TestkingIT

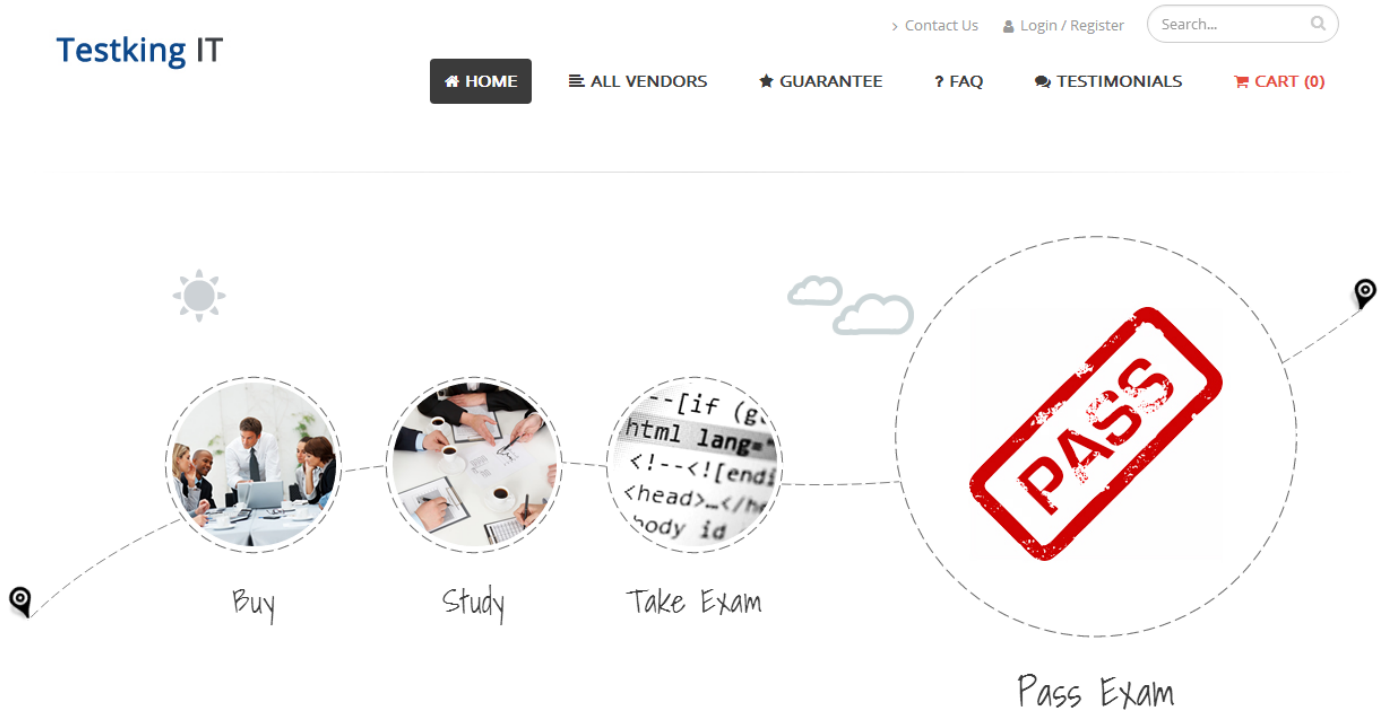

Try **Desktop Test Engine** before you buy

We're not the only ones **happy** about TestKingsIT Practice Material ...

48236+ customers in 100+ countries use TestKingsIT Test Engine. Meet our customers.

iMessenger VOREED GetCustom JET ORANGE iCompany Paradoxx

http://www.testkingit.com/

Latest practice material - Exam Cram - TestKingIT

Exam : **CLF-C01**

Title : Amazon AWS Certified Cloud Practitioner

Vendor : Amazon

Version : DEMO

NO.1 Which AWS service can a company use to perform complex analytical queries?

A. Amazon RDS

B. Amazon DynamoDB

C. Amazon Redshift

D. Amazon ElastiCache

Answer: C

Amazon Redshift is a fully managed, petabyte-scale data warehouse service in the cloud. You can start with just a few hundred gigabytes of data and scale to a petabyte or more. This enables you to use your data to acquire new insights for your business and customers. Amazon Redshift is designed for complex analytical queries that often involve aggregations and joins across very large tables. Amazon Redshift supports standard SQL and integrates with many existing business intelligence tools1.

NO.2 Which Amazon EC2 instance pricing model can provide discounts of up to 90%?

A. Reserved Instances

B. On-Demand

C. Dedicated Hosts

D. Spot Instances

Answer: D

Spot Instances are Amazon EC2 instances that are available at a discounted price compared to On-Demand pricing. Spot Instances use spare EC2 capacity that is not being used by other customers, and the price fluctuates based on supply and demand. Customers can request Spot Instances for their applications and specify the maximum price they are willing to pay per hour. If the Spot price is lower than the customer's bid, the Spot Instance is launched and the customer pays the current Spot price. However, if the Spot price rises above the customer's bid, the Spot Instance is terminated by AWS and the customer is charged for the partial hour of usage. Therefore, Spot Instances can provide discounts of up to 90% or more, but they are not suitable for applications that require continuous or predictable availability. Spot Instances are recommended for applications that are flexible, fault-tolerant, or have low priority, such as batch processing, data analysis, or testing and development.

NO.3 A company needs help managing multiple AWS linked accounts that are reported on a consolidated bill.

Which AWS Support plan includes an AWS concierge whom the company can ask for assistance?

A. AWS Developer Support

B. AWS Enterprise Support

C. AWS Business Support

D. AWS Basic Support

Answer: B

AWS Enterprise Support is the AWS Support plan that includes an AWS concierge whom the company can ask for assistance. According to the AWS Support Plans page, AWS Enterprise Support provides "a dedicated Technical Account Manager (TAM) who provides advocacy and guidance to help plan and build solutions using best practices, coordinate access to subject matter experts, and proactively keep your AWS environment operationally healthy."2 AWS Business Support, AWS Developer Support, and AWS Basic Support do not include a TAM or a concierge service.

NO.4 A company is reviewing the design of an application that will be migrated from on premises to a single Amazon EC2 instance.
What should the company do to make the application highly available?
A. Provision additional EC2 instances in other Availability Zones.
B. Configure an Application Load Balancer (ALB). Assign the EC2 instance as the ALB's target.
C. Use an Amazon Machine Image (AMI) to create the EC2 instance.
D. Provision the application by using an EC2 Spot Instance.
Answer: A
Provisioning additional EC2 instances in other Availability Zones is a way to make the application highly available, as it reduces the impact of failures and increases fault tolerance. Configuring an Application Load Balancer and assigning the EC2 instance as the ALB's target is a way to distribute traffic among multiple instances, but it does not make the application highly available if there is only one instance. Using an Amazon Machine Image to create the EC2 instance is a way to launch a virtual server with a preconfigured operating system and software, but it does not make the application highly available by itself. Provisioning the application by using an EC2 Spot Instance is a way to use spare EC2 capacity at up to 90% off the On-Demand price, but it does not make the application highly available, as Spot Instances can be interrupted by EC2 with a two-minute notification.

NO.5 A company wants to create a set of custom dashboards to collect metrics to monitor its applications.
Which AWS service will meet these requirements?
A. Amazon CloudWatch
B. AWS X-Ray
C. AWS Systems Manager
D. AWS CloudTrail
Answer: A
Amazon CloudWatch is a service that provides monitoring and observability for AWS resources and applications. Users can create custom dashboards to collect and visualize metrics, logs, alarms, and events from different sources5. AWS X-Ray is a service that provides distributed tracing and analysis for applications. AWS Systems Manager is a service that provides operational management for AWS resources and applications. AWS CloudTrail is a service that provides governance, compliance, and auditing for AWS account activity.

NO.6 Which option is a customer responsibility under the AWS shared responsibility model?
A. Maintenance of underlying hardware of Amazon EC2 instances
B. Application data security
C. Physical security of data centers
D. Maintenance of VPC components
Answer: B
The option that is a customer responsibility under the AWS shared responsibility model is B. Application data security.
According to the AWS shared responsibility model, AWS is responsible for the security of the cloud, while the customer is responsible for the security in the cloud. This means that AWS manages the security of the underlying infrastructure, such as the hardware, software, networking, and facilities

that run the AWS services, while the customer manages the security of their applications, data, and resources that they use on top of AWS12.
Application data security is one of the customer responsibilities under the AWS shared responsibility model.
This means that the customer is responsible for protecting their application data from unauthorized access, modification, deletion, or leakage. The customer can use various AWS services and features to help with application data security, such as encryption, key management, access control, logging, and auditing12.
Maintenance of underlying hardware of Amazon EC2 instances is not a customer responsibility under the AWS shared responsibility model. This is part of the AWS responsibility to secure the cloud. AWS manages the physical servers that host the Amazon EC2 instances and ensures that they are updated, patched, and replaced as needed13.
Physical security of data centers is not a customer responsibility under the AWS shared responsibility model.
This is also part of the AWS responsibility to secure the cloud. AWS operates and controls the facilities where the AWS services are hosted and ensures that they are protected from unauthorized access, environmental hazards, fire, and theft14.
Maintenance of VPC components is not a customer responsibility under the AWS shared responsibility model.
This is a shared responsibility between AWS and the customer. AWS provides the VPC service and ensures that it is secure and reliable, while the customer configures and manages their own VPCs and related components, such as subnets, route tables, security groups, network ACLs, gateways, and endpoints15.
References:
1: Shared Responsibility Model - Amazon Web Services (AWS) 2: AWS Cloud Computing - W3Schools 3:
[Amazon EC2 FAQs - Amazon Web Services] 4: [AWS Security - Amazon Web Services] 5: [Amazo n Virtual Private Cloud (VPC) - Amazon Web Services]

NO.7 A security engineer wants a single-tenant AWS solution to create, control, and manage their own cryptographic keys to meet regulatory compliance requirements for data security.
Which AWS service should the engineer use?

A. AWS Key Management Service (AWS KMS)

B. AWS Certificate Manager (ACM)

C. AWS CloudHSM

D. AWS Systems Manager

Answer: C

The correct answer is C because AWS CloudHSM is an AWS service that enables the security engineer to meet the requirements. AWS CloudHSM is a service that provides customers with dedicated hardware security modules (HSMs) to create, control, and manage their own cryptographic keys in the AWS Cloud. AWS CloudHSM allows customers to meet strict regulatory compliance requirements for data security, such as FIPS
140-2 Level 3, PCI-DSS, and HIPAA. The other options are incorrect because they are not AWS services that enable the security engineer to meet the requirements. AWS Key Management Service (AWS KMS) is a service that provides customers with a fully managed, scalable, and integrated key management system to create and control encryption keys for AWS services and applications. AWS

KMS does not provide customers with single-tenant or dedicated HSMs. AWS Certificate Manager (ACM) is a service that provides customers with a simple and secure way to provision, manage, and deploy public and private Secure Sockets Layer/Transport Layer Security (SSL/TLS) certificates for use with AWS services and internal connected resources. ACM does not provide customers with HSMs or cryptographic keys. AWS Systems Manager is a service that provides customers with a unified user interface to view operational data from multiple AWS services and automate operational tasks across their AWS resources. AWS Systems Manager does not provide customers with HSMs or cryptographic keys. Reference: AWS CloudHSM FAQs

NO.8 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 10

D. AWS Enterprise Support

Answer: B

AWS re:Post is a no-cost platform for AWS users to join community groups, ask questions, find answers, and read community-generated articles about best practices. AWS re:Post is a social media platform that connects AWS users with each other and with AWS experts. Users can create posts, comment on posts, follow topics, and join groups related to AWS services, solutions, and use cases. AWS re:Post also features live event feeds, community stories, and AWS Hero profiles. AWS re:Post is a great way to learn from the AWS community, share your knowledge, and get inspired. References:
AWS re:Post
Join the Conversation

NO.9 A company is planning to migrate to the AWS Cloud. The company is conducting organizational transformation and wants to become more responsive to customer inquiries and feedback.
Which tasks should the company perform to meet these requirements, according to the AWS Cloud Adoption Framework (AWS CAF)? (Select TWO.)

A. Realign teams to focus on products and value streams.

B. Create new value propositions with new products and services.

C. Use agile methods to rapidly iterate and evolve.

D. Use a new data and analytics platform to create actionable insights.

E. Migrate and modernize legacy infrastructure.

Answer: A,C

Realigning teams to focus on products and value streams, and using agile methods to rapidly iterate and evolve are tasks that the company should perform to meet the requirements of becoming more responsive to customer inquiries and feedback, according to the AWS Cloud Adoption Framework (AWS CAF). AWS CAF organizes guidance into six areas of focus, called perspectives: business, people, governance, platform, security, and operations. Each perspective is divided into capabilities, which describe the skills and processes to execute the transition effectively. The people perspective helps you prepare your organization for cloud adoption, and includes capabilities such as organizational change management, staff skills and readiness, and organizational alignment. The business perspective helps you align IT strategy with business strategy, and includes capabilities such as business case development, value proposition, and product ownership. Creating new value

propositions with new products and services is a task that belongs to the business perspective, but it is not directly related to the requirement of becoming more responsive to customer inquiries and feedback.
Using a new data and analytics platform to create actionable insights is a task that belongs to the platform perspective, which helps you design, implement, and optimize the architecture of the AWS environment.
However, it is also not directly related to the requirement of becoming more responsive to customer inquiries and feedback. Migrating and modernizing legacy infrastructure is a task that belongs to the operations perspective, which helps you enable, run, use, operate, and recover IT workloads to the level agreed upon with your business stakeholders. However, it is also not directly related to the requirement of becoming more responsive to customer inquiries and feedback.

NO.10 Which AWS service or feature can be used to estimate costs before deployment?
A. AWS Free Tier
B. AWS Pricing Calculator
C. AWS Billing and Cost Management
D. AWS Cost and Usage Report
Answer: B
AWS Pricing Calculator can be used to estimate costs before deployment. AWS Pricing Calculator is a tool that helps the user to compare the cost of AWS services for different use cases and configurations. The user can create estimates for various AWS services, such as Amazon EC2, Amazon S3, Amazon RDS, and more.
The user can also adjust the parameters, such as region, instance type, storage size, and duration, to see how they affect the cost. AWS Pricing Calculator provides a detailed breakdown of the estimated cost, as well as a summary of the key drivers of the cost.

NO.11 A company wants to ensure that two Amazon EC2 instances are in separate data centers with minimal communication latency between the data centers.
How can the company meet this requirement?
A. Place the EC2 instances in two separate AWS Regions connected with a VPC peering connection.
B. Place the EC2 instances in two separate Availability Zones within the same AWS Region.
C. Place one EC2 instance on premises and the other in an AWS Region. Then connect them by using an AWS VPN connection.
D. Place both EC2 instances in a placement group for dedicated bandwidth.
Answer: B
The correct answer is B because placing the EC2 instances in two separate Availability Zones within the same AWS Region is the best way to meet the requirement. Availability Zones are isolated locations within an AWS Region that have independent power, cooling, and networking. Users can launch their resources, such as Amazon EC2 instances, in multiple Availability Zones to increase the fault tolerance and resilience of their applications. Availability Zones within the same AWS Region are connected with low-latency, high-throughput, and highly redundant networking. The other options are incorrect because they are not the best ways to meet the requirement. Placing the EC2 instances in two separate AWS Regions connected with a VPC peering connection is not the best way to meet the requirement because AWS Regions are geographically dispersed and may have higher communication latency between them than Availability Zones within the same AWS Region. VPC

peering connection is a networking connection between two VPCs that enables users to route traffic between them using private IP addresses. Placing one EC2 instance on premises and the other in an AWS Region, and then connecting them by using an AWS VPN connection is not the best way to meet the requirement because on-premises and AWS Region are geographically dispersed and may have higher communication latency between them than Availability Zones within the same AWS Region. AWS VPN connection is a secure and encrypted connection between a user's network and their VPC. Placing both EC2 instances in a placement group for dedicated bandwidth is not the best way to meet the requirement because a placement group is a logical grouping of instances within a single Availability Zone that enables users to launch instances with specific performance characteristics. A placement group does not ensure that the instances are in separate data centers, and it does not provide low-latency communication between instances in different Availability Zones. Reference: [Regions, Availability Zones, and Local Zones], [VPC Peering],
[AWS VPN], [Placement Groups]

NO.12 What does the Amazon S3 Intelligent-Tiering storage class offer?

A. Payment flexibility by reserving storage capacity

B. Long-term retention of data by copying the data to an encrypted Amazon Elastic Block Store (Amazon EBS) volume

C. Automatic cost savings by moving objects between tiers based on access pattern changes

D. Secure, durable, and lowest cost storage for data archival

Answer: C

The Amazon S3 Intelligent-Tiering storage class offers automatic cost savings by moving objects between tiers based on access pattern changes. This storage class is designed for data with unknown or changing access patterns. It has two access tiers: frequent access and infrequent access. Objects are stored in the frequent access tier by default, and are moved to the infrequent access tier after 30 consecutive days of no access. If an object in the infrequent access tier is accessed, it is moved back to the frequent access tier. There are no retrieval fees in S3 Intelligent-Tiering, and no additional tiering fees when objects are moved between access tiers within the S3 Intelligent-Tiering storage class1.

NO.13 A company needs to use standard SQL to query and combine exabytes of structured and semi-structured data across a data warehouse, operational database, and data lake.
Which AWS service meets these requirements?

A. Amazon DynamoDB

B. Amazon Aurora

C. Amazon Athena

D. Amazon Redshift

Answer: D

Amazon Redshift is the service that meets the requirements of using standard SQL to query and combine exabytes of structured and semi-structured data across a data warehouse, operational database, and data lake.
Amazon Redshift is a fully managed, petabyte-scale data warehouse service that allows you to run complex analytic queries using standard SQL and your existing business intelligence tools. Amazon Redshift also supports Redshift Spectrum, a feature that allows you to directly query and join data stored in Amazon S3 using the same SQL syntax. Amazon Redshift can scale up or down to handle any

volume of data and deliver fast query performance5

NO.14 A company wants to set AWS spending targets and track costs against those targets.
Which AWS tool or feature should the company use to meet these requirements?
A. AWS Cost Explorer
B. AWS Budgets
C. AWS Cost and Usage Report
D. Savings Plans
Answer: B
AWS Budgets is a tool that allows users to set AWS spending targets and track costs against those targets.
Users can create budgets for various dimensions, such as service, linked account, tag, and more.
Users can also receive alerts when the actual or forecasted costs exceed or are projected to exceed the budgeted amount.
AWS Cost Explorer, AWS Cost and Usage Report, and Savings Plans are other AWS tools or features that can help users manage and optimize their AWS costs, but they do not enable users to set and track spending targets
.

NO.15 A company has a compute workload that is steady, predictable, and uninterruptible.
Which Amazon EC2 instance purchasing options meet these requirements MOST cost-effectively? (Select TWO.)
A. On-Demand Instances
B. Reserved Instances
C. Spot Instances
D. Saving Plans
E. Dedicated Hosts
Answer: B,D
Reserved Instances and Savings Plans are the most cost-effective purchasing options for a compute workload that is steady, predictable, and uninterruptible. Reserved Instances provide a significant discount compared to On-Demand Instances, and Savings Plans offer flexible and consistent savings on EC2 usage. Both options require a commitment to a consistent amount of usage, in USD per hour, for a term of 1 or 3 years.
On-Demand Instances are suitable for short-term, irregular, or unpredictable workloads, but they are more expensive than Reserved Instances or Savings Plans. Spot Instances are the cheapest option, but they are not suitable for uninterruptible workloads, as they can be reclaimed by AWS at any time.
Dedicated Hosts and Dedicated Instances are designed for compliance and licensing requirements, not for cost optimization. They are more expensive than the other options, as they run on single-tenant hardware. References: Instance purchasing options, Amazon EC2 Pricing, 4 Ways to Purchase Amazon EC2 Instances

NO.16 A company is using AWS Lambda functions to build an application.
Which tasks are the company's responsibility, according to the AWS shared responsibility model? (Select TWO.)
A. Patch the servers where the Lambda functions are deployed.

B. Establish the IAM permissions that define who can run the Lambda functions.
C. Write the code for the Lambda functions to define the application logic.
D. Deploy Amazon EC2 instances to support the Lambda functions.
E. Scale out the Lambda functions when the load increases.
Answer: B,C
According to the AWS shared responsibility model, AWS is responsible for the security of the cloud, while the user is responsible for the security in the cloud. This means that AWS manages the security and maintenance of the underlying infrastructure, such as the servers, networks, and operating systems, while the user manages the security and configuration of the resources and applications that run on AWS. For AWS Lambda functions, the tasks that are the user's responsibility are:
Establish the IAM permissions that define who can run the Lambda functions. IAM is a service that enables users to manage access and permissions for AWS resources and users. Users can create IAM policies, roles, and users to grant or deny permissions to run Lambda functions, invoke other AWS services, or access AWS resources from Lambda functions. [AWS Lambda Permissions] AWS Certified Cloud Practitioner - aws.amazon.com Write the code for the Lambda functions to define the application logic. Lambda functions are units of code that can be written in any supported programming language, such as Python, Node.js, Java, or Go.
Users can write the code for the Lambda functions using the AWS Management Console, the AWS Command Line Interface (AWS CLI), the AWS SDKs, or any code editor of their choice. Users can also use AWS Lambda Layers to share and manage common code and dependencies across multiple functions. [AWS Lambda Overview] AWS Certified Cloud Practitioner - aws.amazon.com

NO.17 What is a characteristic of Convertible Reserved Instances (RIs)?
A. Users can exchange Convertible RIs for other Convertible RIs from a different instance family.
B. Users can exchange Convertible RIs for other Convertible RIs in different AWS Regions.
C. Users can sell and buy Convertible RIs on the AWS Marketplace.
D. Users can shorten the term of their Convertible RIs by merging them with other Convertible RIs.
Answer: A
Convertible Reserved Instances (RIs) are a type of Reserved Instance that allow you to change the attributes of the RI as long as the exchange results in the creation of Reserved Instances of equal or greater value. You can exchange Convertible RIs for other Convertible RIs from a different instance family, size, platform, tenancy, or scope (Region or Availability Zone)3.

NO.18 Which AWS service or feature can the company use to limit the access to AWS services for member accounts?
A. AWS Identity and Access Management (IAM)
B. Service control policies (SCPs)
C. Organizational units (OUs)
D. Access control lists (ACLs)
Answer: B
Service control policies (SCPs) are a type of organization policy that you can use to manage permissions in your organization. SCPs offer central control over the maximum available permissions for all accounts in your organization, allowing you to ensure your accounts stay within your organization's access control guidelines2. SCPs are available only in an organization that has all features enabled2.

NO.19 Which pillar of the AWS Well-Architected Framework includes a design principle about measuring the overall efficiency of workloads in terms of business value?

A. Operational excellence

B. Security

C. Reliability

D. Cost optimization

Answer: A

The operational excellence pillar of the AWS Well-Architected Framework includes a design principle about measuring the overall efficiency of workloads in terms of business value. This principle states that you should monitor and measure key performance indicators (KPIs) and set targets and thresholds that align with your business goals. You should also use feedback loops to continuously improve your processes and procedures1.

NO.20 A company has designed its AWS Cloud infrastructure to run its workloads effectively. The company also has protocols in place to continuously improve supporting processes.
Which pillar of the AWS Well-Architected Framework does this scenario represent?

A. Security

B. Performance efficiency

C. Cost optimization

D. Operational excellence

Answer: D

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

NO.21 What is a benefit of moving to the AWS Cloud in terms of improving time to market?

A. Decreased deployment speed

B. Increased application security

C. Increased business agility

D. Increased backup capabilities

Answer: C

Increased business agility is a benefit of moving to the AWS Cloud in terms of improving time to market.
Business agility refers to the ability of a company to adapt to changing customer needs, market conditions, and competitive pressures. Moving to the AWS Cloud enables business agility by providing faster access to resources, lower upfront costs, and greater scalability and flexibility. By using the AWS Cloud, companies can launch new products and services, experiment with new ideas, and respond to customer feedback more quickly and efficiently. For more information, see [Benefits of Cloud Computing] and [Business Agility].

NO.22 Which AWS service or feature offers security for a VPC by acting as a firewall to control traffic in and out of subnets?

A. AWS Security Hub

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

NO.23 A company is running applications on Amazon EC2 instances in the same AWS account for several different projects. The company wants to track the infrastructure costs for each of the projects separately. The company must conduct this tracking with the least possible impact to the existing infrastructure and with no additional cost.
What should the company do to meet these requirements?
A. Use a different EC2 instance type for each project.
B. Publish project-specific custom Amazon CloudWatch metrics for each application.
C. Deploy EC2 instances for each project in a separate AWS account.
D. Use cost allocation tags with values that are specific to each project.
Answer: D
The correct answer is D because cost allocation tags are a way to track the infrastructure costs for each of the projects separately. Cost allocation tags are key-value pairs that can be attached to AWS resources, such as EC2 instances, and used to categorize and group them for billing purposes. The other options are incorrect because they do not meet the requirements of the question. Use a different EC2 instance type for each project does not help to track the costs for each project, and may impact the performance and compatibility of the applications. Publish project-specific custom Amazon CloudWatch metrics for each application does not help to track the costs for each project, and may incur additional charges for using CloudWatch. Deploy EC2 instances for each project in a separate AWS account does help to track the costs for each project, but it impacts the existing infrastructure and incurs additional charges for using multiple accounts. Reference: Using Cost Allocation Tags

NO.24 According to the AWS shared responsibility model, which of the following are AWS responsibilities? (Select TWO.)
A. Network infrastructure and virtualization of infrastructure
B. Security of application data
C. Guest operating systems
D. Physical security of hardware
E. Credentials and policies
Answer: A,D
The correct answers are A and D because network infrastructure and virtualization of infrastructure and physical security of hardware are AWS responsibilities according to the AWS shared

responsibility model.
The AWS shared responsibility model is a framework that defines the division of responsibilities between AWS and the customer for security and compliance. AWS is responsible for the security of the cloud, which includes the global infrastructure, such as the regions, availability zones, and edge locations; the hardware, software, networking, and facilities that run the AWS services; and the virtualization layer that separates the customer instances and storage. The customer is responsible for the security in the cloud, which includes the customer data, the guest operating systems, the applications, the identity and access management, the firewall configuration, and the encryption. The other options are incorrect because they are not AWS responsibilities according to the AWS shared responsibility model. Security of application data, guest operating systems, and credentials and policies are customer responsibilities according to the AWS shared responsibility model.
Reference: [AWS Shared Responsibility Model]

NO.25 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 asses 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
AWS Cloud Adoption Framework (AWS CAF) is a tool that helps organizations understand how cloud adoption transforms the way they work, and it provides structure to identify and address gaps in skills and processes. Applying the AWS CAF in your organization results in an actionable plan that helps you prepare the cloud environment, enable your staff with new skills, and migrate your applications. AWS Pricing Calculator is a tool that helps you estimate the cost of AWS services for your use cases and compare the cost of different AWS service configurations. AWS Well-Architected Framework is a tool that helps you review and improve your cloud-based architectures and better understand the business impact of your design decisions. AWS Budgets is a tool that helps you plan your service usage, service costs, and instance reservations, and track how close your plan is to your budgeted amount.

NO.26 A company's IT team is managing MySQL database server clusters. The IT team has to patch the database and take backup snapshots of the data in the clusters. The company wants to move this workload to AWS so that these tasks will be completed automatically.
What should the company do to meet these requirements?

A. Deploy MySQL database server clusters on Amazon EC2 instances.

B. Use Amazon RDS with a MySQL database.

C. Use an AWS Cloud Form at ion template to deploy MySQL database servers on Amazon EC2 instances.

D. Migrate all the MySQL database data to Amazon S3.

Answer: B
The company should use Amazon RDS with a MySQL database to meet the requirements of moving its workload to AWS so that the tasks of patching the database and taking backup snapshots of the data in the clusters will be completed automatically. Amazon RDS is a managed service that simplifies

the setup, operation, and scaling of relational databases in the AWS Cloud. Amazon RDS automates common database administration tasks such as patching, backup, and recovery. Amazon RDS also supports MySQL and other popular database engines5

NO.27 Which AWS service or tool should a company use to forecast AWS spending?

A. Amazon DevPay

B. AWS Organizations

C. AWS Trusted Advisor

D. Cost Explorer

Answer: D

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

NO.28 Which controls are the responsibility of both AWS and AWS customers, according to the AWS shared responsibility model? (Select TWO.)

A. Physical and environmental controls

B. Patch management

C. Configuration management

D. Account structures

E. Choice of the AWS Region where data is stored

Answer: B,C

Patch management and configuration management are controls that are the responsibility of both AWS and AWS customers, according to the AWS shared responsibility model. Patch management is the process of applying updates to software and applications to fix vulnerabilities, bugs, or performance issues. Configuration management is the process of defining and maintaining the settings and parameters of systems and applications to ensure their consistency and reliability. AWS is responsible for patching and configuring the software and services that it manages, such as the AWS global infrastructure, the hypervisor, and the AWS managed services. The customer is responsible for patching and configuring the software and services that they manage, such as the guest operating system, the applications, and the AWS customer-managed services.

Physical and environmental controls are the responsibility of AWS, according to the AWS shared responsibility model. Physical and environmental controls are the measures that protect the physical security and availability of the AWS global infrastructure, such as power, cooling, fire suppression, and access control.

AWS is responsible for maintaining these controls and ensuring the resilience and reliability of the AWS Cloud. Account structures are the responsibility of the customer, according to the AWS shared

responsibility model. Account structures are the ways that customers organize and manage their AWS accounts and resources, such as using AWS Organizations, IAM users and roles, resource tagging, and billing preferences.
The customer is responsible for creating and configuring these structures and ensuring the security and governance of their AWS environment. Choice of the AWS Region where data is stored is the responsibility of the customer, according to the AWS shared responsibility model. AWS Regions are geographic areas that consist of multiple isolated Availability Zones. Customers can choose which AWS Region to store their data and run their applications, depending on their latency, compliance, and cost requirements. The customer is responsible for selecting the appropriate AWS Region and ensuring the data sovereignty and regulatory compliance of their data.

NO.29 A company is running its application in the AWS Cloud. The company wants to periodically review its AWS account for cost optimization opportunities.
Which AWS service or tool can the company use to meet these requirements?
A. AWS Cost Explorer
B. AWS Trusted Advisor
C. AWS Pricing Calculator
D. AWS Budgets
Answer: A
AWS Cost Explorer is an AWS service or tool that the company can use to periodically review its AWS account for cost optimization opportunities. AWS Cost Explorer is a tool that enables the company to visualize, understand, and manage their AWS costs and usage over time. The company can use AWS Cost Explorer to access interactive graphs and tables that show the breakdown of their costs and usage by service, region, account, tag, and more. The company can also use AWS Cost Explorer to forecast their future costs, identify trends and anomalies, and discover potential savings by using Reserved Instances or Savings Plans.

NO.30 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
The AWS Well-Architected Framework is a set of best practices and guidelines for designing and operating reliable, secure, efficient, and cost-effective systems in the cloud. The framework consists of five pillars:
operational excellence, performance efficiency, reliability, security, and cost optimization. The security pillar covers the AWS shared responsibility model, which defines the security and compliance responsibilities of AWS and the customers. You can learn more about the AWS Well-Architected Framework from [this whitepaper] or [this digital course].

NO.31 Which AWS service can defend against DDoS attacks?
A. AWS Firewall Manager
B. AWS Shield Standard

C. AWS WAF

D. Amazon Inspector

Answer: B

AWS Shield Standard is a service that provides protection against Distributed Denial of Service (DDoS) attacks for all AWS customers at no additional charge. It automatically detects and mitigates the most common and frequently occurring network and transport layer DDoS attacks that target AWS resources, such as Amazon EC2 instances, Elastic Load Balancers, Amazon CloudFront distributions, and Amazon Route 53 hosted zones. AWS Firewall Manager is a service that allows users to centrally configure and manage firewall rules across their AWS accounts and resources, such as AWS WAF web ACLs, AWS Shield Advanced protections, and Amazon VPC security groups. AWS WAF is a web application firewall that helps protect web applications from common web exploits, such as SQL injection, cross-site scripting, and bot attacks. Amazon Inspector is an automated security assessment service that helps improve the security and compliance of applications deployed on AWS. It analyzes the behavior of the applications and checks for vulnerabilities, exposures, and deviations from best practices.

NO.32 Which AWS service or feature enables users to encrypt data at rest in Amazon S3?

A. IAM policies

B. Server-side encryption

C. Amazon GuardDuty

D. Client-side encryption

Answer: B

Server-side encryption is an encryption option that Amazon S3 provides to encrypt data at rest in Amazon S3.

With server-side encryption, Amazon S3 encrypts an object before saving it to disk in its data centers and decrypts it when you download the objects. You have three server-side encryption options to choose from:

SSE-S3, SSE-C, and SSE-KMS. SSE-S3 uses keys that are managed by Amazon S3. SSE-C allows you to manage your own encryption keys. SSE-KMS uses keys that are managed by AWS Key Management Service (AWS KMS)5.

NO.33 Which design principles should a company apply to AWS Cloud workloads to maximize sustainability and minimize environmental impact? (Select TWO.)

A. Maximize utilization of Amazon EC2 instances.

B. Minimize utilization of Amazon EC2 instances.

C. Minimize usage of managed services.

D. Force frequent application reinstallations by users.

E. Reduce the need for users to reinstall applications.

Answer: A,E

To maximize sustainability and minimize environmental impact, a company should apply the following design principles to AWS Cloud workloads: maximize utilization of Amazon EC2 instances and reduce the need for users to reinstall applications. Maximizing utilization of Amazon EC2 instances means that the company can optimize the performance and efficiency of their compute resources, and avoid wasting energy and money on idle or underutilized instances. The company can use features such as Amazon EC2 Auto Scaling, Amazon EC2 Spot Instances, and AWS Compute

Optimizer to automatically adjust the number and type of instances based on demand, cost, and performance. Reducing the need for users to reinstall applications means that the company can minimize the amount of data and bandwidth required to deliver their applications to users, and avoid unnecessary downloads and updates that consume energy and resources. The company can use services such as Amazon CloudFront, AWS AppStream 2.0, and AWS Amplify to deliver their applications faster, more securely, and more efficiently to users across the globe. Minimizing utilization of Amazon EC2 instances, minimizing usage of managed services, and forcing frequent application reinstallations by users are not design principles that would maximize sustainability and minimize environmental impact. Minimizing utilization of Amazon EC2 instances would reduce the performance and efficiency of the compute resources, and potentially increase the costs and complexity of the cloud workloads. Minimizing usage of managed services would increase the operational overhead and responsibility of the company, and potentially expose them to more security and reliability risks. Forcing frequent application reinstallations by users would increase the amount of data and bandwidth required to deliver the applications to users, and potentially degrade the user experience and satisfaction.

NO.34 Which database engine is compatible with Amazon RDS?

A. Apache Cassandra

B. MongoDB

C. Neo4j

D. PostgreSQL

Answer: D

Amazon RDS supports six database engines: Amazon Aurora, MySQL, MariaDB, PostgreSQL, Oracle, and SQL Server. Apache Cassandra, MongoDB, and Neo4j are not compatible with Amazon RDS. Therefore, the correct answer is D. You can learn more about Amazon RDS and its supported database engines from this page.

NO.35 Which of the following are benefits that a company receives when it moves an on-premises production workload to AWS? (Select TWO.)

A. AWS trains the company's staff on the use of all the AWS services.

B. AWS manages all security in the cloud.

C. AWS offers free support from technical account managers (TAMs).

D. AWS offers high availability.

E. AWS provides economies of scale.

Answer: D,E

The correct answers are D and E because AWS offers high availability and AWS provides economies of scale are benefits that a company receives when it moves an on-premises production workload to AWS. High availability means that AWS has a global infrastructure that allows customers to deploy their applications and data across multiple regions and availability zones. This increases the fault tolerance and resilience of their applications and reduces the impact of failures. Economies of scale means that AWS can achieve lower variable costs than customers can get on their own. This allows customers to pay only for the resources they use and scale up or down as needed. The other options are incorrect because they are not benefits that a company receives when it moves an on-premises production workload to AWS. AWS trains the company's staff on the use of all the AWS services is not a benefit that a company receives when it moves an on-premises production workload to AWS. AWS

does provide various learning resources and training courses for customers, but it does not train the company's staff on the use of all the AWS services. AWS manages all security in the cloud is not a benefit that a company receives when it moves an on-premises production workload to AWS. AWS is responsible for the security of the cloud, but the customer is responsible for the security in the cloud. AWS offers free support from technical account managers (TAMs) is not a benefit that a company receives when it moves an on-premises production workload to AWS. AWS does offer support from TAMs, but only for customers who have the AWS Enterprise Support plan, which is not free. Reference: What is Cloud Computing?, [AWS Shared Responsibility Model], [AWS Support Plans]