the ability to go global in minutes as one of the fundamental advantages of cloud computing.

“Take advantage of economies of scale” is also a key benefit. AWS aggregates demand from a very large customer base and operates infrastructure at massive scale. This scale enables AWS to improve efficiency and offer services on a consumption-based pricing model, so customers can benefit from lower variable costs than they could typically achieve by purchasing and managing equivalent infrastructure independently. These benefits help organizations use cloud resources more efficiently while scaling their workloads as needed. See the AWS overview of cloud computing advantages and AWS documentation on the six advantages of cloud computing.

QUESTION NO: 12

Which actions are best practices for an AWS account root user? (Choose two.)

A. Share root user credentials with team members.

B. Create multiple root users for the account, separated by environment.

C. Enable multi-factor authentication (MFA. on the root user.

D. Create an IAM user with administrator privileges for daily administrative tasks, instead of using the root user.

E. Use programmatic access instead of the root user and password.

ANSWER: C D

Explanation:

Enable multi-factor authentication (MFA) on the root user because the root user has unrestricted access to the AWS account, including account-level settings and sensitive billing or security functions. MFA adds a required second authentication factor, significantly reducing the risk that a compromised password alone could lead to account takeover. AWS specifically recommends enabling MFA for the root user and protecting the root credentials carefully.

Create an IAM user with administrator privileges for daily administrative tasks, instead of using the root user, because routine work should be performed through identities with controlled, auditable access rather than the root identity. In practice,

AWS recommends avoiding root-user access keys and using IAM identities—preferably with appropriate permissions—for everyday administration. This supports least-privilege access management, enables activity tracking by individual identity, and limits use of the root user to tasks that specifically require it. For additional guidance, see the AWS root user best practices and AWS documentation on the account root user.

QUESTION NO: 13

Which actions are examples of a company's effort to rightsize its AWS resources to control cloud costs? (Choose two.)

A. Switch from Amazon RDS to Amazon DynamoDB to accommodate NoSQL datasets.

B. Base the selection of Amazon EC2 instance types on past utilization patterns.

C. Use Amazon S3 Lifecycle policies to move objects that users access infrequently to lower-cost storage tiers.

D. Use Multi-AZ deployments for Amazon RDS.

E. Replace existing Amazon EC2 instances with AWS Elastic Beanstalk.

ANSWER: B C

Explanation:

Rightsizing means aligning AWS resource capacity and configuration with actual workload demand, while avoiding unnecessary cost. "Base the selection of Amazon EC2 instance types on past utilization patterns." is a rightsizing action because historical CPU, memory, network, and other utilization data helps a company choose instance types and sizes that meet performance needs without paying for consistently unused capacity. AWS recommends using metrics and utilization observations to identify underused or oversized resources and select more appropriate configurations.

"Use Amazon S3 Lifecycle policies to move objects that users access infrequently to lower-cost storage tiers." also supports rightsizing and cost control. S3 Lifecycle rules automatically transition objects as their access characteristics change, such as moving eligible data from S3 Standard to S3 Standard-Infrequent Access, S3 Glacier Instant Retrieval, or archival tiers. This matches the storage class and its price to the object's real access pattern without requiring application changes or manual movement of data. Together, these practices continuously align compute and storage consumption with business requirements. See the AWS Well-Architected Framework guidance on compute sizing and Amazon S3 Lifecycle management documentation.

QUESTION NO: 14

Which options are AWS Cloud Adoption Framework (AWS CAF. security perspective capabilities? (Choose two.)

A. Observability

B. Incident and problem management

C. Incident response

D. Infrastructure protection

E. Availability and continuity

ANSWER: C D

Explanation:

Incident response and **Infrastructure protection** are capabilities in the AWS Cloud Adoption Framework (AWS CAF) security perspective. The security perspective helps an organization establish and maintain the safeguards, processes, and technical controls required to protect cloud workloads and data while supporting business objectives and risk-management requirements.

Incident response focuses on preparing for, detecting, containing, investigating, and recovering from security events. This capability includes establishing response procedures, assigning responsibilities, collecting relevant evidence, and using lessons learned to improve security operations. Effective incident response enables teams to minimize the impact of security incidents and restore normal operations efficiently.

Infrastructure protection focuses on securing the underlying cloud environment and workload infrastructure. It covers applying protective controls to networks, compute resources, operating systems, and other infrastructure components, using mechanisms such as segmentation, secure configuration, patching, and monitoring. These controls help reduce exposure to threats and enforce the organization's security requirements across AWS environments.

AWS documents both capabilities within the AWS CAF security perspective. See the AWS CAF Security Perspective and the AWS Cloud Adoption Framework overview.

QUESTION NO: 15

A company wants its Amazon EC2 instances to share the same geographic area but use multiple independent underlying power sources.

Which solution achieves this goal?

A. Use EC2 instances in a single Availability Zone.

B. Use EC2 instances in multiple AWS Regions.

C. Use EC2 instances in multiple Availability Zones in the same AWS Region.

D. Use EC2 instances in the same edge location and the same AWS Region.

ANSWER: C

Explanation:

Use EC2 instances in multiple Availability Zones in the same AWS Region. An AWS Region is a distinct geographic area, while Availability Zones are separate, physically isolated locations within that Region. Each Availability Zone is designed with independent infrastructure, including separate power, cooling, and physical security, so that a disruption affecting one Availability Zone is less likely to affect another. Deploying EC2 instances across multiple Availability Zones therefore keeps the workload in the same regional geographic area while providing resilience through independent underlying power sources and other redundant facilities.

This is a core AWS high-availability design approach. For example, an application can place instances in subnets mapped to different Availability Zones and use an Elastic Load Balancer and an Auto Scaling group spanning those zones. The architecture can continue serving requests when an Availability Zone experiences an infrastructure failure, provided the application itself is designed for multi-AZ operation. AWS describes Regions as consisting of multiple Availability Zones and explains that Availability Zones have independent power, cooling, and physical security. See the AWS Global Infrastructure overview and the Amazon EC2 Regions and Availability Zones documentation.

QUESTION NO: 16

Which actions are best practices for an AWS account root user? (Select TWO.)

A. Share root user credentials with team members.

B. Create multiple root users for the account, separated by environment.

C. Enable multi-factor authentication (MFA) on the root user.

D. Create an IAM user with administrator privileges for daily administrative tasks, instead of using the root user.

ANSWER: C D

Explanation:

Enable multi-factor authentication (MFA) on the root user because the root user has unrestricted access to the account, including access to billing information and account-level settings. MFA requires an additional authentication factor beyond the password, substantially reducing the risk that a compromised password alone could lead to account takeover. AWS specifically recommends enabling MFA for the root user as a foundational account-security measure.

Create an IAM user with administrator privileges for daily administrative tasks, instead of using the root user, because the root user should be reserved for the limited account and service-management operations that require root credentials. Using a separate administrative identity for routine work supports least privilege, provides clearer auditing in AWS CloudTrail, and limits the impact of accidental actions or credential exposure. In modern AWS environments, IAM Identity Center is also commonly used to provide administrative workforce access, but an appropriately protected administrator IAM user remains aligned with this root-user best practice. For AWS guidance, see Root user best practices and Enable MFA for AWS account root users.

QUESTION NO: 17

A company has set up a VPC in its AWS account and has created a subnet in the VPC. The company wants to make the subnet public.

Which AWS features should the company use to meet this requirement? (Select TWO.)

A. Amazon VPC internet gateway

B. Amazon VPC NAT gateway

C. Amazon VPC route tables

D. Amazon VPC network ACL

E. Amazon EC2 security groups

ANSWER: A C

Explanation:

A public subnet is a subnet whose associated route table includes a route to an internet gateway. Therefore, the company needs an **Amazon VPC internet gateway** attached to its VPC and **Amazon VPC route tables** configured with a route that directs internet-bound IPv4 traffic, commonly `0.0.0.0/0`, to that internet gateway. The internet gateway provides the highly available connection point between the VPC and the public internet, while the route table determines that traffic from resources in the subnet should use that connection point. Resources that need direct IPv4 internet communication must also have public IPv4 addresses or Elastic IP addresses, but assigning such an address does not itself make a subnet public. The defining subnet-level configuration is the internet-gateway route in the route table associated with that subnet. AWS documents this setup as the required routing configuration for public subnets. For implementation guidance, see the Amazon VPC internet gateway documentation and the Amazon VPC route table documentation.

QUESTION NO: 18

Which option is a customer responsibility when using Amazon DynamoDB under the AWS Shared Responsibility Model?

A. Physical security of DynamoDB

B. Patching of DynamoDB

C. Access to DynamoDB tables

D. Encryption of data at rest in DynamoDB

ANSWER: C

Explanation:

Access to DynamoDB tables is a customer responsibility because customers control who can perform actions on their DynamoDB resources. Using AWS Identity and Access Management (IAM), customers create and manage identities, permissions policies, roles, and access controls that determine whether users, applications, and AWS services can read, write, update, delete, or administer tables. Customers should apply least-privilege permissions, use roles rather than long-term credentials where appropriate, and regularly review access to ensure that only authorized principals can access sensitive data. DynamoDB also supports resource-based policies and fine-grained access controls in applicable scenarios, allowing customers to further scope access to particular tables, indexes, or items. Under the shared responsibility model, AWS operates and secures the underlying cloud infrastructure and managed DynamoDB service, while the customer remains responsible for security *in* the cloud, including identity and access configuration for their data and resources. AWS explains this division of responsibilities in its Shared Responsibility Model. For details on controlling DynamoDB permissions with IAM, see the Amazon DynamoDB developer guide.

QUESTION NO: 19

A company is running a workload in the AWS Cloud.

Which AWS best practice ensures the MOST cost-effective architecture for the workload?

A. Loose coupling

B. Rightsizing

C. Caching

D. Redundancy

ANSWER: B

Explanation:

Rightsizing is the AWS best practice that most directly supports a cost-effective workload architecture. It means selecting resource types and sizes that match the workload's actual requirements, then continuously reviewing utilization and adjusting capacity as demand changes. For example, a company can use metrics such as CPU utilization, memory use, network throughput, and storage demand to identify Amazon EC2 instances, databases, or other resources that are overprovisioned or underprovisioned. Choosing appropriately sized resources avoids paying for unused capacity while still maintaining the performance the application requires.

Rightsizing is not a one-time activity. Workload patterns can change over time because of growth, seasonality, application releases, or changing user behavior. AWS recommends using services and tools such as AWS Compute Optimizer, Amazon CloudWatch, and AWS Cost Explorer to evaluate utilization, receive recommendations, and make informed adjustments. This ongoing measurement and optimization aligns with the Cost Optimization pillar of the AWS Well-Architected Framework, which emphasizes eliminating unnecessary spending and selecting the most appropriate resources for a workload. See the AWS Well-Architected Cost Optimization Pillar and AWS Compute Optimizer documentation.

QUESTION NO: 20

Which options are AWS Cloud Adoption Framework (AWS CAF. cloud transformation journey recommendations? (Choose two.)

A. Envision phase

B. Align phase

C. Assess phase

D. Mobilize phase

E. Migrate and modernize phase

ANSWER: A B

Explanation:

The AWS Cloud Adoption Framework (AWS CAF) describes a cloud transformation journey that helps organizations progress from defining the intended business outcomes to establishing the organizational capabilities needed to realize them. **Envision phase** is correct because it focuses on identifying and prioritizing transformation opportunities, defining the desired business outcomes, and building a compelling case for change. This gives stakeholders a shared understanding of why the organization is transforming and what value it expects to achieve.

Align phase is also correct because it is where the organization evaluates its current state, identifies capability gaps, and develops an actionable transformation roadmap. Aligning business, people, governance, platform, security, and operations capabilities enables the organization to prepare for execution in a coordinated way. Together, envisioning the target outcomes and aligning the organization around a practical roadmap are foundational recommendations in the AWS CAF transformation journey. AWS describes the AWS CAF transformation journey and its phases on the AWS Cloud Adoption Framework page. Additional AWS CAF guidance is available in the AWS CAF whitepaper.

QUESTION NO: 21

A company wants to visualize and manage AWS Cloud costs and usage for a specific period of time.

Which AWS service or feature will meet these requirements?

A. Cost Explorer

B. Consolidated billing

C. AWS Organizations

D. AWS Budgets

ANSWER: A

Explanation:

Cost Explorer is the AWS cost-management feature designed to visualize, understand, and manage AWS costs and usage over selected time periods. It provides an interactive interface where a company can review historical spending and usage, choose custom date ranges, and group or filter results by dimensions such as AWS service, Region, linked account, usage type, tags, and cost allocation categories. This lets users investigate what drove costs during a particular period rather than viewing only a billing total. Cost Explorer also supplies charts, reports, and cost forecasts, helping teams identify trends and make informed decisions about resource usage and future spending. Data can be viewed at monthly granularity for longer periods and daily granularity for more recent periods, which supports both high-level financial reviews and closer examination of short-term changes. AWS documents Cost Explorer as a tool for analyzing AWS costs and usage, including customizable reports and forecasts. Learn more in the AWS Cost Explorer User Guide and the AWS Cost Explorer product page.

QUESTION NO: 22

A company wants to migrate its application to AWS. The company wants to replace upfront expenses with variable payment that is based on usage.

What should the company do to meet these requirements?

A. Use pay-as-you-go pricing.

B. Purchase Reserved Instances.

C. Pay less by using more.

D. Rightsize instances.

ANSWER: A

Explanation:

Use pay-as-you-go pricing. AWS's pay-as-you-go model lets a customer consume resources as needed and pay for the resources used, rather than making a large upfront capital investment in data center hardware and infrastructure. This converts costs from capital expenditures to variable operational expenses, which directly matches the company's goal of paying according to actual usage. For example, AWS services such as Amazon EC2 can be used on an On-Demand basis, where compute capacity is charged by the time it runs without requiring a long-term commitment. This model also enables the company to scale resource consumption up or down as application demand changes, so its spending can align more closely with workload needs. AWS identifies trading upfront expense for variable expense as a core cloud-computing advantage. The AWS pricing approach is designed so customers can start using services when needed and pay only for their consumption under the applicable service pricing terms. See the AWS Pricing overview and the AWS overview of cloud-computing advantages.

QUESTION NO: 23

In which categories does AWS Trusted Advisor provide recommended actions? (Select TWO.)

A. Operating system patches

B. Cost optimization

C. Repetitive tasks

D. Service quotas

E. Account activity records

ANSWER: B D

Explanation:

AWS Trusted Advisor provides recommendations that help customers align their AWS environments with AWS best practices. Its checks are organized into categories that include cost optimization, performance, security, fault tolerance, and service quotas. Therefore, **Cost optimization** and **Service quotas** are valid categories in which Trusted Advisor provides recommended actions. Cost optimization recommendations identify opportunities to reduce spending, such as resources that are idle, underused, or not configured in the most economical way. Service quotas recommendations help customers monitor usage against AWS service limits, identify resources approaching quota thresholds, and take action such as requesting quota increases when appropriate. These categories support proactive cloud management by helping organizations control costs and avoid operational constraints caused by quota limits. AWS documents the Trusted Advisor check categories and explains that available checks can vary according to the customer's AWS Support plan. See the AWS Trusted Advisor best practice check reference and the AWS Trusted Advisor overview.

QUESTION NO: 24

According to the AWS shared responsibility model, the customer is responsible for applying the latest security updates and patches for which of

the following?

A. Amazon DynamoDB

B. Amazon EC2 instances

C. Amazon RDS instances

D. Amazon S3

ANSWER: B

Explanation:

Amazon EC2 instances are correct because they use the infrastructure-as-a-service model. AWS is responsible for security *of* the cloud, including the physical facilities, hardware, networking, and virtualization layer that support the service. The customer is responsible for security *in* the cloud for the resources they configure and operate. For an Amazon EC2 instance, this includes administering the guest operating system and installing its security updates and patches. Customers must also manage the applications installed on the instance, host-level firewall configuration, identity and access controls, and data protection appropriate to their workload.

This division of responsibility applies even when AWS provides and maintains the underlying infrastructure. Because the customer selects, configures, and controls the operating system environment on an Amazon EC2 instance, the customer must establish a patch-management process that regularly evaluates and deploys operating-system and application updates. AWS summarizes these responsibilities in its shared responsibility model and identifies customer responsibilities for Amazon EC2 security in the service documentation. See the AWS Shared Responsibility Model and the Amazon EC2 security documentation.

QUESTION NO: 25

A company wants to receive a notification when a specific AWS cost threshold is reached.

Which AWS services or tools can the company use to meet this requirement? (Select TWO.)

A. Amazon Simple Queue Service (Amazon SQS)

B. AWS Budgets

C. Cost Explorer

D. Amazon CloudWatch

E. AWS Cost and Usage Report

ANSWER: B D

Explanation:

AWS Budgets lets a company define cost budgets with a specified monetary amount and configure notifications when actual costs or forecasted costs meet or exceed that threshold. Budget alerts can be delivered by email or through Amazon Simple Notification Service (Amazon SNS), allowing the company to notify the appropriate people promptly. Budgets are designed for proactive cost governance because they can evaluate spending against a defined budget amount and can also trigger configured budget actions.

Amazon CloudWatch supports AWS billing alarms that monitor estimated charges and send a notification when the configured charge threshold is crossed. A billing alarm uses the AWS/Billing estimated-charges metric and can publish an alert to an Amazon SNS topic. This is useful when the requirement is to alert on an account's estimated accumulated charges rather than to manage a broader budget plan. To use CloudWatch billing alarms, billing metrics must be enabled for the account, and billing alarms are created in the US East (N. Virginia) Region. Together, AWS Budgets and CloudWatch billing alarms provide AWS-native ways to notify a company when costs reach a chosen threshold. See AWS Budgets documentation and Monitoring estimated charges with CloudWatch.

QUESTION NO: 26

Which AWS services can a company use to translate and analyze text without building and managing machine learning models? (Choose two.)

A. Amazon Translate

B. Amazon Comprehend

C. Amazon Rekognition

D. AWS Batch

E. Amazon EC2 Auto Scaling

ANSWER: A B

Explanation:

Amazon Translate and Amazon Comprehend are AWS artificial intelligence services that provide pre-trained capabilities through APIs. Amazon Translate provides neural machine translation between supported languages. Amazon Comprehend uses natural language processing to identify insights in text, such as entities, key phrases, sentiment, and language.

These managed AI services allow a company to add text translation and analysis capabilities without developing, training, or operating its own machine learning models. For more information, see the Amazon Translate Developer Guide and the Amazon Comprehend Developer Guide.

QUESTION NO: 27

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: A B

Explanation:

Elasticity is the ability to dynamically acquire or release cloud resources so that capacity aligns with changing workload demand. “Scaling the number of Amazon EC2 instances based on traffic” is an elasticity scenario because compute capacity can increase during periods of higher traffic and decrease when demand subsides. AWS Auto Scaling can automate this adjustment using scaling policies and metrics, allowing an application to maintain appropriate capacity without permanently provisioning for peak load.

“Resizing Amazon RDS instances as business needs change” also represents elasticity because a database deployment can be modified to use an instance class and storage configuration that better fits its current workload requirements. As an organization’s database demand grows or changes, it can adjust these resources rather than treating the original capacity selection as fixed for the lifetime of the application.

Both scenarios demonstrate the cloud benefit of matching resource consumption to actual demand. This helps organizations avoid maintaining unnecessary capacity while retaining the ability to accommodate growth. AWS describes elasticity as the ability to scale resources up or down as needed in its overview of cloud-computing advantages: AWS Overview: Six Advantages of Cloud Computing. For RDS instance modifications, see the Amazon RDS User Guide.

QUESTION NO: 28

Which option presents a benefit of using managed services?

A. Eliminate the need to manage underlying infrastructure

B. Provide better performance than self-managed infrastructure

C. Provide unlimited storage capacity

D. Support more applications than self-managed infrastructure

ANSWER: A

Explanation:

Eliminate the need to manage underlying infrastructure is a core benefit of AWS managed services. With a managed service, AWS operates and maintains much of the undifferentiated infrastructure work that customers would otherwise need to perform themselves, such as provisioning and maintaining servers, applying infrastructure-level updates, replacing failed hardware, and supporting service availability. This lets an organization spend more time on its applications, data, and business requirements rather than routine operational activities. The precise division of responsibility depends on the specific AWS service and configuration: customers still retain responsibility for areas such as their data, identities and access permissions, application configuration, and appropriate security settings. For example, AWS explains that its managed database offerings reduce the administrative burden associated with operating database infrastructure, while the customer continues to manage the database content and usage. Using managed services therefore supports the AWS Cloud value proposition of shifting operational responsibilities to AWS where AWS is best positioned to handle them. See the AWS Managed Services overview and the AWS Shared Responsibility Model for further detail.

QUESTION NO: 29

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 the purpose-built AWS time-series database service for storing and analyzing time-stamped data at massive scale. Time-series workloads commonly involve continuously generated events, measurements, and records, such as application metrics, IoT telemetry, clickstreams, and operational monitoring data. Timestream is designed to simplify these workloads by automatically managing the underlying infrastructure, scaling capacity as data volumes and query activity change, and organizing data across memory and magnetic storage tiers to balance query performance and cost. It also provides SQL-based querying and built-in time-series functions, making it practical to analyze recent and historical event data without building and operating a database platform manually. AWS specifically describes Amazon Timestream as a fast, scalable, serverless time-series database service that can store and analyze trillions of time-series data points per day. Therefore, it directly meets the stated requirement for a managed service capable of handling trillions of daily events. See the Amazon Timestream product page and the Amazon Timestream Developer Guide.

QUESTION NO: 30

A company wants to track tags, buckets, and prefixes for its Amazon S3 objects.

Which S3 feature will meet this requirement?

A. S3 Inventory report

B. S3 Lifecycle

C. S3 Versioning

D. S3 ACLs

ANSWER: A

Explanation:

S3 Inventory report is the appropriate feature because it provides scheduled, file-based lists of objects and their metadata for one or more Amazon S3 buckets. An inventory configuration can be scoped to all objects in a bucket or limited to objects

with a specified prefix, allowing the company to report on the relevant object group. The generated inventory includes each object's key name, which identifies its location under a prefix, along with other object properties. The report can also be configured to include object tags as optional metadata fields. This gives the company a scalable way to track object tagging information and object locations across its bucket inventory without requiring individual object-listing requests. Inventory reports are delivered on a daily or weekly schedule in CSV, Apache ORC, or Apache Parquet format to a designated S3 destination bucket. The resulting data can be queried with services such as Amazon Athena or processed through analytics and governance workflows. Amazon documents the available inventory fields, including optional object tags, and explains how inventory configurations can apply a prefix filter: Amazon S3 Inventory and Configuring Amazon S3 Inventory.

QUESTION NO: 31

Which factors affect costs in the AWS Cloud? (Select TWO.)

A. The number of unused AWS Lambda functions

B. The number of configured Amazon S3 buckets

C. Inbound data transfers without acceleration

D. Outbound data transfers without acceleration

E. Compute resources that are currently in use

ANSWER: D E

Explanation:

Outbound data transfers without acceleration affect AWS cost because AWS pricing commonly includes charges for data transferred out from AWS to the internet. The applicable rate depends on factors such as the AWS service, source Region, destination, and total amount of data transferred during the billing period. While services such as Amazon CloudFront can change the architecture and applicable pricing model, outbound transfer is itself a metered consumption dimension that can contribute directly to a customer's bill. AWS provides service-specific pricing and data-transfer details so customers can estimate these charges for their workloads.

Compute resources that are currently in use also affect cost because AWS charges are based on consumption of provisioned or executed compute capacity. For example, Amazon EC2 pricing is associated with instance usage, while serverless services can charge for invocations, execution duration, memory allocation, or other measured usage dimensions. The selected compute service, configuration, pricing model, and period of active use determine the resulting charges. Reviewing actual resource consumption in AWS billing tools and using AWS pricing calculators supports accurate cost estimation and ongoing optimization.

References: AWS Pricing and Amazon EC2 Pricing.

QUESTION NO: 32

Which of the following actions are controlled with AWS Identity and Access Management (IAM)? (Select TWO.)

A. Control access to AWS service APIs and to other specific resources.

B. Provide intelligent threat detection and continuous monitoring.

C. Protect the AWS environment using multi-factor authentication (MFA).

D. Grant users access to AWS data centers.

E. Provide firewall protection for applications from common web attacks.

ANSWER: A C

Explanation:

AWS Identity and Access Management (IAM) controls authorization to AWS services and resources. IAM policies define the permissions that allow or deny actions against AWS service APIs and specific resources. These policies can be attached to IAM identities, such as users, groups, and roles, and can also be used as resource-based policies for supported services. This lets an AWS account implement least-privilege access by specifying who can perform particular actions, on which resources, and under which conditions.

IAM also supports multi-factor authentication (MFA) as an additional sign-in security control. MFA requires a second authentication factor in addition to a password, helping protect the AWS account root user and IAM users if credentials are exposed. AWS recommends enabling MFA, particularly for the root user, because the root user has unrestricted access to account resources. IAM therefore provides both the identity and permission-management capabilities needed to govern API and resource access, along with MFA controls that strengthen authentication security. See the AWS IAM User Guide and AWS guidance on using multi-factor authentication in IAM.

QUESTION NO: 33

Which AWS service or tool gives users the ability to connect with AWS and deploy resources programmatically?

A. Amazon quickSight

B. AWS PrivateLink

C. AWS Direct Connect

D. AWS SDKs

ANSWER: D

Explanation:

AWS SDKs are correct because they provide language-specific libraries and tools that developers use to interact with AWS services through code. An SDK supplies the components needed to make authenticated calls to AWS service APIs from an application, including support for signing requests with AWS credentials, request serialization, response handling, retries, and error handling. AWS offers SDKs for popular languages and platforms, such as Python, JavaScript, Java, .NET, Go, Ruby, PHP, C++, and Rust. For example, a developer can use an AWS SDK in Python to create an Amazon S3 bucket, launch Amazon EC2 instances, invoke an AWS Lambda function, or configure other AWS resources without manually interacting with the AWS Management Console. SDKs therefore enable programmatic access to AWS and are a foundational way for applications and automation scripts to call AWS APIs. Although infrastructure can also be deployed using infrastructure-as-code tools, the question asks for the AWS tool that lets users connect to AWS and work with resources programmatically; AWS SDKs directly fulfill that role by exposing AWS APIs through idiomatic language libraries. AWS documents the available SDKs and their purpose at AWS Developer Tools. Details about making API requests and using AWS SDKs are available in the AWS SDKs and Tools Reference Guide.

QUESTION NO: 34

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

A. Server-side encryption with Amazon S3 managed encryption keys (SSE-S3)

B. Server-side encryption with AWS KMS managed keys (SSE-KMS)

C. TLS

D. SSL

E. Transparent Data Encryption (TDE)

ANSWER: A B

Explanation:

Server-side encryption with Amazon S3 managed encryption keys (SSE-S3) and server-side encryption with AWS KMS managed keys (SSE-KMS) both protect Amazon S3 objects while they are stored at rest. With SSE-S3, Amazon S3 manages the encryption keys and uses strong server-side encryption automatically when the object is written to the bucket. This is the default encryption baseline for newly uploaded S3 objects unless another encryption method is explicitly requested. It is appropriate when an organization wants S3 to handle key management without requiring direct control over keys.

Server-side encryption with AWS KMS managed keys (SSE-KMS) encrypts objects through AWS Key Management Service. It provides additional control and visibility over key usage, including AWS KMS permissions, key policies, and audit records in AWS CloudTrail for KMS API activity. Organizations commonly select SSE-KMS when governance, access-control requirements, or auditing requirements call for using a KMS key. Both methods encrypt data at rest on S3 storage infrastructure and decrypt it transparently for authorized requests. AWS documents the available S3 server-side encryption choices at Protecting data with server-side encryption and describes the KMS integration at Using SSE-KMS encryption.

QUESTION NO: 35

Which of the following actions are controlled with AWS Identity and Access Management (1AM)? (Select TWO.)

A. Control access to AWS service APIs and to other specific resources.

B. Provide intelligent threat detection and continuous monitoring.

C. Protect the AWS environment using multi-factor authentication (MFA).

D. Grant users access to AWS data centers.

E. Provide firewall protection for applications from common web attacks.

ANSWER: A C

Explanation:

AWS Identity and Access Management (IAM) controls authentication and authorization for AWS accounts. “Control access to AWS service APIs and to other specific resources” is correct because IAM lets account administrators define who or what can make requests to AWS services and resources. IAM identities, including users and roles, receive permissions through policies. These policies can allow or deny actions on specified resources, enabling least-privilege access for people, workloads, and applications.

“Protect the AWS environment using multi-factor authentication (MFA)” is also correct because IAM supports MFA as an additional authentication factor for AWS account access. MFA requires a user to provide a second factor in addition to their password when signing in, helping protect IAM users and the AWS account root user from unauthorized use of compromised credentials. IAM can enforce MFA-related requirements through policy conditions for sensitive actions, strengthening account security. For details, see the AWS IAM User Guide introduction and AWS guidance on using multi-factor authentication in AWS.

QUESTION NO: 36

A company is running and managing its own Docker environment on Amazon EC2 instances. The company wants an alternative to help manage

cluster size, scheduling, and environment maintenance.

Which AWS service meets these requirements?

A. AWS Lambda

B. Amazon RDS

C. AWS Fargate

D. Amazon Athena

ANSWER: C

Explanation:

AWS Fargate is the appropriate service because it provides serverless compute capacity for containerized applications. Rather than operating a self-managed Docker environment on Amazon EC2 instances, the company can run container workloads using Fargate with Amazon Elastic Container Service or Amazon Elastic Kubernetes Service. AWS provisions and manages the underlying compute infrastructure, including the servers, operating systems, and capacity required to run the containers. This removes the operational responsibility of selecting, patching, scaling, and maintaining EC2 worker instances.

Fargate is designed to let teams define the resources required by each containerized task or pod, such as CPU and memory, while the AWS container orchestration service handles placement and scheduling. The company can therefore focus on its application configuration and desired scaling behavior instead of maintaining the cluster infrastructure. This directly addresses the need to reduce administration associated with cluster capacity, workload scheduling, and environment maintenance. For additional details, see the AWS Fargate documentation for Amazon ECS and the AWS Fargate product page.

QUESTION NO: 37

A company is migrating an application that includes an Oracle database to AWS. The company cannot rewrite the application.

To which AWS service could the company migrate the database?

A. Amazon Athena

B. Amazon DynamoDB®

C. Amazon DocumentDB (with MongoDB compatibility)

D. Amazon RDS for Oracle

ANSWER: D

Explanation:

Amazon RDS for Oracle is the appropriate destination because it provides a managed relational database service that supports Oracle Database engines. Since the application cannot be rewritten, retaining Oracle database compatibility is essential: the application can continue using its existing Oracle-specific SQL, schemas, database features, and connectivity patterns rather than requiring conversion to a different database model or engine. Amazon RDS reduces operational effort by managing routine database administration activities such as provisioning, backups, software patching, monitoring, and failover capabilities, while the customer retains responsibility for database design, application integration, and access configuration. The company can choose licensing approaches that fit its Oracle licensing situation, including License Included for eligible Oracle Database editions or Bring Your Own License where applicable. This makes Amazon RDS for Oracle a practical migration target for an Oracle-backed application when minimizing application changes is the requirement. AWS documents the supported Oracle versions, editions, licensing models, and managed-service capabilities in the Amazon RDS User Guide for Oracle. For an overview of the operational features provided by the service, see Amazon RDS for Oracle.

QUESTION NO: 38

Which of the following can be components of a VPC in the AWS Cloud? (Select TWO.)

A. Amazon API Gateway

B. Amazon S3 buckets and objects

C. AWS Storage Gateway

D. Internet gateway

E. Subnet

ANSWER: D E

Explanation:

An Internet gateway and a subnet can both be components of an Amazon Virtual Private Cloud (Amazon VPC). A subnet is a logical partition of a VPC IP address range that AWS resources, such as Amazon EC2 instances, can be placed into. Each subnet resides in one Availability Zone, enabling customers to organize workloads by availability, routing, and access requirements. A VPC can contain one or more subnets, and route tables determine how traffic leaves or moves within those subnets.

An internet gateway is a VPC component that enables communication between resources in a VPC and the internet. It is attached to a VPC, and routes in a subnet route table can direct internet-bound traffic to it. For applicable resources to communicate publicly, they also need an appropriate public IPv4 address or IPv6 address and security rules that permit the desired traffic. Together, subnets provide the VPC's internal network segmentation, while an internet gateway can provide a path for authorized external connectivity. AWS documents VPC components and their networking roles in the Amazon VPC User Guide and explains internet gateway behavior in its Internet gateways documentation.

QUESTION NO: 39

Which option is an AWS Cloud Adoption Framework (AWS CAF. foundational capability for the operations perspective?

A. Performance and capacity management

B. Application portfolio management

C. Identity and access management

D. Product management

ANSWER: A

Explanation:

Performance and capacity management is a foundational capability in the operations perspective of the AWS Cloud Adoption Framework (AWS CAF). The operations perspective focuses on operating, monitoring, maintaining, and continually improving cloud workloads so that they deliver agreed business outcomes. Performance and capacity management helps an organization establish processes to measure workload performance, forecast demand, identify resource constraints, and ensure that sufficient capacity is available as requirements change.

In practice, this capability supports the ongoing operational use of metrics, monitoring, scaling plans, and capacity forecasts. Teams use it to understand whether applications are meeting performance expectations and to make informed decisions about resource allocation before demand affects users. This aligns with AWS CAF's emphasis on defining operational practices that support reliability, efficiency, and continuous improvement throughout a workload's lifecycle.

AWS identifies performance and capacity management among the capabilities of the operations perspective. For the AWS CAF perspective model and its operations guidance, see the AWS CAF Operations Perspective and the AWS Cloud Adoption Framework overview.

QUESTION NO: 40

A user has been granted permission to change their own IAM user password.

Which AWS services can the user use to change the password? (Select TWO.)

A. AWS Command Line Interface (AWS CLI)

B. AWS Key Management Service (AWS KMS)

C. AWS Management Console

D. AWS Resource Access Manager (AWS RAM)

E. AWS Secrets Manager

ANSWER: A C

Explanation:

A user who has been granted permission to change their own IAM password can do so through the **AWS Management Console** or the **AWS Command Line Interface (AWS CLI)**. In the console, the IAM user signs in, opens their security credentials, and chooses the password-change action. This is the standard browser-based method for users who need to manage their own sign-in credentials. AWS documents this workflow as part of IAM user password management.

Using the AWS CLI, the signed-in IAM user can call the `aws iam change-password` command. The command changes the password for the IAM user whose access keys are used to make the request; it does not require specifying a user name. The user must have an IAM policy permitting the `iam:ChangePassword` action, and any configured account password policy still applies to the new password. These two interfaces directly invoke IAM password-management functionality, so they are the applicable choices. For implementation details, see the AWS IAM User Guide: Change an IAM user password and the AWS CLI change-password command reference.

QUESTION NO: 41

A company is releasing a business-critical application. Before the release, the company needs strategic planning assistance from AWS. During the release, the company needs AWS infrastructure event management and real-time support.

What should the company do to meet these requirement?

A. Access AWS Trusted Advisor.

B. Contact the AWS Partner Network (APN).

C. Sign up for AWS Enterprise Support.

D. Contact AWS Professional Services.

ANSWER: C

Explanation:

Sign up for AWS Enterprise Support. is correct because it is designed for organizations operating business- and mission-critical workloads that need proactive, personalized AWS guidance as well as support during important operational events. Enterprise Support includes access to a designated Technical Account Manager (TAM), who provides strategic guidance, helps the company plan and optimize its AWS environment, and coordinates resources around the organization's business objectives. This directly supports the need for strategic planning before the application release.

For the release itself, Enterprise Support provides Infrastructure Event Management (IEM). IEM offers planning and operational guidance for planned events, such as launches and migrations, and can include real-time support and coordination from AWS specialists during the event. This helps the company prepare its architecture and operational processes, monitor the release appropriately, and obtain rapid assistance if issues arise. Enterprise Support also provides 24/7 access to cloud support engineers with the fastest response targets for critical production-impacting cases. AWS describes Enterprise Support features at AWS Enterprise Support and explains event support at AWS Infrastructure Event Management.

QUESTION NO: 42

What is a benefit of using an Elastic Load Balancing (ELB) load balancer with applications running in the AWS Cloud?

A. An ELB will automatically scale resources to meet capacity needs.

B. An ELB can balance traffic across multiple compute resources.

C. An ELB can span multiple AWS Regions.

D. An ELB can balance traffic between multiple internet gateways.

ANSWER: B

Explanation:

An ELB can balance traffic across multiple compute resources. Elastic Load Balancing distributes incoming application traffic among multiple targets, such as Amazon EC2 instances, containers, IP addresses, and AWS Lambda functions, depending on the load balancer type and configured target group. By routing requests across healthy, registered targets, ELB helps an application handle varying request volumes while improving availability and fault tolerance. For example, an Application Load Balancer can direct HTTP/HTTPS requests to multiple application targets in more than one Availability Zone. ELB continuously performs health checks and sends traffic only to targets that are considered healthy, helping maintain reliable application access when an individual target becomes unavailable. A load balancer is therefore a core component of a scalable, highly available AWS architecture: it provides one endpoint for clients and intelligently distributes the workload behind that endpoint. Capacity changes to the underlying compute fleet can be handled separately through services such as EC2 Auto Scaling, while ELB provides the traffic-distribution layer. For additional details, see the AWS Elastic Load Balancing User Guide and the Elastic Load Balancing product page.

QUESTION NO: 43

A company wants to implement controls (guardrails) in a newly created AWS Control Tower landing zone.

Which AWS services or features can the company use to create and define these controls (guardrails)? (Select TWO.)

A. AWS Config

B. Service control policies (SCPs)

C. Amazon GuardDuty

D. AWS Identity and Access Management (1AM)

E. Security groups

ANSWER: A B

Explanation:

AWS Config and Service control policies (SCPs) are the AWS Control Tower mechanisms used to implement landing-zone controls, historically called guardrails. AWS Config supports detective controls. These controls continuously evaluate resource configurations against AWS Config rules and identify resources that do not meet an organization's governance requirements. This enables centralized visibility into compliance across accounts enrolled in the Control Tower landing zone.

Service control policies (SCPs) support preventive controls. SCPs are policies managed through AWS Organizations that establish permission guardrails for accounts and organizational units. A preventive control can restrict actions before they occur, such as preventing use of services or AWS Regions that the organization has not approved. Control Tower applies its managed controls to the appropriate organizational units and accounts, helping maintain consistent governance as accounts are provisioned and enrolled.

Together, AWS Config-based detective controls and SCP-based preventive controls provide complementary governance: one evaluates and reports configuration compliance, while the other establishes organization-wide permission boundaries. AWS describes the control types and their underlying implementation in the AWS Control Tower controls documentation. Details about how SCPs set maximum available permissions are available in the AWS Organizations SCP documentation.

QUESTION NO: 44

A company needs to migrate a PostgreSQL database from on-premises to Amazon RDS.

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

A. Cloud Adoption Readiness Tool

B. AWS Migration Hub

C. AWS Database Migration Service (AWS DMS)

D. AWS Application Migration Service

ANSWER: C

Explanation:

AWS Database Migration Service (AWS DMS) is the appropriate service for migrating an on-premises PostgreSQL database to Amazon RDS. AWS DMS is designed to migrate relational databases, data warehouses, NoSQL databases, and other supported data stores to AWS with minimal downtime. For this migration, the on-premises PostgreSQL instance is configured as the source endpoint and an Amazon RDS for PostgreSQL DB instance is configured as the target endpoint. AWS DMS can perform an initial full load of the existing database data and then use change data capture (CDC) to replicate ongoing changes from the source while the application remains in use. This approach helps reduce the cutover window and supports a controlled migration to the RDS environment. The service also provides migration tasks and monitoring capabilities so the company can track progress and validate the movement of data. PostgreSQL is a supported source and target engine for AWS DMS, including migrations from self-managed PostgreSQL databases to Amazon RDS for PostgreSQL. For a detailed list of supported source-target combinations and migration planning guidance, see the AWS DMS PostgreSQL source documentation and the AWS Database Migration Service User Guide.

QUESTION NO: 45

Which AWS services are connectivity services for a VPC? (Select TWO.)

A. AWS Site-to-Site VPN

B. AWS Direct Connect

C. Amazon Connect

D. AWS Key Management Service (AWS KMS)

E. AWS Identity and Access Management (IAM)

ANSWER: A B

Explanation:

AWS Site-to-Site VPN and **AWS Direct Connect** are VPC connectivity services that link external networks to AWS resources hosted in a VPC. AWS Site-to-Site VPN creates encrypted IPsec VPN tunnels between an on-premises network and AWS. A VPN connection terminates at a virtual private gateway or transit gateway on the AWS side and typically uses a customer gateway device on the external-network side. It is useful for securely extending an on-premises network to VPC workloads over the internet.

AWS Direct Connect provides a dedicated private physical network connection between a customer network and AWS through a Direct Connect location. Private virtual interfaces can be used to access VPC resources through a virtual private gateway or Direct Connect gateway. This approach can provide more consistent network performance than internet-based connectivity and can support hybrid architectures. Site-to-Site VPN can also be used with Direct Connect to add encryption when required. Together, these services are core AWS offerings for connecting on-premises environments to VPCs. See the AWS Site-to-Site VPN documentation and the AWS Direct Connect User Guide.

QUESTION NO: 46

A company needs to plan, schedule, and run hundreds of thousands of computing jobs on AWS.

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

A. AWS Step Functions

B. AWS Service Catalog

C. Amazon Simple Queue Service (Amazon SQS)

D. AWS Batch

ANSWER: D

Explanation:

AWS Batch is designed to plan, schedule, and run batch computing workloads at any scale, including hundreds of thousands of jobs. A batch workload consists of discrete jobs that can run independently or as part of a dependency workflow, such as simulations, media transcoding, financial analysis, scientific processing, or large-scale data processing. With AWS Batch, the company defines jobs, job queues, and compute environments rather than having to build and operate its own scheduler and compute fleet.

AWS Batch evaluates queued jobs, their resource requirements, priorities, and dependencies, then provisions and scales the appropriate compute capacity automatically. It can run jobs on Amazon EC2, AWS Fargate, or Amazon EKS compute environments. This enables efficient execution of very large job volumes while reducing operational overhead: AWS Batch handles scheduling, queue management, retries, and integration with AWS compute resources. The service is specifically documented as supporting batch computing planning, scheduling, and execution across a range from a few jobs to hundreds of thousands of jobs. For more information, see the AWS Batch product page and the AWS Batch User Guide.

QUESTION NO: 47

A company needs to create and publish interactive business intelligence dashboards. The dashboards require insights that are powered by machine learning.

Which AWS service or tool will meet these requirements?

A. AWS Glue Studio

B. Amazon QuickSight

C. Amazon Redshift

D. Amazon Athena

ANSWER: B

Explanation:

Amazon QuickSight is the appropriate service because it is AWS's cloud-native, serverless business intelligence service for analyzing data and creating, publishing, and sharing interactive dashboards. It provides visual authoring capabilities that let business users build dashboards from connected data sources without managing BI infrastructure. QuickSight also includes ML Insights, which applies AWS machine learning capabilities to surface useful findings directly in analyses and dashboards. These capabilities include anomaly detection, forecasting, automated insights, and natural-language features, enabling users to identify trends and unusual changes without having to develop and operate their own machine learning models. A dashboard can be published for other users and can present interactive visualizations, filters, drill-downs, and embedded insights. This directly satisfies both parts of the requirement: interactive BI dashboard creation and machine-learning-powered insights. AWS describes QuickSight as a unified BI service that provides interactive dashboards and ML-powered insights at cloud scale. For details, see the Amazon QuickSight product page and the Amazon QuickSight ML Insights documentation.

QUESTION NO: 48

Which AWS service provides a scalable data warehouse solution?

A. Amazon S3

B. Amazon DynamoDB

C. Amazon Kinesis Data Streams

D. Amazon Redshift

ANSWER: D

Explanation:

Amazon Redshift is AWS's fully managed, scalable cloud data warehouse service. It is designed to store and analyze large volumes of structured and semi-structured data using SQL, supporting analytical workloads such as business intelligence reporting, dashboards, and complex aggregations. Redshift uses a columnar storage architecture and massively parallel processing to distribute query execution across compute resources, helping organizations obtain high performance as their data volumes and analytical needs grow. AWS manages much of the operational work involved in running the warehouse, including infrastructure provisioning, backups, software patching, and monitoring capabilities. Capacity can be adjusted by changing node configurations or using Amazon Redshift Serverless, which automatically provisions and scales resources to meet workload demands. Redshift also integrates with common data and analytics services, including Amazon S3, AWS Glue, Amazon QuickSight, and data ingestion tools, enabling organizations to build end-to-end analytics solutions in AWS. AWS explicitly describes Amazon Redshift as a fully managed, petabyte-scale cloud data warehouse service. See the Amazon Redshift product page and the Amazon Redshift Management Guide for service details and architecture information.

QUESTION NO: 49

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

Which benefit of the AWS Cloud does this scenario demonstrate?

A. Cost savings

B. Improved operational resilience

C. Increased business agility

D. Enhanced security

ANSWER: C

Explanation:

Increased business agility is demonstrated because the company can obtain the compute capacity it needs on demand and turn an idea into a new marketing campaign much faster. AWS enables teams to provision Amazon EC2 instances in minutes rather than waiting through the traditional procurement, installation, and configuration process for on-premises hardware. This speed lets the business experiment, respond to market opportunities, and adapt to changing customer or competitive needs with shorter delivery cycles.

AWS identifies increased agility as a core cloud benefit: cloud computing allows organizations to innovate more quickly by providing access to a broad range of technology resources when needed. In this scenario, reducing the time to launch campaigns from three weeks to three days is direct evidence that the organization can move faster operationally and commercially. The ability to scale compute resources up when demand requires it supports timely execution without requiring the company to predict and purchase all future capacity in advance. This creates flexibility for teams to act on new opportunities while maintaining the pace needed by the business. See the AWS overview of cloud computing benefits and the AWS Cloud Adoption Framework guidance on business outcomes and organizational agility.

QUESTION NO: 50

Which task does AWS perform automatically?

A. Encrypt data that is stored in Amazon DynamoDB.

B. Patch Amazon EC2 instances.

C. Encrypt user network traffic.

D. Create TLS certificates for users' websites.

ANSWER: A

Explanation:

AWS automatically encrypts data stored in Amazon DynamoDB at rest. DynamoDB provides encryption at rest for all tables by default, helping protect table data without requiring the customer to implement or operate encryption infrastructure. By default, DynamoDB uses an AWS owned key, which AWS creates, manages, and uses on the customer's behalf. Customers can alternatively select an AWS KMS key that they manage when they need greater control over key policies, access permissions, rotation, or auditing, but encryption at rest remains an inherent DynamoDB capability.

This automatic protection applies to the underlying storage used by DynamoDB, including table data and related artifacts such as local and global secondary indexes. It is an example of AWS handling security *of* the cloud for a managed service: AWS operates the storage systems and provides built-in encryption at rest, while customers remain responsible for configuring their data access controls and choosing key-management settings when required. AWS documents DynamoDB's default encryption behavior in the DynamoDB Developer Guide. The AWS shared responsibility model provides additional context on how responsibilities differ by service and deployment model: AWS Shared Responsibility Model.

QUESTION NO: 51

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

A. Security inside of code

B. Selection of CPU resources

C. Patching of operating system

D. Writing and updating of code

E. Security of underlying infrastructure

ANSWER: A D

Explanation:

With AWS Lambda, the company remains responsible for **Security inside of code** and **Writing and updating of code**. Lambda is a serverless service, meaning AWS operates the infrastructure that runs the functions. However, the shared responsibility model still assigns customers responsibility for what they deploy and configure in their workloads.

Security inside of code includes implementing secure application behavior, such as validating untrusted input, managing errors safely, protecting secrets, using encryption appropriately, and ensuring dependencies and function logic do not introduce vulnerabilities. The company must also configure the function and its permissions securely, including following least-privilege practices for the function's execution role.

Writing and updating of code is also a customer duty because the company creates the Lambda function's business logic, tests it, deploys revised versions, and maintains the application code and its dependencies over time. AWS provides the managed runtime environment and service platform, but it does not own or maintain the customer's function implementation. AWS documentation describes Lambda as running code without requiring customers to administer servers, while customers retain responsibility for their application code and configuration. See the AWS Lambda Developer Guide introduction and the AWS Shared Responsibility Model.

QUESTION NO: 52

A company needs to store data from a recommendation engine in a database.

A. Amazon RDS for PostgreSQL

B. Amazon DynamoDB

C. Amazon Neptune

D. Amazon Aurora

ANSWER: B

Explanation:

Amazon DynamoDB is the best fit because it is a fully managed, serverless NoSQL database service designed to store and retrieve application data with consistently low latency at virtually any scale. A recommendation engine commonly needs a highly scalable data store for items such as recommendation results, user-to-item mappings, preference attributes, or cached recommendation records. DynamoDB supports key-value and document data models, allowing an application to access these records efficiently by using a primary key and, where appropriate, secondary indexes.

It provides the least operational overhead because AWS manages the underlying infrastructure, including server provisioning, hardware maintenance, software patching, replication, durability, and availability. With on-demand capacity mode, the company can also avoid capacity planning and automatically accommodate changing request volumes while paying for requests actually used. DynamoDB includes features such as point-in-time recovery, backups, encryption, and global tables that can be enabled without administering database servers. These managed capabilities let the team focus on the recommendation application and its data model rather than database operations. See the Amazon DynamoDB product page and the DynamoDB Developer Guide.

QUESTION NO: 53

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

A. Pay for compute resources in advance.

B. Make data-driven decisions to determine cloud architectural design.

C. Emphasize manual processes to allow for changes.

D. Test systems at production scale.

E. Refine operational procedures infrequently.

ANSWER: B D

Explanation:

Make data-driven decisions to determine cloud architectural design. and **Test systems at production scale.** are AWS Cloud design principles described in the AWS Well-Architected Framework. Data-driven architecture means using measurable evidence—including workload metrics, operational data, business requirements, and experiment results—to guide design choices. AWS services make it practical to collect telemetry, monitor workload behavior, and adjust architecture as usage patterns and requirements evolve. This supports informed trade-offs involving performance, resilience, security, and cost rather than relying solely on assumptions.

Testing at production scale is also fundamental to cloud design because cloud resources can be provisioned on demand. Teams can create realistic environments and run experiments that reflect actual traffic volumes, failure conditions, and scaling behavior. Testing under these conditions provides confidence that an architecture will operate as intended when it serves real users. It also enables teams to validate changes, identify operational limits, and improve workload resilience before an issue affects customers. These practices align with AWS guidance to use data to drive architectural decisions and to test systems at production scale as part of continuous architectural improvement. See the AWS Well-Architected general design principles and the AWS Well-Architected Framework.

QUESTION NO: 54

Which AWS service can a company use to directly query and analyze AWS Cost and Usage Reports?

A. Amazon OpenSearch Service

B. Amazon Athena

C. Amazon Aurora

D. AWS Glue

ANSWER: B

Explanation:

Amazon Athena is the correct service because AWS Cost and Usage Reports are delivered to an Amazon S3 bucket, and Athena can run standard SQL queries directly against data stored in S3. This enables a company to explore detailed cost, usage, pricing, reservation, savings plan, and resource-level information without first loading the report into a database or provisioning query infrastructure. Athena is serverless, so there are no servers to manage; charges are generally based on the amount of data scanned by queries. AWS provides integration guidance and a CloudFormation template that can create the Athena table and related resources for Cost and Usage Report data, making it easier to begin analyzing the report. For efficient and current querying, AWS recommends configuring the Cost and Usage Report with Amazon Athena integration and report overwrite enabled. This use case is especially valuable for creating custom cost analyses, allocating costs, and investigating usage patterns beyond the standard views available in AWS billing tools. See the AWS Cost and Usage Report documentation for querying with Athena and the Amazon Athena User Guide.

QUESTION NO: 55

Elasticity in the AWS Cloud refers to which of the following? (Choose two.)

A. How quickly an Amazon EC2 instance can be restarted

B. The ability to rightsize resources as demand shifts

C. The maximum amount of RAM an Amazon EC2 instance can use

D. The pay-as-you-go billing model

E. How easily resources can be procured when they are needed

ANSWER: B E

Explanation:

Elasticity is the AWS Cloud capability to adjust the amount of computing capacity used in response to changing workload demand. “The ability to rightsize resources as demand shifts” describes this because an elastic architecture can add resources when demand grows and reduce them when demand declines, helping maintain appropriate performance without continuously operating unnecessary capacity. This adjustment can be automated through services such as Amazon EC2 Auto Scaling, which monitors conditions and changes capacity to match demand.

“How easily resources can be procured when they are needed” is also part of elasticity. AWS resources can be provisioned rapidly, enabling an organization to obtain capacity when a workload requires it and release that capacity when it no longer does. The AWS Well-Architected Framework describes elasticity as dynamically acquiring and releasing resources to match demand, while avoiding overprovisioning. This capability supports efficient operation of variable workloads and lets organizations align infrastructure capacity with actual usage over time. Learn more in the AWS Well-Architected Framework documentation and the Amazon EC2 Auto Scaling User Guide.

QUESTION NO: 56

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

A. AmazonKendra

B. Amazon Lex

C. AmazonTextract

D. AmazonPolly

ANSWER: B

Explanation:

Amazon Lex is the appropriate service because it is purpose-built for creating conversational interfaces, including chatbots that interact with users through text or voice. A developer can define the chatbot's intents, sample utterances, slots, prompts, and fulfillment behavior, allowing the application to recognize user requests and guide conversations toward the required outcome. Amazon Lex uses the same core technologies associated with Amazon Alexa, including automatic speech recognition for voice interactions and natural language understanding to interpret user intent.

For integration with an existing web application, Amazon Lex provides runtime APIs that let the application send user messages to a bot and receive the bot's responses. It also supports embedding a web user interface through integrations such as the Amazon Lex Web UI, and bot logic can invoke AWS Lambda when dynamic fulfillment or back-end processing is needed. This combination enables a company to add a managed, scalable conversational experience to its application without building natural-language processing capabilities from scratch. AWS documentation describes Amazon Lex as a service for building conversational interfaces using voice and text and documents its web deployment and API integration capabilities.

References: Amazon Lex product page and Amazon Lex Developer Guide.

QUESTION NO: 57

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

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

A. Elastic Load Balancing

B. Amazon EC2 Auto Scaling

C. AWS App2Container (A2C)

D. AWS Systems Manager

ANSWER: B

Explanation:

Amazon EC2 Auto Scaling is the appropriate service because it automatically adjusts the number of Amazon EC2 instances in response to application demand. The company creates an Auto Scaling group, which defines the desired, minimum, and maximum instance capacity for the application. Scaling policies can then add instances when metrics such as average CPU utilization, request count, or custom Amazon CloudWatch metrics indicate increased workload. When demand decreases, the group can remove unneeded instances while respecting the configured minimum capacity. This gives the application sufficient compute capacity during periods of higher demand without requiring administrators to manually launch every additional instance. Auto Scaling groups can also maintain the desired number of healthy instances by replacing instances that become unhealthy, supporting application availability. For predictable traffic patterns, scheduled scaling can adjust capacity at specific times, while dynamic scaling responds to real-time changes in monitored demand. This elasticity helps align EC2 capacity with workload requirements and can reduce costs by avoiding unnecessary continuously running instances. See the Amazon EC2 Auto Scaling User Guide and the AWS documentation on scaling based on demand.

QUESTION NO: 58

A company wants to protect a public web application from common web exploits and distributed denial of service (DDoS) attacks.

Which AWS services should the company use to meet these requirements? (Choose two.)

A. AWS WAF

B. AWS Shield

C. Amazon Inspector

D. AWS Artifact

E. AWS Secrets Manager

ANSWER: A B

Explanation:

AWS WAF helps protect web applications from common web exploits by allowing the company to create web access control rules. AWS WAF can inspect HTTP(S) requests and block, allow, or count requests based on conditions such as IP addresses, HTTP headers, URI paths, SQL injection patterns, and cross-site scripting patterns. AWS WAF can be associated with resources such as an Application Load Balancer, Amazon CloudFront distribution, or Amazon API Gateway REST API.

AWS Shield provides managed protection against DDoS attacks. AWS Shield Standard is automatically included at no additional cost and provides protection against common, frequently occurring network and transport layer DDoS attacks. AWS Shield Advanced provides additional detection, mitigation, and response capabilities for workloads that require enhanced protection. For more information, see What is AWS WAF? and AWS Shield documentation.

QUESTION NO: 59

A company is planning to migrate its application to the AWS Cloud.

Which AWS tool or set of resources should the company use to analyze and assess its readiness for migration?

A. AWS Cloud Adoption Framework (AWS CAF)

B. AWS Pricing Calculator

C. AWS Well-Architected Framework

D. AWS Budgets

ANSWER: A

Explanation:

AWS Cloud Adoption Framework (AWS CAF) is the appropriate set of resources because it helps organizations assess their cloud readiness and plan the organizational changes needed for a successful AWS migration. Rather than focusing only on technical infrastructure, AWS CAF provides guidance across six perspectives: business, people, governance, platform, security, and operations. Reviewing these perspectives helps a company identify capability gaps, define stakeholders and responsibilities, establish governance and security practices, and develop an actionable cloud transformation and migration plan.

AWS CAF also supports a Migration Readiness Assessment, which evaluates an organization's current maturity and identifies areas to address before or during migration. This makes it well suited to the question's need to analyze and assess readiness, including people, processes, technology, and business considerations. AWS describes CAF as guidance for developing efficient and effective plans for cloud adoption, and its migration guidance explains how readiness assessments help organizations prepare for large-scale migration. See the AWS Cloud Adoption Framework overview and the AWS Migration Readiness Assessment guidance.

QUESTION NO: 60

Which of the following is the customer's responsibility under the AWS shared responsibility model? (Choose two.)

A. Maintain the configuration of infrastructure devices.

B. Maintain patching and updates within the hardware infrastructure.

C. Maintain the configuration of guest operating systems and applications.

D. Manage decisions involving encryption options.

E. Maintain infrastructure hardware.

ANSWER: C D

Explanation:

Under the AWS shared responsibility model, customers are responsible for security *in* the cloud, including the configuration and management of the operating systems and software that they deploy. Therefore, **Maintain the configuration of guest operating systems and applications.** is a customer responsibility. For example, when using Amazon EC2, the customer administers the guest operating system, applies appropriate operating-system and application configuration, and manages the security of workloads running on the instance.

Manage decisions involving encryption options. is also a customer responsibility. AWS provides encryption capabilities and services, but customers decide whether encryption is needed for their data, select suitable encryption settings, manage access to encryption keys according to the service and configuration used, and apply encryption to meet their security and compliance requirements. The exact responsibility can vary with the AWS service and whether it is managed or customer-controlled, but deciding how to protect customer data remains the customer's responsibility. AWS documents this division as AWS securing the underlying cloud infrastructure while customers secure their data, identities, configurations, and workloads. See the AWS Shared Responsibility Model and the AWS encryption guidance.

QUESTION NO: 61

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

A. Chief financial officers (CFOs)

B. IT architects

C. Chief information officers (CIOs)

D. Chief data officers (CDOs)

E. Engineers

ANSWER: B E

Explanation:

The correct stakeholders are **IT architects** and **Engineers**. In the AWS Cloud Adoption Framework (AWS CAF), the platform perspective addresses the design, implementation, and ongoing evolution of an organization's cloud platform. Its goal is to provide a scalable, secure, and well-managed technology foundation on which workloads can run. IT architects are key stakeholders because they translate business and technical requirements into cloud architecture designs, including account structures, networking, compute, storage, and integration patterns. Engineers are also central stakeholders because they build, automate, deploy, operate, and optimize the platform capabilities that architects design. Their work commonly includes infrastructure as code, CI/CD enablement, platform automation, observability, and service integration. AWS identifies architects and engineers among the technology-focused stakeholders responsible for establishing and evolving the cloud platform. These roles collaborate to ensure that the platform can support application teams while meeting organizational requirements for reliability, scalability, and operational efficiency. For additional detail, see the AWS CAF Platform perspective introduction and the AWS Cloud Adoption Framework overview.

QUESTION NO: 62

What is the total amount of storage offered by Amazon S3?

A. WOMB

B. 5 GB

C. 5 TB

D. Unlimited

ANSWER: D

Explanation:

Unlimited is correct because Amazon S3 is designed to provide virtually unlimited storage capacity. Customers can create buckets and store any number of objects as their data requirements grow, without needing to provision a fixed total amount of storage in advance. This elasticity is a fundamental cloud characteristic: storage scales on demand, and customers are charged for the storage and related requests they actually use rather than purchasing a predefined storage pool.

Although S3 has limits that apply to an individual object, those limits do not define the total amount of data a customer can store. For example, a single object stored in S3 can be up to 5 TB in size, but an account can store an effectively unlimited number of objects. This distinction is important: 5 TB is an object-size limit, not the service's aggregate storage capacity. AWS documentation describes S3 as offering virtually unlimited storage and explains its object storage model, while the S3 user guide documents the maximum object size. See the Amazon S3 service page and the Amazon S3 User Guide.

QUESTION NO: 63

Which options are AWS Cloud Adoption Framework (AWS CAF. governance perspective capabilities? (Choose two.)

A. Identity and access management

B. Cloud financial management

C. Application portfolio management

D. Innovation management

E. Product management

ANSWER: B C

Explanation:

Cloud financial management and **Application portfolio management** are capabilities in the AWS Cloud Adoption Framework (AWS CAF) governance perspective. The governance perspective helps an organization govern its cloud environment in a way that aligns cloud investments, decisions, and workloads with business objectives while managing value, cost, risk, and compliance.

Cloud financial management establishes practices for planning, tracking, allocating, optimizing, and governing cloud spending. This capability supports informed investment decisions and helps ensure that cloud consumption delivers measurable business value. Application portfolio management provides a structured approach for assessing and managing the organization’s application estate. It helps teams understand application dependencies, business value, lifecycle status, and modernization priorities, which supports effective migration and ongoing governance decisions.

These capabilities enable leadership and stakeholders to prioritize cloud initiatives, manage financial accountability, and make portfolio-level decisions throughout a cloud transformation. AWS lists both capabilities under the governance perspective in its AWS CAF guidance. See the AWS CAF Governance Perspective and the AWS Cloud Adoption Framework overview.

QUESTION NO: 64

A retail company is migrating its IT infrastructure applications from on premises to the AWS Cloud. Which costs will the company eliminate with this migration? (Select TWO.)

A. Cost of data center operations

B. Cost of application licensing

C. Cost of marketing campaigns

D. Cost of physical server hardware

E. Cost of network management

ANSWER: A D

Explanation:

Moving infrastructure applications from an on-premises environment to AWS eliminates the company's responsibility for operating its own data center facilities. This includes expenses associated with running and maintaining the physical location that houses servers, such as facility operations, power, cooling, equipment space, and the operational effort required to support that environment. AWS operates and maintains the underlying global infrastructure, allowing customers to consume resources without owning or operating a traditional data center.

The migration also eliminates the cost of purchasing and maintaining physical server hardware for the workloads moved to AWS. Instead of making capital investments in servers and periodically replacing or upgrading them, the company uses AWS compute services on a consumption-based pricing model. AWS is responsible for the physical hosts, their lifecycle, and the infrastructure required to run them. This shift from customer-owned hardware to AWS-managed infrastructure is a core cloud economic benefit, reducing upfront capital expenditure and avoiding hardware refresh planning. AWS describes these benefits as trading fixed costs for variable costs and stopping the need to guess capacity in advance. See the AWS Well-Architected Framework and AWS Pricing.

QUESTION NO: 65

A company needs to determine whether its use of AWS services meets regulatory requirements. The company also needs to download AWS compliance reports and agreements.

Which AWS services or resources can the company use to meet these requirements? (Choose two.)

A. AWS Artifact

B. AWS Compliance Center

C. AWS Trusted Advisor

D. Amazon Macie

E. AWS Health Dashboard

ANSWER: A B

Explanation:

AWS Artifact provides on-demand access to AWS compliance reports, certifications, and agreements. The company can use AWS Artifact Reports to download documents such as AWS ISO, SOC, and PCI reports that can support its compliance assessment activities.

AWS Compliance Center provides information about AWS compliance programs, security and compliance resources, and AWS services that can help customers meet regulatory requirements. The company can use these materials to understand AWS compliance capabilities and the shared responsibility model. For more information, see AWS Artifact and the AWS Compliance Center.

QUESTION NO: 66

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

A. Patch AWS network devices.

B. Set user password rules.

C. Provide physical security for compute resources.

D. Configure security groups.

E. Patch the operating system of an Amazon EC2 instance.

ANSWER: A C

Explanation:

AWS is responsible for "security *of* the cloud," meaning it operates and protects the underlying global infrastructure that delivers AWS services. This responsibility includes the facilities, physical hardware, networking equipment, and the foundational software and virtualization components used to run the AWS Cloud. Therefore, **Patch AWS network devices.** is an AWS responsibility: AWS maintains the networking infrastructure that supports its services, including applying maintenance and security patches to AWS-managed network hardware and software.

Provide physical security for compute resources. is also AWS's responsibility. AWS secures its data centers through physical safeguards such as controlled facility access, monitoring, environmental protections, and resilient infrastructure operations. Customers do not manage or physically access the servers, networking devices, or facilities hosting AWS services. This division lets customers focus on securing their workloads and configurations while AWS protects the underlying cloud infrastructure. The exact customer responsibilities can vary by service type, but AWS always retains responsibility for the physical infrastructure and core networking that make up the cloud environment. See the AWS Shared Responsibility Model and the AWS Overview of Security of the Cloud.

QUESTION NO: 67

Which AWS services or features provide disaster recovery solutions for Amazon EC2 instances? (Select TWO.)

A. EC2 Reserved Instances

B. EC2 Amazon Machine Images (AMIs)

C. Amazon Elastic Block Store (Amazon EBS) snapshots

D. AWS Shield

E. Amazon GuardDuty

ANSWER: B C

Explanation:

EC2 Amazon Machine Images (AMIs) and Amazon Elastic Block Store (Amazon EBS) snapshots provide foundational recovery capabilities for Amazon EC2 workloads. An AMI captures the configuration needed to launch an instance, including its operating system, application server, and applications; it can also include block device mappings and, for Amazon EBS-backed instances, snapshots of attached EBS volumes. Maintaining current AMIs enables an organization to rapidly launch replacement instances when an Availability Zone, Region, or individual instance failure affects the original workload. AMIs can be copied to another AWS Region to support regional disaster recovery planning.

Amazon EBS snapshots create point-in-time backups of EBS volumes. They are incremental after the first snapshot and are stored durably by AWS. During recovery, a snapshot can be used to create a new EBS volume, which can then be attached to a recovered or newly launched EC2 instance. Snapshot copies can also be created in a separate Region, making the underlying data available for a cross-Region recovery strategy. Together, AMIs restore the compute environment while EBS

snapshots restore persistent block-storage data, supporting practical EC2 disaster recovery procedures. See the Amazon EC2 AMI documentation and the Amazon EBS snapshot documentation.

QUESTION NO: 68

A company is running an order processing system on Amazon EC2 instances. The company wants to migrate microservices-based application.

Which combination of AWS services can the application use to meet these requirements? (Select TWO.)

A. Amazon Simple Queue Service (Amazon SQS)

B. AWS Lambda

C. AWS Migration Hub

D. AWS AppSync

E. AWS Application Migration Service

ANSWER: A B

Explanation:

Amazon Simple Queue Service (Amazon SQS) and **AWS Lambda** provide a common serverless pattern for building or modernizing an order-processing application into independently deployable microservices. Amazon SQS provides a fully managed queue that decouples service components: an order-facing service can place work messages on a queue, while downstream processing components consume those messages asynchronously. This separation allows each component to scale according to its own workload and helps the overall system tolerate temporary spikes or a temporarily unavailable consumer. For order workloads that require message ordering and deduplication, Amazon SQS FIFO queues support ordered processing and exactly-once processing semantics.

AWS Lambda supplies the compute layer for event-driven microservices without requiring the company to provision, patch, or operate EC2 servers for each processing function. Lambda can be configured to poll an Amazon SQS queue and invoke a function with queued messages, allowing functions such as payment validation, inventory updates, and notification processing to be implemented and scaled independently. Together, SQS and Lambda support loose coupling, asynchronous processing, and managed operational scalability. See the Amazon SQS Developer Guide and Using Lambda with Amazon SQS.

QUESTION NO: 69

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

A. EC2 includes operating system patch management

B. EC2 integrates with Amazon VPC, AWS CloudTrail, and AWS Identity and Access Management (IAM).

C. EC2 has a 100% service level agreement (SLA).

D. EC2 has a flexible, pay-as-you-go pricing model.

E. EC2 has automatic storage cost optimization.

ANSWER: B D

Explanation:

"EC2 integrates with Amazon VPC, AWS CloudTrail, and AWS Identity and Access Management (IAM)." is an advantage because EC2 instances work with core AWS services that support network isolation, access control, and account activity auditing. Amazon VPC enables customers to define virtual networks for workloads, including IP address ranges, subnets,

route tables, and security controls. IAM enables customers to manage identities and permissions according to least-privilege practices, while CloudTrail records API activity for governance, compliance, and operational auditing. This integrated cloud service model helps organizations establish and operate controlled application environments without building all of those supporting capabilities from scratch in an on-premises data center.

"EC2 has a flexible, pay-as-you-go pricing model." is also an advantage because AWS offers multiple ways to pay for EC2 capacity based on workload needs. On-Demand Instances provide flexibility without long-term commitments, Savings Plans can reduce cost for predictable usage commitments, and Spot Instances can provide significant savings when workloads can tolerate interruptions. Customers can scale capacity as requirements change and pay for the resources selected rather than making large up-front hardware investments. AWS documents EC2 pricing options at Amazon EC2 pricing and describes EC2 networking and security integrations in the Amazon EC2 User Guide.

QUESTION NO: 70

Which of the following is a recommended design principle for AWS Cloud architecture?

A. Design tightly coupled components.

B. Build a single application component that can handle all the application functionality.

C. Make large changes on fewer iterations to reduce chances of failure.

D. Avoid monolithic architecture by segmenting workloads.

ANSWER: D

Explanation:

Avoid monolithic architecture by segmenting workloads is a recommended AWS Cloud architecture principle because it promotes loosely coupled, modular components that can be developed, deployed, scaled, and recovered independently. Segmenting a workload into appropriate components or services reduces the blast radius of a failure: an issue in one component is less likely to interrupt every capability of the application. It also enables teams to make smaller, reversible changes more frequently, supporting automation and continuous improvement. For example, separating a web tier, application services, and data services allows each tier to use AWS services and scaling policies suited to its own requirements. This design can improve resilience, operational agility, and the efficient use of resources while allowing the architecture to evolve as demand changes. AWS Well-Architected guidance identifies using loosely coupled dependencies as a reliability design principle and emphasizes designing workloads so that failures can be isolated and handled effectively. The AWS architecture guidance also describes decomposing workloads into components with clear responsibilities as an approach that supports scalable and resilient cloud systems. See the AWS Well-Architected Reliability Pillar design principles and the AWS Well-Architected Framework.

QUESTION NO: 71

A company is migrating to the AWS Cloud and plans to run experimental workloads for 3 to 6 months on AWS.

Which pricing model will meet these requirements?

A. Use Savings Plans for a 3-year term.

B. Use Dedicated Hosts.

C. Buy Reserved Instances.

D. Use On-Demand Instances.

ANSWER: D

Explanation:

Use On-Demand Instances is correct because the company needs capacity for a short, experimental period of only 3 to 6 months and may need flexibility as the workloads change. Amazon EC2 On-Demand pricing allows customers to pay for

compute capacity by the second or hour, depending on the operating system, with no upfront payment and no long-term commitment. This lets the company launch instances when experiments require them and stop or terminate them when testing is complete, avoiding an obligation beyond the temporary workload duration. This pricing approach is particularly appropriate during a cloud migration when usage patterns, instance requirements, and the eventual production architecture may not yet be known. The company can evaluate performance and cost during the experimental phase, then choose a commitment-based savings option later if it establishes stable, predictable usage. AWS describes On-Demand Instances as suitable for workloads with short-term, irregular, or unpredictable usage and for applications that cannot be interrupted. See the Amazon EC2 On-Demand Pricing page and the Amazon EC2 User Guide for On-Demand Instances.

QUESTION NO: 72

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

A. Observability

B. Data governance

C. Platform architecture

D. Portfolio management

E. Configuration management

ANSWER: A E

Explanation:

Observability and **Configuration management** are capabilities in the Operations perspective of the AWS Cloud Adoption Framework (AWS CAF). The Operations perspective helps an organization ensure that cloud workloads are delivered, operated, and improved at agreed service levels. It focuses on the day-to-day practices required to maintain workload health, reliability, and operational effectiveness.

Observability provides the operational insight needed to understand the state and behavior of applications and infrastructure. Teams use metrics, logs, traces, monitoring, and alerting to detect issues, investigate their causes, and make informed operational decisions. This visibility supports reliable service delivery and continual improvement.

Configuration management establishes and maintains accurate information about cloud resources, workload components, their relationships, and their configurations. It supports controlled change processes, helps identify configuration drift, and enables operators to assess the potential impact of changes or incidents. Maintaining this dependable configuration record is fundamental to operating workloads consistently at scale.

AWS identifies both capabilities within the Operations perspective in its CAF guidance. See the AWS CAF Operations perspective and the AWS Cloud Adoption Framework overview.

QUESTION NO: 73

A company stores data in Amazon S3.

According to the AWS shared responsibility model, which task is the company’s responsibility?

A. Maintaining the physical storage infrastructure in AWS data centers

B. Securing the network infrastructure between AWS Regions

C. Configuring bucket policies to control access to S3 objects

D. Patching the underlying S3 service platform

ANSWER: C

Explanation:

Configuring bucket policies to control access to S3 objects is the company's responsibility because Amazon S3 is a managed AWS service and the customer remains accountable for security *in* the cloud. This includes deciding which users, roles, accounts, or services are permitted to access buckets and objects, and then implementing those decisions through S3 bucket policies, IAM policies, access point policies, and S3 Block Public Access settings. A bucket policy is a resource-based policy attached to an S3 bucket that can grant or restrict permissions for actions such as reading, writing, and deleting objects. The company must configure these permissions according to its business, privacy, and compliance requirements and regularly review them as access needs change. AWS provides the S3 service and the security features used to manage access, but it cannot determine which principals should be authorized to access a customer's data. This division reflects the AWS shared responsibility model: AWS protects the infrastructure that runs cloud services, while customers configure identity, permissions, and protection of the data they store. For further details, see the AWS Shared Responsibility Model and the Amazon S3 bucket policy documentation.

QUESTION NO: 74

Which of the following is a managed AWS service that is used specifically for extract, transform, and load (ETL) data?

A. Amazon Athena

B. AWS Glue

C. Amazon S3

D. AWS Snowball Edge

ANSWER: B

Explanation:

AWS Glue is the correct answer because it is AWS's fully managed service designed specifically for extract, transform, and load (ETL) workloads. It helps organizations discover, prepare, clean, transform, and move data from one or more sources into destinations such as data lakes and data warehouses. AWS Glue includes a Data Catalog for storing metadata about datasets, crawlers that can infer schemas and populate that catalog, and serverless ETL jobs that can run transformation code without requiring customers to provision or manage servers. This makes it useful for building and operating repeatable data-integration pipelines at scale. AWS Glue can work with data in services including Amazon S3, Amazon Redshift, and relational databases, allowing data to be converted into analytics-ready formats and organized for querying. Its managed and serverless design means AWS handles much of the underlying infrastructure, scaling, and operational work while customers focus on defining the data transformations and workflow requirements. For more information, see the AWS Glue product page and the AWS Glue Developer Guide.

QUESTION NO: 75

A company is using Amazon EC2 instances.

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

A. Maintain the network infrastructure.

B. Patch the guest operating system.

C. Configure a security group on deployed EC2 instances.

D. Provide physical security for the underlying hardware of the EC2 instances.

E. Manage the underlying hypervisor.

ANSWER: B C

Explanation:

For Amazon EC2, the customer is responsible for security *in* the cloud, including the guest operating system and the configuration of network controls for their workloads. Therefore, **Patch the guest operating system.** is a customer

responsibility. The company chooses the operating system running on its EC2 instances and must perform activities such as applying security updates, managing installed software, and maintaining the instance-level configuration.

Configure a security group on deployed EC2 instances. is also a customer responsibility. Security groups act as virtual firewalls for EC2 instances, and customers define and maintain their inbound and outbound rules to permit only the required network traffic. Properly limiting these rules is an important part of protecting an EC2 workload.

AWS's shared responsibility model distinguishes these customer-managed workload controls from AWS management of the cloud's underlying facilities and infrastructure. AWS documentation specifically identifies customer responsibilities for EC2 that include guest OS patching and security-group configuration. See the AWS Shared Responsibility Model and the Amazon EC2 security documentation.

QUESTION NO: 76

Which of the following are the customer's responsibilities under the AWS shared responsibility model?(Select TWO.)

A. Maintain the configuration of infrastructure devices.

B. Maintain patching and updates within the hardware infrastructure.

C. Maintain the configuration of guest operating systems and applications.

D. Manage decisions involving encryption options.

E. Maintain infrastructure hardware.

ANSWER: C D

Explanation:

"Maintain the configuration of guest operating systems and applications" is a customer responsibility when using Amazon EC2 and other infrastructure services where the customer administers the guest operating system. The customer is responsible for securing and configuring the OS, installing and maintaining application software, and applying relevant patches and updates within that guest environment. This reflects the shared responsibility model's principle that customers are responsible for security *in* the cloud, with the exact scope varying by AWS service.

"Manage decisions involving encryption options" is also a customer responsibility. Customers classify their data and choose whether and how to use AWS encryption features, including selecting AWS managed keys or customer managed keys where applicable, configuring encryption for supported services, and controlling access to keys and encrypted data through IAM policies and permissions. AWS provides the cloud infrastructure and managed encryption capabilities, but the customer retains responsibility for appropriate data-protection configuration and access management. AWS describes these customer duties in its shared responsibility model documentation and explains how responsibility shifts according to the service category. See the AWS Shared Responsibility Model and the AWS Risk and Compliance whitepaper.

QUESTION NO: 77

Which benefits can customers gain by using AWS Marketplace? (Select TWO.)

A. Speed of business

B. Fewer legal objections

C. Ability to pay with credit cards

D. No requirement for product licenses for any products

E. Free use of all services for the first hour

ANSWER: A B

Explanation:

Speed of business and **Fewer legal objections** are benefits customers can gain through AWS Marketplace. AWS Marketplace provides a curated digital catalog where customers can discover, evaluate, procure, and deploy third-party software, data products, and professional services that run on AWS. This reduces the time traditionally spent locating suppliers, conducting separate purchasing processes, and coordinating deployment, helping teams obtain approved solutions more quickly.

AWS Marketplace also helps streamline commercial and legal processes. Customers can use standardized AWS Marketplace agreements and, where applicable, private offers and negotiated terms. Centralizing procurement through an AWS account and consolidating eligible charges onto the AWS bill can reduce the administrative effort associated with separate vendor relationships, invoices, and contract reviews. These capabilities support faster adoption of software while maintaining an organization's procurement controls and governance practices.

AWS describes Marketplace as a digital catalog designed to make it easier to find, buy, deploy, and manage third-party offerings: AWS Marketplace overview. Additional details on buyer procurement and subscription workflows are available in the AWS Marketplace Buyer Guide.

QUESTION NO: 78

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

A. How quickly an Amazon EC2 instance can be restarted

B. The ability to rightsized resources as demand shifts

C. The maximum amount of RAM an Amazon EC2 instance can use

D. The pay-as-you-go billing model

E. How easily resources can be procured when they are needed

ANSWER: B E

Explanation:

Elasticity is the capability to adjust cloud resource capacity in response to changing demand. "The ability to rightsized resources as demand shifts" describes this directly: workloads can scale capacity out when demand rises and scale it back when demand falls, so the resources provisioned remain aligned with the current workload requirement. In AWS, this adjustment can be automated through services and features such as AWS Auto Scaling, reducing the need for manual infrastructure provisioning while helping maintain workload performance and efficient utilization.

"How easily resources can be procured when they are needed" is also a core part of elasticity. AWS provides on-demand access to compute, storage, and other services, allowing organizations to obtain capacity quickly rather than purchasing, installing, and maintaining physical hardware in advance. Elasticity includes both acquiring additional resources when needed and releasing unneeded resources when demand declines. This ability supports variable and unpredictable workloads by making capacity responsive rather than fixed. AWS describes elasticity as adapting resources to demand, and the AWS Well-Architected Framework recommends automatically scaling infrastructure to meet changes in demand. See AWS: What is cloud elasticity? and AWS Well-Architected Reliability Pillar design principles.

QUESTION NO: 79

Which design principle is achieved by following the reliability pillar of the AWS Well-Architected Framework?

A. Vertical scaling

B. Manual failure recovery

C. Testing recovery procedures

D. Changing infrastructure manually

ANSWER: C

Explanation:

Testing recovery procedures is a design principle of the Reliability Pillar in the AWS Well-Architected Framework. Reliability is concerned with a workload's ability to perform its intended function consistently and correctly, including recovering from disruptions such as infrastructure failures, application defects, configuration issues, or increased demand. Recovery capabilities should not merely be documented; they must be exercised regularly to confirm that the workload can meet its defined recovery objectives.

By testing recovery procedures, teams validate that backups can be restored, failover mechanisms operate as intended, alarms and runbooks support effective response, and recovery processes work under realistic failure conditions. These tests also expose gaps before an actual incident occurs, allowing the organization to improve automation, operational readiness, and resilience. AWS specifically identifies testing recovery procedures as a Reliability Pillar design principle, alongside practices such as automatically recovering from failure, scaling horizontally, and managing change through automation. See the AWS Well-Architected Reliability Pillar and the Reliability design principles documentation.

QUESTION NO: 80

Which task is a responsibility of AWS, according to the AWS shared responsibility model?

A. Enable client-side encryption for objects that are stored in Amazon S3.

B. Configure IAM security policies to comply with the principle of least privilege.

C. Patch the guest operating system on an Amazon EC2 instance.

D. Apply updates to the Nitro Hypervisor.

ANSWER: D

Explanation:

Apply updates to the Nitro Hypervisor is correct because AWS is responsible for security *of* the cloud. This includes operating, maintaining, and protecting the foundational infrastructure that delivers AWS services, such as AWS data centers, physical networking, hardware, and the virtualization layer. The AWS Nitro System provides the underlying technology used by Amazon EC2, including the Nitro Hypervisor, which securely isolates workloads running on EC2 instances. Because this layer is operated by AWS rather than by an individual customer, AWS performs its maintenance and security updates, including hypervisor updates.

The shared responsibility model assigns responsibilities according to the service being used, but the infrastructure supporting EC2 remains AWS-managed. Customers retain responsibility for security *in* the cloud, including configuration and management of resources and workloads they deploy. Keeping the virtualization platform current is essential to maintaining isolation between instances and protecting the broader AWS environment, so AWS centrally manages those updates without customers needing access to, or administrative control over, the Nitro Hypervisor. See the AWS Shared Responsibility Model and the AWS Nitro System Hypervisor documentation.

QUESTION NO: 81

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: C D

Explanation:

Reserved Instances and Savings Plans are the appropriate pricing models when a company has steady production EC2 usage and is prepared to make a one- or three-year commitment to lower its costs. Reserved Instances provide a discount relative to On-Demand pricing in exchange for committing to a specific EC2 instance configuration and term. For eligible workloads, selecting an upfront payment option can further increase the discount. This model is well suited to stable, predictable instance requirements.

Savings Plans similarly provide discounted pricing in exchange for a commitment to a consistent amount of compute spend per hour over a one- or three-year term. They can be purchased with All Upfront, Partial Upfront, or No Upfront payment options. EC2 Instance Savings Plans apply to a chosen instance family in a specific AWS Region and generally provide the greatest Savings Plans discount, while Compute Savings Plans offer more flexibility as usage needs change. Both models reduce the cost of continued production usage while retaining the reliability needed for workloads that are expected to run continuously. AWS describes these commitment-based models in its Reserved Instances documentation and Savings Plans User Guide.

QUESTION NO: 82

A company needs to identify who accessed an AWS service and what action was performed for a given time period.

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

A. Amazon CloudWatch

B. AWS CloudTrail

C. AWS Security Hub

D. Amazon Inspector

ANSWER: B

Explanation:

AWS CloudTrail is the appropriate service because it records activity in an AWS account, including actions performed through the AWS Management Console, AWS CLI, SDKs, and other AWS services. Its event records provide key audit details such as the identity that made a request, the time of the request, the AWS service and API action involved, the resources affected, and the source IP address. A company can review CloudTrail event history for recent management events or configure a trail to deliver ongoing event logs to Amazon S3 for retention, analysis, and compliance requirements. CloudTrail Lake can also be used to retain and query event data for a chosen time range. These capabilities directly support determining who accessed an AWS service and precisely what action they performed during a specified period. AWS documents CloudTrail as a service for governance, compliance, operational auditing, and risk auditing because it provides a record of actions taken by users, roles, and AWS services. For more information, see the AWS CloudTrail User Guide and the documentation for viewing CloudTrail event history.

QUESTION NO: 83

A company has many developers who need programmatic access to AWS services. The company must provide the access in compliance with AWS security best practices.

Which solution will meet these requirements?

A. Require multi-factor authentication (MFA) for the AWS account root user and all IAM users. Rotate access keys.

B. Create a single shared IAM user account for all the developers.

C. Use the AWS account root user for programmatic access. Rotate access keys.

D. Create IAM permissions boundaries. Require multi-factor authentication (MFA) for the AWS account root user.

E. Use AWS IAM Identity Center or federation to provide each developer with an individual identity and temporary credentials through IAM roles.

ANSWER: E

Explanation:

Use AWS IAM Identity Center (or an external identity provider integrated through federation) to give each developer an individual workforce identity, and grant access through IAM roles that provide temporary credentials. This approach follows AWS guidance to use temporary credentials rather than long-term IAM user access keys whenever possible. Developers can obtain short-lived, automatically refreshed credentials for AWS CLI and SDK access, while permissions are assigned centrally through permission sets and roles according to least privilege. Individual identities also provide clear accountability in AWS logging and auditing tools because actions can be traced to a specific authenticated user or assumed role session. MFA can be enforced through the organization's identity provider or IAM Identity Center, strengthening access to the credential-issuing process without relying on shared, static secrets. Temporary credentials reduce operational overhead and security exposure because they expire automatically and do not require developers to create, store, distribute, or periodically rotate long-term access keys. AWS recommends federation and temporary credentials for human access, including programmatic access from developer tools. See the IAM security best practices and IAM Identity Center temporary credential guidance.

QUESTION NO: 84

A retail company is building a new mobile app. The company is evaluating whether to build the app at an on-premises data center or in the AWS Cloud.

responsibility model?

A. Amazon FSx for Windows File Server

B. Amazon Workspaces virtual Windows desktop

C. AWS Directory Service for Microsoft Active Directory

D. Amazon RDS for Microsoft SQL Server

E. AWS is responsible for security of the cloud, and the customer is responsible for security in the cloud.

ANSWER: E

Explanation:

The AWS shared responsibility model is the applicable concept when a company evaluates hosting an application in AWS rather than in an on-premises data center. Under this model, AWS is responsible for security *of* the cloud, including the physical facilities, hardware, networking infrastructure, and foundational services that run the AWS Cloud. The customer remains responsible for security *in* the cloud. For a mobile app, this includes configuring identity and access management, protecting application data, securely configuring AWS services, maintaining application code, and applying the appropriate controls for the services the company uses.

The exact customer responsibility varies by service type. For example, AWS operates the underlying infrastructure for a managed service, while the customer must still manage such items as data classification, user permissions, and configuration. This division allows the retail company to reduce the operational burden of owning and maintaining physical data-center infrastructure while retaining control of its application, data, and security settings. AWS documents this allocation as a fundamental part of using AWS services and designing workloads securely. See the AWS Shared Responsibility Model and the AWS Risk and Compliance whitepaper.

QUESTION NO: 85

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

A. AWS Storage Gateway

B. Virtual private gateway

C. NAT gateway

D. Customer gateway

E. Internet gateway

ANSWER: B D

Explanation:

An AWS Site-to-Site VPN connection includes a customer gateway and a virtual private gateway as the endpoints that establish encrypted IPsec tunnels between an on-premises network and AWS. The customer gateway represents the customer-side device and configuration: typically a physical or software VPN appliance in the on-premises network with a static public IP address. It supplies the external endpoint from which the encrypted VPN tunnels originate.

The virtual private gateway is the AWS-side VPN endpoint. It is attached to a VPC and provides the AWS endpoint for the Site-to-Site VPN connection. After the connection is configured, route tables and routing—using static routes or Border Gateway Protocol—direct traffic between the VPC and the customer network through the VPN tunnels. AWS provides two tunnels per Site-to-Site VPN connection to support high availability, and the customer gateway device should be configured to use both tunnels where possible.

For AWS documentation describing customer gateways, virtual private gateways, and Site-to-Site VPN connection architecture, see the AWS Site-to-Site VPN User Guide and the AWS guide to creating a Site-to-Site VPN connection.

QUESTION NO: 86

Which of the following are advantages of the AWS Cloud? (Select TWO.)

A. Trade variable expenses for capital expenses

B. High economies of scale

C. Launch globally in minutes

D. Focus on managing hardware infrastructure

E. Overprovision to ensure capacity

ANSWER: B C

Explanation:

High economies of scale and **Launch globally in minutes** are core advantages of AWS Cloud computing. AWS aggregates demand from a very large customer base and operates infrastructure at significant scale. This scale helps AWS achieve lower variable costs for computing, storage, and other infrastructure services than most individual organizations could achieve by building and operating comparable data centers themselves. Those efficiencies can be reflected in the pay-as-you-go cloud model.

AWS also enables organizations to launch workloads globally quickly. Its global infrastructure includes AWS Regions and Availability Zones in geographic locations around the world. Customers can provision services in a selected Region through the AWS Management Console, APIs, or infrastructure-as-code tools, without first building physical facilities or deploying their own hardware in that location. This supports faster expansion, lower time to market, and architectures designed for geographic reach and resilience. AWS identifies both economies of scale and the ability to go global in minutes among the principal benefits of cloud computing. See AWS: What is cloud computing? and AWS Global Infrastructure.

QUESTION NO: 87

A company stores data in Amazon S3.

According to the AWS shared responsibility model, which task is the companyâ€™s responsibility?

A. Maintaining the physical storage infrastructure in AWS data centers

B. Securing the network infrastructure between AWS Regions

C. Configuring bucket policies to control access to S3 objects

D. Patching the underlying S3 service platform

ANSWER: C

Explanation:

Configuring bucket policies to control access to S3 objects is the company's responsibility because Amazon S3 is a managed AWS service, but customers retain control of their data and the permissions that govern access to it. An S3 bucket policy is a resource-based AWS Identity and Access Management (IAM) policy attached to a bucket. It can define which principals are allowed or denied particular S3 actions, such as listing a bucket, reading objects, or uploading objects. The company must configure these permissions to meet its own business, security, and compliance requirements, including applying least-privilege access and preventing unintended public access. Under the AWS shared responsibility model, AWS is responsible for security *of* the cloud, including the infrastructure and managed service components used to deliver Amazon S3. The customer is responsible for security *in* the cloud, which includes its content, identity and access configuration, and service-level security settings. AWS provides the mechanisms for bucket policies and access controls, but it is the company that decides and configures who should be able to access its buckets and objects. See the AWS Shared Responsibility Model and the Amazon S3 bucket policy documentation.

QUESTION NO: 88

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

A. Amazon EC2

B. Amazon RDS

C. Amazon Lightsail

D. AWS Step Functions

E. AWS Cost Explorer

ANSWER: E

Explanation:

AWS Cost Explorer is correct because its rightsizing recommendations feature uses AWS Compute Optimizer to analyze Amazon EC2 workload metrics and identify instances that may be overprovisioned. Compute Optimizer evaluates utilization data, including CPU, memory where supported, network, and other workload characteristics over time, then recommends more suitable instance types or sizes. Cost Explorer presents these recommendations in a cost-management context, helping customers estimate potential savings from adjusting the compute capacity of eligible EC2 instances. This integration is useful because it connects technical resource-sizing analysis with financial visibility, allowing users to prioritize optimization opportunities according to expected cost impact. The recommendations are based on observed workload behavior rather than a fixed rule, so they can better align resources with actual demand. AWS notes that Cost Explorer's rightsizing recommendations are powered by Compute Optimizer and are intended to help identify EC2 resources that can be resized to reduce costs while continuing to meet workload needs. For details, see the AWS Cost Explorer rightsizing recommendations documentation and the AWS Compute Optimizer User Guide.

QUESTION NO: 89

Which AWS feature provides a no-cost platform for AWS users to join community groups, ask questions, find answers, and read community-

generated articles about best practices?

A. AWS Knowledge Center

B. AWS re:Post

C. AWS IQ

D. AWS Enterprise Support

ANSWER: B

Explanation:

AWS re:Post is the correct answer because it is AWS's free, community-driven knowledge service where AWS users can ask technical questions, receive answers, participate in topic-focused community groups, and access knowledge content. It is designed to help customers learn from AWS experts and other community members while finding practical guidance for using AWS services and following recommended practices. AWS re:Post also includes AWS-managed knowledge content and community articles, making it a central place to discover troubleshooting information, architectural guidance, and service-specific discussions.

The service is available at no additional cost to AWS users, aligning directly with the question's emphasis on a no-cost platform. Its community spaces support discussions around AWS topics and allow users to engage with peers who have similar technical interests or implementation needs. This makes AWS re:Post a community and knowledge-sharing destination rather than merely a support-plan benefit or a professional-services marketplace. AWS describes re:Post as a place to get answers to AWS questions from AWS experts and the broader AWS community. For details, see the AWS re:Post website and the AWS re:Post User Guide.

QUESTION NO: 90

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 the appropriate service because it delivers fully managed Windows file storage built on Windows Server. It provides native support for the Server Message Block (SMB) protocol, which Microsoft Windows workloads commonly use to access shared files. The service is designed for Windows-based applications and integrates with Microsoft Active Directory for identity, authentication, and access control. It also supports Windows file-system capabilities such as NTFS permissions, access control lists, Distributed File System (DFS) namespaces, and shadow copies, allowing organizations to migrate or run Microsoft-oriented file workloads while retaining familiar file-sharing behavior. AWS manages the underlying file servers, including hardware provisioning, software updates, backups, and availability features, reducing the operational work required to operate a traditional Windows file server. This directly satisfies the requirement for a managed file server with native Microsoft workload and file-system support, accessed through SMB. For more detail, see the Amazon FSx for Windows File Server User Guide and the Amazon FSx for Windows File Server product page.

QUESTION NO: 91

Which of the following is a way to use Amazon EC2 Auto Scaling groups to scale capacity in the AWS Cloud?

A. Scale the number of EC2 instances in or out automatically, based on demand.

B. Use serverless EC2 instances.

C. Scale the size of EC2 instances up or down automatically, based on demand.

D. Transfer unused CPU resources between EC2 instances.

ANSWER: A

Explanation:

Amazon EC2 Auto Scaling groups scale capacity by automatically adjusting the number of EC2 instances available to an application in response to demand or a defined scaling policy. An Auto Scaling group specifies minimum, maximum, and desired instance capacity. AWS can then launch additional instances when metrics such as average CPU utilization, request count, or a custom Amazon CloudWatch metric indicate that more capacity is needed. When demand declines, the group can terminate instances while respecting the configured minimum capacity. This is horizontal scaling: increasing or decreasing the number of compute instances rather than changing the hardware characteristics of an existing instance. Auto Scaling groups also help maintain availability by replacing unhealthy instances and maintaining the configured desired capacity. For example, an application can run with a small baseline fleet during normal traffic and automatically add instances during a peak period, helping align resource capacity with workload needs and cost. For details, see the Amazon EC2 Auto Scaling User Guide and the AWS documentation on scaling based on demand.

QUESTION NO: 92

A company wants to store compliance records for 7 years at the lowest possible Amazon S3 storage cost. The records will be retrieved only in rare audit situations, and retrieval can take up to 12 hours.

Which Amazon S3 storage class should the company use?

A. S3 Standard

B. S3 Standard-Infrequent Access (S3 Standard-IA)

C. S3 Glacier Instant Retrieval

D. S3 Glacier Deep Archive

ANSWER: D

Explanation:

Amazon S3 Glacier Deep Archive is designed for long-term data retention and digital preservation at the lowest Amazon S3 storage cost. It is appropriate for data that is accessed very rarely and can tolerate retrieval times of hours. Standard retrievals from S3 Glacier Deep Archive typically complete within 12 hours.

The company can use an S3 Lifecycle rule to transition compliance records to S3 Glacier Deep Archive automatically after the records are stored in another S3 storage class. This supports retention requirements while reducing the cost of data that does not require immediate access. For more information, see the Amazon S3 storage classes documentation.

QUESTION NO: 93

A company runs a web application on Amazon EC2 instances. The company wants to distribute incoming traffic across healthy instances and automatically adjust the number of instances as demand changes.

Which AWS services should the company use to meet these requirements? (Choose two.)

A. Elastic Load Balancing

B. Amazon EC2 Auto Scaling

C. Amazon Route 53 Resolver

D. AWS Direct Connect

E. Amazon EBS

ANSWER: A B

Explanation:

Elastic Load Balancing distributes incoming application traffic across multiple targets, such as Amazon EC2 instances, in one or more Availability Zones. A load balancer performs health checks and routes traffic to healthy registered targets, improving application availability and fault tolerance.

Amazon EC2 Auto Scaling helps maintain application availability and adjust Amazon EC2 capacity automatically. An Auto Scaling group can launch replacement instances when existing instances become unhealthy and can add or remove instances in response to demand or scheduled scaling policies. Together, Elastic Load Balancing and Amazon EC2 Auto Scaling distribute traffic and align compute capacity with workload demand. For more information, see the Elastic Load Balancing User Guide and the Amazon EC2 Auto Scaling User Guide.

QUESTION NO: 94

Which of the following are customer responsibilities under the AWS shared responsibility model? (Select TWO.)

A. Physical security of AWS facilities

B. Configuration of security groups

C. Encryption of customer data on AWS

ANSWER: B C

Explanation:

Under the AWS shared responsibility model, customers are responsible for security *in* the cloud. This includes configuring the AWS services and resources they use to protect their own workloads, applications, and data. **Configuration of security groups** is a customer responsibility because security groups act as virtual firewalls for AWS resources, and customers define the inbound and outbound rules that permit or restrict traffic. Customers must apply rules that meet their application's connectivity and security requirements.

Encryption of customer data on AWS is also a customer responsibility. Customers determine whether encryption is needed for their data, select and configure the relevant encryption features, and manage access to encryption keys according to the AWS service and key-management approach in use. AWS provides encryption capabilities and underlying cloud infrastructure, but the customer remains accountable for using those capabilities appropriately to protect its data. The precise responsibility can vary by AWS service, but data classification, access controls, and encryption configuration remain key customer obligations. See the AWS Shared Responsibility Model and the Amazon VPC security groups documentation.

QUESTION NO: 95

Which AWS service or tool can provide rightsizing recommendations for Amazon EC2 resources at no additional cost?

A. AWS Well-Architected Tool

B. Amazon CloudWatch

C. AWS Cost Explorer

D. Amazon S3 analytics

ANSWER: C

Explanation:

AWS Cost Explorer is correct because it includes rightsizing recommendations that help customers identify Amazon EC2 instances that may be oversized or underutilized. These recommendations use historical instance utilization data, including CPU utilization and memory utilization when Amazon CloudWatch Agent metrics are available, to suggest instance types or sizes that can better match workload requirements and reduce costs. Cost Explorer presents the estimated savings opportunity and can help customers evaluate the potential impact before making changes to their environments.

The AWS Cost Explorer service itself is available at no additional charge; AWS customers can use it to view and analyze cost and usage data, including its optimization recommendations. Rightsizing recommendations are particularly useful as part of an ongoing cost-management practice because workloads can change over time, making a previously appropriate instance size less efficient. AWS documents the available Amazon EC2 rightsizing recommendation types, utilization inputs, and savings estimates in its Cost Management guide. See AWS Cost Explorer rightsizing recommendations and AWS Cost Explorer.

QUESTION NO: 96

How does AWS Cloud computing help businesses reduce costs? (Choose two.)

A. AWS charges the same prices for services in every AWS Region.

B. AWS enables capacity to be adjusted on demand.

C. AWS offers discounts for Amazon EC2 instances that remain idle for more than 1 week.

D. AWS does not charge for data sent from the AWS Cloud to the internet.

E. AWS eliminates many of the costs of building and maintaining on-premises data centers.

ANSWER: B E

Explanation:

AWS enables capacity to be adjusted on demand because cloud resources can scale up when demand rises and scale down when demand falls. This elasticity helps a business avoid purchasing and operating infrastructure sized for peak capacity at all times. Instead, the business can provision resources when needed and pay only for the capacity it uses, reducing waste from underutilized servers and associated operational overhead.

AWS eliminates many of the costs of building and maintaining on-premises data centers because AWS operates the physical facilities and underlying infrastructure. Customers do not need to make large upfront capital investments in data center space, hardware, power, cooling, physical security, or ongoing equipment maintenance. The AWS consumption-based model also shifts many costs from fixed, long-term investments to variable expenses aligned with actual usage. These cloud economics allow organizations to benefit from AWS's scale while focusing spending and staff effort on business workloads rather than operating physical infrastructure. AWS describes these benefits as trading upfront expense for variable expense and avoiding the cost and complexity of running data centers. See the AWS overview of the advantages of cloud computing and the AWS pricing page.

QUESTION NO: 97

A company has a managed IAM policy that does not grant the necessary permissions for users to accomplish required tasks.

How can this be resolved?

A. Enable AWS Shield Advanced.

B. Create a custom IAM policy.

C. Use a third-party web application firewall (WAF. managed rule from the AWS Marketplace.

D. Use AWS Key Management Service (AWS KMS) to create a customer-managed key.

ANSWER: B

Explanation:

Creating a custom IAM policy resolves this issue because a customer-managed policy can define the precise AWS actions, resources, and conditions that users need for their required work. Unlike an AWS managed policy, whose permission set is maintained by AWS and cannot be edited by the customer, a customer-managed policy is created and maintained within the company's AWS account. The company can add the missing permissions while applying least-privilege principles, limiting access to only the relevant services and resources. The resulting policy can then be attached to the appropriate IAM users, groups, or roles. Customer-managed policies are reusable, so the company can consistently apply the same tailored permission set wherever it is needed, and can revise it later if job requirements change. AWS recommends using IAM policies to explicitly grant only the permissions necessary for a workload or user role. For more information, see the AWS IAM documentation on managed and inline policies and creating IAM policies.

QUESTION NO: 98

Which AWS service should a company use to organize characterize, and search large numbers of images?

A. Amazon Transcribe

B. Amazon Rekognition

C. Amazon Aurora

D. Amazon QuickSight

ANSWER: B

Explanation:

Amazon Rekognition is the appropriate service because it uses machine learning to analyze images and extract useful metadata at scale. Its image-analysis capabilities can identify labels and objects, detect faces and face attributes, recognize text in images, identify unsafe content, and detect image properties. A company can use these analysis results to characterize an image collection with consistent tags or attributes, store the resulting metadata, and then use that metadata to organize and search a large image library. For example, an application could submit images to the Rekognition DetectLabels API, receive labels such as objects, scenes, and concepts with confidence scores, and use those labels to support cataloging and search experiences. Rekognition is designed to provide pre-trained computer vision capabilities through APIs, so the company does not need to build, train, or operate its own image-recognition models for common analysis tasks. AWS describes Amazon Rekognition as a service that adds image and video analysis to applications and can identify objects, people, text, scenes, and activities. See the Amazon Rekognition product page and the Amazon Rekognition Developer Guide.

QUESTION NO: 99

A company wants to migrate its on-premises workloads to the AWS Cloud. The company wants to separate workloads for chargeback to different

departments.

Which AWS services or features will meet these requirements? (Choose two.)

A. Placement groups

B. Consolidated billing

C. Edge locations

D. AWS Config

E. Multiple AWS accounts

ANSWER: B E

Explanation:

Multiple AWS accounts provide a practical boundary for separating each department's workloads. Each account has its own AWS resources, permissions, quotas, and cost data, enabling the company to allocate workloads to departments and identify the costs incurred by each one. This multi-account approach is a core AWS recommendation for isolating business units, environments, and billing responsibilities while retaining centralized governance.

Consolidated billing, available through AWS Organizations, lets the company combine the charges from its member accounts into one payer account while preserving account-level cost visibility. The organization can use AWS Cost Explorer and cost and usage reports to review charges by linked account, which supports internal chargeback processes. Consolidated billing can also aggregate usage across eligible accounts for volume-based pricing benefits and provides a single bill for payment administration. Together, separate AWS accounts establish departmental workload and cost boundaries, and consolidated billing centralizes payment while making it possible to attribute costs to the appropriate department.

For details, see AWS Organizations and consolidated billing and Organizing Your AWS Environment.

QUESTION NO: 100

A company wants a report that lists the status of multi-factor authentication (MFA) devices that all users in the company ' s AWS account use.

Which AWS feature or service will meet this requirement?

A. AWS Cost and Usage Reports

B. IAM credential reports

C. Detailed Billing Reports

D. AWS Cost Explorer reports

ANSWER: B

Explanation:

IAM credential reports provide an account-level report containing credential-related information for every IAM user in an AWS account. The report includes fields that identify whether MFA is enabled for each IAM user, including the `mfa_active` column. This lets the company review MFA status across its users in a single downloadable report and use that information to support access-control reviews, security audits, and compliance checks. An administrator generates the credential report in IAM and can download it as a CSV file for filtering, analysis, or archival. IAM credential reports also include other useful credential status data, such as password usage, access key status, and access key rotation information, but the MFA status field directly addresses the stated requirement. AWS documents that credential reports list all IAM users in an account and the status of their credentials, including MFA device information. For details about generating and interpreting the report, see the AWS IAM User Guide: Getting credential reports for your AWS account. The available report fields, including `mfa_active`, are documented in the credential report fields reference.

QUESTION NO: 101

A company is migrating its public website to AWS. The company wants to host the domain name for the website on AWS.

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

A. AWS Lambda

B. Amazon Route 53

C. Amazon CloudFront

D. AWS Direct Connect

ANSWER: B

Explanation:

Amazon Route 53 is the appropriate service because it provides AWS domain registration and authoritative Domain Name System (DNS) hosting. The company can register a new domain name with Route 53 or transfer registration for an existing domain to Route 53 Registrar. It can then create a public hosted zone containing DNS records that direct users of the website's domain to the selected AWS endpoint, such as an Application Load Balancer, an Amazon S3 static website endpoint, an Amazon CloudFront distribution, or another public resource. Route 53 is designed as a highly available and scalable DNS service, so it resolves domain queries globally and supports common DNS record types. It also offers routing capabilities such as simple, weighted, latency-based, failover, geolocation, and multivalue answer routing, allowing a public website's DNS configuration to evolve with its availability and traffic-delivery requirements. Domain registration and DNS hosting are related but distinct functions; Route 53 supports both, which directly meets the requirement to host the website's domain name on AWS. For details, see the Amazon Route 53 Developer Guide and Registering domain names with Amazon Route 53.

QUESTION NO: 102

Which AWS service provides command line access to AWS tools and resources directly from a web browser?

A. AWS CloudHSM

B. AWS CloudShell

C. Amazon WorkSpaces

D. AWS Cloud Map

ANSWER: B

Explanation:

AWS CloudShell provides browser-based command-line access for managing AWS services and resources. It is available directly within the AWS Management Console, so a user can open a shell without installing command-line software locally, configuring a terminal, or managing underlying compute infrastructure. CloudShell includes commonly used tools such as the AWS Command Line Interface (AWS CLI), enabling users to run AWS commands against resources in their account from the browser. It also provides persistent storage for files in the user's home directory, which is useful for scripts, configuration files, and other command-line work. The environment is pre-authenticated with the AWS console credentials of the signed-in user, and its permissions are governed by that user's AWS Identity and Access Management (IAM) permissions. This makes CloudShell a convenient console-integrated environment for operational tasks, automation, and learning AWS commands while retaining normal AWS access controls. For more details, see the AWS CloudShell User Guide and the AWS CloudShell product page.

QUESTION NO: 103

Which of the following is a cloud benefit that AWS offers to its users?

A. The ability to configure AWS data center hypervisors

B. The ability to purchase hardware in advance of increased traffic

C. The ability to deploy to AWS on a global scale

D. Compliance audits for user IT environments

ANSWER: C

Explanation:

The ability to deploy to AWS on a global scale is a core cloud benefit. AWS operates a global infrastructure organized into Regions and Availability Zones, allowing customers to deploy workloads in geographic locations that are appropriate for their users, latency objectives, business-continuity needs, and data-residency requirements. Rather than building and operating physical data centers around the world, an organization can use AWS infrastructure in multiple locations and expand its application footprint as its needs evolve. This global reach supports delivering applications closer to end users and designing architectures that can remain available even if an individual facility experiences a disruption. AWS services can also help route users and distribute content globally, while customers retain control over where they deploy their resources. The AWS Cloud Practitioner curriculum identifies going global in minutes as one of the fundamental advantages of cloud computing: AWS enables customers to deploy applications in multiple Regions without the time and capital expense of establishing their own global infrastructure. AWS publishes the locations and structure of its worldwide infrastructure at AWS Global Infrastructure. For an overview of the cloud advantage of global reach and rapid deployment, see Six Advantages of Cloud Computing.

QUESTION NO: 104

Which design principles support the reliability pillar of the AWS Well-Architected Framework? (Choose two.)

A. Perform operations as code.

B. Enable traceability.

C. Automatically scale to meet demand.

D. Deploy resources globally to improve response time.

E. Automatically recover from failure.

ANSWER: C E

Explanation:

Automatically scale to meet demand supports reliability because workloads must continue to perform correctly as usage changes. AWS recommends scaling horizontally and using automation to add or remove resources in response to demand. This reduces dependence on manually estimating capacity, helps prevent resource exhaustion during peak traffic, and enables the workload to maintain its intended level of service as load varies. Services such as Amazon EC2 Auto Scaling and AWS Lambda can provide this adaptive capacity.

Automatically recover from failure is a core Reliability Pillar design principle. A reliable workload detects failures and uses automated mechanisms to restore service, such as health checks, replacement of unhealthy instances, failover, and self-healing infrastructure. Automating recovery reduces recovery time and limits the operational impact of component failures, enabling the workload to withstand disruptions without requiring immediate human intervention. These practices are foundational to designing systems that meet availability and resiliency objectives. See the AWS Well-Architected Reliability Pillar design principles and the Amazon EC2 Auto Scaling documentation.

QUESTION NO: 105

A company wants to migrate its database to a managed AWS service that is compatible with PostgreSQL.

Which AWS services will meet these requirements? (Choose two.)

A. Amazon Athena

B. Amazon RDS

C. Amazon EC2

D. Amazon DynamoDB

E. Amazon Aurora

ANSWER: B E

Explanation:

Amazon RDS is a managed relational database service that supports PostgreSQL as a database engine. It handles much of the undifferentiated operational work involved in running a relational database, including provisioning, backups, software patching, monitoring, and scaling capabilities. A company can therefore migrate a PostgreSQL workload to Amazon RDS for PostgreSQL while retaining PostgreSQL compatibility and using a fully managed AWS database offering.

Amazon Aurora also meets the requirement because Amazon Aurora PostgreSQL-Compatible Edition is designed to be compatible with PostgreSQL. Aurora is a fully managed relational database service that automates key administration tasks and provides Aurora's cloud-native architecture for availability, performance, and storage scalability. This makes it an appropriate managed destination for a database migration when PostgreSQL compatibility is required. The choice between Amazon RDS for PostgreSQL and Aurora PostgreSQL-Compatible can depend on the application's performance, availability, scaling, and cost needs. AWS documents supported engines for Amazon RDS at Amazon RDS User Guide and describes Aurora PostgreSQL compatibility at Amazon Aurora User Guide.

QUESTION NO: 106

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

A. AWS Basic Support

B. AWS Developer Support

C. AWS Business Support

D. AWS Enterprise Support

ANSWER: D

Explanation:

AWS Enterprise Support is correct because it includes access to a designated Technical Account Manager (TAM). The TAM serves as a customer's primary technical point of contact within AWS, working proactively to understand the customer's business goals, AWS environment, and operational requirements. This role provides strategic technical guidance, helps customers plan and optimize their AWS usage, coordinates access to specialized AWS resources when appropriate, and supports operational readiness through activities such as reviews and recommendations.

Enterprise Support is intended for organizations with business- or mission-critical workloads that need a highly personalized support relationship in addition to 24/7 access to cloud support engineers. AWS describes the TAM as a designated technical point of contact who provides proactive guidance and assists customers in achieving their cloud objectives. The AWS Support plan comparison also identifies the Technical Account Manager as an Enterprise Support feature. For current plan features and details, see the AWS Support plans page and the AWS Enterprise Support documentation.

QUESTION NO: 107

Which tasks are customer responsibilities, according to the AWS shared responsibility model? (Choose two.)

A. Configure the AWS provided security group firewall.

B. Classify company assets in the AWS Cloud.

C. Determine which Availability Zones to use for Amazon S3 buckets.

D. Patch or upgrade Amazon DynamoDB.

E. Select Amazon EC2 instances to run AWS Lambda on.

ANSWER: A B

Explanation:

“Configure the AWS provided security group firewall.” is a customer responsibility because security groups are virtual firewalls that control inbound and outbound network traffic for AWS resources. AWS supplies and operates the underlying infrastructure, but the customer defines the security-group rules appropriate for its workloads, such as permitted protocols, ports, and source or destination ranges. This is part of the customer’s responsibility for security *in* the cloud.

“Classify company assets in the AWS Cloud.” is also a customer responsibility. Customers retain ownership and control of their data and are responsible for understanding its sensitivity, applying suitable handling requirements, and selecting appropriate access controls and protections. Asset and data classification informs decisions such as encryption, retention, monitoring, and which identities should be authorized to access data. AWS’s shared responsibility model assigns customers responsibility for their data, identity and access management, and configuration choices, while AWS is responsible for protecting the infrastructure that runs AWS services. The exact division can vary by service, but these governance and configuration activities remain customer-managed responsibilities. See the AWS Shared Responsibility Model and the Amazon VPC security group documentation.

QUESTION NO: 108

A company needs a managed NFS file system that the company can use with its AWS compute....

Which AWS service or feature will meet these requirements?

A. Amazon Elastic Block Store (Amazon EBS)

B. AWS Storage Gateway Tape Gateway

C. Amazon S3 Glacier Flexible Retrieval

D. Amazon Elastic File System (Amazon EFS)

ANSWER: D

Explanation:

Amazon Elastic File System (Amazon EFS) is the appropriate service because it provides a fully managed, elastic file system that is accessed using the Network File System (NFS) protocol. AWS compute resources, including Amazon EC2 instances, can mount an EFS file system concurrently, allowing multiple instances to access the same shared files and directories. This makes it well suited to workloads such as web serving, content management, development environments, containerized applications, and shared application data that require standard file-system semantics rather than object or block storage.

EFS automatically grows and shrinks as files are added or removed, so the company does not need to provision file-system capacity in advance or manage underlying storage infrastructure. It also integrates with AWS networking and security controls: mount targets provide network access within selected Availability Zones, while security groups and IAM-related controls help govern access. EFS is designed for high availability and durability across multiple Availability Zones in an AWS Region, making it a managed shared storage choice for AWS compute workloads. See the Amazon EFS User Guide overview and the Amazon EFS product page for service details.

QUESTION NO: 109

A company needs to migrate its website from on premises to the AWS Cloud. The website must be hosted on hardware that is not shared with other companies. The company wants to use its existing per-socket, per-core software licenses.

Which Amazon EC2 instance purchasing option will meet these requirements?

A. Dedicated Instance

B. Reserved Instance

C. Dedicated Host

D. On-Demand Instance

ANSWER: C

Explanation:

Dedicated Host is the appropriate Amazon EC2 purchasing option because it provides an entire physical server dedicated to a single AWS customer. This satisfies the requirement that the website run on hardware that is not shared with other companies. Unlike a tenancy setting that only ensures single-customer hardware use, a Dedicated Host also exposes host-level attributes and gives the customer control over instance placement on that physical server.

Those host-level capabilities are important for bring-your-own-license scenarios involving software licensed per physical socket or per physical core. AWS identifies Dedicated Hosts as supporting eligible existing server-bound licenses, including per-socket and per-core licensing, provided that the customer's licensing agreement permits use in AWS. Customers can allocate a host and launch supported EC2 instances onto that specific host, helping them maintain the physical-server relationship needed to apply qualifying licenses and meet software-audit requirements.

This combination of dedicated physical infrastructure, visibility into sockets and cores, and instance-placement control directly addresses both the isolation requirement and the existing license model. AWS documentation describes Dedicated Hosts and their licensing use cases at Amazon EC2 Dedicated Hosts overview and provides licensing guidance at Amazon EC2 Dedicated Hosts.

QUESTION NO: 110

Which of the following are customer responsibilities under the AWS shared responsibility model? (Select TWO.)

A. Physical security of AWS facilities

B. Configuration of security groups

C. Encryption of customer data on AWS

D. Management of AWS Lambda infrastructure

E. Management of network throughput of each AWS Region

ANSWER: B C

Explanation:

Under the AWS shared responsibility model, customers are responsible for security *in* the cloud. This includes making security choices for the AWS resources and data they deploy. **Configuration of security groups** is a customer responsibility because security groups are virtual firewalls that control inbound and outbound traffic to customer workloads. Customers must define and maintain rules that permit only the network access required by their applications.

Encryption of customer data on AWS is also a customer responsibility. AWS provides encryption capabilities and services, but customers decide whether to encrypt their data, select appropriate encryption settings, and manage access to encryption keys according to the AWS service and key-management approach they use. These responsibilities reflect the customer's control over data, identity and access management, and workload configuration. AWS explains the division of duties in its Shared Responsibility Model. Security group configuration is further described in the Amazon VPC security group documentation.

QUESTION NO: 111

A company wants to provide one of its employees with access to Amazon RDS. The company also wants to limit the interaction to only the AWS CLI and AWS software development kits (SDKs).

Which combination of actions should the company take to meet these requirements while following the principles of least privilege? (Select TWO)

A. Create an IAM user and provide AWS Management Console access only.

B. Create an IAM user and provide programmatic access only.

C. Create an IAM role and provide AWS Management Console access only.

D. Create an IAM policy with administrator access and attach it to the IAM user.

E. Create an IAM policy with Amazon RDS access and attach it to the IAM user.

ANSWER: B E

Explanation:

Creating an IAM user and providing programmatic access only meets the requirement to use the AWS CLI and AWS SDKs without granting access to the AWS Management Console. Programmatic access for an IAM user is provided through access keys, which can be used to sign AWS CLI and SDK requests. The access keys should be securely managed and used only by the employee who requires the access.

Creating an IAM policy with Amazon RDS access and attaching it to the IAM user applies least privilege when the policy is scoped to only the RDS API actions, database resources, and conditions the employee needs. For example, the policy can allow read-only RDS actions or a limited set of management actions rather than broad permissions across AWS services. AWS IAM policies can control access through allowed actions, resources, and request conditions, enabling the company to grant the minimum necessary permissions for the employee's RDS responsibilities. For implementation guidance, see the AWS IAM access keys documentation and the Amazon RDS identity-based policy examples.

QUESTION NO: 112

A company needs a repository that stores source code. The company needs a way to update the running software when the code changes.

Which combination of AWS services will meet these requirements? (Select TWO.)

A. AWS CodeCommit

B. AWS CodeDeploy

C. Amazon DynamoDB

D. Amazon S3

E. Amazon Elastic Container Service (Amazon ECS)

ANSWER: A B

Explanation:

AWS CodeCommit and AWS CodeDeploy meet the two stated requirements. AWS CodeCommit provides managed, Git-based source repositories, enabling a company to store application source code and maintain source-control history in AWS. Development teams can push changes to a repository and use standard Git workflows to collaborate on code.

AWS CodeDeploy automates deployments of application revisions to compute environments, including Amazon EC2 instances, on-premises servers, AWS Lambda functions, and Amazon ECS services. A deployment workflow can use a code revision produced from repository changes, allowing the running application to be updated in a controlled and repeatable manner when new code is ready for release. CodeDeploy supports deployment configurations and lifecycle hooks, which help teams roll out updated software consistently.

Together, AWS CodeCommit serves as the managed source-code repository and AWS CodeDeploy provides the automated mechanism for releasing updated application code to the runtime environment. See the AWS CodeCommit User Guide and the AWS CodeDeploy User Guide.

QUESTION NO: 113

A company has a compute workload that is steady, predictable, and uninterruptible.

Which Amazon EC2 instance purchasing options meet these requirements MOST cost-effectively? (Choose two.)

A. On-Demand Instances

B. Reserved Instances

C. Spot Instances

D. Saving Plans

E. Dedicated Hosts

ANSWER: B D

Explanation:

Reserved Instances and Saving Plans are the most cost-effective choices for a workload with stable, predictable, continuous compute usage. Both pricing models provide substantial discounts compared with On-Demand pricing in exchange for a one- or three-year commitment. This matches a workload that is expected to run consistently rather than intermittently.

Reserved Instances provide a discounted billing rate for a specified EC2 instance configuration and term. They are well suited when the company can predict its instance family, Region, operating system, tenancy, and expected duration. Savings Plans provide similar commitment-based savings while offering additional flexibility. Compute Savings Plans can apply to eligible EC2 usage across instance families, sizes, Regions, and also to AWS Fargate and AWS Lambda usage, making them valuable if the workload may evolve while remaining consistently active.

Because the workload is uninterruptible, the company should use commitment-based discounts for dependable capacity rather than an interruptible pricing approach. AWS describes Reserved Instances and Savings Plans as ways to reduce costs for steady-state usage. See the Amazon EC2 Reserved Instances documentation and the AWS Savings Plans User Guide.

QUESTION NO: 114

Which AWS services or features enable users to connect on-premises networks to a VPC? (Choose two.)

A. AWS VPN

B. Elastic Load Balancing

C. AWS Direct Connect

D. VPC peering

E. Amazon CloudFront

ANSWER: A C

Explanation:

AWS VPN and AWS Direct Connect are services designed to establish connectivity between an on-premises network and Amazon Virtual Private Cloud (Amazon VPC). AWS VPN provides encrypted connectivity over the public internet. In particular, a Site-to-Site VPN connects an on-premises network, through a customer gateway device, to a virtual private gateway or transit gateway associated with AWS networking resources. This is commonly used when encrypted hybrid connectivity is needed without provisioning a dedicated physical connection.

AWS Direct Connect provides a dedicated private network connection from an on-premises environment to AWS at a Direct Connect location. It can provide more consistent network performance and can support private connectivity to VPC resources through a private virtual interface and a Direct Connect gateway or virtual private gateway. Organizations can use Direct Connect alone or combine it with AWS Site-to-Site VPN to add encryption or resiliency. Both services are core AWS

hybrid-networking options for extending on-premises networks into a VPC. See the AWS Site-to-Site VPN documentation and the AWS Direct Connect User Guide.

QUESTION NO: 115

An ecommerce company has been monitoring usage of its online store that is hosted on a fleet of Amazon EC2 instances. Surges in traffic occur every weekend day at the same time and last for approximately 4 hours.

Which AWS service should the company use to ensure that there are enough instances to meet the surges in demand?

A. AWS Lambda

B. Amazon EventBridge

C. Elastic Load Balancing (ELB)

D. Amazon EC2 Auto Scaling

ANSWER: D

Explanation:

Amazon EC2 Auto Scaling is the appropriate service because it automatically adjusts the number of EC2 instances in response to defined demand requirements. Since the traffic increase occurs predictably at the same time on weekend days and lasts about four hours, the company can configure an Auto Scaling scheduled scaling action. This increases the Auto Scaling group's desired capacity before the expected surge begins, giving instances time to launch and become healthy, and then reduces capacity after the surge has ended. The Auto Scaling group also maintains the configured desired number of healthy instances, replacing instances if they become unhealthy. This provides capacity for recurring demand while avoiding the cost of continuously operating peak-level capacity throughout the week. Scheduled scaling is particularly suitable when demand patterns are known in advance, as in this scenario. The company can additionally use metrics-based dynamic scaling if traffic occasionally differs from the expected pattern, but scheduled scaling directly addresses the predictable weekend schedule. AWS documents scheduled scaling as a way to set scaling actions for recurring schedules, and EC2 Auto Scaling as the service that automatically adds or removes EC2 instances to maintain application availability and optimize cost. See the EC2 Auto Scaling scheduled scaling documentation and the Amazon EC2 Auto Scaling overview.

QUESTION NO: 116

A company needs to run code in response to an event notification that occurs when objects are uploaded to an Amazon S3 bucket.

Which AWS service will integrate directly with the event notification?

A. AWS Lambda

B. Amazon EC2

C. Amazon Elastic Container Registry (Amazon ECR)

D. AWS Elastic Beanstalk

ANSWER: A

Explanation:

AWS Lambda is the correct service because Amazon S3 can be configured to invoke a Lambda function directly when a specified bucket event occurs, including the creation of an object through an upload. The company places its event-driven processing logic in a Lambda function and configures an S3 event notification for an object-created event. When S3 receives an uploaded object that matches the configured event type and any optional prefix or suffix filters, S3 invokes the function asynchronously and supplies event details, such as the bucket name and object key. The function can then run processing tasks such as validating the upload, extracting metadata, generating thumbnails, or initiating downstream workflows.

This integration is serverless: AWS manages the infrastructure needed to run the code, while the company configures the function, permissions, and S3 notification. The Lambda function's resource-based policy must permit Amazon S3 to invoke it, and the S3 bucket and function must be in the same AWS Region for direct S3-to-Lambda event notifications. For implementation details, see the Amazon S3 Event Notifications documentation and the AWS Lambda guide for processing Amazon S3 events.

QUESTION NO: 117

Which benefits does a company gain when the company moves from on-premises IT architecture to the AWS Cloud? (Select TWO.)

A. Reduced or eliminated tasks for hardware troubleshooting, capacity planning, and procurement

B. Elimination of the need for trained IT staff

C. Automatic security configuration of all applications that are migrated to the cloud

D. Elimination of the need for disaster recovery planning

E. Faster deployment of new features and applications

ANSWER: A E

Explanation:

Moving to AWS reduces or eliminates tasks for hardware troubleshooting, capacity planning, and procurement because AWS operates the physical data centers, servers, storage hardware, and networking infrastructure. Instead of purchasing equipment, forecasting data-center capacity far in advance, and replacing failed physical components, a company can provision resources on demand through AWS services. This shifts effort away from maintaining undifferentiated infrastructure and toward work that supports business applications and customer outcomes.

AWS also enables faster deployment of new features and applications. Teams can rapidly create and adjust computing, storage, database, and networking resources through the AWS Management Console, APIs, and infrastructure as code. On-demand provisioning and automation let organizations experiment, deploy, scale, and retire resources much more quickly than traditional procurement and data-center setup processes. These cloud characteristics support greater agility and faster time to market while allowing capacity to align with current demand.

AWS describes these cloud advantages as trading capital expense for variable expense, benefiting from massive economies of scale, stopping guessing capacity, increasing speed and agility, and avoiding spending on data-center operations. See the AWS overview of the six advantages of cloud computing and the AWS Well-Architected Framework.

QUESTION NO: 118

An ecommerce company plans to move its data center workload to the AWS Cloud to support highly dynamic usage patterns. Which benefits make the AWS Cloud cost-effective for the migration of this type of workload? (Select TWO.)

A. Reliability

B. Security

C. Elasticity

D. Pay-as-you-go resource pricing

E. High availability

ANSWER: C D

Explanation:

Elasticity makes AWS cost-effective for highly dynamic ecommerce usage because resources can automatically increase when customer demand rises and decrease when demand subsides. For example, compute capacity can scale out during a promotion or seasonal peak and scale back afterward. This avoids the traditional data center approach of purchasing and maintaining enough hardware for maximum anticipated demand, even when that capacity sits idle for much of the year. AWS describes elasticity as the ability to acquire resources as needed and release them when they are no longer needed.

Pay-as-you-go resource pricing complements elasticity by charging customers for the resources they consume rather than requiring large upfront investments in physical servers, facilities, and excess capacity. This consumption-based approach aligns infrastructure spending with actual business activity. Together, elastic provisioning and usage-based pricing enable the company to respond to changing workload levels while reducing the cost of underused infrastructure. For more information, see the AWS overview of cloud-computing advantages and the AWS Pricing page.

QUESTION NO: 119

What does the concept of agility mean in AWS Cloud computing? (Choose two.)

A. The speed at which AWS resources are implemented

B. The speed at which AWS creates new AWS Regions

C. The ability to experiment quickly

D. The elimination of wasted capacity

E. The low cost of entry into cloud computing

ANSWER: A C

Explanation:

In AWS Cloud computing, agility is the ability to obtain and use technology resources quickly so that an organization can respond rapidly to changing business needs. “The speed at which AWS resources are implemented” reflects this because users can provision compute, storage, databases, networking, and other services on demand rather than waiting for lengthy purchasing, installation, and data-center setup processes. This allows teams to move from an idea to a working environment in minutes or hours.

“The ability to experiment quickly” is also central to agility. AWS enables teams to create temporary environments, test an idea or workload, evaluate results, and then scale, modify, or remove the resources as needed. Because resources can be provisioned programmatically and paid for based on use, experimentation can occur with much shorter feedback cycles. This supports innovation and lets organizations adapt services and products faster. AWS identifies increased agility as a core cloud advantage, emphasizing rapid access to a broad range of technologies and faster experimentation. See the AWS Overview: Six Advantages of Cloud Computing and AWS: What is Cloud Computing?.

QUESTION NO: 120

A company is migrating its on-premises server to an Amazon EC2 instance. The server must stay active at all times for the next 12 months.

Which EC2 pricing option is the MOST cost-effective for the company's workload?

A. On-Demand

B. Dedicated Hosts

C. Spot Instances

D. Reserved Instances

ANSWER: D

Explanation:

Reserved Instances are the most cost-effective choice for this predictable, continuously running workload. The company knows that it needs EC2 capacity for a full 12-month period, so it can make a one-year commitment rather than paying the standard hourly On-Demand rate throughout the year. Amazon EC2 Reserved Instances offer a billing discount in exchange for committing to a specific instance configuration for either a one-year or three-year term. This model is designed for steady-state usage, where an instance is expected to run consistently over time.

For this scenario, a one-year Reserved Instance aligns directly with the stated duration. The organization can select payment terms that suit its budget, including All Upfront, Partial Upfront, or No Upfront; greater upfront payment generally provides a larger discount. Reserved Instances can be either Standard or Convertible, depending on how much flexibility the company requires to change instance attributes during the commitment. If capacity assurance in a particular Availability Zone is also required, a zonal Reserved Instance can provide a capacity reservation in addition to the pricing benefit.

See the Amazon EC2 Reserved Instances pricing page and the Amazon EC2 Reserved Instances documentation.

QUESTION NO: 121

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

A. Amazon S3

B. AWS Lambda

C. Amazon Elastic Block Store (Amazon EBS)

D. Amazon SageMaker

E. AWS Storage Gateway

ANSWER: A C

Explanation:

Amazon S3 is correct because it is AWS object storage and stores data as objects in buckets. Files of virtually any type can be uploaded to and retrieved from S3, and the service is designed for massive scale, high durability, and a wide range of storage classes. This makes it a core AWS service for storing and serving files such as documents, images, backups, logs, and media content.

Amazon Elastic Block Store (Amazon EBS) is also correct because it provides persistent block storage volumes for Amazon EC2 instances. An EBS volume can be attached to an instance, formatted with a file system, and then used by the operating system to create, read, modify, and retain files. EBS is especially appropriate when an application running on EC2 needs low-latency, persistent storage that behaves like a disk volume. AWS documents S3 as an object storage service for storing and retrieving data, and documents EBS as durable block-level storage for use with EC2 instances. See the Amazon S3 service page and the Amazon EBS User Guide.

QUESTION NO: 122

A company needs access to checks and recommendations that help the company follow AWS best practices for cost optimization, security, fault tolerance, performance, and service quotas.

Which combination of an AWS service and AWS Support plan on the AWS account will meet these requirements?

A. AWS Trusted Advisor with AWS Developer Support

B. AWS Health Dashboard with AWS Enterprise Support

C. AWS Trusted Advisor with AWS Business Support

D. AWS Health Dashboard with AWS Enterprise On-Ramp Support

ANSWER: C

Explanation:

AWS Trusted Advisor with AWS Business Support meets the requirement because AWS Trusted Advisor evaluates an AWS environment against AWS best practices and provides actionable checks and recommendations across the relevant categories: cost optimization, security, fault tolerance, performance, and service quotas. These checks help a company identify opportunities such as reducing unnecessary spending, addressing security-related configuration concerns, improving resilience, enhancing workload performance, and monitoring resource limits that could affect deployments or scaling. Although a limited set of Trusted Advisor checks is available to all AWS customers, AWS Business Support provides access to the full Trusted Advisor check set and its broader recommendation capabilities. This support-plan level is therefore sufficient for an account that needs comprehensive best-practice guidance in every category named in the question. AWS documents Trusted Advisor's categories and checks at AWS Trusted Advisor User Guide. The available support plans and their included Trusted Advisor access are described on the AWS Support Plans page. Using this combination gives the company a centralized AWS-native source of recommendations aligned with operational best practices.