ATMIC NETWORKS
YOUR GUIDE TO GLOBAL CERTIFIATION

"Please note that these files may not be up to date. However, the questions will help you understand the exam format and typical question patterns."

www.atmicnetworks .com

Warning: Keep connected with our support team for latest updates

Question: 1

Which database engine is compatible with Amazon RDS?

A. Apache Cassandra
B. MongoDB
C. Neo4j
D. PostgreSQL

Answer: D

Explanation:
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

Question: 2

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 a service that lets you run code without provisioning or managing servers. You can use Lambda to process event notifications from Amazon S3 when objects are uploaded or deleted. Lambda integrates directly with the event notification and invokes your code automatically. Therefore, the correct answer is A.

Question: 3

A company wants to centrally manage security policies and billing services within a multi-account AWS environment. Which AWS service should the company use to meet these requirements?

A. AWS Identity and Access Management (IAM)
B. AWS Organizations
C. AWS Resource Access Manager (AWS RAM)
D. AWS Config

Answer: B

Explanation:
AWS Organizations is a service that helps you centrally manage and govern your environment as you grow and scale your AWS resources. You can use AWS Organizations to create groups of accounts and apply policies to them. You can also use AWS Organizations to consolidate billing for multiple accounts. Therefore, the correct answer is B. You can learn more about AWS Organizations and its features .

Question: 4

What are the characteristics of Availability Zones? (Select TWO.)

A. All Availability Zones in an AWS Region are interconnected with high-bandwidth, low-latency networking
B. Availability Zones are physically separated by a minimum of distance of 150 km (100 miles).
C. All traffic between Availability Zones is encrypted.
D. Availability Zones within an AWS Region share redundant power, networking, and connectivity.
E. Every Availability Zone contains a single data center.

Answer: A,D

Explanation:
Availability Zones are physically separate locations within an AWS Region that are engineered to be isolated from failures. Each Availability Zone has independent power, cooling, and physical security, and is connected to other Availability Zones in the same Region by a low-latency network. Therefore, the correct answers are A and D. You can learn more about Availability Zones and their characteristics

Question: 5

Which AWS Well-Architected Framework concept represents a system's ability to remain functional when the system encounters operational problems?

A. Consistency
B. Elasticity

C. Durability
D. Latency

Answer: B

Explanation:
The AWS Well-Architected Framework is a set of best practices and guidelines for designing and operating systems in the cloud. The framework consists of five pillars: operational excellence, security, reliability, performance efficiency, and cost optimization. The concept of elasticity represents a system's ability to adapt to changes in demand by scaling resources up or down automatically. Therefore, the correct answer is B. You can learn more about the AWS Well- Architected Framework and its pillars

Question: 6

Which AWS service or tool does AWS Control Tower use to create resources?

A. AWS CloudFormation
B. AWS Trusted Advisor
C. AWS Directory Service
D. AWS Cost Explorer

Answer: A

Explanation:
AWS Control Tower uses AWS CloudFormation to create resources in your landing zone. AWS CloudFormation is a service that helps you model and set up your AWS resources using templates. AWS Control Tower supports creating AWS::ControlTower::EnabledControl resources in AWS CloudFormation. Therefore, the correct answer is A. You can learn more about AWS Control Tower and AWS CloudFormation

Question: 7

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:

Some of the advantages of using Amazon EC2 instances to host applications in the AWS Cloud instead of on premises are:

EC2 integrates with Amazon VPC, AWS CloudTrail, and AWS Identity and Access Management (IAM). Amazon VPC lets you provision a logically isolated section of the AWS Cloud where you can launch AWS resources in a virtual network that you define. AWS CloudTrail enables governance, compliance, operational auditing, and risk auditing of your AWS account. AWS IAM enables you to manage access to AWS services and resources securely. Therefore, the correct answer is B. You can learn more about Amazon EC2 and its integration with other AWS services

EC2 has a flexible, pay-as-you-go pricing model. You only pay for the compute capacity you use, and you can scale up and down as needed. You can also choose from different pricing options, such as On-Demand, Savings Plans, Reserved Instances, and Spot Instances, to optimize your costs. Therefore, the correct answer is D. You can learn more about Amazon EC2 pricing

The other options are incorrect because:

EC2 does not include operating system patch management. You are responsible for managing and maintaining your own operating systems on EC2 instances. You can use AWS Systems Manager to automate common maintenance tasks, such as applying patches, or use Amazon EC2 Image Builder to create and maintain secure images. Therefore, the incorrect answer is A.

EC2 does not have a 100% service level agreement (SLA). The EC2 SLA guarantees 99.99% availability for each EC2 Region, not for each individual instance. Therefore, the incorrect answer is C.

EC2 does not have automatic storage cost optimization. You are responsible for choosing the right storage option for your EC2 instances, such as Amazon Elastic Block Store (EBS) or Amazon Elastic File System (EFS), and monitoring and optimizing your storage costs. You can use AWS Cost Explorer or AWS Trusted Advisor to analyze and reduce your storage spending. Therefore, the incorrect answer is E.

Question: 8

Which option is an advantage of AWS Cloud computing that minimizes variable costs?

A. High availability
B. Economies of scale
C. Global reach
D. Agility

Answer: B

Explanation:

One of the advantages of AWS Cloud computing is that it minimizes variable costs by leveraging economies of scale. This means that AWS can achieve lower costs per unit of computing resources by spreading the fixed costs of building and maintaining data centers over a large number of customers. As a result, AWS can offer lower and more predictable prices to its customers, who only pay for the resources they consume. Therefore, the correct answer is B. You can learn more about AWS pricing and economies of scale

Question: 9

Which pillar of the AWS Well-Architected Framework focuses on the ability to run workloads effectively, gain insight into operations, and continuously improve supporting processes and procedures?

A. Cost optimization
B. Reliability
C. Operational excellence
D. Performance efficiency

Answer: C

Explanation:
The AWS Well-Architected Framework is a set of best practices and guidelines for designing and operating systems in the cloud. The framework consists of five pillars: operational excellence, security, reliability, performance efficiency, and cost optimization. The operational excellence pillar focuses on the ability to run workloads effectively, gain insight into operations, and continuously improve supporting processes and procedures. Therefore, the correct answer is C. You can learn more about the AWS Well-Architected Framework and its pillars.

Question: 10

Which benefit is included with an AWS Enterprise Support plan?

A. AWS Partner Network (APN) support at no cost
B. Designated support from an AWS technical account manager (TAM)
C. On-site support from AWS engineers
D. AWS managed compliance as code with AWS Config

Answer: B

Explanation:
AWS offers different support plans to meet the needs of different customers. The AWS Enterprise Support plan is the highest level of support that provides customers with concierge-like service, where the main focus is helping them achieve their outcomes and find success in the cloud. One of the benefits of the AWS Enterprise Support plan is that customers get designated support from an AWS technical account manager (TAM), who provides consultative architectural and operational guidance based on their applications and use cases. Therefore, the correct answer is B. You can learn more about AWS support plans and their benefits .

Question: 11

A company plans to migrate to AWS and wants to create cost estimates for its AWS use cases. Which AWS service or tool can the company use to meet these requirements?

A. AWS Pricing Calculator
B. Amazon CloudWatch
C. AWS Cost Explorer
D. AWS Budgets

Answer: A

Explanation:
AWS Pricing Calculator is a web-based planning tool that customers can use to create estimates for their AWS use cases. They can use it to model their solutions before building them, explore the AWS service price points, and review the calculations behind their estimates. Therefore, the correct answer is A. You can learn more about AWS Pricing Calculator and how it works.

Question: 12

A developer needs to build an application for a retail company. The application must provide real-time product recommendations that are based on machine learning.
Which AWS service should the developer use to meet this requirement?

A. AWS Health Dashboard
B. Amazon Personalize
C. Amazon Forecast
D. Amazon Transcribe

Answer: B

Explanation:
Amazon Personalize is a fully managed machine learning service that customers can use to generate personalized recommendations for their users. It can also generate user segments based on the users' affinity for certain items or item metadata. Amazon Personalize uses the customers' data to train and deploy custom recommendation models that can be integrated into their applications. Therefore, the correct answer is B. You can learn more about Amazon Personalize and its use case.

Question: 13

A company deploys its application on Amazon EC2 instances. The application occasionally experiences sudden increases in demand. The company wants to ensure that its application can respond to changes in demand at the lowest possible cost.
Which AWS service or tool will meet these requirements?

A. AWS Auto Scaling
B. AWS Compute Optimizer
C. AWS Cost Explorer
D. AWS Well-Architected Framework

Answer: A

Explanation:
AWS Auto Scaling is the AWS service or tool that will meet the requirements of ensuring that the application can respond to changes in demand at the lowest possible cost. AWS Auto Scaling allows users to automatically adjust the number of Amazon EC2 instances based on the application's performance and availability needs. AWS Auto Scaling can also optimize costs by helping users select the most cost-effective EC2 instances for their application1

Question: 14

A company wants to use a managed service to simplify the setup, operation, and scaling of its MySQL database in the AWS Cloud.
Which AWS service will meet these requirements?

A. Amazon EMR
B. Amazon RDS
C. Amazon Redshift
D. Amazon DynamoDB

Answer: B

Explanation:
Amazon RDS is the AWS service that will meet the requirements of using a managed service to simplify the setup, operation, and scaling of a MySQL database in the AWS Cloud. Amazon RDS is a relational database service that supports MySQL and other popular database engines. Amazon RDS handles routine database tasks such as provisioning, patching, backup, recovery, and scaling. Amazon RDS also offers high availability, security, and compatibility features3

Question: 15

A company deploys its application to multiple AWS Regions and configures automatic failover between those Regions.
Which cloud concept does this architecture represent?

A. Security
B. Reliability
C. Scalability
D. Cost optimization

Answer: B

Explanation:
Reliability is the cloud concept that this architecture represents. Reliability is the ability of a system to recover from infrastructure or service disruptions, dynamically acquire computing resources to meet demand, and mitigate disruptions such as misconfigurations or transient network issues. Deploying an application to multiple AWS Regions and configuring automatic failoverbetween those Regions enhances the reliability of the application by reducing the impact of regional failures and increasing the availability of the application4

Question: 16

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

Explanation:
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

Question: 17

A company recently migrated to the AWS Cloud. The company needs to determine whether its newly imported Amazon EC2 instances are the appropriate size and type.
Which AWS services can provide this information to the company? {Select TWO.)

A. AWS Auto Scaling
B. AWS Control Tower
C. AWS Trusted Advisor

D. AWS Compute Optimizer
E. Amazon Forecast

Answer: C,D

Explanation:
AWS Trusted Advisor and AWS Compute Optimizer are the AWS services that can provide information to the company about whether its newly imported Amazon EC2 instances are theappropriate size and type. AWS Trusted Advisor is an online tool that provides best practices recommendations in five categories: cost optimization, performance, security, fault tolerance, and service limits. AWS Trusted Advisor can help users identify underutilized or idle EC2 instances, and suggest ways to reduce costs and improve performance. AWS Compute Optimizer is a service that analyzes the configuration and utilization metrics of EC2 instances and delivers recommendations for optimal instance types, sizes, and configurations. AWS Compute Optimizer helps users improve performance, reduce costs, and eliminate underutilized resources

Question: 18

A company has a social media platform in which users upload and share photos with other users. The company wants to identify and remove inappropriate photos. The company has no machine learning (ML) scientists and must build this detection capability with no ML expertise.
Which AWS service should the company use to build this capability?

A. Amazon SageMaker
B. Amazon Textract
C. Amazon Rekognition
D. Amazon Comprehend

Answer: C

Explanation:
Amazon Rekognition is the AWS service that the company should use to build the capability of identifying and removing inappropriate photos. Amazon Rekognition is a service that uses deep learning technology to analyze images and videos for various purposes, such as face detection, object recognition, text extraction, and content moderation. Amazon Rekognition can help users detect unsafe or inappropriate content in images and videos, such as nudity, violence, or drugs, and provide confidence scores for each label. Amazon Rekognition does not require any machine learning expertise, and users can easily integrate it with other AWS services

Question: 19

A company's user base needs to remotely access virtual desktop computers from the internet Which AWS service provides this functionality?

A. Amazon Connect
B. Amazon Cognito
C. Amazon Workspaces
D. Amazon Upstream 2.0

Answer: C

Explanation:

Amazon Workspaces is the AWS service that provides the functionality of remotely accessing virtual desktop computers from the internet. Amazon Workspaces is a fully managed, secure desktop-as-a- service (DaaS) solution that allows users to provision cloud-based virtual desktops and access them from anywhere, using any supported device. Amazon Workspaces helps users reduce the complexity and cost of managing and maintaining physical desktops, and provides a consistent and secure user experience

Question: 20

Amazon Elastic File System (Amazon EFS) and Amazon FSx offer which type of storage?

A. File storage
B. Object storage
C. Block storage
D. Instance store

Answer: A

Explanation:

Amazon Elastic File System (Amazon EFS) and Amazon FSx offer file storage. File storage is a type of storage that organizes data into files and folders, and allows multiple users or applications to access and share the same files over a network. Amazon EFS is a fully managed, scalable, and elastic file system that supports the Network File System (NFS) protocol and can be used with Amazon EC2 instances and AWS Lambda functions. Amazon FSx is a fully managed service that provides two file system options: Amazon FSx for Windows File Server, which supports the Server Message Block(SMB) protocol and is compatible with Microsoft Windows applications; and Amazon FSx for Lustre, which is a high-performance file system that is optimized for compute-intensive workloads

Question: 21

Which AWS service or feature is used to Troubleshoot network connectivity issues between Amazon EC2 instances?

A. AWS Certificate Manager (ACM)
B. Internet gateway
C. VPC Flow Logs
D. AWS CloudHSM

Answer: C

Explanation:

VPC Flow Logs is the AWS service or feature that is used to troubleshoot network connectivity issues between Amazon EC2 instances. VPC Flow Logs is a feature that enables users to captureinformation about the IP traffic going to and from network interfaces in their VPC. VPC Flow Logs can help users monitor and diagnose network-related issues, such as traffic not reaching an instance, or an instance not responding to requests. VPC Flow Logs can be published to Amazon CloudWatch Logs, Amazon S3, or Amazon Kinesis Data Firehose for analysis and storage.

Question: 22

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 and compute resources that are currently in use are the factors that affect costs in the AWS Cloud. Outbound data transfers without acceleration refer to the amount of data that is transferred from AWS to the internet, without using any service that can optimize the speed and cost of the data transfer, such as AWS Global Accelerator or Amazon CloudFront. Outbound data transfers are charged at different rates depending on the source and destination AWS Regions, and the volume of data transferred. Compute resources that are currently in use refer to the AWS services and resources that provide computing capacity, such as Amazon EC2 instances, AWS Lambda functions, or Amazon ECS tasks. Compute resources are charged based on the type, size, and configuration of the resources, and the duration and frequency of their usage.

Question: 23

Which design principles support the reliability pillar of the AWS Well-Architected Framework? (Select 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:
The design principles that support the reliability pillar of the AWS Well-Architected Framework are: automatically scale to meet demand, and automatically recover from failure. These principles help users design systems that can handle changes in load, avoid disruptions, and resume normal operations quickly. Automatically scaling to meet demand means adjusting the capacity of the system based on the current and anticipated workload, using services such as AWS Auto Scaling, Amazon EC2, and AWS Lambda. Automatically recovering from failure means detecting and resolving issues, using services such as Amazon CloudWatch, AWS CloudFormation, and AWS CloudTrail

Question: 24

Which of the following are user authentication services managed by AWS? (Select TWO.)

A. Amazon Cognito
B. AWS Lambda
C. AWS License Manager

D. AWS Identity and Access Management (IAM)
E. AWS CodeStar

Answer: A,D

Explanation:
The user authentication services managed by AWS are: Amazon Cognito and AWS Identity and Access Management (IAM). These services help users securely manage and control access to their AWS resources and applications. Amazon Cognito is a service that provides user sign-up, sign-in, and access control for web and mobile applications. Amazon Cognito supports various identity providers, such as Facebook, Google, and Amazon, as well as custom user pools. AWS IAM is a service that enables users to create and manage users, groups, roles, and permissions for AWS services and resources. AWS IAM supports various authentication methods, such as passwords, access keys, and multi-factorauthentication (MFA)

Question: 25

company wants to protect its AWS Cloud information, systems, and assets while performing risk assessment and mitigation tasks. Which pillar of the AWS Well-Architected Framework is supported by these goals?

A. Reliability
B. Security
C. Operational excellence
D. Performance efficiency

Answer: B

Explanation:
The pillar of the AWS Well-Architected Framework that is supported by the goals of protecting AWS Cloud information, systems, and assets while performing risk assessment and mitigation tasks is security. Security is the ability to protect information, systems, and assets while delivering business value through risk assessments and mitigation strategies. The security pillar covers topics such as identity and access management, data protection, infrastructure protection, detective controls, incident response, and compliance

Question: 26

A company is configuring its AWS Cloud environment. The company's administrators need to group users together and apply permissions to the group.
Which AWS service or feature can the company use to meet these requirements?

A. AWS Organizations
B. Resource groups
C. Resource tagging
D. AWS Identity and Access Management (IAM)

Answer: D

Explanation:
The AWS service or feature that the company can use to group users together and apply permissions to the group is AWS Identity and Access Management (IAM). AWS IAM is a service that enables users to create and manage users, groups, roles, and permissions for AWS services and resources. Users can use IAM groups to organize multiple users that have similar access requirements, and attach policies to the groups that define the permissions for the users in the group. This simplifies the management and administration of user access

Question: 27

A company has two AWS accounts in an organization in AWS Organizations for consolidated billing. All of the company's AWS resources are hosted in one AWS Region.
Account A has purchased five Amazon EC2 Standard Reserved Instances (RIs) and has four EC2 instances running. Account B has not purchased any RIs and also has four EC2 instances running.
Which statement is true regarding pricing for these eight instances?

A. The eight instances will be charged as regular instances.
B. Four instances will be charged as RIs, and four will be charged as regular instances.
C. Five instances will be charged as RIs, and three will be charged as regular instances.
D. The eight instances will be charged as RIs.

Answer: B

Explanation:
The statement that is true regarding pricing for these eight instances is: four instances will be charged as RIs, and four will be charged as regular instances. Amazon EC2 Reserved Instances (RIs) are a pricing model that allows users to reserve EC2 instances for a specific term and benefit from discounted hourly rates and capacity reservation. RIs are purchased for a specific AWS Region, and can be shared across multiple accounts in an organization in AWS Organizations for consolidated billing. However, RIs are applied on a first-come, first-served basis, and there is no guarantee that all instances in the organization will be charged at the RI rate. In this case, Account A has purchased five RIs and has four instances running, so all four instances will be charged at the RI rate. Account B has not purchased any RIs and also has four instances running, so all four instances will be charged at the regular rate. The remaining RI in Account A will not be applied to any instance in Account B, and will be wasted.

Question: 28

Which of the following is an advantage that users experience when they move on-premises workloads to the AWS Cloud?

A. Elimination of expenses for running and maintaining data centers
B. Price discounts that are identical to discounts from hardware providers
C. Distribution of all operational controls to AWS
D. Elimination of operational expenses

Answer: A

Explanation:
The advantage that users experience when they move on-premises workloads to the AWS Cloud is: elimination of expenses for running and maintaining data centers. By moving on-premises workloads to the AWS Cloud, users can reduce or eliminate the costs associated with owning and operating physical servers, storage, network equipment, and facilities. These costs include hardware purchase, maintenance, repair, power, cooling, security, and staff. Users can also benefit from the pay-as-you- go pricing model of AWS, which allows them to pay only for the resources they use, and scale up or down as needed.

Question: 29

Which of the following is a cost efficiency principle related to the AWS Cloud?

A. Right-size services based on capacity requirements.
B. Use the Billing Dashboard to access information about monthly bills.
C. Use AWS Organizations to combine the expenses of multiple accounts into a single bill.
D. Tag all AWS resources.

Answer: A

Explanation:
One of the cost efficiency principles related to the AWS Cloud is to right-size services based on capacity requirements. This means choosing the most appropriate type and size of AWS resources to meet the performance and scalability needs of the applications, while avoiding over-provisioning or under-provisioning. By right-sizing services, users can optimize the costs and benefits of using the AWS Cloud1

Question: 30

A cloud engineer needs to download AWS security and compliance documents for an upcoming audit.
Which AWS service can provide the documents?

A. AWS Trusted Advisor
B. AWS Artifact
C. AWS Well-Architected Tool
D. AWS Systems Manager

Answer: B

Explanation:
AWS Artifact is the AWS service that can provide security and compliance documents for an upcoming audit. AWS Artifact is a self-service portal that allows users to access and download AWS compliance reports and agreements. These documents provide evidence of AWS's compliance with global, regional, and industry-specific security standards and regulations

Question: 31

A company has been storing monthly reports in an Amazon S3 bucket. The company exports the report data into comma-separated values (.csv) files. A developer wants to write a simple query that can read all of these files and generate a summary report. Which AWS service or feature should the developer use to meet these requirements with the LEAST amount of operational overhead?

A. Amazon S3 Select
B. Amazon Athena
C. Amazon Redshift
D. Amazon EC2

Answer: C

Explanation:
Amazon Athena is the AWS service that the developer should use to write a simple query that can read all of the .csv files stored in an Amazon S3 bucket and generate a summary report. Amazon

Athena is an interactive query service that allows users to analyze data in Amazon S3 using standard SQL. Amazon Athena does not require any server setup or management, and users only pay for the queries they run. Amazon Athena can handle various data formats, including .csv, and can integrate with other AWS services such as Amazon QuickSight for data visualization

Question: 32

Which task requires the use of AWS account root user credentials?

A. The deletion of IAM users
B. The change to a different AWS Support plan
C. The creation of an organization in AWS Organizations
D. The deletion of Amazon EC2 instances

Answer: C

Explanation:
The creation of an organization in AWS Organizations requires the use of AWS account root user credentials. The AWS account root user is the email address that was used to create the AWS account. The root user has complete access to all AWS services and resources in the account, and can perform sensitive tasks such as changing the account settings, closing the account, or creating an organization. The root user credentials should be used sparingly and securely, and only for tasks that cannot be performed by IAM users or roles4

Question: 33

Which feature of the AWS Cloud gives users the ability to pay based on current needs rather than forecasted needs?

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

Answer: B

Explanation:

Pay-as-you-go pricing is the feature of the AWS Cloud that gives users the ability to pay based on current needs rather than forecasted needs. Pay-as-you-go pricing means that users only pay for theAWS services and resources they use, without any upfront or long-term commitments. This allows users to scale up or down their usage depending on their changing business requirements, and avoid paying for idle or unused capacity. Pay-as-you-go pricing also enables users to benefit from the economies of scale and lower costs of AWS as they grow their business5

Question: 34

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 (AmazonEBS) 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

Explanation:

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.

Question: 35

Which AWS service gives users the ability to provision a dedicated and private network connection from their internal network to AWS?

A. AWS CloudHSM
B. AWS Direct Connect
C. AWS VPN
D. Amazon Connect

Answer: B

Explanation:

AWS Direct Connect gives users the ability to provision a dedicated and private network connection from their internal network to AWS. AWS Direct Connect links the user's internal network to an AWSDirect Connect location over a standard Ethernet fiber-optic cable. One end of the cable is connected to the user's router, the other to an AWS Direct Connect router. With this connection in place, the user can create virtual interfaces directly to the AWS cloud and Amazon Virtual Private Cloud (Amazon VPC), bypassing internet service providers in the network path2.

Question: 36

A company is hosting a web application in a Docker container on Amazon EC2.

AWS is responsible for which of the following tasks?

A. Scaling the web application and services developed with Docker
B. Provisioning or scheduling containers to run on clusters and maintain their availability
C. Performing hardware maintenance in the AWS facilities that run the AWS Cloud
D. Managing the guest operating system, including updates and security patches

Answer: C

Explanation:

AWS is responsible for performing hardware maintenance in the AWS facilities that run the AWS Cloud. This is part of the shared responsibility model, where AWS is responsible for the security of the cloud, and the customer is responsible for security in the cloud. AWS is also responsible for the global infrastructure that runs all of the services offered in the AWS Cloud, including the hardware, software, networking, and facilities that run AWS Cloud services3. The customer is responsible for the guest operating system, including updates and security patches, as well as the web application and services developed with Docker4.

Question: 37

Which design principle should be considered when architecting in the AWS Cloud?

A. Think of servers as non-disposable resources.
B. Use synchronous integration of services.
C. Design loosely coupled components.
D. Implement the least permissive rules for security groups.

Answer: C

Explanation:

Designing loosely coupled components is a design principle that should be considered when architecting in the AWS Cloud. Loose coupling is a way of designing systems to reduce interdependencies and minimize the impact of changes. Loose coupling allows components to interact with each other through well-defined interfaces, rather than direct references. This reducesthe risk of failures and errors propagating across the system, and enables greater scalability, availability, and maintainability5.

Question: 38

Which AWS service or tool helps to centrally manage billing and allow controlled access to resources across AWS accounts?

A. AWS Identity and Access Management (IAM)
B. AWS Organizations
C. AWS Cost Explorer
D. AWS Budgets

Answer: B

Explanation:
AWS Organizations helps to centrally manage billing and allow controlled access to resources across AWS accounts. AWS Organizations is a service that enables the user to consolidate multiple AWS accounts into an organization that can be managed as a single unit. AWS Organizations allows the user to create groups of accounts and apply policies to them, such as service control policies (SCPs) that specify the services and actions that users and roles can access in the accounts. AWS Organizations also enables the user to use consolidated billing, which combines the usage and charges from all the accounts in the organization into a single bill.

Question: 39

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

Explanation:
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.

Question: 40

Which of the following promotes AWS Cloud architectural best practices for designing and operating reliable, secure, efficient, and cost-effective systems?

A. AWS Serverless Application Model framework
B. AWS Business Support

C. Principle of least privilege
D. AWS Well-Architected Framework

Answer: D

Explanation:
AWS Well-Architected Framework promotes AWS Cloud architectural best practices for designing and operating reliable, secure, efficient, and cost-effective systems. AWS Well-Architected Framework is a set of guidelines and best practices that help the user to evaluate and improve the architecture of their applications and workloads on AWS. AWS Well-Architected Framework consists of five pillars: operational excellence, security, reliability, performance efficiency, and cost optimization. Each pillar provides a set of design principles, questions, and best practices that help the user to achieve the desired outcomes for their systems.

Question: 41

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

A. Management of the guest operating systems
B. Maintenance of the configuration of infrastructure devices
C. Management of the host operating systems and virtualization
D. Maintenance of the software that powers Availability ZonesA company has refined its workload to use specific AWS services to improve efficiency and reduce cost.

Answer: A

Explanation:
Management of the guest operating systems is a customer's responsibility, according to the AWS shared responsibility model. The AWS shared responsibility model defines the different security and compliance responsibilities of AWS and the customer. AWS is responsible for the security of the cloud, which includes the physical infrastructure, hardware, software, and facilities that run the AWS Cloud. The customer is responsible for security in the cloud, which includes the configuration and management of the guest operating systems, applications, data, and network traffic protection

Question: 42

Which best practice for cost governance does this example show?

A. Resource controls
B. Cost allocation
C. Architecture optimization
D. Tagging enforcement

Answer: C

Explanation:
Architecture optimization is the best practice for cost governance that this example shows. Architecture optimization is the process of

designing and implementing AWS solutions that are efficient, scalable, and cost-effective. By using specific AWS services to improve efficiency and reduce cost, the company is following the architecture optimization best practice. Some of the techniques for architecture optimization include using the right size and type of resources, leveraging elasticity and scalability, choosing the most suitable storage class, and using serverless and managed services2.

Question: 43

Which activity can companies complete by using AWS Organizations?

A. Troubleshoot the performance of applications.
B. Manage service control policies (SCPs).
C. Migrate applications to microservices.
D. Monitor the performance of applications.

Answer: B

Explanation:
Managing service control policies (SCPs) is an activity that companies can complete by using AWS Organizations. AWS Organizations is a service that enables the user to consolidate multiple AWS accounts into an organization that can be managed as a single unit. AWS Organizations allows the user to create groups of accounts and apply policies to them, such as service control policies (SCPs) that specify the services and actions that users and roles can access in the accounts. AWS Organizations also enables the user to use consolidated billing, which combines the usage and charges from all the accounts in the organization into a single bill3.

Question: 44

Which AWS service or feature is used to send both text and email messages from distributed applications?

A. Amazon Simple Notification Service (Amazon SNS)
B. Amazon Simple Email Service (Amazon SES)
C. Amazon CloudWatch alerts
D. Amazon Simple Queue Service (Amazon SQS)

Answer: A

Explanation:
Amazon Simple Notification Service (Amazon SNS) is the AWS service or feature that is used to send both text and email messages from distributed applications. Amazon SNS is a fully managed pub/sub messaging service that enables the user to send messages to multiple subscribers or endpoints, such as email addresses, phone numbers, HTTP endpoints, AWS Lambda functions, and more. Amazon SNS can be used to send notifications, alerts, confirmations, and reminders from applications to users or other applications4.

Question: 45

Which of the following is a benefit of decoupling an AWS Cloud architecture?

A. Reduced latency

B. Ability to upgrade components independently
C. Decreased costs
D. Fewer components to manage

Answer: B

Explanation:

A benefit of decoupling an AWS Cloud architecture is the ability to upgrade components independently. Decoupling is a way of designing systems to reduce interdependencies and minimize the impact of changes. Decoupling allows components to interact with each other through well- defined interfaces, rather than direct references. This reduces the risk of failures and errors propagating across the system, and enables greater scalability, availability, and maintainability. By decoupling an AWS Cloud architecture, the user can upgrade or modify one component without affecting the other components5.

Question: 46

Which of the following describes an AWS Region?

A. A specific location within a geographic area that provides high availability
B. A set of data centers spanning multiple countries
C. A global picture of a user's cloud computing environment
D. A collection of databases that can be accessed from a specific geographic area only

Answer: A

Explanation:

An AWS Region is a specific location within a geographic area that provides high availability. An AWS Region consists of two or more Availability Zones, which are isolated locations within the sameRegion. Each Availability Zone has independent power, cooling, and physical security, and is connected to the other Availability Zones in the same Region by low-latency, high-throughput, and highly redundant networking. AWS services are available in multiple Regions around the world, allowing the user to choose where to run their applications and store their data1.

Question: 47

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.

Which of the following are benefits of building this app in the AWS Cloud? (Select TWO.)

A. A large upfront capital expense and low variable expenses
B. Increased speed for trying out new projects
C. Complete control over the physical security of the infrastructure
D. Flexibility to scale up in minutes as the application becomes popular
E. Ability to pick the specific data centers that will host the application servers

Answer: A,D

Explanation:

AWS Directory Service for Microsoft Active Directory is the AWS service that provides a managed Microsoft Active Directory in the AWS Cloud. It enables the user to use their existing Active Directory users, groups, and policies to access AWS resources, such as Amazon EC2 instances, Amazon S3 buckets, and AWS Single Sign-On. It also integrates with other Microsoft applications and services, such as Microsoft SQL Server, Microsoft Office 365, and Microsoft SharePoint

Question: 48

Which AWS service should a cloud practitioner use to receive real-time guidance for provisioning resources, based on AWS best practices related to security, cost optimization, and service limits?

A. AWS Trusted Advisor
B. AWS Config
C. AWS Security Hub
D. AWS Systems Manager

Answer: A

Explanation:

AWS Trusted Advisor is the AWS service that provides real-time guidance for provisioning resources, based on AWS best practices related to security, cost optimization, and service limits. AWS Trusted Advisor inspects the user's AWS environment and provides recommendations for improving performance, security, and reliability, reducing costs, and following best practices. AWS TrustedAdvisor also alerts the user when they are approaching or exceeding their service limits, and helps them request limit increases3.

Question: 49

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

A. The ability to turn over the responsibility for all security to AWS.
B. The ability to use the pay-as-you-go model.
C. The ability to have full control over the physical infrastructure.
D. No longer having to guess what capacity will be required.
E. No longer worrying about users access controls.

Answer: B,D

Explanation:

The advantages of moving to the AWS Cloud are the ability to use the pay-as-you-go model and no longer having to guess what capacity will be required. The pay-as-you-go model allows the user to pay only for the resources they use, without any upfront or long-term commitments. This reduces the cost and risk of over-provisioning or under-provisioning resources. No longer having to guess what capacity will be required means that the user can scale their resources up or down according to the demand, without wasting money on idle resources or losing customers due to insufficient capacity4.

Question: 50

A company is migrating a relational database server to the AWS Cloud. The company wants to minimize administrative overhead of database maintenance tasks.

Which AWS service will meet these requirements?

A. Amazon DynamoDB
B. Amazon EC2
C. Amazon Redshift
D. Amazon RDS

Answer: D

Explanation:

Amazon RDS is the AWS service that will meet the requirements of migrating a relational database server to the AWS Cloud and minimizing administrative overhead of database maintenance tasks.

Amazon RDS is a fully managed relational database service that handles routine database tasks, such as provisioning, patching, backup, recovery, failure detection, and repair. Amazon RDS supports several database engines, such as MySQL, PostgreSQL, Oracle, SQL Server, and Amazon Aurora5.

Question: 51

A company is reviewing its operating policies.

Which policy complies with guidance in the security pillar of the AWS Well-Architected Framework?

A. Ensure that employees have access to all company data.
B. Expand employees' permissions as they gain more experience.
C. Grant all privileges and access to all users.
D. Apply security requirements at all layers of a process.

Answer: D

Explanation:

Applying security requirements at all layers of a process is a policy that complies with guidance in the security pillar of the AWS Well-Architected Framework. The security pillar of the AWS Well- Architected Framework provides best practices for securing the user's data and systems in the AWS Cloud. One of the design principles of the security pillar is to apply security at all layers, which means that the user should implement defense-in-depth strategies and avoid relying on a single security mechanism. For example, the user should use multiple security controls, such as encryption, firewalls, identity and access management, and logging and monitoring, to protect their data and resources at different layers.

Question: 52

Which task is the responsibility of a company that is using Amazon RDS?

A. Provision the underlying infrastructure.
B. Create IAM policies to control administrative access to the service.
C. Install the cables to connect the hardware for compute and storage.
D. Install and patch the RDS operating system.

Answer: B

Explanation:
The correct answer is B because AWS IAM policies can be used to control administrative access to the Amazon RDS service. The other options are incorrect because they are the responsibilities of AWS, not the company that is using Amazon RDS. AWS manages the provisioning, cabling, installation, and patching of the underlying infrastructure for Amazon RDS. Reference: Amazon RDS FAQs

Question: 53

A company is designing an identity access management solution for an application. The company wants users to be able to use their social media, email, or online shopping accounts to access the application.
Which AWS service provides this functionality?

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

Answer: C

Explanation:
The correct answer is C because Amazon Cognito provides identity federation and user authentication for web and mobile applications. Amazon Cognito allows users to sign in with their social media, email, or online shopping accounts. The other options are incorrect because they do not provide identity federation or user authentication. AWS IAM Identity Center (AWS Single Sign- On) is a service that enables users to access multiple AWS accounts and applications with a single sign-on experience. AWS Config is a service that enables users to assess, audit, and evaluate the configurations of their AWS resources. AWS Identity and Access Management (IAM) is a service that enables users to manage access to AWS resources using users, groups, roles, and policies. Reference: Amazon Cognito FAQs

Question: 54

Which AWS service aggregates, organizes, and prioritizes security alerts and findings from multiple AWS services?

A. Amazon Detective
B. Amazon Inspector
C. Amazon Macie
D. AWS Security Hub

Answer: D

Explanation:

The correct answer is D because AWS Security Hub is a service that aggregates, organizes, and prioritizes security alerts and findings from multiple AWS services, such as Amazon GuardDuty, Amazon Inspector, Amazon Macie, AWS Firewall Manager, and AWS IAM Access Analyzer. The other options are incorrect because they are not services that aggregate security alerts and findings from multiple AWS services. Amazon Detective is a service that helps users analyze and visualize security data to investigate and remediate potential issues. Amazon Inspector is a service that helps users find security vulnerabilities and deviations from best practices in their Amazon EC2 instances. Amazon Macie is a service that helps users discover, classify, and protect sensitive data stored in Amazon S3. Reference: AWS Security Hub FAQs

Question: 55

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:

The correct answers are B and C because they are advantages of the AWS Cloud. High economies of scale means that AWS can achieve lower variable costs than customers can get on their own. Launch globally in minutes means that AWS has a global infrastructure that allows customers to deploy their applications and data across multiple regions and availability zones. The other options are incorrect because they are not advantages of the AWS Cloud. Trade variable expenses for capital expenses means that customers have to invest heavily in data centers and servers before they know how they will use them. Focus on managing hardware infrastructure means that customers have to spend time and money on maintaining and upgrading their physical resources. Overprovision to ensure capacity means that customers have to pay for more resources than they actually need to avoid performance issues. Reference: What is Cloud Computing?

Question: 56

Which AWS service is a key-value database that provides sub-millisecond latency on a large scale?

A. Amazon DynamoDB
B. Amazon Aurora
C. Amazon DocumentDB (with MongoDB compatibility)
D. Amazon Neptune

Answer: A

Explanation:

The correct answer is A because Amazon DynamoDB is a key-value database that provides submillisecond latency on a large scale. Amazon DynamoDB is a fully managed, serverless, and scalable NoSQL database service that supports both key-value and document data models. The other options are incorrect because they are not key-value databases. Amazon Aurora is a relational database that is compatible with MySQL and PostgreSQL. Amazon DocumentDB (with MongoDB compatibility) is a document database that is compatible with MongoDB. Amazon Neptune is a graph database that supports property graph and RDF models. Reference: Amazon DynamoDB FAQs

Question: 57

Which AWS service or tool provides users with the ability to monitor AWS service quotas?

A. AWS CloudTrail
B. AWS Cost and Usage Reports
C. AWS Trusted Advisor
D. AWS Budgets

Answer: C

Explanation:

The correct answer is C because AWS Trusted Advisor is an AWS service or tool that provides users with the ability to monitor AWS service quotas. AWS Trusted Advisor is an online tool that provides users with real-time guidance to help them provision their resources following AWS best practices. One of the categories of checks that AWS Trusted Advisor performs is service limits, which monitors the usage of each AWS service and alerts users when they are close to reaching the default limit. The other options are incorrect because they are not AWS services or tools that provide users with the ability to monitor AWS service quotas. AWS CloudTrail is a service that enables users to track user activity and API usage across their AWS account. AWS Cost and Usage Reports is a tool that enables users to access comprehensive information about their AWS costs and usage. AWS Budgets is a tool that enables users to plan their service usage, costs, and reservations. Reference: [AWS Trusted Advisor FAQs]

Question: 58

Which of the following is an advantage of AWS Cloud computing?

A. Trade security for elasticity.
B. Trade operational excellence for agility.
C. Trade fixed expenses for variable expenses.
D. Trade elasticity for performance.

Answer: C

Explanation:

The correct answer is C because AWS Cloud computing allows customers to trade fixed expenses for variable expenses. This means that customers only pay for the resources they use, and can scale up or down as needed. The other options are incorrect because they are not advantages of AWS Cloud computing. Trade security for elasticity means that customers have to compromise on the protection of their

data and applications in order to adjust their capacity quickly. Trade operational excellence for agility means that customers have to sacrifice the quality and reliability of their operations in order to respond to changing needs faster. Trade elasticity for performance means that customers have to limit their ability to scale up or down in order to achieve higher speed and efficiency. Reference: What is Cloud Computing?

Question: 59

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

Explanation:
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

Question: 60

A company has an online shopping website and wants to store customers' credit card dat a. The company must meet Payment Card Industry (PCI) standards.

Which service can the company use to access AWS compliance documentation?

A. Amazon Cloud Directory
B. AWS Artifact
C. AWS Trusted Advisor
D. Amazon Inspector

Answer: B

Explanation:
The correct answer is B because AWS Artifact is a service that provides access to AWS compliance documentation, such as audit reports, security certifications, and agreements. AWS Artifact allows customers to download, review, and accept the documents that are relevant to their use of AWS services. The other options are incorrect because they are not services that provide access to AWS compliance

documentation. Amazon Cloud Directory is a service that enables customers to create flexible cloud-native directories for organizing hierarchies of data. AWS Trusted Advisor is a service that provides real-time guidance to help customers follow AWS best practices for security, performance, cost optimization, and fault tolerance. Amazon Inspector is a service that helps customers find security vulnerabilities and deviations from best practices in their Amazon EC2 instances. Reference: [AWS Artifact FAQs]

Question: 61

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:
The correct answers are B and D because a virtual private gateway and a customer gateway are components of an AWS Site-to-Site VPN connection. A virtual private gateway is the AWS side of the VPN connection that attaches to the customer's VPC. A customer gateway is the customer side of the VPN connection that resides in the customer's network. The other options are incorrect because they are not components of an AWS Site-to-Site VPN connection. AWS Storage Gateway is a service that connects on-premises software applications with cloud-based storage. NAT gateway is a service that enables instances in a private subnet to connect to the internet or other AWS services, but prevents the internet from initiating a connection with those instances. Internet gateway is a service that enables communication between instances in a VPC and the internet. Reference: [What is AWS Site- to-Site VPN?]

Question: 62

A company runs thousands of simultaneous simul-ations using AWS Batch. Each simul-ation is stateless, is fault tolerant, and runs for up to 3 hours.
Which pricing model enables the company to optimize costs and meet these requirements?

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

Answer: B

Explanation:
The correct answer is B because Spot Instances enable the company to optimize costs and meet the requirements. Spot Instances are spare EC2 instances that are available at up to 90% discount compared to On-Demand prices. Spot Instances are suitable for stateless, fault-tolerant, and flexible applications that can run for any duration. The other options are incorrect because they do not enable the company to optimize costs and meet the requirements. Reserved Instances are EC2 instances that are reserved for a specific period of time (one or three years) in exchange for a lower hourly rate. Reserved Instances are suitable for steady-state or predictable workloads that run for a long duration. On-Demand Instances are EC2 instances that are launched and billed at a fixed hourly rate. On-Demand Instances are suitable for short-term, irregular, or unpredictable workloads that cannot be interrupted. Dedicated Instances are EC2 instances that run

on hardware that is dedicated to a single customer. Dedicated Instances are suitable for workloads that require regulatory compliance or data isolation. Reference: [Amazon EC2 Instance Purchasing Options]

Question: 63

A company has an application with robust hardware requirements. The application must be accessed by students who are using lightweight, low-cost laptops.
Which AWS service will help the company deploy the application without investing in backend infrastructure or high end client hardware?

A. Amazon AppStream 2.0
B. AWS AppSync
C. Amazon WorkLink
D. AWS Elastic Beanstalk

Answer: A

Explanation:
The correct answer is A because Amazon AppStream 2.0 is a service that will help the company deploy the application without investing in backend infrastructure or high end client hardware. Amazon AppStream 2.0 is a fully managed, secure application streaming service that allows customers to stream desktop applications from AWS to any device running a web browser. Amazon AppStream 2.0 handles the provisioning, scaling, patching, and maintenance of the backend infrastructure, and delivers high performance and responsive user experience. The other options are incorrect because they are not services that will help the company deploy the application without investing in backend infrastructure or high end client hardware. AWS AppSync is a service that enables customers to create flexible APIs for synchronizing data across multiple data sources. Amazon WorkLink is a service that enables customers to provide secure, one-click access to internalwebsites and web apps from mobile devices. AWS Elastic Beanstalk is a service that enables customers to deploy and manage web applications using popular platforms such as Java, .NET, PHP, and Node.js. Reference: [Amazon AppStream 2.0 FAQs]

Question: 64

Which AWS service will help a company identify the user who deleted an Amazon EC2 instance yesterday?

A. Amazon CloudWatch
B. AWS Trusted Advisor
C. AWS CloudTrail
D. Amazon Inspector

Answer: C

Explanation:
The correct answer is C because AWS CloudTrail is a service that will help a company identify the user who deleted an Amazon EC2 instance yesterday. AWS CloudTrail is a service that enables users to track user activity and API usage across their AWS account. AWS CloudTrail records the details of every API call made to AWS services, such as the identity of the caller, the time of the call, the source IP address of the caller, the parameters and responses of the call, and more. Users can use AWS CloudTrail to audit, monitor, and troubleshoot their AWS resources and actions. The other options

are incorrect because they are not services that will help a company identify the user who deleted an Amazon EC2 instance yesterday. Amazon CloudWatch is a service that enables users to collect, analyze, and visualize metrics, logs, and events from their AWS resources and applications. AWS Trusted Advisor is a service that provides real-time guidance to help users follow AWS best practices for security, performance, cost optimization, and fault tolerance. Amazon Inspector is a service that helps users find security vulnerabilities and deviations from best practices in their Amazon EC2 instances. Reference: AWS CloudTrail FAQs

Question: 65

Which AWS database service provides in-memory data storage?

A. Amazon DynamoDB
B. Amazon ElastiCache
C. Amazon RDS
D. Amazon Timestream

Answer: B

Explanation:
The correct answer is B because Amazon ElastiCache is a service that provides in-memory data storage. Amazon ElastiCache is a fully managed, scalable, and high-performance service that supports two popular open-source in-memory engines: Redis and Memcached. Amazon ElastiCache allows users to store and retrieve data from fast, low-latency, and high-throughput in-memory systems. Users can use Amazon ElastiCache to improve the performance of their applications by caching frequently accessed data, reducing database load, and enabling real-time data processing. The other options are incorrect because they are not services that provide in-memory data storage. Amazon DynamoDB is a service that provides key-value and document data storage. Amazon RDS is a service that provides relational data storage. Amazon Timestream is a service that provides time series data storage. Reference: Amazon ElastiCache FAQs

Question: 66

A company is using a third-party service to back up 10 TB of data to a tape library. The on-premises backup server is running out of space. The company wants to use AWS services for the backups without changing its existing backup workflows.
Which AWS service should the company use to meet these requirements?

A. Amazon Elastic Block Store (Amazon EBS)
B. AWS Storage Gateway
C. Amazon Elastic Container Service (Amazon ECS)
D. AWS Lambda

Answer: B

Explanation:
The correct answer is B because AWS Storage Gateway is a service that should be used by the company to meet the requirements. AWS Storage Gateway is a service that connects on-premises software applications with cloud-based storage. AWS Storage Gateway supports three types of gateways: file gateway, volume gateway, and tape gateway. The tape gateway type enables users to back up and archive data to virtual tapes in AWS without changing their existing backup workflows. Users can use their existing backup applications and tape libraries to store data on virtual tapes in Amazon S3 or Amazon S3 Glacier. The other options are incorrect because they are not services

that should be used by the company to meet the requirements. Amazon Elastic Block Store (Amazon EBS) is a service that provides block-level storage volumes for Amazon EC2 instances. Amazon Elastic Container Service (Amazon ECS) is a service that enables users to run, scale, and secure containerized applications on AWS. AWS Lambda is a service that enables users to run code without provisioning or managing servers. Reference: AWS Storage Gateway FAQs

Question: 67

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:
The correct answer is D because AWS Enterprise Support is the support plan that provides customers with access to an AWS technical account manager (TAM). AWS Enterprise Support is the highest level of support plan offered by AWS, and it provides customers with the most comprehensive and personalized support experience. An AWS TAM is a dedicated technical resource who works closely with customers to understand their business and technical needs, provide proactive guidance, and coordinate support across AWS teams. The other options are incorrect because they are not support plans that provide customers with access to an AWS TAM. AWS Basic Support is the default and free support plan that provides customers with access to online documentation, forums, and account information. AWS Developer Support is the lowest level of paid support plan that provides customers with access to technical support during business hours, general guidance, and best practice recommendations. AWS Business Support is the intermediate level of paid support plan that provides customers with access to technical support 24/7, system health checks, architectural guidance, and case management. Reference: AWS Support Plans

Question: 68

A company is designing a web application that will run on Amazon EC2 instances.
Which AWS services and features will improve availability and reduce the impact of failures for this application?
(Select TWO.)

A. Amazon EC2 Auto Scaling for the EC2 instances
B. VPC subnet ACLs to check the health of a service
C. Resources that are distributed across multiple Availability Zones
D. Configuration of AWS Server Migration Service (AWS SMS) to move the EC2 instances to a differentAWS Region
E. Resources that are distributed across multiple AWS points of presence

Answer: A,C

Explanation:
The correct answers are A and C because Amazon EC2 Auto Scaling and resources that are distributed across multiple Availability Zones are AWS services and features that will improve availability and reduce the impact of failures for the web application. Amazon EC2 Auto Scaling

is a service that enables users to automatically adjust the number of Amazon EC2 instances in response to changes in demand or performance. Amazon EC2 Auto Scaling helps users to maintain optimal availability and performance of their applications by adding or removing instances as needed. Resources that are distributed across multiple Availability Zones are AWS features that enable users to increase the fault tolerance and resilience of their applications. 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 protect their applications from the failure of a single location. The other options are incorrect because they are not AWS services and features that will improve availability and reduce the impact of failures for the web application. VPC subnet ACLs are AWS features that enable users to control the inbound and outbound traffic to and from their subnets within a VPC. VPC subnet ACLs do not check the health of a service, but rather filter the network traffic based on rules. Configuration of AWS Server Migration Service (AWS SMS) is an AWS service that enables users to migrate their on-premises servers to AWS. Configuration of AWS SMS does not help to move the Amazon EC2 instances to a different AWS Region, but rather to migrate the servers from the source environment to AWS. Resources that are distributed across multiple AWS points of presence are AWS features that enable users to deliver content to their end users with low latency and high performance. AWS points of presence are edge locations that are part of the AWS Global Infrastructure. Users can use services such as Amazon CloudFront and AWS Global Accelerator to distribute their content across multiple AWS points of presence. Reference: Amazon EC2 Auto Scaling, [Regions, Availability Zones, and Local Zones]

Question: 69

An Availability Zone consists of:

A. one or more data centers in a single location.
B. two or more data centers in multiple locations.
C. one or more physical hosts in a single data center.
D. two or more physical hosts in multiple data centers.

Answer: A

Explanation:
The correct answer is A because an Availability Zone consists of one or more data centers in a single location. An Availability Zone is an isolated location within an AWS Region that has independent power, cooling, and networking. Each Availability Zone has one or more data centers that host the physical servers and storage devices that run the AWS services. The other options are incorrect because they are not accurate descriptions of an Availability Zone. Two or more data centers in multiple locations are not an Availability Zone, but rather multiple Availability Zones within an AWS Region. One or more physical hosts in a single data center are not an Availability Zone, butrather the components of a data center within an Availability Zone. Two or more physical hosts in multiple data centers are not an Availability Zone, but rather the components of multiple data centers within one or more Availability Zones. Reference: [Regions, Availability Zones, and Local Zones]

Question: 70

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 anAWS VPN connection.
D. Place both EC2 instances in a placement group for dedicated bandwidth.

Answer: B

Explanation:

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 AvailabilityZones. Reference: [Regions, Availability Zones, and Local Zones], [VPC Peering], [AWS VPN], [Placement Groups]

Question: 71

Which of the following acts as an instance-level firewall to control inbound and outbound access?

A. Network access control list
B. Security groups
C. AWS Trusted Advisor
D. Virtual private gateways

Answer: B

Explanation:

The correct answer is B because security groups are AWS features that act as instance-level firewalls to control inbound and outbound access. Security groups are virtual firewalls that can be attached to one or more Amazon EC2 instances. Users can configure rules for security groups to allow or deny traffic based on protocols, ports, and source or destination IP addresses. The other options are incorrect because they are not AWS features that act as instance-level firewalls to control inbound and outbound access. Network access control list is an AWS feature that acts as a subnet-level firewall to control inbound and outbound access. AWS Trusted Advisor is an AWS service that provides real-time guidance to help users follow AWS best practices for security, performance, cost optimization, and fault tolerance. Virtual private gateways are AWS features that enable users to create a secure and encrypted connection between their VPC and their on-premises network. Reference: Security Groups for Your VPC

Question: 72

A company has an application that uses AWS services. During scaling events, the company wants to keep

application usage within AWS service quotas.

Which AWS services or tools can report on the quotas so that the company can improve the reliability of the application? (Select TWO.)

A. Service Quotas console
B. AWS Trusted Advisor
C. AWS Systems Manager
D. AWS Shield
E. AWS Cost Explorer

Answer: A,B

Explanation:

The correct answers are A and B because Service Quotas console and AWS Trusted Advisor are AWS services or tools that can report on the quotas so that the company can improve the reliability of the application. Service Quotas console is an AWS tool that enables users to view and manage their quotas for AWS services from a central location. Users can use Service Quotas console to request quota increases, track quota usage, and set up alarms for approaching quota limits. AWS Trusted Advisor is an AWS service that provides real-time guidance to help users follow AWS best practices for security, performance, cost optimization, and fault tolerance. One of the categories of checks that AWS Trusted Advisor performs is service limits, which monitors the usage of each AWS service and alerts users when they are close to reaching the default limit. The other options are incorrect because they are not AWS services or tools that can report on the quotas so that the company can improve the reliability of the application. AWS Systems Manager is an AWS service that enables users to automate operational tasks, manage configuration and compliance, and monitor system health and performance. AWS Shield is an AWS service that protects users from distributed denial of service (DDoS) attacks. AWS Cost Explorer is an AWS tool that enables users to visualize, understand, and manage their AWS costs and usage. Reference: Service Quotas, AWS Trusted Advisor FAQs

Question: 73

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:

The correct answers are B and D because making data-driven decisions to determine cloud architectural design and testing systems at production scale are AWS Cloud design principles. Making data-driven decisions to determine cloud architectural design means that users should collect and analyze data from their AWS resources and applications to optimize their performance, availability, security, and cost. Testing systems at production scale means that users should simulate real-world scenarios and load conditions to validate the functionality, reliability, and scalability of their systems. The other options are incorrect because they are not AWS Cloud design principles. Paying for compute resources in advance means that users have to invest heavily in data centers and servers before they know how they will use them. This is not a cloud design principle, but rather a traditional IT model. Emphasizing manual processes to allow for changes means that users have to rely on human intervention and coordination to perform operational tasks and updates. This is not a cloud design principle, but rather a source of inefficiency and error. Refining operational procedures infrequently means that users have to stick to the same

methods andpractices without adapting to the changing needs and feedback. This is not a cloud design principle, but rather a hindrance to innovation and improvement. Reference: AWS Well-Architected Framework

Question: 74

A company needs to migrate all of its development teams to a cloud-based integrated development environment (IDE). Which AWS service should the company use?

A. AWS CodeBuild
B. AWS Cloud9
C. AWS OpsWorks
D. AWS Cloud Development Kit (AWS CDK)

Answer: B

Explanation:
The correct answer is B because AWS Cloud9 is an AWS service that enables users to run their existing custom, nonproduction workloads in the AWS Cloud quickly and cost-effectively. AWS Cloud9 is a cloud-based integrated development environment (IDE) that allows users to write, run, and debug code from a web browser. AWS Cloud9 supports multiple programming languages, such as Python, Java, Node.js, and more. AWS Cloud9 also provides users with a terminal that can access AWS services and resources, such as Amazon EC2 instances, AWS Lambda functions, and AWS CloudFormation stacks. The other options are incorrect because they are not AWS services that enable users to run their existing custom, nonproduction workloads in the AWS Cloud quickly and cost-effectively. AWS CodeBuild is an AWS service that enables users to compile, test, and package their code for deployment. AWS OpsWorks is an AWS service that enables users to configure and manage their applications using Chef or Puppet. AWS Cloud Development Kit (AWS CDK) is an AWS service that enables users to define and provision their cloud infrastructure using familiar programming languages, such as TypeScript, Python, Java, and C#. Reference: AWS Cloud9 FAQs

Question: 75

A company needs to run its existing custom, nonproduction workloads in the AWS Cloud quickly and cost-effectively. The workloads can recover from interruptions easily.
Which pricing model should the company use?

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

Answer: C

Explanation:

The correct answer is C because Spot Instances are the pricing model that enables the company to run its existing custom, nonproduction workloads in the AWS Cloud quickly and cost-effectively. Spot Instances are spare Amazon EC2 instances that are available at up to 90%

discount compared to On- Demand prices. Spot Instances are suitable for stateless, fault-tolerant, and flexible workloads that can recover from interruptions easily. The other options are incorrect because they are not the pricing model that enables the company to run its existing custom, nonproduction workloads in the AWS Cloud quickly and cost-effectively. Reserved Instances are Amazon EC2 instances that are reserved for a specific period of time (one or three years) in exchange for a lower hourly rate. Reserved Instances are suitable for steady-state or predictable workloads that run for a long duration. On-Demand Instances are Amazon EC2 instances that are launched and billed at a fixed hourly rate. On-Demand Instances are suitable for short-term, irregular, or unpredictable workloads that cannot be interrupted. Dedicated Hosts are physical servers that are dedicated to a single customer. Dedicated Hosts are suitable for workloads that require regulatory compliance or data isolation. Reference: Amazon EC2 Instance Purchasing Options

Question: 76

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

Explanation:
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 operatingsystems, and credentials and policies are customer responsibilities according to the AWS shared responsibility model. Reference: [AWS Shared Responsibility Model]

Question: 77

Which options does AWS make available for customers who want to learn about security in the cloud in an instructor-led setting? (Select TWO.)

A. AWS Trusted Advisor
B. AWS Online Tech Talks
C. AWS Blog
D. AWS Forums
E. AWS Classroom Training

Answer: B,E

Explanation:

The correct answers are B and E because AWS Online Tech Talks and AWS Classroom Training are options that AWS makes available for customers who want to learn about security in the cloud in an instructor-led setting. AWS Online Tech Talks are live, online presentations that cover a broad range of topics at varying technical levels. AWS Online Tech Talks are delivered by AWS experts and feature live Q&A sessions with the audience. AWS Classroom Training are in-person or virtual courses that are led by accredited AWS instructors. AWS Classroom Training offer hands-on labs, exercises, and best practices to help customers gain confidence and skills on AWS. The other options are incorrect because they are not options that AWS makes available for customers who want to learn about security in the cloud in an instructor-led setting. AWS Trusted Advisor is an AWS service that provides real-time guidance to help customers follow AWS best practices for security, performance, cost optimization, and fault tolerance. AWS Blog is an AWS resource that provides news, announcements, and insights from AWS experts and customers. AWS Forums are AWS resources that enable customers to interact with other AWS users and get feedback and support. Reference: AWS Online Tech Talks, AWS Classroom Training

Question: 78

A company wants to host its relational databases on AWS. The databases have predefined schemas that the company needs to replicate on AWS.

Which AWS services could the company use for the databases? (Select TWO.)

A. Amazon Aurora
B. Amazon RDS
C. Amazon DocumentDB (with MongoDB compatibility)
D. Amazon Neptune
E. Amazon DynamoDB

Answer: A,B

Explanation:

The correct answers are A and B because Amazon Aurora and Amazon RDS are AWS services that the company could use for the relational databases. Amazon Aurora is a relational database that is compatible with MySQL and PostgreSQL. Amazon Aurora is a fully managed, scalable, and high- performance service that offers up to five times the throughput of standard MySQL and up to three times the throughput of standard PostgreSQL. Amazon RDS is a service that enables users to set up, operate, and scale relational databases in the cloud. Amazon RDS supports six popular database engines: MySQL, PostgreSQL, Oracle, SQL Server, MariaDB, and Amazon Aurora. The other options are incorrect because they are not AWS services that the company could use for the relational databases. Amazon DocumentDB (with MongoDB compatibility) is a document database that is compatible with MongoDB. Amazon Neptune is a graph database that supports property graph and RDF models. Amazon DynamoDB is a key-value and document database. Reference: Amazon Aurora, Amazon RDS

Question: 79

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

Explanation:
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 nottrain 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]

Question: 80

A company needs a content delivery network that provides secure delivery of data, videos, applications, and APIs to users globally with low latency and high transfer speeds.

Which AWS service meets these requirements?

A. Amazon CloudFront
B. Elastic Load Balancing
C. Amazon S3
D. Amazon Elastic Transcoder

Answer: A

Explanation:
The correct answer is A because Amazon CloudFront is an AWS service that provides secure delivery of data, videos, applications, and APIs to users globally with low latency and high transfer speeds. Amazon CloudFront is a fast content delivery network (CDN) that integrates with other AWS services, such as Amazon S3, Amazon EC2, AWS Lambda, and AWS Shield. Amazon CloudFront delivers content through a worldwide network of edge locations that are located close to the end users. The other options are incorrect because they are not AWS services that provide secure delivery of data, videos, applications, and APIs to users globally with low latency and high transfer speeds. Elastic Load Balancing is an AWS service that distributes incoming traffic across multiple targets, such as Amazon EC2 instances, containers, and IP addresses. Amazon S3 is an AWS service that provides object storage for data of any size and type. Amazon Elastic Transcoder is an AWS service that converts media files from their original source format into different formats that will play on various devices.

Reference: Amazon CloudFront FAQs

Question: 81

An application is running on multiple Amazon EC2 instances. The company wants to make the application highly available by configuring a load balancer with requests forwarded to the EC2 instances based on URL paths.
Which AWS load balancer will meet these requirements and take the LEAST amount of effort to deploy?

A. Network Load Balancer
B. Application Load Balancer
C. AWS OpsWorks Load Balancer
D. Custom Load Balancer on Amazon EC2

Answer: B

Explanation:
The correct answer is B because Application Load Balancer is an AWS load balancer that will meet the requirements and take the least amount of effort to deploy. Application Load Balancer is a type of Elastic Load Balancing that operates at the application layer (layer 7) of the OSI model and routes requests to targets based on the content of the request. Application Load Balancer supports advanced features, such as path-based routing, host-based routing, and HTTP header-based routing. The other options are incorrect because they are not AWS load balancers that will meet the requirements and take the least amount of effort to deploy. Network Load Balancer is a type of Elastic Load Balancing that operates at the transport layer (layer 4) of the OSI model and routes requests to targets based on the destination IP address and port. Network Load Balancer does not support path-based routing. AWS OpsWorks Load Balancer is not an AWS load balancer, but rather a feature of AWS OpsWorks that enables users to attach an Elastic Load Balancing load balancer to a layer of their stack. Custom Load Balancer on Amazon EC2 is not an AWS load balancer, but rather a user-defined load balancer that runs on an Amazon EC2 instance. Custom Load Balancer on Amazon EC2 requires more effort to deploy and maintain than an AWS load balancer. Reference: Elastic Load Balancing

Question: 82

A company needs to use dashboards and charts to analyze insights from business data. Which AWS service will provide the dashboards and charts for these insights?

A. Amazon Macie
B. Amazon Aurora
C. Amazon QuickSight
D. AWS CloudTrail

Answer: C

Explanation:
The correct answer is C because Amazon QuickSight is an AWS service that will provide the dashboards and charts for the insights from business data. Amazon QuickSight is a fully managed, scalable, and serverless business intelligence service that enables users to create and share interactive dashboards and charts. Amazon QuickSight can connect to various data sources, such as Amazon S3, Amazon RDS,

Amazon Redshift, and more. Amazon QuickSight also provides users with machine learning insights, such as anomaly detection, forecasting, and natural languagenarratives. The other options are incorrect because they are not AWS services that will provide the dashboards and charts for the insights from business data. Amazon Macie is an AWS service that helps users discover, classify, and protect sensitive data stored in Amazon S3. Amazon Aurora is an AWS service that provides a relational database that is compatible with MySQL and PostgreSQL. AWS CloudTrail is an AWS service that enables users to track user activity and API usage across their AWS account. Reference: Amazon QuickSight FAQs

Question: 83

A large company has a workload that requires hardware to remain on premises. The company wants to use the same management and control plane services that it currently uses on AWS.
Which AWS service should the company use to meet these requirements?

A. AWS Device Farm
B. AWS Fargate
C. AWS Outposts
D. AWS Ground Station

Answer: C

Explanation:
The correct answer is C because AWS Outposts is an AWS service that enables the company to meet the requirements. AWS Outposts is a fully managed service that extends AWS infrastructure, services, APIs, and tools to virtually any datacenter, co-location space, or on-premises facility. AWS Outposts allows customers to run their workloads on the same hardware and software that AWS uses in its cloud, while maintaining local access and control. The other options are incorrect because they are not AWS services that enable the company to meet the requirements. AWS Device Farm is an AWS service that enables customers to test their mobile and web applications on real devices in the AWS Cloud. AWS Fargate is an AWS service that enables customers to run containers without having to manage servers or clusters. AWS Ground Station is an AWS service that enables customers to communicate with satellites and downlink data from orbit. Reference: AWS Outposts FAQs

Question: 84

When a user wants to utilize their existing per-socket, per-core, or per-virtual machine software licenses for a Microsoft Windows server running on AWS, which Amazon EC2 instance type is required?

A. Spot Instances
B. Dedicated Instances
C. Dedicated Hosts
D. Reserved Instances

Answer: C

Explanation:
The correct answer is C because Dedicated Hosts are Amazon EC2 instances that are required when a user wants to utilize their existing per-socket, per-core, or per-virtual machine software licenses for a Microsoft Windows server running on AWS. Dedicated Hosts are physical servers that are dedicated to a single customer. Dedicated Hosts allow customers to use their existing server-bound software

licenses, such as Windows Server, SQL Server, and SUSE Linux Enterprise Server, subject to their license terms. The other options are incorrect because they are not Amazon EC2 instances that are required when a user wants to utilize their existing per-socket, per-core, or per-virtual machine software licenses for a Microsoft Windows server running on AWS. Spot Instances are spare Amazon EC2 instances that are available at up to 90% discount compared to On-Demand prices. Spot Instances are suitable for stateless, fault-tolerant, and flexible workloads that can recover from interruptions easily. Dedicated Instances are Amazon EC2 instances that run on hardware that is dedicated to a single customer, but not to a specific physical server. Dedicated Instances do not allow customers to use their existing server-bound software licenses. Reserved Instances are Amazon EC2 instances that are reserved for a specific period of time (one or three years) in exchange for a lower hourly rate. Reserved Instances are suitable for steady-state or predictable workloads that run for a long duration. Reserved Instances do not allow customers to use their existing server-bound software licenses. Reference: Dedicated Hosts, Amazon EC2 Instance Purchasing Options

Question: 85

Which AWS service should a cloud engineer use to view API calls to AWS services?

A. Amazon CloudWatch
B. AWS CloudTrail
C. AWS Config
D. AWS Artifact

Answer: B

Explanation:
The correct answer is B because AWS CloudTrail is an AWS service that a cloud engineer can use to view API calls to AWS services. AWS CloudTrail is a service that enables customers to track user activity and API usage across their AWS account. AWS CloudTrail records the details of every API call made to AWS services, such as the identity of the caller, the time of the call, the source IP address of the caller, the parameters and responses of the call, and more. Customers can use AWS CloudTrail to audit, monitor, and troubleshoot their AWS resources and actions. The other options are incorrect because they are not AWS services that a cloud engineer can use to view API calls to AWS services. Amazon CloudWatch is an AWS service that enables customers to collect, analyze, and visualize metrics, logs, and events from their AWS resources and applications. AWS Config is an AWS service that enables customers to assess, audit, and evaluate the configurations of their AWS resources. AWS Artifact is an AWS service that provides customers with on-demand access to AWS compliance reports and select online agreements. Reference: AWS CloudTrail FAQs

Question: 86

What can a user accomplish using AWS CloudTrail?

A. Generate an IAM user credentials report.
B. Record API calls made to AWS services.
C. Assess the compliance of AWS resource configurations with policies and guidelines.
D. Ensure that Amazon EC2 instances are patched with the latest security updates.A company uses Amazon Workspaces.

Answer: B

Explanation:

AWS CloudTrail is an AWS service that enables users to accomplish the task of recording API calls made to AWS services. AWS CloudTrail is a service that tracks user activity and API usage across the AWS account. AWS CloudTrail records the details of every API call made to AWS services, such as the identity of the caller, the time of the call, the source IP address of the caller, the parameters and responses of the call, and more. Users can use AWS CloudTrail to audit, monitor, and troubleshoot their AWS resources and actions. The other options are incorrect because they are not tasks that users can accomplish using AWS CloudTrail. Generating an IAM user credentials report is a task that users can accomplish using IAM, which is an AWS service that enables users to manage access and permissions to AWS resources and services. Assessing the compliance of AWS resource configurations with policies and guidelines is a task that users can accomplish using AWS Config, which is an AWS service that enables users to assess, audit, and evaluate the configurations of their AWS resources. Ensuring that Amazon EC2 instances are patched with the latest security updates is a task that users can accomplish using AWS Systems Manager, which is an AWS service that enables users to automate operational tasks, manage configuration and compliance, and monitor system health and performance. Reference: AWS CloudTrail FAQs

Question: 87

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

A. Set up multi-factor authentication (MFA) for each Workspaces user account.
B. Ensure the environmental safety and security of the AWS infrastructure that hosts Workspaces.
C. Provide security for Workspaces user accounts through AWS Identity and Access Management(IAM).
D. Configure AWS CloudTrail to log API calls and user activity.A company stores data in an Amazon S3 bucket. The company must control who has permission to read, write,or delete objects that the company stores in the S3 bucket.

Answer: B

Explanation:
The correct answer is B because ensuring the environmental safety and security of the AWS infrastructure that hosts Workspaces is the responsibility of AWS, 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 the responsibility of the customer, according to the AWS shared responsibility model. Setting up multi-factor authentication (MFA) for each Workspaces user account, providing security for Workspaces user accounts through AWS Identity and Access Management (IAM), configuring AWS CloudTrail to log API calls and user activity, and encrypting data at rest and in transit are all tasks that the customer has to perform to secure their Workspaces environment. Reference: AWS Shared Responsibility Model, Amazon WorkSpaces Security

Question: 88

Which AWS features will meet these requirements? (Select TWO.)

A. Security groups
B. Network ACLs
C. S3 bucket policies
D. IAM user policies
E. S3 bucket versioning

Answer: C,D

Explanation:

The correct answers are C and D because S3 bucket policies and IAM user policies are AWS features that will meet the requirements. S3 bucket policies are access policies that can be attached to Amazon S3 buckets to grant or deny permissions to the bucket and the objects it contains. S3 bucket policies can be used to control who has permission to read, write, or delete objects that the company stores in the S3 bucket. IAM user policies are access policies that can be attached to IAM users to grant or deny permissions to AWS resources and actions. IAM user policies can be used to control who has permission to read, write, or delete objects that the company stores in the S3 bucket. The other options are incorrect because they are not AWS features that will meet the requirements. Security groups and network ACLs are AWS features that act as firewalls to control inbound and outbound traffic to and from Amazon EC2 instances and subnets. Security groups andnetwork ACLs do not control who has permission to read, write, or delete objects that the company stores in the S3 bucket. S3 bucket versioning is an AWS feature that enables users to keep multiple versions of the same object in the same bucket. S3 bucket versioning can be used to recover from accidental overwrites or deletions of objects, but it does not control who has permission to read, write, or delete objects that the company stores in the S3 bucket. Reference: Using Bucket Policies and User Policies, Security Groups for Your VPC, Network ACLs, [Using Versioning]

Question: 89

Which of the following is a recommended design principle of the AWS Well-Architected Framework?

A. Reduce downtime by making infrastructure changes infrequently and in large increments.
B. Invest the time to configure infrastructure manually.
C. Learn to improve from operational failures.
D. Use monolithic application design for centralization.

Answer: C

Explanation:

The correct answer is C because learning to improve from operational failures is a recommended design principle of the AWS Well-Architected Framework. The AWS Well-Architected Framework is a set of best practices and guidelines for designing and operating reliable, secure, efficient, and costeffective systems in the cloud. The AWS Well-Architected Framework consists of five pillars: operational excellence, security, reliability, performance efficiency, and cost optimization. Each pillar has a set of design principles that describe the characteristics of a well-architected system. Learning to improve from operational failures is a design principle of the operational excellence pillar, which focuses on running and monitoring systems to deliver business value and continually improve supporting processes and procedures. The other options are incorrect because they are not recommended design principles of the AWS Well-Architected Framework. Reducing downtime by making infrastructure changes infrequently and in large increments is not a design principle of the AWS Well-Architected Framework, but rather a source of risk and inefficiency. A well-architected system should implement changes frequently and in small increments to minimize the impact and scope of failures. Investing the time to configure infrastructure manually is not a design principle of the AWS Well-Architected Framework, but rather a source of human error and inconsistency. A well- architected system should automate manual tasks to improve the speed and accuracy of operations. Using monolithic application design for centralization is not a design principle of the AWS Well- Architected Framework, but rather a source of complexity and rigidity. A well-architected system should use loosely coupled and distributed components to enable scalability and resilience. Reference: [AWS Well-Architected Framework]

Question: 90

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

Explanation:
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

Question: 91

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:
The correct answers are A and C because patching AWS network devices and providing physical security for compute resources are tasks that are the responsibility of AWS, 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 tasks that are the responsibility of the customer, according to the AWS shared responsibility model. Setting user password rules, configuring security groups, and patching the operating system of an Amazon EC2 instance are all tasks that the customer has to perform to secure

their AWS environment. Reference: AWS Shared Responsibility Model

Question: 92

Which AWS service or feature captures information about the network traffic to and from an Amazon EC2 instance?

A. VPC Reachability Analyzer
B. Amazon Athena
C. VPC Flow Logs
D. AWS X-Ray

Answer: C

Explanation:
The correct answer is C because VPC Flow Logs is an AWS service or feature that captures information about the network traffic to and from an Amazon EC2 instance. VPC Flow Logs is a feature that enables customers to capture information about the IP traffic going to and from network interfaces in their VPC. VPC Flow Logs can help customers to monitor and troubleshoot connectivity issues, such as traffic not reaching an instance or traffic being rejected by a security group. The other options are incorrect because they are not AWS services or features that capture information about the network traffic to and from an Amazon EC2 instance. VPC Reachability Analyzer is an AWS service or feature that enables customers to perform connectivity testing between resources in their VPC and identify configuration issues that prevent connectivity. Amazon Athena is an AWS service that enables customers to query data stored in Amazon S3 using standard SQL. AWS X-Ray is an AWS service that enables customers to analyze and debug distributed applications, such as those built using a microservices architecture. Reference: VPC Flow Logs

Question: 93

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

A. Availability
B. Reliability
C. Scalability
D. Responsive design
E. Operational excellence

Answer: B,E

Explanation:
The correct answers to the questions are B and E because reliability and operational excellence are pillars of the AWS Well-Architected Framework. 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 AWS Well-Architected Framework consists of five pillars: operational excellence, security, reliability, performance efficiency, and cost optimization. Each pillar has a set of design principles that describe the characteristics of a well-architected system. Reliability is the pillar that focuses on the ability of a system to recover from failures and meet business and customer demand. Operational excellence is the pillar that focuses on the ability of a system to run and monitor processes that support business outcomes and continually improve. The other options are incorrect because they are not pillars of the AWS Well-Architected Framework. Availability, scalability, and responsive design are important aspects of cloud architecture, but they are not separate pillars in

the framework. Availability and scalability are related to the reliability and performance efficiency pillars, while responsive design is related to the customer experience and user interface. Reference: AWS Well-Architected Framework

Question: 94

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

A. Determine application dependencies with operating systems.
B. Provide user access with AWS Identity and Access Management (IAM).
C. Secure the data center in an Availability Zone.
D. Patch the hypervisor.
E. Provide network availability in Availability Zones.

Answer: B

Explanation:
The correct answer to the question is B because providing user access with AWS Identity and Access Management (IAM) is a customer responsibility according to the AWS shared responsibilitymodel. 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. IAM is an AWS service that enables customers to manage access and permissions to AWS resources and services. Customers are responsible for creating and managing IAM users, groups, roles, and policies, and ensuring that they follow the principle of least privilege. Reference: AWS Shared Responsibility Model

Question: 95

A user wants to identify any security group that is allowing unrestricted incoming SSH traffic. Which AWS service can be used to accomplish this goal?

A. Amazon Cognito
B. AWS Shield
C. Amazon Macie
D. AWS Trusted Advisor

Answer: D

Explanation:
The correct answer to the question is D because AWS Trusted Advisor is an AWS service that can be used to accomplish the goal of identifying any security group that is allowing unrestricted incoming SSH traffic. AWS Trusted Advisor is a service that provides customers with recommendations that help them follow AWS best practices. Trusted Advisor evaluates the customer's AWS environment and identifies ways to optimize their AWS infrastructure, improve security and performance, reduce costs, and monitor service quotas. One of

the checks that Trusted Advisor performs is the Security Groups - Specific Ports Unrestricted check, which flags security groups that allow unrestricted access to specific ports, such as port 22 for SSH. Customers can use this check to review and modify their security group rules to restrict SSH access to only authorized sources. Reference: Security Groups - Specific Ports Unrestricted

Question: 96 Which AWS feature or resource is a deployable Amazon EC2 instance template that is prepackaged with software and security requirements?

A. Amazon Elastic Block Store (Amazon EBS) volume
B. AWS CloudFormation template
C. Amazon Elastic Block Store (Amazon EBS) snapshot
D. Amazon Machine Image (AMI)

Answer: D

Explanation:
An Amazon Machine Image (AMI) is a deployable Amazon EC2 instance template that is prepackaged with software and security requirements. It provides the information required to launch an instance, which is a virtual server in the cloud. You can use an AMI to launch as many instances as you need. You can also create your own custom AMIs or use AMIs shared by other AWS users1.

Question: 97

Which AWS service is a highly available and scalable DNS web service?

A. Amazon VPC
B. Amazon CloudFront
C. Amazon Route 53
D. Amazon Connect

Answer: C

Explanation:
Amazon Route 53 is a highly available and scalable DNS web service. It is designed to give developers and businesses an extremely reliable and cost-effective way to route end users to Internet applications by translating domain names into the numeric IP addresses that computers use to connect to each other2. Amazon Route 53 also offers other features such as health checks, traffic management, domain name registration, and DNSSEC3.

Question: 98

Which of the following is a characteristic of the AWS account root user?

A. The root user is the only user that can be configured with multi-factor authentication (MFA).
B. The root user is the only user that can access the AWS Management Console.
C. The root user is the first sign-in identity that is available when an AWS account is created.
D. The root user has a password that cannot be changed.

Answer: C

Explanation:

The AWS account root user is the first sign-in identity that is available when an AWS account is created. It has complete access to all AWS services and resources in the account. The root user email address and password are the same credentials that are used to sign in to the AWS Management Console4. The root user should be used only to perform a few account and servicemanagement tasks. For day-to-day tasks, it is recommended to use AWS Identity and Access Management (IAM) users or roles instead.

Question: 99

Which AWS service provides the ability to host a NoSQL database in the AWS Cloud?

A. Amazon Aurora
B. Amazon DynamoDB
C. Amazon RDS
D. Amazon Redshift

Answer: B

Explanation:

Amazon DynamoDB is a fully managed NoSQL database service that provides fast and predictable performance with seamless scalability. It supports both key-value and document data models, and allows you to create tables that can store and retrieve any amount of data, and serve any level of request traffic. You can also use DynamoDB Streams to capture data modification events in DynamoDB tables.

Question: 100

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

A. WOMB
B. 5 GB
C. 5 TB
D. Unlimited

Answer: D

Explanation:

Amazon S3 offers unlimited storage for any amount of data. You can store as many objects as you want, and each object can be as large as 5 terabytes. You pay only for the storage space that you actually use, and there are no minimum commitments or upfront fees. Amazon S3 also provides high durability, availability, scalability, and security for your data.

Question: 101

Which AWS network services or features allow Cl DR block notation when providing an IP address range?

(Select TWO.)

A. Security groups
B. Amazon Machine Image (AMI)
C. Network access control list (network ACL)
D. AWS Budgets
E. Amazon Elastic Block Store (Amazon EBS)

Answer: A,C

Explanation:
Security groups and network access control lists (network ACLs) are two AWS network services or features that allow CIDR block notation when providing an IP address range. Security groups act as a firewall for associated Amazon EC2 instances, controlling both inbound and outbound traffic at the instance level. Network ACLs act as a firewall for associated subnets, controlling both inbound and outbound traffic at the subnet level. Both security groups and network ACLs use CIDR block notation to specify the IP address ranges that are allowed or denied

Question: 102

A company has a workload that requires data to be collected, analyzed, and stored on premises. The company wants to extend the use of AWS services to run on premises with access to the company network and the company's VPC.
Which AWS service meets this requirement?

A. AWS Outposts
B. AWS Storage Gateway
C. AWS Direct Connect
D. AWS Snowball

Answer: A

Explanation:
AWS Outposts is an AWS service that meets the requirement of running AWS services on premises with access to the company network and the company's VPC. AWS Outposts is a fully managed service that extends AWS infrastructure, AWS services, APIs, and tools to virtually any datacenter, colocation space, or on-premises facility for a truly consistent hybrid experience. AWS Outposts is ideal for workloads that require low latency access to on-premises systems, local data processing, or local data storage2.

Question: 103

A company wants to deploy and manage a Docker-based application on AWS.
Which solution meets these requirements with the LEAST amount of operational overhead?

A. An open-source Docker orchestrator on Amazon EC2 instances
B. AWS AppSync
C. Amazon Elastic Container Registry (Amazon ECR)
D. Amazon Elastic Container Service (Amazon ECS)

Answer: D

Explanation:
Amazon Elastic Container Service (Amazon ECS) is a solution that meets the requirements of deploying and managing a Docker-based application on AWS with the least amount of operational overhead. Amazon ECS is a fully managed container orchestration service that makes it easy to run, scale, and secure Docker container applications on AWS. Amazon ECS eliminates the need for you to install, operate, and scale your own cluster management infrastructure. With simple API calls, you can launch and stop container-enabled applications, query the complete state of your cluster, and access many familiar features likesecurity groups, Elastic Load Balancing, EBS volumes, and IAM roles3.

Question: 104

When designing AWS workloads to be operational even when there are component failures, what is an AWS best practice?

A. Perform quarterly disaster recovery tests.
B. Place the main component on the us-east-1 Region.
C. Design for automatic failover to healthy resources.
D. Design workloads to fit on a single Amazon EC2 instance.

Answer: C

Explanation:
Designing for automatic failover to healthy resources is an AWS best practice when designing AWS workloads to be operational even when there are component failures. This means that you should architect your system to handle the loss of one or more components without impacting the availability or performance of your application. You can use various AWS services and features to achieve this, such as Auto Scaling, Elastic Load Balancing, Amazon Route 53, Amazon CloudFormation, and AWS CloudFormation4.

Question: 105

Which AWS service provides highly durable object storage?

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

Answer: A

Explanation:
Amazon S3 is the AWS service that provides highly durable object storage. Amazon S3 is designed to provide 99.999999999% durability of objects over a given year. This means that you can store your data with high confidence that it will not be lost. Amazon S3 also provides high availability, scalability, security, and performance for your data. You can use Amazon S3 to store and retrieve any amount of data, at any time, from anywhere on the web5.

Question: 106

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

Explanation:
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.

Question: 107

Who enables encryption of data at rest for Amazon Elastic Block Store (Amazon EBS)?

A. AWS Support
B. AWS customers
C. AWS Key Management Service (AWS KMS)
D. AWS Trusted Advisor

Answer: B

Explanation:
AWS customers are responsible for enabling encryption of data at rest for Amazon Elastic Block Store (Amazon EBS). Amazon EBS encryption offers a simple encryption solution for your EBSvolumes that does not require you to build, maintain, and secure your own key management infrastructure. You can encrypt both the boot and data volumes of your EC2 instances. You can use AWS Key Management Service (AWS KMS) customer master keys (CMKs) or your own CMKs to encrypt your volumes2.

Question: 108

Who is responsible for decommissioning end-of-life underlying storage devices that are used to host data on AWS?

A. Customer
B. AWS
C. Account creator
D. Auditing team

Answer: B

Explanation:
AWS is responsible for decommissioning end-of-life underlying storage devices that are used to host data on AWS. AWS follows strict and audited data destruction processes to ensure that customer data is not exposed to unauthorized individuals or devices when an AWS storage device reaches the end of its useful life. AWS uses techniques detailed in DoD 5220.22-M ("National Industrial Security Program Operating Manual") or NIST 800-88 ("Guidelines for Media Sanitization") to destroy data as part of the decommissioning process3.

Question: 109

A company wants to manage access and permissions for its third-party software as a service (SaaS) applications. The company wants to use a portal where end users can access assigned AWS accounts and AWS Cloud applications.
Which AWS service should the company use to meet these requirements?

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

Answer: B

Explanation:
AWS IAM Identity Center (AWS Single Sign-On) is the AWS service that the company should use to meet the requirements of managing access and permissions for its third-party SaaS applications. AWS Single Sign-On is a cloud-based service that makes it easy to centrally manage single sign- on(SSO) access to multiple AWS accounts and business applications. You can use AWS Single Sign-On to enable your users to sign in to a user portal with their existing corporate credentials and access all of their assigned accounts and applications from one place4.

Question: 110

A large company wants to track the combined AWS usage costs of all of its linked accounts. How can this be accomplished?

A. Use AWS Trusted Advisor to generate customized summary reports.
B. Use AWS Organizations to generate consolidated billing reports.
C. Use AWS Budgets to set utilization targets and receive summary reports.
D. Use the AWS Control Tower dashboard to get a summary report of all linked account costs.

Answer: B

Explanation:
The company can use AWS Organizations to track the combined AWS usage costs of all of its linked accounts. AWS Organizations is a service that enables you to consolidate multiple AWS accounts into an organization that you can manage centrally. You can use AWS Organizations to create a consolidated billing report that shows the charges incurred by each account in your organization as well as the total charges across all accounts. You can also use AWS Organizations to apply policies and controls to your accounts to help you manage costs and security5.

Question: 111

A company wants its Amazon EC2 instances to operate in a highly available environment, even if there is a natural disaster in a particular geographic area.

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 edge locations.
D. Use Amazon CloudFront with the EC2 instances configured as the source.

Answer: B

Explanation:

To achieve high availability in the event of a natural disaster, the company should use EC2 instances in multiple AWS Regions. AWS Regions are geographically isolated areas that consist of multiple Availability Zones. Availability Zones are physically separate locations within an AWS Region that areengineered to be isolated from failures. By using EC2 instances in multiple AWS Regions, the company can ensure that its applications can continue to run even if one Region is affected by a disaster. AWS Global InfrastructureAWS Well-Architected Framework

Question: 112

Using Amazon Elastic Container Service (Amazon ECS) to break down a monolithic architecture into microservices is an example of:

A. a loosely coupled architecture.
B. a tightly coupled architecture.
C. a stateless architecture.
D. a stateful architecture.

Answer: A

Explanation:

Using Amazon Elastic Container Service (Amazon ECS) to break down a monolithic architecture into microservices is an example of a loosely coupled architecture. A loosely coupled architecture is one where the components are independent and can communicate with each other through well-defined interfaces. This allows for greater scalability, flexibility, and resilience. A tightly coupled architecture is one where the components are interdependent and rely on each other for functionality. This can lead to increased complexity, fragility, and difficulty in changing or scaling the system. Amazon ECS OverviewAWS Well-Architected Framework

Question: 113

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

A. Build architectures with tightly coupled resources.
B. Use AWS Trusted Advisor to meet security best practices.

C. Use automation to recover immediately from failure.
D. Rightsize Amazon EC2 instances to ensure optimal performance.
E. Simulate failures to test recovery processes.

Answer: C,E

Explanation:
The design principles for reliability in the AWS Cloud are:
Test recovery procedures. The best way to ensure that systems can recover from failures is to regularly test them using simulated scenarios. This can help identify gaps and improve the recovery process.
Automatically recover from failure. By using automation, systems can detect and correct failures without human intervention. This can reduce the impact and duration of failures and improve the availability of the system.
Scale horizontally to increase aggregate system availability. By adding more redundant resources to the system, the impact of individual resource failures can be reduced. This can also improve the performance and scalability of the system.
Stop guessing capacity. By using monitoring and automation, systems can adjust the capacity based on the demand and performance metrics. This can prevent failures due to insufficient or excessive capacity and optimize the cost and efficiency of the system.
Manage change in automation. By using automation, changes to the system can be applied in a consistent and controlled manner. This can reduce the risk of human errors and configuration drifts that can cause failures. AWS Well-Architected Framework

Question: 114

Which statements represent the cost-effectiveness of the AWS Cloud? (Select TWO.)

A. Users can trade fixed expenses for variable expenses.
B. Users can deploy all over the world in minutes.
C. AWS offers increased speed and agility.
D. AWS is responsible for patching the infrastructure.
E. Users benefit from economies of scale.

Answer: D,E

Explanation:
The statements that represent the cost-effectiveness of the AWS Cloud are:
Users can trade fixed expenses for variable expenses. By using the AWS Cloud, users can pay only for the resources they use, instead of investing in fixed and upfront costs for hardware and software. This can lower the total cost of ownership and increase the return on investment.
Users benefit from economies of scale. By using the AWS Cloud, users can leverage the massive scale and efficiency of AWS to access lower prices and higher performance. AWS passes the cost savings to the users through price reductions and innovations. AWS Cloud Value Framework

Question: 115

A company wants to migrate its on-premises data warehouse to AWS. The information in the data warehouse is used to populate analytics dashboards.
Which AWS service should the company use for the data warehouse?

A. Amazon ElastiCache
B. Amazon Aurora
C. Amazon RDS
D. Amazon Redshift

Answer: D

Explanation:

The AWS service that the company should use for the data warehouse is Amazon Redshift. Amazon Redshift is a fully managed, petabyte-scale data warehouse service that is optimized for analytical queries. It can integrate with various data sources and business intelligence tools to provide fast and cost-effective insights. Amazon Redshift also offers high availability, scalability, security, and compliance features. [Amazon Redshift Overview]

Question: 116

Which benefit does Amazon Rekognition provide?

A. The ability to place watermarks on images
B. The ability to detect objects that appear in pictures
C. The ability to resize millions of images automatically
D. The ability to bid on object detection jobs

Answer: B

Explanation:

Amazon Rekognition is a service that provides deep learning-based image and video analysis. One of the benefits of Amazon Rekognition is the ability to detect objects that appear in pictures, such as faces, landmarks, animals, text, and scenes. This can enable applications to perform tasks such as face recognition, face verification, face comparison, face search, celebrity recognition, emotion detection, age range estimation, gender identification, facial analysis, facial expression recognition, and more. Amazon Rekognition OverviewAWS Certified Cloud Practitioner - aws.amazon.com

Question: 117

Which AWS service uses a combination of publishers and subscribers?

A. AWS Lambda
B. Amazon Simple Notification Service (Amazon SNS)
C. Amazon CloudWatch
D. AWS CloudFormation

Answer: B

Explanation:

Amazon Simple Notification Service (Amazon SNS) is a service that provides fully managed pub/sub messaging. Pub/sub messaging is a pattern that uses a combination of publishers and subscribers. Publishers are entities that produce messages and send them to topics. Subscribers are entities that receive messages from topics. Topics are logical access points that act as communication channels between publishers and subscribers. Amazon SNS enables applications to decouple, scale, and coordinate the delivery of messages to multiple endpoints, such as email, SMS, mobile push notifications, Lambda functions, SQS queues, and HTTP/S endpoints. Amazon SNS OverviewAWS Certified Cloud Practitioner - aws.amazon.com

Question: 118

A company is developing an application that uses multiple AWS services. The application needs to use temporary, limited-privilege credentials for authentication with other AWS APIs.
Which AWS service or feature should the company use to meet these authentication requirements?

A. Amazon API Gateway
B. IAM users
C. AWS Security Token Service (AWS STS)
D. IAM instance profiles

Answer: C

Explanation:
AWS Security Token Service (AWS STS) is a service that enables applications to request temporary, limited-privilege credentials for authentication with other AWS APIs. AWS STS can be used to grant access to AWS resources to users who are federated (using IAM roles), switched (using IAM users), or cross-account (using IAM roles). AWS STS can also be used to assume a role within the same account or a different account. The credentials issued by AWS STS are short-term and have a limited scope, which can enhance the security and compliance of the application. AWS STS OverviewAWS Certified Cloud Practitioner - aws.amazon.com

Question: 119

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 RDS
C. Amazon DocumentDB (with MongoDB compatibility)

Answer: C

Explanation:
Amazon Relational Database Service (Amazon RDS) is a service that provides fully managed relational database engines. Amazon RDS supports several database engines, including Oracle, MySQL, PostgreSQL, MariaDB, SQL Server, and Amazon Aurora. Amazon RDS can be used to migrate an application that includes an Oracle database to AWS without rewriting the application, as long as the application is compatible with the Oracle version and edition supported by Amazon RDS. Amazon RDS can also provide benefits such as high availability, scalability, security, backup and restore, and performance optimization. [Amazon RDS Overview] AWS Certified Cloud Practitioner - aws.amazon.com

Question: 120

Which of the following is an AWS value proposition that describes a user's ability to scale infrastructure based on demand?

A. Speed of innovation
B. Resource elasticity
C. Decoupled architecture
D. Global deployment

Answer: B

Explanation:
Resource elasticity is an AWS value proposition that describes a user's ability to scale infrastructure based on demand. Resource elasticity means that the user can provision or deprovision resources quickly and easily, without any upfront commitment or long-term contract. Resource elasticity can help the user optimize the cost and performance of the application, as well as respond to changing business needs and customer expectations. Resource elasticity can be achieved by using services such as Amazon EC2, Amazon S3, Amazon RDS, Amazon DynamoDB, Amazon ECS, and AWS Lambda. [AWS Cloud Value Framework] AWS Certified Cloud Practitioner - aws.amazon.com

Question: 121

A company needs to continuously monitor its environment to analyze network and account activity and identify potential security threats.
Which AWS service should the company use to meet these requirements?

A. AWS Artifact
B. Amazon Macie
C. AWS Identity and Access Management (IAM)
D. Amazon GuardDuty

Answer: D

Explanation:
Amazon GuardDuty is a service that provides intelligent threat detection and continuous monitoring for the AWS environment. It analyzes network and account activity using machine learning and threat intelligence to identify potential security threats, such as unauthorized access, compromised credentials, malicious hosts, and reconnaissance activities. It also generates detailed and actionable findings that can be viewed on the AWS Management Console or sent to other AWS services, such as Amazon CloudWatch Events and AWS Lambda, for further analysis or remediation. Amazon GuardDuty OverviewAWS Certified Cloud Practitioner - aws.amazon.com

Question: 122

Which AWS service can report how AWS resource configurations have changed over time?

A. AWS CloudTrail
B. Amazon CloudWatch
C. AWS Config

D. Amazon Inspector

Answer: C

Explanation:
AWS Config is a service that enables users to assess, audit, and evaluate the configurations of AWS resources. It continuously monitors and records the configuration changes of the resources and evaluates them against desired configurations and best practices. It also provides a detailed view of the resource configuration history and relationships, as well as compliance reports and notifications. AWS Config can help users maintain consistent and secure configurations, troubleshoot issues, and simplify compliance auditing. AWS Config OverviewAWS Certified Cloud Practitioner - aws.amazon.com

Question: 123

Which AWS benefit is demonstrated by on-demand technology services that enable companies to replace upfront fixed expenses with variable expenses?

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

Answer: B

Explanation:
Pay-as-you-go pricing is an AWS benefit that demonstrates the ability of users to replace upfront fixed expenses with variable expenses. With pay-as-you-go pricing, users only pay for the resources they consume, without any long-term contracts or commitments. This can lower the total cost of ownership and increase the return on investment. Pay-as-you-go pricing also provides flexibility and scalability, as users can adjust their resource usage according to their changing needs and demands. AWS Cloud Value FrameworkAWS Certified Cloud Practitioner - aws.amazon.com

Question: 124

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

Explanation:

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

Question: 125

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

A. AWS Elastic Beanstalk
B. AWS Config
C. AWS OpsWorksQ D. AWS Application Discovery Service
D. Amazon Kinesis

Answer: A,C

Explanation:
The services that can be used to deploy applications on AWS are:
AWS Elastic Beanstalk. This is a service that simplifies the deployment and management of web applications on AWS. Users can upload their application code and Elastic Beanstalk automatically handles the provisioning, scaling, load balancing, monitoring, and health checking of the resources needed to run the application. Users can also retain full control and access to the underlying resources and customize their configuration settings. Elastic Beanstalk supports multiple platforms, such as Java, .NET, PHP, Node.js, Python, Ruby, Go, and Docker. [AWS Elastic Beanstalk
Overview] AWS Certified Cloud Practitioner - aws.amazon.com
AWS OpsWorks. This is a service that provides configuration management and automation for AWS resources. Users can define the application architecture and the configuration of each resource using Chef or Puppet, which are popular open-source automation platforms. OpsWorks then automatically creates and configures the resources according to the user's specifications. OpsWorks also provides features such as auto scaling, monitoring, and integration with other AWS services. OpsWorks has two offerings: OpsWorks for Chef Automate and OpsWorks for Puppet Enterprise. [AWS OpsWorks Overview] AWS Certified Cloud Practitioner - aws.amazon.com

Question: 126

Which statement describes a characteristic of the AWS global infrastructure?

A. Edge locations contain multiple AWS Regions.
B. AWS Regions contain multiple Regional edge caches.
C. Availability Zones contain multiple data centers.
D. Each data center contains multiple edge locations.

Answer: C

Explanation:

Availability Zones contain multiple data centers. This is a characteristic of the AWS global infrastructure, which consists of AWS Regions, Availability Zones, and edge locations. AWS Regions are geographically isolated areas that contain multiple Availability Zones. Availability Zones are physically separate locations within an AWS Region that are engineered to be isolated from failures and connected by low-latency, high-throughput, and highly redundant networking. Each Availability Zone contains one or more data centers that house the servers and storage devices that run AWS services. Edge locations are sites that are located closer to the end users and provide caching and content delivery services. AWS Global InfrastructureAWS Certified Cloud Practitioner - aws.amazon.com

Question: 127

Which of the following is available to a company that has an AWS Business Support plan?

A. AWS Support concierge
B. AWS DDoS Response Team (DRT)
C. AWS technical account manager (TAM)
D. AWS Health API

Answer: D

Explanation:

AWS Health API is available to a company that has an AWS Business Support plan. The AWS Health API provides programmatic access to the AWS Health information that is presented in the AWS

Personal Health Dashboard. The AWS Health API can help users get timely and personalized information about events that can affect the availability and performance of their AWS resources, such as scheduled maintenance, network issues, or service disruptions. The AWS Health API can also integrate with other AWS services, such as Amazon CloudWatch Events and AWS Lambda, to enable automated actions and notifications. AWS Health API OverviewAWS Support Plans

Question: 128

Which pillar of the AWS Well-Architected Framework focuses on the return on investment of moving into the AWS Cloud?

A. Sustainability
B. Cost optimization
C. Operational excellence
D. Reliability

Answer: B

Explanation:

Cost optimization is the pillar of the AWS Well-Architected Framework that focuses on the return on investment of moving into the AWS Cloud. Cost optimization means that users can achieve the desired business outcomes at the lowest possible price point, while maintaining

high performance and reliability. Cost optimization can be achieved by using various AWS features and best practices, such as pay-as-you-go pricing, right-sizing, elasticity, reserved instances, spot instances, cost allocation tags, cost and usage reports, and AWS Trusted Advisor. [AWS Well-Architected Framework] AWS Certified Cloud Practitioner - aws.amazon.com

Question: 129

Which AWS service or feature offers HTTP attack protection to users running public-facing web applications?

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

Answer: D

Explanation:

AWS WAF is the AWS service or feature that offers HTTP attack protection to users running publicfacing web applications. AWS WAF is a web application firewall that helps users protect their web applications from common web exploits, such as SQL injection, cross-site scripting, and bot attacks. Users can create custom rules to define the web traffic that they want to allow, block, or count. Users can also use AWS Managed Rules, which are pre-configured rules that are curated and maintained by

AWS or AWS Marketplace Sellers. AWS WAF can be integrated with other AWS services, such as Amazon CloudFront, Amazon API Gateway, and Application Load Balancer, to provide comprehensive security for web applications. [AWS WAF Overview] AWS Certified Cloud Practitioner - aws.amazon.com

Question: 130

What is an Availability Zone?

A. A location where users can deploy compute, storage, database, and other select AWS serviceswhere no AWS Region currently exists
B. One or more discrete data centers with redundant power, networking, and connectivity
C. One or more clusters of servers where new workloads can be deployed
D. A fast content delivery network (CDN) service that securely delivers data, videos, applications, andAPIs to users globally

Answer: B

Explanation:

An Availability Zone is one or more discrete data centers with redundant power, networking, and connectivity. Availability Zones are part of the AWS global infrastructure, which consists of AWS Regions, Availability Zones, and edge locations. Availability Zones are physically separate locations within an AWS Region that are engineered to be isolated from failures and connected by low-latency, high-throughput, and highly redundant networking. Each Availability Zone contains one or more data centers that house the servers and storage devices that run AWS services. Availability Zones enable users to design and operate fault-tolerant and high-availability applications on AWS. AWS Global InfrastructureAWS Certified Cloud Practitioner - aws.amazon.com

Question: 131

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 cloud benefit that AWS offers to its users. AWS has a global infrastructure that consists of AWS Regions, Availability Zones, and edge locations. Users can choose from multiple AWS Regions around the world to deploy their applications and data closer to their end users, while also meeting their compliance and regulatory requirements. Users can also leverage AWS services, such as Amazon CloudFront, Amazon Route 53, and AWS Global Accelerator, to improve the performance and availability of their global applications. AWS also provides tools and guidance to help users optimize their global deployments, such as AWS Well-Architected Framework, AWS CloudFormation, and AWS Migration Hub. AWS Global Infrastructure [AWS Cloud Value Framework] AWS Certified Cloud Practitioner - aws.amazon.com

Question: 132

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

Explanation:
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.

Question: 133

A company wants to track its AWS account's service costs. The company also wants to receive notifications when costs are forecasted to reach a specific level.
Which AWS service or tool provides this functionality?

A. AWS Budgets
B. AWS Cost Explorer

C. Savings Plans
D. AWS Billing Conductor

Answer: A

Explanation:
AWS Budgets gives you the ability to set custom budgets that alert you when your costs or usage exceed (or are forecasted to exceed) your budgeted amount. You can also use AWS Budgets to set reservation utilization or coverage targets and receive alerts when your utilization drops below the threshold you define2.

Question: 134 An ecommerce company has migrated its IT infrastructure from an on-premises data center to the AWS Cloud.
Which AWS service is used to track, record, and audit configuration changes made to AWS resources?

A. AWS Shield
B. AWS Config
C. AWS IAM
D. Amazon Inspector

Answer: B

Explanation:
AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. AWS Config continuously monitors and records your AWS resource configurations and allows you to automate the evaluation of recorded configurations against desired configurations. With AWS Config, you can review changes in configurations and relationships between AWS resources, dive into detailed resource configuration histories, and determine your overall compliance against the configurations specified in your internal guidelines3.

Question: 135

A company needs to test a new application that was written in Python. The code will activate when new images are stored in an Amazon S3 bucket. The application will put a watermark on each image and then will store the images in a different S3 bucket.

Which AWS service should the company use to conduct the test with the LEAST amount of operational overhead?

A. Amazon EC2
B. AWS CodeDeploy
C. AWS Lambda
D. Amazon Lightsail

Answer: C

Explanation:
AWS Lambda is a compute service that lets you run code without provisioning or managing servers. AWS Lambda executes your code only when needed and scales automatically, from a few requests per day to thousands per second. You pay only for the compute time you

consume - there is no charge when your code is not running. With AWS Lambda, you can run code for virtually any type of application or backend service - all with zero administration. AWS Lambda runs your code on a high- availability compute infrastructure and performs all of the administration of the compute resources, including server and operating system maintenance, capacity provisioning and automatic scaling, code monitoring and logging

Question: 136

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 groupsQ C. Encryption of customer data on AWS
C. Management of AWS Lambda infrastructureQ E. Management of network throughput of each AWS Region

Answer: B,C

Explanation:
The AWS shared responsibility model describes how AWS and the customer share responsibility for security and compliance of the AWS environment. AWS is responsible for the security of the cloud, which includes the physical security of AWS facilities, the infrastructure, hardware, software, and networking that run AWS services. The customer is responsible for security in the cloud, which includes the configuration of security groups, the encryption of customer data on AWS, the management of AWS Lambda infrastructure, and the management of network throughput of each AWS Region.

Question: 137

Which AWS service or tool can be used to consolidate payments for a company with multiple AWS accounts?

A. AWS Cost and Usage Report
B. AWS Organizations
C. Cost Explorer
D. AWS Budgets

Answer: B

Explanation:
AWS Organizations is an account management service that enables you to consolidate multiple AWS accounts into an organization that you create and centrally manage. AWS Organizationsincludes consolidated billing and account management capabilities that enable you to better meet the budgetary, security, and compliance needs of your business1.

Question: 138

How can an AWS user conduct security assessments of Amazon EC2 instances, NAT gateways, and Elastic Load Balancers in a way that is approved by AWS?

A. Flood a target with requests.
B. Use Amazon Inspector.
C. Perform penetration testing.
D. Use the AWS Service Health Dashboard.

Answer: B

Explanation:
Amazon Inspector is an automated security assessment service that helps improve the security and compliance of applications deployed on AWS. Amazon Inspector automatically assesses applications for exposure, vulnerabilities, and deviations from best practices. After performing an assessment, Amazon Inspector produces a detailed list of security findings prioritized by level of severity2.

Question: 139

Which AWS service will help protect applications running on AWS from DDoS attacks?

A. Amazon GuardDuty
B. AWS WAF
C. AWS Shield
D. Amazon Inspector

Answer: C

Explanation:
AWS Shield is a managed Distributed Denial of Service (DDoS) protection service that safeguards applications running on AWS. AWS Shield provides always-on detection and automatic inline mitigations that minimize application downtime and latency, so there is no need to engage AWS Support to benefit from DDoS protection3.

Question: 140

A cloud engineer wants to know the percentage of the allocated compute units that are in use for a specific Amazon EC2 instance. Which AWS service can provide this information?

A. AWS CloudTrail
B. AWS Config
C. Amazon CloudWatch
D. AWS Artifact

Answer: C

Explanation:
Amazon CloudWatch is a monitoring and observability service built for DevOps engineers, developers, site reliability engineers (SREs), and IT managers. CloudWatch provides you with data and actionable insights to monitor your applications, respond to system-wide

performance changes, optimize resource utilization, and get a unified view of operational health. CloudWatch collects monitoring and operational data in the form of logs, metrics, and events, providing you with a unified view of AWS resources, applications, and services that run on AWS and on-premises servers

Question: 141

Which activity is a customer responsibility in the AWS Cloud according to the AWS shared responsibility model?

A. Ensuring network connectivity from AWS to the internet
B. Patching and fixing flaws within the AWS Cloud infrastructure
C. Ensuring the physical security of cloud data centers
D. Ensuring Amazon EBS volumes are backed up

Answer: D

Explanation:
The AWS shared responsibility model describes how AWS and the customer share responsibility for security and compliance of the AWS environment. AWS is responsible for the security of the cloud, which includes the physical security of AWS facilities, the infrastructure, hardware, software, and networking that run AWS services. The customer is responsible for security in the cloud, which includes the configuration of security groups, the encryption of customer data on AWS, the management of AWS Lambda infrastructure, and the management of network throughput of each AWS Region. One of the customer responsibilities is to ensure that Amazon EBS volumes are backed up.

Question: 142

Which AWS service meets this requirement?

A. AWS CloudFormation
B. AWS Elastic Beanstalk
C. AWS Cloud9
D. AWS CloudShell

Answer: A

Explanation:
AWS CloudFormation is a service that gives developers and businesses an easy way to create a collection of related AWS and third-party resources, and provision and manage them in an orderly and predictable fashion. You can use AWS CloudFormation's sample templates or create your own templates to describe the AWS and third-party resources, and any associated dependencies or runtime parameters, required to run your application.

Question: 143

A company wants to use the AWS Cloud as an offsite backup location for its on-premises infrastructure.
Which AWS service will meet this requirement MOST cost-effectively?

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

Answer: A

Explanation:
Amazon S3 is the most cost-effective service for storing offsite backups of on-premises infrastructure. Amazon S3 offers low-cost, durable, and scalable storage that can be accessed from anywhere over the internet. Amazon S3 also supports lifecycle policies, versioning, encryption, and cross-region replication to optimize the backup and recovery process. Amazon EFS, Amazon FSx, and Amazon EBS are more suitable for storing data that requires high performance, low latency, and frequent access12

Question: 144

A company is building a serverless architecture that connects application data from multiple data sources. The company needs a solution that does not require additional code.
Which AWS service meets these requirements?

A. AWS Lambda
B. Amazon Simple Queue Service (Amazon SQS)
C. Amazon CloudWatch
D. Amazon EventBridge

Answer: D

Explanation:
Amazon EventBridge is the service that meets the requirements of building a serverless architecture that connects application data from multiple data sources without requiring additional code. Amazon EventBridge is a serverless event bus service that allows you to easily connect your applications with data from AWS services, SaaS applications, and your own applications. You can use Amazon EventBridge to create rules that match events and route them to targets such as AWS Lambda functions, Amazon SNS topics, Amazon SQS queues, or other AWS services. Amazon EventBridge handles the event ingestion, delivery, security, authorization, and error handling for you34

Question: 145

A company needs to use standard SQL to query and combine exabytes of structured and semistructured 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

Explanation:
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

Question: 146

A company's information security manager is supervising a move to AWS and wants to ensure that AWS best practices are followed. The manager has concerns about the potential misuse of AWS account root user credentials.
Which of the following is an AWS best practice for using the AWS account root user credentials?

A. Allow only the manager to use the account root user credentials for normal activities.
B. Use the account root user credentials only for Amazon EC2 instances from the AWS Free Tier.
C. Use the account root user credentials only when they alone must be used to perform a requiredfunction.
D. Use the account root user credentials only for the creation of private VPC subnets.

Answer: C

Explanation:
The AWS best practice for using the AWS account root user credentials is to use them only when they alone must be used to perform a required function. The AWS account root user credentials have full access to all the resources in the account, and therefore pose a security risk if compromised or misused. You should create individual IAM users with the minimum necessary permissions for everyday tasks, and use AWS Organizations to manage multiple accounts. You should also enable multi-factor authentication (MFA) and rotate the password for the root user regularly. Some of the functions that require the root user credentials are changing the account name, closing the account, changing the support plan, and restoring an IAM user's access.

Question: 147

A company needs to store data across multiple Availability Zones in an AWS Region. The data will not be accessed regularly but must be immediately retrievable.
Which Amazon Elastic File System (Amazon EFS) storage class meets these requirements MOST cost effectively?

A. EFS Standard
B. EFS Standard-Infrequent Access(EFS Standard-IA)
C. EFS One Zone
D. EFS One Zone-Infrequent Access (EFS One Zone-IA)

Answer: B

Explanation:
EFS Standard-Infrequent Access (EFS Standard-IA) is the storage class that meets the requirements of storing data across multiple Availability Zones in an AWS Region, that will not be accessed regularly but must be immediately retrievable, most cost-effectively. EFS Standard-IA is designed for files that are accessed less frequently, but still require the same high performance, low latency, and high availability as EFS Standard. EFS Standard-IA has a lower storage cost than EFS Standard, but charges a small additional fee for each access. EFS One Zone and EFS One Zone-IA store data in a single Availability Zone, which reduces the availability and durability compared to EFS Standard and EFS Standard-IA.

Question: 148

A company wants to establish a security layer in its VPC that will act as a firewall to control subnet traffic.
Which AWS service or feature will meet this requirement?

A. Routing tables
B. Network access control lists (network ACLs)
C. Security groups
D. Amazon GuardDuty

Answer: C

Explanation:
Security groups are the service or feature that meets the requirement of establishing a security layer in a VPC that will act as a firewall to control subnet traffic. Security groups are stateful firewalls that control the inbound and outbound traffic at the instance level. You can assign one or more security groups to each instance in a VPC, and specify the rules that allow or deny traffic based on the protocol, port, and source or destination. Security groups are associated with network interfaces, and therefore apply to all the instances in the subnets that use those network interfaces. Routing tables are used to direct traffic between subnets and gateways, not to filter traffic. Network ACLs are stateless firewalls that control the inbound and outbound traffic at the subnet level, but they are less granular and more cumbersome to manage than security groups. Amazon GuardDuty is a threat detection service that monitors your AWS account and workloads for malicious or unauthorized activity, not a firewall service.

Question: 149

A newly created IAM user has no IAM policy attached.
What will happen when the user logs in and attempts to view the AWS resources in the account?

A. All AWS services will be read-only access by default.
B. Access to all AWS resources will be denied.
C. Access to the AWS billing services will be allowed.
D. Access to AWS resources will be allowed through the AWS CLL

Answer: B

Explanation:
Access to all AWS resources will be denied if a newly created IAM user has no IAM policy attached and logs in and attempts to view the AWS resources in the account. IAM policies are the way to grant permissions to IAM users, groups, and roles to access and manage AWS resources. By default, IAM users have no permissions, unless they are explicitly granted by an IAM policy. Therefore, a newly created IAM user without any IAM policy attached will not be able to view or perform anyactions on the AWS resources in the account. Access to the AWS billing services and AWS CLI will also be denied, unless the user has the necessary permissions.

Question: 150

A cloud practitioner is analyzing Amazon EC2 instance performance and usage to provide recommendations for potential cost savings.
Which cloud concept does this analysis demonstrate?

A. Auto scaling
B. Rightsizing
C. Load balancing
D. High availability

Answer: B

Explanation:
Rightsizing is the cloud concept that this analysis demonstrates. Rightsizing is the process of optimizing the performance and cost of your AWS resources by selecting the most appropriate type, size, and configuration based on your workload requirements and usage patterns. Rightsizing can help you achieve potential cost savings by reducing the over-provisioning or under-utilization of your resources. You can use various AWS tools and services, such as AWS Cost Explorer, AWS Compute Optimizer, and AWS Trusted Advisor, to analyze your resource utilization and performance metrics, and receive recommendations for rightsizing.

Question: 151

An auditor needs to find out whether a specific AWS service is compliant with specific compliance frameworks.
Which AWS service will provide this information?

A. AWS Artifact
B. AWS Trusted Advisor
C. Amazon GuardDuty
D. AWS Certificate Manager (ACM)

Answer: A

Explanation:
AWS Artifact is the service that will provide the information about whether a specific AWS service is compliant with specific compliance frameworks. AWS Artifact is a self-service portal that allows you to access, review, and download AWS security and compliance reports and

agreements. You canuse AWS Artifact to verify the compliance status of AWS services across various regions and compliance programs, such as ISO, PCI, SOC, FedRAMP, HIPAA, and more12

Question: 152

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:

The duties that are the responsibility of a company that is using AWS Lambda are security inside of code and writing and updating of code. AWS Lambda is a serverless compute service that allows you to run code without provisioning or managing servers, scaling, or patching. AWS Lambda takes care of the security of the underlying infrastructure, such as the operating system, the network, and the firewall. However, the company is still responsible for the security of the code itself, such as encrypting sensitive data, validating input, and handling errors. The company is also responsible for writing and updating the code that defines the Lambda function, and choosing the runtime environment, such as Node.js, Python, or Java. AWS Lambda does not require the selection of CPU resources, as it automatically allocates them based on the memory configuration34

Question: 153

Which AWS services and features are provided to all customers at no charge? (Select TWO.)

A. Amazon Aurora
B. VPC
C. Amazon SageMaker
D. AWS Identity and Access Management (IAM)
E. Amazon Polly

Answer: B,D

Explanation:

The AWS services and features that are provided to all customers at no charge are VPC and AWS Identity and Access Management (IAM). VPC is a service that allows you to launch AWS resources in a logically isolated virtual network that you define. You can create and use a VPC at no additional charge, and you only pay for the resources that you launch in the VPC, such as EC2 instances or EBSvolumes. IAM is a service that allows you to manage access and permissions to AWS resources. You can create and use IAM users, groups, roles, and policies at no additional charge, and you only pay for the AWS resources that the IAM entities access. Amazon Aurora, Amazon SageMaker, andAmazon Polly are not free services, and they charge based on the usage and features that you choose5

Question: 154

Which AWS services or features can control VPC traffic? (Select TWO.)

A. Security groups
B. AWS Direct Connect
C. Amazon GuardDuty
D. Network ACLs
E. Amazon Connect

Answer: A,D

Explanation:
The AWS services or features that can control VPC traffic are security groups and network ACLs. Security groups are stateful firewalls that control the inbound and outbound traffic at the instance level. You can assign one or more security groups to each instance in a VPC, and specify the rules that allow or deny traffic based on the protocol, port, and source or destination. Network ACLs are stateless firewalls that control the inbound and outbound traffic at the subnet level. You can associate one network ACL with each subnet in a VPC, and specify the rules that allow or deny traffic based on the protocol, port, and source or destination. AWS Direct Connect, Amazon GuardDuty, and Amazon Connect are not services or features that can control VPC traffic. AWS Direct Connect is a service that establishes a dedicated network connection between your premises and AWS. Amazon GuardDuty is a service that monitors your AWS account and workloads for malicious or unauthorized activity. Amazon Connect is a service that provides a cloud-based contact center solution.

Question: 155

A company needs to identify the last time that a specific user accessed the AWS Management Console.
Which AWS service will provide this information?

A. Amazon Cognito
B. AWS CloudTrail
C. Amazon Inspector
D. Amazon GuardDuty

Answer: B

Explanation:
AWS CloudTrail is the service that will provide the information about the last time that a specific user accessed the AWS Management Console. AWS CloudTrail is a service that records the API calls and events made by or on behalf of your AWS account. You can use AWS CloudTrail to view, search, and download the history of AWS console sign-in events, which include the user name, date, time, source IP address, and other details of the sign-in activity. Amazon Cognito, Amazon Inspector, and Amazon GuardDuty are not services that will provide this information. Amazon Cognito is a service that provides user authentication and authorization for web and mobile applications. Amazon Inspector is a service that assesses the security and compliance of your applications running on AWS. Amazon GuardDuty is a service that monitors your AWS account and workloads for malicious or unauthorized activity.

Question: 156

A company's application stores data in an Amazon S3 bucket. The company has an AWS Lambda function that processes data in the S3 bucket. The company needs to invoke the function once a day at a specific time. Which AWS service should the company use to meet this requirement?

A. AWS Managed Services (AMS)
B. AWS CodeStar
C. Amazon EventBridge
D. AWS Step Functions

Answer: C

Explanation:
Amazon EventBridge is the service that the company should use to meet the requirement of invoking the Lambda function once a day at a specific time. Amazon EventBridge is a serverless event bus service that allows you to easily connect your applications with data from AWS services, SaaS applications, and your own applications. You can use Amazon EventBridge to create rules that match events and route them to targets such as AWS Lambda functions, Amazon SNS topics, Amazon SQS queues, or other AWS services. You can also use Amazon EventBridge to create scheduled rules that trigger your targets at a specific time or interval, such as once a day. AWS Managed Services (AMS), AWS CodeStar, and AWS Step Functions are not services that the company should use to meet this requirement. AMS is a service that provides operational management for your AWS infrastructure and applications. AWS CodeStar is a service that provides a unified user interface for managing software development projects on AWS. AWS Step Functions is a service that coordinates multiple AWS services into serverless workflows.

Question: 157

A company uses Amazon Aurora as its database service. The company wants to encrypt its databases and database backups. Which party manages the encryption of the database clusters and database snapshots, according to the AWS shared responsibility model?

A. AWS
B. The company
C. AWS Marketplace partners
D. Third-party partners

Answer: A

Explanation:
AWS manages the encryption of the database clusters and database snapshots for Amazon Aurora, as well as the encryption keys. This is part of the AWS shared responsibility model, where AWS is responsible for the security of the cloud, and the customer is responsible for the security in the cloud. Encryption is one of the security features that AWS provides to protect the data at rest and in transit. For more information, see Amazon Aurora FAQs and AWS Shared Responsibility Model.

Question: 158

Which AWS solution gives companies the ability to use protocols such as NFS to store and retrieve objects in Amazon S3?

A. Amazon FSx for Lustre
B. AWS Storage Gateway volume gateway
C. AWS Storage Gateway file gateway
D. Amazon Elastic File System (Amazon EFS)

Answer: C

Explanation:
AWS Storage Gateway file gateway allows companies to use protocols such as NFS and SMB to store and retrieve objects in Amazon S3. File gateway provides a seamless integration between onpremises applications and Amazon S3, and enables low-latency access to data through local caching. File gateway also supports encryption, compression, and lifecycle management of the objects in Amazon S3. For more information, see What is AWS Storage Gateway? and File Gateway.

Question: 159

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 AWS service or tool that can help the company launch the number of EC2 instances that will be needed to handle the workload. Amazon EC2 Auto Scaling automatically adjusts the capacity of the EC2 instances based on the demand and the predefined scaling policies. Amazon EC2 Auto Scaling also helps to improve availability and reduce costs by scaling in and out as needed. For more information, see What is Amazon EC2 Auto Scaling? and [Getting Started with Amazon EC2 Auto Scaling].

Question: 160

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 the design principle that is achieved by following the reliability pillar of the AWS Well-Architected Framework. The reliability pillar focuses on the ability of a system to recover from failures and prevent disruptions. Testing recovery procedures helps to ensure that the system can handle different failure scenarios and restore normal operations as quickly as possible. Testing recovery procedures also helps to identify and mitigate any risks or gaps in the system design and implementation. For more information, see [Reliability Pillar] and [Testing for Reliability].

Question: 161

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

Explanation:

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].

Question: 162

In which of the following AWS services should database credentials be stored for maximum security?

A. AWS Identity and Access Management (IAM)
B. AWS Secrets Manager
C. Amazon S3
D. AWS Key Management Service (AWS KMS)

Answer: B

Explanation:

AWS Secrets Manager is the AWS service where database credentials should be stored for maximum security. AWS Secrets Manager helps to protect the secrets, such as database credentials, passwords, API keys, and tokens, that are used to access applications, services, and

resources. AWS Secrets Manager enables secure storage, encryption, rotation, and retrieval of the secrets. AWS Secrets

Manager also integrates with other AWS services, such as AWS Identity and Access Management (IAM), AWS Key Management Service (AWS KMS), and AWS Lambda. For more information, see [What is AWS Secrets Manager?] and [Getting Started with AWS Secrets Manager].

Question: 163

A company needs to configure rules to identify threats and protect applications from malicious network access.
Which AWS service should the company use to meet these requirements?

A. AWS Identity and Access Management (IAM)
B. Amazon QuickSight
C. AWS WAF
D. Amazon Detective

Answer: C

Explanation:
AWS WAF is the AWS service that the company should use to configure rules to identify threats and protect applications from malicious network access. AWS WAF is a web application firewall that helps to filter, monitor, and block malicious web requests based on customizable rules. AWS WAF can be integrated with other AWS services, such as Amazon CloudFront, Amazon API Gateway, and Application Load Balancer. For more information, see What is AWS WAF? and How AWS WAF Works.

Question: 164

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 the benefit of the AWS Cloud that this scenario demonstrates. 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, the company can launch new marketing campaigns in 3 days instead of 3 weeks, which shows that it can respond to customer feedback more quickly and efficiently. For more information,

see Benefits of Cloud Computing and [Business Agility].

Question: 165

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:
The costs that the company will eliminate with this migration are the cost of application licensing and the cost of physical server hardware. The cost of application licensing is the fee that the company has to pay to use the software applications on its on-premises servers. The cost of physical server hardware is the expense that the company has to incur to purchase, maintain, and upgrade the servers and related equipment. By migrating to the AWS Cloud, the company can avoid these costs by using the AWS services and resources that are already licensed and managed by AWS. For more information, see [Cloud Economics] and [AWS Total Cost of Ownership (TCO) Calculator].

Question: 166

Which AWS Support plan assigns an AWS concierge agent to a company's account?

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

Answer: D

Explanation:
AWS Enterprise Support is the AWS Support plan that assigns an AWS concierge agent to a company's account. AWS Enterprise Support is the highest level of support that AWS offers, and it provides the most comprehensive and personalized assistance. An AWS concierge agent is a dedicated technical account manager who acts as a single point of contact for the company and helps to optimize the AWS environment, resolve issues, and access AWS experts. For more information, see [AWS Support Plans] and [AWS Concierge Support].

Question: 167 A company hosts an application on an Amazon EC2 instance. The EC2 instance needs to access several AWS resources, including Amazon S3 and Amazon DynamoDB.

What is the MOST operationally efficient solution to delegate permissions?

A. Create an IAM role with the required permissions. Attach the role to the EC2 instance.

B. Create an IAM user and use its access key and secret access key in the application.
C. Create an IAM user and use its access key and secret access key to create a CLI profile in the EC2 instance.
D. Create an IAM role with the required permissions. Attach the role to the administrativeIAM user.

Answer: A

Explanation:
Creating an IAM role with the required permissions and attaching the role to the EC2 instance is the most operationally efficient solution to delegate permissions. An IAM role is an entity that defines a set of permissions for making AWS service requests. An IAM role can be assumed by an EC2 instance to access other AWS resources, such as Amazon S3 and Amazon DynamoDB, without having to store any credentials on the instance. This solution is more secure and scalable than using IAM users and their access keys. For more information, see [IAM Roles for Amazon EC2] and [Using an IAM Role to Grant Permissions to Applications Running on Amazon EC2 Instances].

Question: 168

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

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

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

Question: 169 A company is building an application that will receive millions of database queries each second. The company needs the data store for the application to scale to meet these needs.

Which AWS service will meet this requirement?

A. Amazon DynamoDB
B. AWS Cloud9
C. Amazon ElastiCache for Memcached
D. Amazon Neptune

Answer: A

Explanation:
Amazon DynamoDB is the AWS service that will meet the requirement of building an application that will receive millions of database queries each second. Amazon DynamoDB is a fully managed NoSQL database service that provides fast and consistent performance, scalability, and durability. Amazon DynamoDB can handle any level of request traffic and automatically scale up or down the capacity based on the demand. Amazon DynamoDB also supports in-memory caching with Amazon DynamoDB Accelerator (DAX) to improve the response time and reduce the cost. For more information, see What is Amazon DynamoDB? and Amazon DynamoDB Features.

Question: 170

An application runs on multiple Amazon EC2 instances that access a shared file system simultaneously.
Which AWS storage service should be used?

A. Amazon EBS
B. Amazon EFS
C. Amazon S3
D. AWS Artifact

Answer: B

Explanation:
Amazon Elastic File System (Amazon EFS) is the AWS storage service that should be used for an application that runs on multiple Amazon EC2 instances that access a shared file system simultaneously. Amazon EFS is a fully managed service that provides a scalable, elastic, and highly available file system for Linux-based workloads. Amazon EFS supports the Network File System version 4 (NFSv4) protocol and allows multiple EC2 instances to read and write data to the same file system concurrently. Amazon EFS also integrates with other AWS services, such as AWS Backup, AWS CloudFormation, and AWS CloudTrail. For more information, see What is Amazon Elastic File System? and [Amazon EFS Use Cases].

Question: 171

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

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

Answer: D

Explanation:
Physical and environmental controls are entirely the responsibility of AWS, according to the AWS shared responsibility model. The AWS shared responsibility model defines the division of responsibilities between AWS and the customer for security and compliance. AWS is responsible for the security of the cloud, which includes the physical and environmental controls of the AWS global infrastructure, such as power, cooling, fire suppression, and physical access. The customer is responsible for the security in the cloud, which includes the

configuration and management of the AWS resources and applications. For more information, see [AWS Shared Responsibility Model] and [AWS Cloud Security].

Question: 172

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

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

Answer: D

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

Question: 173

A company wants to use Amazon EC2 instances to run a stateless and restartable process after business hours.
Which AWS service provides DNS resolution?

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

Answer: C

Explanation:
Amazon Route 53 is the AWS service that provides DNS resolution. DNS (Domain Name System) is a service that translates domain names into IP addresses. Amazon Route 53 is a highly available and scalable cloud DNS service that offers domain name registration, DNS routing, and health checking. Amazon Route 53 can route the traffic to various AWS services, such as Amazon EC2, Amazon S3, and Amazon CloudFront. Amazon Route 53 can also integrate with other AWS services, such as AWS Certificate Manager, AWS Shield, and AWS WAF. For more information, see [What is Amazon Route 53?] and [Amazon Route 53 Features].

Question: 174

Which group shares responsibility with AWS for security and compliance of AWS accounts and resources?

A. Third-party vendors
B. Customers
C. Reseller partners
D. Internet providers

Answer: B

Explanation:
Customers share responsibility with AWS for security and compliance of AWS accounts and resources. This is part of the AWS shared responsibility model, which defines the division of responsibilities between AWS and the customer for security and compliance. AWS is responsible for the security of the cloud, which includes the physical and environmental controls of the AWS global infrastructure, such as power, cooling, fire suppression, and physical access. The customer is responsible for the security in the cloud, which includes the configuration and management of the AWS resources and applications, such as identity and access management, encryption, firewall, and backup. For more information, see AWS Shared Responsibility Model and AWS Cloud Security.

Question: 175

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

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

Answer: C

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

Question: 176

A company moves a workload to AWS to run on Amazon EC2 instances. The company needs to run the workload in the most cost-effective way.
What can the company do to meet this requirement?

A. Use AWS Key Management Service (AWS KMS).
B. Use multiple AWS accounts and consolidated billing.
C. Use AWS CloudFormation to deploy the infrastructure.
D. Rightsized all the EC2 instances that are used in the deployment.

Answer: D

Explanation:
Rightsizing all the EC2 instances that are used in the deployment is the best way to run the workload in the most cost-effective way. Rightsizing means choosing the optimal instance type and size for the workload based on the performance and capacity requirements. Rightsizing helps to avoid overprovisioning or under-provisioning of the EC2 instances, which can result in wasted resources or poor performance. Rightsizing also helps to take advantage of the different pricing models and features that AWS offers, such as On-Demand, Reserved, and Spot Instances, and Auto Scaling. For more information, see Rightsizing Your Instances and [Cost Optimization with AWS].

Question: 177

A company needs to launch an Amazon EC2 instance.

Which of the following can the company use during the launch process to configure the root volume of the EC2 instance?

A. Amazon EC2 Auto Scaling
B. Amazon Data Lifecycle Manager (Amazon DLM)
C. Amazon Machine Image (AMI)
D. Amazon Elastic Block Store (Amazon EBS) volume

Answer: C

Explanation:
Amazon Machine Image (AMI) is the option that the company can use during the launch process to configure the root volume of the EC2 instance. An AMI is a template that contains the software configuration, such as the operating system, applications, and settings, required to launch an EC2 instance. An AMI also specifies the volume size and type of the root device for the instance. Thecompany can choose an AMI provided by AWS, the AWS Marketplace, or the AWS community, or create a custom AMI. For more information, see [Amazon Machine Images (AMI)] and [Launching an Instance Using the Launch Instance Wizard].

Question: 178

A company plans to migrate its on-premises workload to AWS. Before the migration, the company needs to estimate its future AWS service costs.
Which AWS service or tool should the company use to meet this requirement?

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

Answer: C

Explanation:
AWS Pricing Calculator is the AWS service or tool that the company should use to estimate its future AWS service costs before the

migration. AWS Pricing Calculator is a web-based tool that allows the company to create cost estimates for various AWS services and scenarios. AWS Pricing Calculator helps the company to compare the costs of running the workload on premises versus on AWS, and to optimize the costs by choosing the best options for the workload. AWS Pricing Calculator also provides a detailed breakdown of the cost components and a downloadable report. For more information, see [AWS Pricing Calculator] and [Getting Started with AWS Pricing Calculator].

Question: 179

A company suspects that its AWS resources are being used for illegal activities. Which AWS group or team should the company notify?

A. AWS Abuse team
B. AWS Support team
C. AWS technical account managers
D. AWS Professional Services team

Answer: A

Explanation:
AWS Abuse team is the AWS group or team that the company should notify if it suspects that its AWS resources are being used for illegal activities. AWS Abuse team is a dedicated team that handles reports of abuse, such as spam, phishing, malware, denial-of-service attacks, and unauthorized access, involving AWS resources. The company can contact the AWS Abuse team byfilling out the [Report Abuse of AWS Resources form] or sending an email to abuse@amazonaws.com. The company should provide as much information as possible, such as the source and destination IP addresses, timestamps, log files, and screenshots, to help the AWS Abuse team investigate and take appropriate actions. For more information, see [Reporting Abuse] and [AWS Acceptable Use Policy].

Question: 180

A company wants an in-memory data store that is compatible with open source in the cloud. Which AWS service should the company use?

A. Amazon DynamoDB
B. Amazon ElastiCache
C. Amazon Elastic Block Store (Amazon EBS)
D. Amazon Redshift

Answer: B

Explanation:
Amazon ElastiCache is a fully managed in-memory data store service that is compatible with open source engines such as Redis and Memcached1. It provides fast and scalable performance for applications that require high throughput and low latency1. Amazon DynamoDB is a fully managed NoSQL database service that provides consistent and single-digit millisecond latency at any scale2. Amazon EBS is a block storage service that provides persistent and durable storage volumes for Amazon EC2 instances3. Amazon Redshift is a fully managed data warehouse service that allows users to run complex analytic queries using SQL4.

Question: 181

A company wants to improve its security and audit posture by limiting Amazon EC2 inbound access. According to the AWS shared responsibility model, which task is the responsibility of the customer?

A. Protect the global infrastructure that runs all of the services offered in the AWS Cloud.
B. Configure logical access controls for resources, and protect account credentials.
C. Configure the security used by managed services.
D. Patch and back up Amazon Aurora.

Answer: B

Explanation:
According to the AWS shared responsibility model, the customer is responsible for configuring logical access controls for resources, and protecting account credentials. This includes managing IAM user permissions, security group rules, network ACLs, encryption keys, and other aspects of access management1. AWS is responsible for protecting the global infrastructure that runs all of the services offered in the AWS Cloud, such as the hardware, software, networking, and facilities. AWS isalso responsible for configuring the security used by managed services, such as Amazon RDS, Amazon DynamoDB, and Amazon Aurora2.

Question: 182

Which task is the responsibility of AWS when using AWS services?

A. Management of IAM user permissions
B. Creation of security group rules for outbound access
C. Maintenance of physical and environmental controls
D. Application of Amazon EC2 operating system patches

Answer: C

Explanation:
AWS is responsible for maintaining the physical and environmental controls of the AWS Cloud, such as power, cooling, fire suppression, and physical security1. The customer is responsible for managing the IAM user permissions, creating security group rules for outbound access, applying Amazon EC2 operating system patches, and other aspects of security in the cloud1.

Question: 183

A company wants to push VPC Flow Logs to an Amazon S3 bucket.

A company wants to optimize long-term compute costs of AWS Lambda functions and Amazon EC2 instances.
Which AWS purchasing option should the company choose to meet these requirements?

A. Dedicated Hosts
B. Compute Savings Plans
C. Reserved Instances

D. Spot Instances

Answer: B

Explanation:
Compute Savings Plans are a flexible and cost-effective way to optimize long-term compute costs of AWS Lambda functions and Amazon EC2 instances. With Compute Savings Plans, customers can commit to a consistent amount of compute usage (measured in $/hour) for a 1-year or 3-year termand receive a discount of up to 66% compared to On-Demand prices3. Dedicated Hosts are physical servers with EC2 instance capacity fully dedicated to the customer's use. They are suitable for customers who have specific server-bound software licenses or compliance requirements4. Reserved Instances are a pricing model that provides a significant discount (up to 75%) compared to On-Demand pricing and a capacity reservation for EC2 instances. They are available in 1-year or 3year terms and different payment options5. Spot Instances are spare EC2 instances that are available at up to 90% discount compared to On-Demand prices. They are suitable for customers who have flexible start and end times, can withstand interruptions, and can handle excess capacity.

Question: 184

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

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

Answer: A

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

Question: 185

A company needs to centralize its operational dat

a. The company also needs to automate tasks across all of its Amazon EC2 instances. Which AWS service can the company use to meet these requirements?

A. AWS Trusted Advisor
B. AWS Systems Manager
C. AWS CodeDeploy
D. AWS Elastic Beanstalk

Answer: B

Explanation:
AWS Systems Manager is a service that enables users to centralize and automate the management of their AWS resources. It provides a unified user interface to view operational data, such as inventory, patch compliance, and performance metrics. It also allows users to automate common and repetitive tasks, such as patching, backup, and configuration management, across all of their Amazon EC2 instances1. AWS Trusted Advisor is a service that provides best practices and recommendations to optimize the performance, security, and cost of AWS resources2. AWS CodeDeploy is a service that automates the deployment of code and applications to Amazon EC2 instances or other compute services3. AWS Elastic Beanstalk is a service that simplifies the deployment andmanagement of web applications using popular platforms, such as Java, PHP, and Node.js4.

Question: 186

A company needs Amazon EC2 instances for a workload that can tolerate interruptions.
Which EC2 instance purchasing option meets this requirement with the LARGEST discount compared to On-Demand prices?

A. Spot Instances
B. Convertible Reserved Instances
C. Standard Reserved Instances
D. Dedicated Hosts

Answer: A

Explanation:
Spot Instances are spare Amazon EC2 instances that are available at up to 90% discount compared to On-Demand prices. They are suitable for workloads that can tolerate interruptions, such as batch processing, data analysis, and testing. Spot Instances are allocated based on the current supply and demand, and can be reclaimed by AWS with a two-minute notice when the demand exceeds the supply5. Convertible Reserved Instances are a type of Reserved Instances that provide a significant discount (up to 54%) compared to On-Demand prices and a capacity reservation for Amazon EC2 instances. They are available in 1-year or 3-year terms and allow users to change the instance family, size, operating system, or tenancy during the term. Standard Reserved Instances are another type of Reserved Instances that provide a larger discount (up to 75%) compared to On-Demand prices and a capacity reservation for Amazon EC2 instances. They are available in 1-year or 3-year terms and do not allow users to change the instance attributes during the term. Dedicated Hosts are physical servers with Amazon EC2 instance capacity fully dedicated to the user's use. They are suitable for users who have specific server-bound software licenses or compliance requirements.

Question: 187

Which AWS service can defend against DDoS attacks?

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

Answer: B

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

Question: 188

A company wants its Amazon EC2 instances to share the same geographic area but use redundant underlying power sources. Which solution will meet these requirements?

A. Use EC2 instances across multiple Availability Zones in the same AWS Region.
B. Use Amazon CloudFront as the database for the EC2 instances.
C. Use EC2 instances in the same edge location and the same Availability Zone.
D. Use EC2 instances in AWS OpsWorks stacks in different AWS Regions.

Answer: A

Explanation:
Using EC2 instances across multiple Availability Zones in the same AWS Region is a solution that meets the requirements of sharing the same geographic area but using redundant underlying power sources. Availability Zones are isolated locations within an AWS Region that have independent power, cooling, and physical security. They are connected through low-latency, high-throughput, and highly redundant networking. By launching EC2 instances in different Availability Zones, users can increase the fault tolerance and availability of their applications. Amazon CloudFront is a contentdelivery network (CDN) service that speeds up the delivery of web content and media to end users by caching it at the edge locations closer to them. It is not a database service and cannot be used to store operational data for EC2 instances. Edge locations are sites that are part of the Amazon CloudFront network and are located in many cities around the world. They are not the same as Availability Zones and do not provide redundancy for EC2 instances. AWS OpsWorks is a configuration management service that allows users to automate the deployment and management of applications using Chef or Puppet. It can be used to create stacks that span multiple AWS Regions, but this would not meet the requirement of sharing the same geographic area.

Question: 189

A company needs to design a solution for the efficient use of compute resources for an enterprise workload. The company needs to make informed decisions as its technology needs evolve. Which pillar of the AWS Well-Architected Framework do these requirements represent?

A. Operational excellence
B. Performance efficiency
C. Cost optimization

D. Reliability

Answer: B

Explanation:
Performance efficiency is the pillar of the AWS Well-Architected Framework that represents the requirements of designing a solution for the efficient use of compute resources for an enterprise workload and making informed decisions as the technology needs evolve. It focuses on using the right resources and services for the workload, monitoring performance, and continuously improving the efficiency of the solution. Operational excellence is the pillar of the AWS Well-Architected Framework that represents the ability to run and monitor systems to deliver business value and to continually improve supporting processes and procedures. Cost optimization is the pillar of the AWS Well-Architected Framework that represents the ability to run systems to deliver business value at the lowest price point. Reliability is the pillar of the AWS Well-Architected Framework that represents the ability of a system to recover from infrastructure or service disruptions, dynamically acquire computing resources to meet demand, and mitigate disruptions such as misconfigurations or transient network issues.

Question: 190

What does "security of the cloud" refer to in the AWS shared responsibility model?

A. Availability of AWS services such as Amazon EC2
B. Security of the cloud infrastructure that runs all the AWS services
C. Implementation of password policies for IAM users
D. Security of customer environments by using AWS Network Firewall partners

Answer: B

Explanation:
Security of the cloud refers to the security of the cloud infrastructure that runs all the AWS services. This includes the hardware, software, networking, and facilities that AWS operates and manages. AWS is responsible for protecting the security of the cloud as part of the AWS shared responsibility model. Availability of AWS services such as Amazon EC2 refers to the ability of the services to be up and running and to meet the expected performance. Availability is part of the reliability pillar of the AWS Well-Architected Framework and is a shared responsibility between AWS and the customer . Implementation of password policies for IAM users refers to the security of the customer data and applications in the cloud. This includes the configuration and management of IAM user permissions, encryption keys, security group rules, network ACLs, and other aspects of access management. The customer is responsible for protecting the security in the cloud as part of the AWS shared responsibility model. Security of customer environments by using AWS Network Firewall partners refers to the security of the customer data and applications in the cloud. AWS Network Firewall is a managed service that provides network protection for Amazon VPCs. It allows customers to use AWS Marketplace partners to implement firewall rules and policies. The customer is responsible for protecting the security in the cloud as part of the AWS shared responsibility model .

Question: 191

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

A. Amazon DevPay
B. AWS Organizations

C. AWS Trusted Advisor
D. Cost Explorer

Answer: D

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

Question: 192

Which AWS service is always free of charge for users?

A. Amazon S3
B. Amazon Aurora
C. Amazon EC2
D. AWS Identity and Access Management (IAM)

Answer: D

Explanation:
AWS Identity and Access Management (IAM) is a service that allows users to manage access to AWS resources and services. It enables users to create and manage users, groups, roles, and policies that control who can do what in AWS. IAM is always free of charge for users, as there is no additional cost for using IAM with any AWS service1. Amazon S3 is a storage service that provides scalable, durable, and secure object storage. Amazon S3 has a free tier that offers 5 GB of storage, 20,000 GET requests, and 2,000 PUT requests per month for one year. However, users are charged for any additional usage beyond the free tier limits2. Amazon Aurora is a relational database service that is compatible with MySQL and PostgreSQL. Amazon Aurora has a free tier that offers 750 hours of Aurora Single-AZ db.t2.small database usage and 20 GB of storage per month for one year. However, users are charged for any additional usage beyond the free tier limits3. Amazon EC2 is a compute service that provides resizable virtual servers. Amazon EC2 has a free tier that offers 750 hours of Linux and Windows t2.micro instances per month for one year. However, users are charged for any additional usage beyond the free tier limits4.

Question: 193

A company has multiple AWS accounts that include compute workloads that cannot be interrupted. The company wants to obtain billing discounts that are based on the company's use of AWS services. Which AWS feature or purchasing option will meet these requirements?

A. Resource tagging

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

Answer: B

Explanation:
Consolidated billing is an AWS feature that allows users to combine the usage and costs of multiple AWS accounts into a single bill. This enables users to obtain billing discounts that are based on the company's use of AWS services, such as volume pricing tiers, Reserved Instance discounts, andSavings Plans discounts5. Resource tagging is an AWS feature that allows users to assign metadata to AWS resources, such as EC2 instances, S3 buckets, and Lambda functions. This enables users to organize, track, and manage their AWS resources, such as filtering, grouping, and reporting. Pay-as-you-go pricing is an AWS pricing model that allows users to pay only for the resources and services they use, without any upfront or long-term commitments. This enables users to lower their costs by scaling up or down as needed, and avoiding over-provisioning or under-utilization. Spot

Instances are spare EC2 instances that are available at up to 90% discount compared to On-Demand prices. They are suitable for workloads that can tolerate interruptions, such as batch processing, data analysis, and testing. Spot Instances are allocated based on the current supply and demand, and can be reclaimed by AWS with a two-minute notice when the demand exceeds the supply.

Question: 194

A company has an environment that includes Amazon EC2 instances, Amazon Lightsail, and onpremises servers. The company wants to automate the security updates for its operating systems and applications.
Which solution will meet these requirements with the LEAST operational effort?

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

Answer: C

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

Question: 195

A company that is planning to migrate to the AWS Cloud is based in an isolated area that has limited internet connectivity. The company needs to perform local data processing on premises. The company needs a solution that can operate without a stable internet connection. Which AWS service will meet these requirements?

A. Amazon S3
B. AWS Snowball Edge
C. AWS StorageGateway
D. AWS Backup

Answer: B

Explanation:
AWS Snowball Edge is a service that provides a physical device that can store up to 100 TB of data and perform local data processing on premises. It enables users to transfer data to and from the AWS Cloud in areas with limited or no internet connectivity. It also supports AWS Greengrass, which allows users to run AWS Lambda functions and other AWS services locally without a stable internet connection. Amazon S3 is a storage service that provides scalable, durable, and secure object storage. It requires a stable internet connection to transfer data to and from the AWS Cloud. AWS Storage Gateway is a service that provides a hybrid storage solution that connects on-premises applications to AWS Cloud storage services, such as Amazon S3, Amazon S3 Glacier, and Amazon EBS. It requires a stable internet connection to synchronize data between the on-premises and cloud storage. AWS Backup is a service that provides a centralized and automated solution to back up data across AWS services and on-premises resources. It requires a stable internet connection to transfer data to and from the AWS Cloud.

Question: 196

A company wants to migrate its applications to the AWS Cloud. The company plans to identify and prioritize any business transformation opportunities and evaluate its AWS Cloud readiness. Which AWS service or tool should the company use to meet these requirements?

A. AWS Cloud Adoption Framework (AWS CAF)
B. AWS Managed Services (AMS)
C. AWS Well-Architected Framework
D. AWS Migration Hub

Answer: A

Explanation:
AWS Cloud Adoption Framework (AWS CAF) is a service or tool that helps users migrate their applications to the AWS Cloud. It provides guidance and best practices to identify and prioritize any business transformation opportunities and evaluate their AWS Cloud readiness. It also helps users align their business and technical perspectives, create an actionable roadmap, and measure their progress. AWS Managed Services (AMS) is a service that provides operational services for AWS infrastructure and applications. It helps users reduce their operational overhead and risk, and focuson their core business. It does not help users identify and prioritize any business transformation opportunities and evaluate their AWS Cloud readiness. AWS Well-Architected Framework is a tool that helps users design and implement secure, high-performing, resilient, and efficient solutions on AWS. It provides a set of questions and best practices across five pillars: operational excellence, security, reliability, performance efficiency, and cost optimization. It does not help users identify and prioritize any business transformation opportunities and evaluate their AWS Cloud readiness. AWS Migration Hub is a service that provides a single location to track and manage the migration of applications to AWS. It helps users discover their on-premises servers, group them into

applications, and choose the right migration tools. It does not help users identify and prioritize any business transformation opportunities and evaluate their AWS Cloud readiness.

Question: 197

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

Explanation:
Patch management and configuration management are controls that are the responsibility of both AWS and AWS customers, according to the AWS shared responsibility model. Patch management is the process of applying updates to software and applications to fix vulnerabilities, bugs, or performance issues. Configuration management is the process of defining and maintaining the settings and parameters of systems and applications to ensure their consistency and reliability. AWS is responsible for patching and configuring the software and services that it manages, such as the AWS global infrastructure, the hypervisor, and the AWS managed services. The customer is responsible for patching and configuring the software and services that they manage, such as the guest operating system, the applications, and the AWS customer-managed services. Physical and environmental controls are the responsibility of AWS, according to the AWS shared responsibility model. Physical and environmental controls are the measures that protect the physical security and availability of the AWS global infrastructure, such as power, cooling, fire suppression, and access control. AWS is responsible for maintaining these controls and ensuring the resilience and reliability of the AWS Cloud. Account structures are the responsibility of the customer, according to the AWS shared responsibility model. Account structures are the ways that customers organize and manage their AWS accounts and resources, such as using AWS Organizations, IAM users and roles, resource tagging, and billing preferences. The customer is responsible for creating and configuring these structures and ensuring the security and governance of their AWS environment. Choice of the AWS Region where data is stored is the responsibility of the customer, according to the AWS sharedresponsibility 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.

Question: 198

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 (IAM)
E. Security groups

Answer: A,B

Explanation:
AWS Config and service control policies (SCPs) are AWS services or features that the company can use to create and define controls (guardrails) in a newly created AWS Control Tower landing zone. AWS Config is a service that enables users to assess, audit, and evaluate the configurations of their AWS resources. It can be used to create rules that check for compliance with the desired configurations and report any deviations. AWS Control Tower provides a set of predefined AWS Config rules that can be enabled as guardrails to enforce compliance across the landing zone1. Service control policies (SCPs) are a type of policy that can be used to manage permissions in AWS Organizations. They can be used to restrict the actions that the users and roles in the member accounts can perform on the AWS resources. AWS Control Tower provides a set of predefined SCPsthat can be enabled as guardrails to prevent access to certain services or regions across the landing zone2. Amazon GuardDuty is a service that provides intelligent threat detection and continuous monitoring for AWS accounts and resources. It is not a feature that can be used to create and define controls (guardrails) in a landing zone. AWS Identity and Access Management (IAM) is a service that allows users to manage access to AWS resources and services. It can be used to create users, groups, roles, and policies that control who can do what in AWS. It is not a feature that can be used to create and define controls (guardrails) in a landing zone. Security groups are virtual firewalls that control the inbound and outbound traffic for Amazon EC2 instances. They can be used to allow or deny access to an EC2 instance based on the port, protocol, and source or destination. They are not a feature that can be used to create and define controls (guardrails) in a landing zone.

Question: 199

A developer wants to use an Amazon S3 bucket to store application logs that contain sensitive data. Which AWS service or feature should the developer use to restrict read and write access to the S3 bucket?

A. Security groups
B. Amazon CloudWatch
C. AWS CloudTrail
D. ACLs

Answer: D

Explanation:
ACLs are an AWS service or feature that the developer can use to restrict read and write access to the S3 bucket. ACLs are access control lists that grant basic permissions to other AWS accounts or predefined groups. They can be used to grant read or write access to an S3 bucket or an object3. Security groups are virtual firewalls that control the inbound and outbound traffic for Amazon EC2 instances. They are not a service or feature that can be used to restrict access to an S3 bucket. Amazon CloudWatch is a service that provides monitoring and observability for AWS resources and applications. It can be used to collect and analyze metrics, logs, events, and alarms. It is not a service or feature that can be used to restrict access to an S3 bucket. AWS CloudTrail is a service that provides governance, compliance, and audit for AWS accounts and resources. It can be used to track and record the API calls and user activity in AWS. It is not a service or feature that can be used to restrict access to an S3 bucket.

Question: 200

Which AWS service or tool helps companies measure the environmental impact of their AWS usage?

A. AWS customer carbon footprint tool
B. AWS Compute Optimizer

C. Sustainability pillar

D. OS-Climate (Open Source Climate Data Commons)

Answer: A

Explanation:

AWS customer carbon footprint tool is an AWS service or tool that helps companies measure the environmental impact of their AWS usage. It allows users to estimate the carbon emissions associated with their AWS resources and services, such as EC2, S3, and Lambda. It also provides recommendations and best practices to reduce the carbon footprint and improve the sustainability of their AWS workloads4. AWS Compute Optimizer is an AWS service that helps users optimize the performance and cost of their EC2 instances and Auto Scaling groups. It provides recommendations for optimal instance types, sizes, and configurations based on the workload characteristics andutilization metrics. It does not help users measure the environmental impact of their AWS usage. Sustainability pillar is a concept that refers to the ability of a system to operate in an environmentally friendly and socially responsible manner. It is not an AWS service or tool that helps users measure the environmental impact of their AWS usage. OS-Climate (Open Source Climate Data Commons) is an initiative that aims to provide open source data, tools, and platforms to accelerate climate action and innovation. It is not an AWS service or tool that helps users measure the environmental impact of their AWS usage.

Question: 201

Which option is a perspective that includes foundational capabilities of the AWS Cloud Adoption Framework (AWS CAF)?

A. Sustainability

B. Operations

C. Performance efficiency

D. Reliability

Answer: B

Explanation:

Operations is an option that is a perspective that includes foundational capabilities of the AWS Cloud Adoption Framework (AWS CAF). Operations is one of the six perspectives of the AWS CAF, along with business, people, governance, platform, and security. Operations focuses on the processes and procedures to support the ongoing management and maintenance of the cloud-based IT assets. It covers topics such as monitoring, backup and recovery, change management, incident management, and automation5. Sustainability is not a perspective of the AWS CAF, but a concept that refers to the ability of a system to operate in an environmentally friendly and socially responsible manner. Performance efficiency is not a perspective of the AWS CAF, but a pillar of the AWS Well-Architected Framework. It focuses on using the right resources and services for the workload, monitoring performance, and continuously improving the efficiency of the solution. Reliability is not a perspective of the AWS CAF, but a pillar of the AWS Well-Architected Framework. It focuses on the ability of a system to recover from infrastructure or service disruptions, dynamically acquire computing resources to meet demand, and mitigate disruptions such as misconfigurations or transient network issues.

Question: 202

Which AWS service can a company use to securely store and encrypt passwords for a database?

A. AWS Shield
B. AWS Secrets Manager
C. AWS Identity and Access Management (IAM)
D. Amazon Cognito

Answer: B

Explanation:
AWS Secrets Manager is an AWS service that can be used to securely store and encrypt passwords for a database. It allows users to manage secrets, such as database credentials, API keys, and tokens, in a centralized and secure way. It also provides features such as automatic rotation, fine-grained access control, and auditing. AWS Shield is an AWS service that provides protection against Distributed Denial of Service (DDoS) attacks for AWS resources and services. It does not store or encrypt passwords for a database. AWS Identity and Access Management (IAM) is an AWS service that allows users to manage access to AWS resources and services. It can be used to create users, groups, roles, and policies that control who can do what in AWS. It does not store or encrypt passwords for a database. Amazon Cognito is an AWS service that provides user identity and data synchronization for web and mobile applications. It can be used to authenticate and authorize users, manage user profiles, and sync user data across devices. It does not store or encrypt passwords for a database.

Question: 203

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

A. Identity and access management
B. Hard drive initialization
C. Protection of data center hardware
D. Security of Availability Zones

Answer: A

Explanation:
Identity and access management is the customer's responsibility, according to the AWS shared responsibility model. This means that the customer is responsible for managing user access to the AWS resources, using tools such as AWS Identity and Access Management (IAM), AWS Single Sign-On (SSO), and AWS Organizations. The customer is also responsible for securing their data in transit and at rest, using encryption, key management, and other methods. Hard drive initialization, protection of data center hardware, and security of Availability Zones are AWS's responsibility, as they are part of the infrastructure, physical security, and network security that AWS provides to the customer12

Question: 204

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

A. AWS Transit Gateway
B. Internet gateway
C. Amazon VPC

D. Amazon EC2

Answer: C

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

Question: 205

Which AWS service offers a global content delivery network (CDN) that helps companies securely deliver websites, videos, applications, and APIs at high speeds with low latency?

A. Amazon EC2
B. Amazon CloudFront
C. Amazon CloudWatch
D. AWS CloudFormation

Answer: B

Explanation:
Amazon CloudFront is the AWS service that offers a global content delivery network (CDN) that helps companies securely deliver websites, videos, applications, and APIs at high speeds with low latency. Amazon CloudFront is a web service that speeds up distribution of static and dynamic web content, such as HTML, CSS, JavaScript, and image files, to users. Amazon CloudFront uses a global network of edge locations, located near users' geographic locations, to cache and serve content with high availability and performance. Amazon CloudFront also provides features such as AWS Shield for DDoS protection, AWS Certificate Manager for SSL/TLS encryption, AWS WAF for web application firewall, and AWS Lambda@Edge for customizing content delivery with serverless code. Amazon EC2, Amazon CloudWatch, and AWS CloudFormation are not services that offer a global CDN.
Amazon EC2 is a service that provides scalable compute capacity in the cloud. Amazon CloudWatch is a service that provides monitoring and observability for AWS resources andapplications. AWS CloudFormation is a service that provides a common language to model and provision AWS resources and their dependencies.

Question: 206

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

A. Agility
B. Economies of scale
C. Global reach

D. Pay-as-you-go pricing

Answer: C

Explanation:

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

Question: 207

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

Explanation:

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

Question: 208

An ecommerce company wants to design a highly available application that will be hosted on multiple Amazon EC2 instances. How should the company deploy the EC2 instances to meet these requirements?

A. Across multiple edge locations
B. Across multiple VPCs
C. Across multiple Availability Zones
D. Across multiple AWS accounts

Answer: C

Explanation:
The company should deploy the EC2 instances across multiple Availability Zones to design a highly available application. Availability Zones are isolated locations within an AWS Region that are engineered to be fault-tolerant and operate independently of each other. By deploying the EC2 instances across multiple Availability Zones, the company can ensure that their application can withstand the failure of an entire Availability Zone and continue to operate with minimal disruption. Deploying the EC2 instances across multiple edge locations, VPCs, or AWS accounts will not provide the same level of availability and fault tolerance as Availability Zones. Edge locations are part of the Amazon CloudFront service, which is a content delivery network (CDN) that caches and serves webcontent to users. VPCs are virtual networks that isolate the AWS resources within an AWS Region. AWS accounts are the primary units of ownership and access control for AWS resources12

Question: 209

Which AWS Cloud design principle does a company follow by using AWS CloudTrail?

A. Recover automatically.
B. Perform operations as code.
C. Measure efficiency.
D. Ensure traceability.

Answer: D

Explanation:
The company follows the AWS Cloud design principle of ensuring traceability by using AWS CloudTrail. AWS CloudTrail is a service that records the API calls and events made by or on behalf of the AWS account. The company can use AWS CloudTrail to monitor, audit, and analyze the activity and changes in their AWS resources and applications. AWS CloudTrail helps the company to achieve compliance, security, governance, and operational efficiency. Recovering automatically, performing operations as code, and measuring efficiency are other AWS Cloud design principles, but they are not directly related to using AWS CloudTrail. Recovering automatically means that the company can design their cloud workloads to handle failures gracefully and resume normal operations without manual intervention. Performing operations as code means that the company can automate the creation, configuration, and management of their cloud resources using scripts or templates. Measuring efficiency means that the company can monitor and optimize the performance and utilization of their cloud resources and applications34

Question: 210

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

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

Answer: B

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

Question: 211

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

A. Configure AWS Identity and Access Management (IAM).
B. Configure security groups on Amazon EC2 instances.
C. Secure the access of physical AWS facilities.
D. Patch applications that run on Amazon EC2 instances.
E. Perform infrastructure patching and maintenance.

Answer: C,E

Explanation:
The tasks that are the responsibility of AWS according to the AWS shared responsibility model are securing the access of physical AWS facilities and performing infrastructure patching and maintenance. The AWS shared responsibility model defines the division of

responsibilities between AWS and the customer for security and compliance. AWS is responsible for the security of the cloud, which includes the physical security of the hardware, software, networking, and facilities that run the AWS services. AWS is also responsible for the maintenance and patching of the infrastructure that supports the AWS services. The customer is responsible for the security in the cloud, which includes the configuration and management of the AWS resources and applications that they use. Configuring AWS Identity and Access Management (IAM), configuring security groups on Amazon EC2 instances, and patching applications that run on Amazon EC2 instances are tasks that are the responsibility of the customer, not AWS.

Question: 212

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:
The combination of AWS services that the application can use to migrate to a microservices-based application are Amazon Simple Queue Service (Amazon SQS) and AWS Lambda. Amazon SQS is a fully managed message queuing service that enables customers to decouple and scale microservices, distributed systems, and serverless applications. The application can use Amazon SQS to send, store, and receive messages between the microservices, ensuring that each message is processed only once and in the right order. AWS Lambda is a serverless compute service that allows customers to run code without provisioning or managing servers. The application can use AWS Lambda to create and deploy microservices as functions that are triggered by events, such as messages from Amazon SQS. AWS Migration Hub, AWS AppSync, and AWS Application Migration Service are not the best services to use for migrating to a microservices-based application. AWS Migration Hub is a service that provides a single location to track the progress of application migrations across multiple AWS and partner solutions. AWS AppSync is a service that simplifies the development of GraphQL APIs for real-time and ofline data synchronization. AWS Application Migration Service is a service that enables customers to migrate their on-premises applications to AWS without making any changes to the applications, servers, or databases.

Question: 213

A company wants to access a report about the estimated environmental impact of the company's AWS usage.
Which AWS service or feature should the company use to meet this requirement?

A. AWS Organizations
B. IAM policy
C. AWS Billing console
D. Amazon Simple Notification Service (Amazon SNS)

Answer: C

Explanation:
The company should use the AWS Billing console to access a report about the estimated environmental impact of the company's AWS usage. The AWS Billing console provides customers with various tools and reports to manage and monitor their AWS costs and usage. One of the reports available in the AWS Billing console is the AWS Sustainability Dashboard, which shows the estimated carbon footprint and energy mix of the customer's AWS usage. The company can use this dashboard to measure and improve the sustainability of their cloud workloads. AWS Organizations, IAM policy, and Amazon Simple Notification Service (Amazon SNS) are not services or features that can provide a report about the estimated environmental impact of the company's AWS usage. AWS Organizations is a service that enables customers to centrally manage and govern their AWS accounts. IAM policy is a document that defines the permissions for an IAM identity (user, group, or role) or an AWS resource. Amazon SNS is a fully managed pub/sub messaging service that enables customers to send messages to subscribers or other AWS services.

Question: 214

A company has an AWS-hosted website located behind an Application Load Balancer. The company wants to safeguard the website from SQL injection or cross-site scripting.
Which AWS service should the company use?

A. Amazon GuardDuty
B. AWS WAF
C. AWS Trusted Advisor
D. Amazon Inspector

Answer: B

Explanation:
The company should use AWS WAF to safeguard the website from SQL injection or cross-site scripting. AWS WAF is a web application firewall that helps protect web applications from common web exploits that could affect availability, compromise security, or consume excessive resources. The company can use AWS WAF to create custom rules that block malicious requests that match certain patterns, such as SQL injection or cross-site scripting. AWS WAF can be applied to web applications that are behind an Application Load Balancer, Amazon CloudFront, or Amazon API Gateway. Amazon GuardDuty, AWS Trusted Advisor, and Amazon Inspector are not the best services to use for this purpose. Amazon GuardDuty is a threat detection service that monitors for malicious activity and unauthorized behavior across the AWS accounts and resources. AWS Trusted Advisor is a service that provides best practice recommendations for cost optimization, performance, security, and fault tolerance. Amazon Inspector is a service that assesses the security and compliance of applications running on Amazon EC2 instances12

Question: 215

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

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

Answer: B

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

Question: 216

A company runs a database on Amazon Aurora in the us-east-1 Region. The company has a disaster recovery requirement that the database be available in another Region.
Which solution meets this requirement with minimal disruption to the database operations?

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

Answer: B

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

Question: 217

Which AWS service requires the customer to patch the guest operating system?

A. AWS Lambda
B. Amazon OpenSearch Service
C. Amazon EC2
D. Amazon ElastiCache

Answer: C

Explanation:
The AWS service that requires the customer to patch the guest operating system is Amazon EC2. Amazon EC2 is a service that provides scalable compute capacity in the cloud, and allows customers to launch and run virtual servers, called instances, with a variety of operating systems, configurations, and specifications. The customer is responsible for patching and updating the guest operating system and any applications that run on the EC2 instances, as part of the security in the cloud. AWS Lambda, Amazon OpenSearch Service, and Amazon ElastiCache are not services that require the customer to patch the guest operating system. AWS Lambda is a serverless compute service that allows customers to run code without provisioning or managing servers. Amazon OpenSearch Service is a fully managed service that makes it easy to deploy, operate, and scale OpenSearch clusters in the AWS Cloud. Amazon ElastiCache is a fully managed service that provides in-memory data store and cache solutions, such as Redis and Memcached. These services are managed by AWS, and AWS is responsible for patching and updating the underlying infrastructure and software.

Question: 218

Which benefit of the AWS Cloud helps companies achieve lower usage costs because of the aggregate usage of all AWS users?

A. No need to guess capacity
B. Ability to go global in minutes
C. Economies of scale
D. Increased speed and agility

Answer: A

Explanation:
The benefit of the AWS Cloud that helps companies achieve lower usage costs because of the aggregate usage of all AWS users is economies of scale. Economies of scale means that AWS can achieve lower costs and higher efficiency by operating at a massive scale and passing the savings to the customers. AWS leverages the aggregate usage of all AWS users to negotiate better prices with hardware vendors, optimize power consumption, and improve operational processes. As a result, AWS can offer lower and more flexible pricing options to the customers, such as pay-as-you-go, reserved, and spot pricing models. No need to guess capacity, ability to go global in minutes, and increased speed and agility are other benefits of the AWS Cloud, but they are not directly related to the aggregate usage of all AWS users. No need to guess capacity means that AWS customers can avoid the risk of over-provisioning or under-provisioning resources, and scale up or down as needed. Ability to go global in minutes means that AWS customers can deploy their applications and data in multiple regions around the world, and deliver them to users with high performance and availability. Increased speed and agility means that AWS customers can quickly and easily provision and access AWS resources, and accelerate their innovation and time to market.

Question: 219

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

Question: 220

A company wants to migrate to the AWS Cloud. The company needs the ability to acquire resources when the resources are necessary.
The company also needs the ability to release those resources when the resources are no longer necessary.
Which architecture concept of the AWS Cloud meets these requirements?

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

Answer: A

Explanation:
The architecture concept of the AWS Cloud that meets the requirements of the company that wants to migrate to the AWS Cloud and needs the ability to acquire and release resources as needed is elasticity. Elasticity means that AWS customers can quickly and easily provision and scale up or down AWS resources as their demand changes, without any upfront costs or long-term commitments. AWS

provides various tools and services that enable customers to achieve elasticity, such as Amazon EC2 Auto Scaling, Amazon CloudWatch, and AWS CloudFormation. Elasticity helps customers optimize their performance, availability, and cost efficiency. Availability, reliability, and durability are other architecture concepts of the AWS Cloud, but they are not directly related to the ability to acquire and release resources as needed. Availability means that AWS customers can access their AWS resources and applications whenever and wherever they need them. Reliability means that AWS customers can depend on their AWS resources and applications to functioncorrectly and consistently. Durability means that AWS customers can preserve their data and objects for long periods of time without loss or corruption12

Question: 221

Which AWS service or tool provides recommendations to help users get rightsized Amazon EC2 instances based on historical workload usage data?

A. AWS Pricing Calculator
B. AWS Compute Optimizer
C. AWS App Runner
D. AWS Systems Manager

Answer: B

Explanation:
The AWS service or tool that provides recommendations to help users get rightsized Amazon EC2 instances based on historical workload usage data is AWS Compute Optimizer. AWS Compute Optimizer is a service that analyzes the configuration and performance of the AWS resources, such as Amazon EC2 instances, and provides recommendations for optimal resource types and sizes based on the workload patterns and metrics. AWS Compute Optimizer helps users improve the performance, availability, and cost efficiency of their AWS resources. AWS Pricing Calculator, AWS App Runner, and AWS Systems Manager are not the best services or tools to use for this purpose. AWS Pricing Calculator is a tool that helps users estimate the cost of using AWS services based on their requirements and preferences. AWS App Runner is a service that helps users easily and quickly deploy web applications and APIs without managing any infrastructure. AWS Systems Manager is a service that helps users automate and manage the configuration and operation of their AWS resources and applications34

Question: 222

Which AWS service is designed to help users orchestrate a workflow process for a set of AWS Lambda functions?

A. Amazon DynamoDB
B. AWS CodePipeline
C. AWS Batch
D. AWS Step Functions

Answer: D

Explanation:
The AWS service that is designed to help users orchestrate a workflow process for a set of AWS Lambda functions is AWS Step Functions. AWS Step Functions is a service that helps userscoordinate multiple AWS services into serverless workflows that can be triggered by events, such as messages, API calls, or schedules. AWS Step Functions allows users to create and visualize complex workflows that can

include branching, parallel execution, error handling, retries, and timeouts. AWS Step Functions can integrate with AWS Lambda to orchestrate a sequence of Lambda functions that perform different tasks or logic. Amazon DynamoDB, AWS CodePipeline, and AWS Batch are not the best services to use for orchestrating a workflow process for a set of AWS Lambda functions. Amazon DynamoDB is a fully managed NoSQL database service that provides fast and consistent performance, scalability, and flexibility. AWS CodePipeline is a fully managed continuous delivery service that helps users automate the release process of their applications. AWS Batch is a fully managed service that helps users run batch computing workloads on the AWS Cloud.

Question: 223

Which options are perspectives that include foundational capabilities of the AWS Cloud Adoption Framework (AWS CAF)? (Select TWO.)

A. Sustainability
B. Security
C. Operations
D. Performance efficiency
E. Reliability

Answer: C,D

Explanation:

The options that are perspectives that include foundational capabilities of the AWS Cloud Adoption Framework (AWS CAF) are operations and performance efficiency. The AWS CAF is a guidance that helps organizations design and travel an accelerated path to successful cloud adoption. The AWS CAF organizes the cloud adoption process into six areas of focus, called perspectives, which are business, people, governance, platform, security, and operations. Each perspective is divided into capabilities, which are further divided into skills and responsibilities. The operations perspective focuses on the management and monitoring of the cloud resources and applications, as well as the automation and optimization of the operational processes. The operations perspective capabilities are operations support, operations integration, and service management. The performance efficiency perspective focuses on the selection and configuration of the right cloud resources and services to meet the performance requirements of the applications, as well as the continuous improvement and innovation of the cloud solutions. The performance efficiency perspective capabilities are selection, review, and monitoring. Sustainability, security, and reliability are not perspectives of the AWS CAF, but they are aspects of the AWS Well-Architected Framework. The AWS Well-Architected Framework is a guidance that helps users build and operate secure, reliable, efficient, and cost-effective systems in the cloud. The AWS Well-Architected Framework consists of five pillars, which are operational excellence, security, reliability, performance efficiency, and cost optimization. Sustainability is a cross-cutting theme that applies to all the pillars, and refers to the environmental and social impact of the cloud solutions.

Question: 224

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

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

Answer: D

Explanation:

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

Question: 225

A company needs to host a highly available application in the AWS Cloud. The application runs infrequently for short periods of time.

Which AWS service will meet these requirements with the LEAST amount of operational overhead?

A. Amazon EC2
B. AWS Fargate
C. AWS Lambda
D. Amazon Aurora

Answer: C

Explanation:

The AWS service that will meet the requirements of the company that needs to host a highly available application in the AWS Cloud that runs infrequently for short periods of time with the least amount of operational overhead is AWS Lambda. AWS Lambda is a serverless compute service that allows customers to run code without provisioning or managing servers. The company can use AWS Lambda to create and deploy their application as functions that are triggered by events, such as API calls, messages, or schedules. AWS Lambda automatically scales the compute resources based on the demand, and customers only pay for the compute time they consume. AWS Lambda also simplifies the management and maintenance of the application, as customers do not need to worry about the underlying infrastructure, security, or availability. Amazon EC2, AWS Fargate, and Amazon Aurora are not the best services to use for this purpose. Amazon EC2 is a service that provides scalable compute capacity in the cloud, and allows customers to launch and run virtual servers, called instances, with a variety of operating systems, configurations, and specifications. Amazon EC2 requires customers to provision and manage the instances, and pay for the instance hours they use, regardless of the application usage. AWS Fargate is a serverless compute engine for containers that allows customers to run containerized applications without managing servers or clusters. AWS Fargate requires customers to specify the amount of CPU and memory resources for each container, and pay for the resources they allocate, regardless of the application usage. Amazon Aurora is a fully managed relational database service that provides high performance, availability, and compatibility. Amazon Aurora is not a compute service, and it is not suitable for hosting an application that runs infrequently for short periods of time12

Question: 226

A company is planning a migration to the AWS Cloud and wants to examine the costs that are associated with different workloads.
Which AWS tool will meet these requirements?

A. AWS Budgets
B. AWS Cost Explorer
C. AWS Pricing Calculator
D. AWS Cost and Usage Report

Answer: C

Explanation:
The AWS tool that will meet the requirements of the company that is planning a migration to the AWS Cloud and wants to examine the costs that are associated with different workloads is AWS Pricing Calculator. AWS Pricing Calculator is a tool that helps customers estimate the cost of using AWS services based on their requirements and preferences. The company can use AWS Pricing Calculator to compare the costs of different AWS services and configurations, such as Amazon EC2, Amazon S3, Amazon RDS, and more. AWS Pricing Calculator also provides detailed breakdowns of the cost components, such as compute, storage, network, and data transfer. AWS Pricing Calculator helps customers plan and optimize their cloud budget and migration strategy. AWS Budgets, AWS Cost Explorer, and AWS Cost and Usage Report are not the best tools to use for this purpose. AWSBudgets is a tool that helps customers monitor and manage their AWS spending and usage against predefined budget limits and thresholds. AWS Cost Explorer is a tool that helps customers analyze and visualize their AWS spending and usage trends over time. AWS Cost and Usage Report is a tool that helps customers access comprehensive and granular information about their AWS costs and usage in a CSV or Parquet file. These tools are more useful for tracking and optimizing the existing AWS costs and usage, rather than estimating the costs of different workloads34

Question: 227

A company is hosting a web application on Amazon EC2 instances. The company wants to implement custom conditions to filter and control inbound web traffic.
Which AWS service will meet these requirements?

A. Amazon GuardDuty
B. AWSWAF
C. Amazon Macie
D. AWS Shield

Answer: B

Explanation:
The AWS service that will meet the requirements of the company that is hosting a web application on Amazon EC2 instances and wants to implement custom conditions to filter and control inbound web traffic is AWS WAF. AWS WAF is a web application firewall that helps protect web applications from common web exploits that could affect availability, compromise security, or consume excessive resources. The company can use AWS WAF to create custom rules that block malicious requests that match certain patterns, such as SQL injection or cross-site scripting. AWS WAF can be applied to web applications that are behind an Application Load Balancer, Amazon CloudFront, or Amazon API Gateway. Amazon GuardDuty, Amazon Macie, and AWS Shield are not the best services to use for this purpose. Amazon

GuardDuty is a threat detection service that monitors for malicious activity and unauthorized behavior across the AWS accounts and resources. Amazon Macie is a data security and data privacy service that uses machine learning and pattern matching to discover, classify, and protect sensitive data stored in Amazon S3. AWS Shield is a managed distributed denial of service (DDoS) protection service that safeguards web applications running on AWS. These services are more useful for detecting and preventing different types of threats and attacks, rather than filtering and controlling inbound web traffic based on custom conditions.

Question: 228

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

A. AmazonKendra
B. Amazon Lex
C. AmazonTextract
D. AmazonPolly

Answer: B

Explanation:
The AWS service that will meet the requirements of the company that wants to create a chatbot and integrate the chatbot with its current web application is Amazon Lex. Amazon Lex is a service that helps customers build conversational interfaces using voice and text. The company can use Amazon Lex to create a chatbot that can understand natural language and respond to user requests, using the same deep learning technologies that power Amazon Alexa. Amazon Lex also provides easy integration with other AWS services, such as Amazon Comprehend, Amazon Polly, and AWS Lambda, as well as popular platforms, such as Facebook Messenger, Slack, and Twilio. Amazon Lex helps customers create engaging and interactive chatbots for their web applications. Amazon Kendra, Amazon Textract, and Amazon Polly are not the best services to use for this purpose. Amazon Kendra is a service that helps customers provide accurate and natural answers to natural language queries using machine learning. Amazon Textract is a service that helps customers extract text and data from scanned documents using optical character recognition (OCR) and machine learning. Amazon Polly is a service that helps customers convert text into lifelike speech using deep learning. These services are more useful for different types of natural language processing and generation tasks, rather than creating and integrating chatbots.

Question: 229

Which AWS service is used to temporarily provide federated security credentials to a

A. Amazon GuardDuty
B. AWS Simple Token Service (AWS STS)
C. AWS Secrets Manager
D. AWS Certificate Manager

Answer: B

Explanation:
The AWS service that is used to temporarily provide federated security credentials to a user is AWS Security Token Service (AWS STS). AWS STS is a service that enables customers to request temporary, limited-privilege credentials for AWS Identity and Access Management (IAM)

users or for users that they authenticate (federated users). The company can use AWS STS to grant federated users access to AWS resources without creating permanent IAM users or sharing long-term credentials. AWS STS helps customers manage and secure access to their AWS resources for federated users. Amazon GuardDuty, AWS Secrets Manager, and AWS Certificate Manager are not the best services to use for this purpose. Amazon GuardDuty is a threat detection service that monitors for malicious activity and unauthorized behavior across the AWS accounts and resources.AWS Secrets Manager is a service that helps customers manage and rotate secrets, such as database credentials, API keys, and passwords. AWS Certificate Manager is a service that helps customers 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. These services are more useful for different types of security and compliance tasks, rather than providing temporary federated security credentials to a user.

Question: 230

A company wants to securely store Amazon RDS database credentials and automatically rotate user passwords periodically.
Which AWS service or capability will meet these requirements?

A. Amazon S3
B. AWS Systems Manager Parameter Store
C. AWS Secrets Manager
D. AWS CloudTrail

Answer: B

Explanation:
AWS Secrets Manager is a service that helps you protect access to your applications, services, and IT resources. This service enables you to easily rotate, manage, and retrieve database credentials, API keys, and other secrets throughout their lifecycle1. Amazon S3 is a storage service that does not offer automatic rotation of credentials. AWS Systems Manager Parameter Store is a service that provides secure, hierarchical storage for configuration data management and secrets management2, but it does not offer automatic rotation of credentials. AWS CloudTrail is a service that enables governance, compliance, operational auditing, and risk auditing of your AWS account3, but it does not store or rotate credentials.

Question: 231

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

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

Answer: A

Explanation:
Provisioning additional EC2 instances in other Availability Zones is a way to make the application highly available, as it reduces the impact of failures and increases fault tolerance. Configuring an Application Load Balancer and assigning the EC2 instance as the ALB's target is a

way to distribute traffic among multiple instances, but it does not make the application highly available if there is only one instance. Using an Amazon Machine Image to create the EC2 instance is a way to launch a virtual server with a preconfigured operating system and software, but it does not make the application highly available by itself. Provisioning the application by using an EC2 Spot Instance is a way to use spare EC2 capacity at up to 90% off the On-Demand price, but it does not make the application highly available, as Spot Instances can be interrupted by EC2 with a two-minute notification.

Question: 232

Which AWS service provides a highly accurate and easy-to-use enterprise search service that is powered by machine learning (ML)?

A. Amazon Kendra
B. Amazon SageMaker
C. Amazon Augmented AI (Amazon A2I)
D. Amazon Polly

Answer: A

Explanation:
Amazon Kendra is a service that provides a highly accurate and easy-to-use enterprise search service that is powered by machine learning. Kendra delivers powerful natural language search capabilities to your websites and applications so your end users can more easily find the information they need within the vast amount of content spread across your company. Amazon SageMaker is a service that provides a fully managed platform for data scientists and developers to quickly and easily build, train, and deploy machine learning models at any scale. Amazon Augmented AI (Amazon A2I) is a service that makes it easy to build the workflows required for human review of ML predictions. Amazon A2I brings human review to all developers, removing the undifferentiated heavy lifting associated with building human review systems or managing large numbers of human reviewers. Amazon Polly is a service that turns text into lifelike speech, allowing you to create applications that talk, and build entirely new categories of speech-enabled products. None of these services provide an enterprise search service that is powered by machine learning.

Question: 233

A company provides a software as a service (SaaS) application. The company has a new customer that is based in a different country.
The new customer's data needs to be hosted in that country.
Which AWS service or infrastructure component should the company use to meet this requirement?

A. AWS Shield
B. Amazon S3 Object Lock
C. AWS Regions
D. Placement groups

Answer: C

Explanation:
AWS Regions are geographic areas around the world where AWS has clusters of data centers. Each AWS Region consists of multiple, isolated, and physically separate AZ's within a geographic area. By hosting the customer's data in a specific AWS Region, the company can

meet the requirement of hosting the data in the customer's country. AWS Shield is a service that provides always-on detection and automatic inline mitigations that minimize application downtime and latency, so there is no need to engage AWS Support to benefit from DDoS protection. Amazon S3 Object Lock is afeature that allows you to store objects using a write-once-read-many (WORM) model. You can use it to prevent an object from being deleted or overwritten for a fixed amount of time or indefinitely. Placement groups are logical grouping of instances within a single Availability Zone. Placement groups enable applications to participate in a low-latency, 10 Gbps network. None of these services or infrastructure components can help the company host the customer's data in a different country.

Question: 234

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

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

Answer: B

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

Question: 235

A company has developed a distributed application that recovers gracefully from interruptions. The application periodically processes large volumes of data by using multiple Amazon EC2 instances. The application is sometimes idle for months.
Which EC2 instance purchasing option is MOST cost-effective for this use case?

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

Answer: B

Explanation:
Spot Instances are instances that use spare EC2 capacity that is available for up to 90% off the On- Demand price. Because Spot Instances can be interrupted by EC2 with two minutes of notification when EC2 needs the capacity back, you can use them for applications that have flexible start and end times, or that can withstand interruptions5. This option is most cost-effective for the use case described in the question. Reserved Instances are instances that you purchase for a one-year or three-year term, and pay a lower hourly rate compared to On-Demand Instances. This option is suitable for applications that have steady state or predictable usage. Dedicated Instances are instances that run on hardware that's dedicated to a single customer within an Amazon VPC. This option is suitable for applications that have stringent regulatory or compliance requirements. On- Demand Instances are instances that you pay for by the second, with no long-

term commitments or upfront payments. This option is suitable for applications that have unpredictable or intermittent workloads.

Question: 236

A company is preparing to launch a redesigned website on AWS. Users from around the world will download digital handbooks from the website.
Which AWS solution should the company use to provide these static files securely?

A. Amazon Kinesis Data Streams
B. Amazon CloudFront with Amazon S3
C. Amazon EC2 instances with an Application Load Balancer
D. Amazon Elastic File System (Amazon EFS)

Answer: B

Explanation:
Amazon CloudFront with Amazon S3 is a solution that allows you to provide static files securely to users from around the world. Amazon CloudFront is a fast content delivery network (CDN) service that securely delivers data, videos, applications, and APIs to customers globally with low latency, high transfer speeds, all within a developer-friendly environment. Amazon S3 is an object storage service that offers industry-leading scalability, data availability, security, and performance. You can use Amazon S3 to store and retrieve any amount of data from anywhere. You can also configure Amazon S3 to work with Amazon CloudFront to distribute your content to edge locations near your users for faster delivery and lower latency. Amazon Kinesis Data Streams is a service that enables you to build custom applications that process or analyze streaming data for specialized needs. This option is not relevant for providing static files securely. Amazon EC2 instances with an Application Load Balancer is a solution that allows you to distribute incoming traffic across multiple targets, such as EC2 instances, in multiple Availability Zones. This option is suitable for dynamic web applications, but not necessary for static files. Amazon Elastic File System (Amazon EFS) is a service that provides a simple, scalable, fully managed elastic NFS file system for use with AWS Cloud services and onpremises resources. This option is not relevant for providing static files securely.

Question: 237

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

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

Answer: D

Explanation:
Amazon ElastiCache is a service that offers fully managed in-memory data store and cache services that deliver sub-millisecond response times to applications. You can use Amazon ElastiCache to improve the performance of your applications by retrieving data from fast, managed, in-memory data stores, instead of relying entirely on slower disk-based databases. Amazon Aurora is a relational database service that combines the performance and availability of high-end commercial databases with the simplicity and cost-effectiveness of open source databases. Amazon RDS is a service that makes it easy to set up, operate, and scale a relational database in the cloud. Amazon

DynamoDB is a key-value and document database that delivers single-digit millisecond performance at any scale. None of these services are in-memory data store services.

Question: 238

Which AWS service or tool offers consolidated billing?

A. AWS Artifact
B. AWS Budgets
C. AWS Organizations
D. AWS Trusted AdvisorA company wants to limit its employees' AWS access to a portfolio of predefined AWS resources.

Answer: C

Explanation:
AWS Organizations is a service that enables you to consolidate multiple AWS accounts into an organization that you create and centrally manage. With AWS Organizations, you can create a single payment method for all the AWS accounts in your organization through consolidated billing. Consolidated billing enables you to see a combined view of AWS charges incurred by all accounts in your organization, as well as get a detailed cost report for each individual AWS account associated with your organization. AWS Artifact is a service that provides on-demand access to AWS' security and compliance reports and select online agreements. AWS Budgets is a service that enables you to plan your service usage, service costs, and instance reservations. AWS Trusted Advisor is a servicethat provides real-time guidance to help you provision your resources following AWS best practices. None of these services or tools offer consolidated billing.

Question: 239

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

A. AWS Config
B. AWS software development kits (SDKs)
C. AWS Service Catalog
D. AWS AppSync

_____ Answer: C

Explanation:
AWS Service Catalog is a service that allows you to create and manage catalogs of IT services that are approved for use on AWS. You can use AWS Service Catalog to centrally manage commonly deployed IT services and help your organization achieve consistent governance and meet your compliance requirements, while enabling users to quickly deploy only the approved IT services they need1. AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. AWS software development kits (SDKs) are tools that enable you to easily integrate your applications with AWS services using your preferred programming language. AWS AppSync is a service that simplifies application development by letting you create a flexible API to securely access, manipulate, and combine data from one or more data sources. None of these services can help you limit your employees' AWS access to a portfolio of predefined AWS resources.

Question: 240

A company is running workloads for multiple departments within a single VPC. The company needs to be able to bill each department for its resource usage.
Which action should the company take to accomplish this goal with the LEAST operational overhead?

A. Add a department tag to each resource and configure cost allocation tags.
B. Move each department resource to its own VPC.
C. Move each department resource to its own AWS account.
D. Use AWS Organizations to get a billing report for each department.

Answer: A

Explanation:
Adding a department tag to each resource and configuring cost allocation tags is an action that can help you accomplish the goal of billing each department for its resource usage with the least operational overhead. Tags are simple labels consisting of a key and an optional value that you can assign to AWS resources. You can use tags to organize your resources and track your AWS costs ona detailed level. Cost allocation tags enable you to track your AWS costs on a detailed level. After you activate cost allocation tags, AWS uses the cost allocation tags to organize your resource costs on your cost allocation report, to make it easier for you to categorize and track your AWS costs2. Moving each department resource to its own VPC or its own AWS account is an action that can help you isolate and control the resources for each department, but it would incur more operational overhead than using tags. Using AWS Organizations to get a billing report for each department is an action that can help you consolidate billing and payment across multiple AWS accounts, but it would not help you bill each department for its resource usage within a single VPC.

Question: 241

A large company has multiple departments. Each department has its own AWS account. Each department has purchased Amazon EC2 Reserved Instances. Some departments do not use all the Reserved Instances that they purchased, and other departments need more Reserved Instances than they purchased.
The company needs to manage the AWS accounts for all the departments so that the departments can share the Reserved Instances.
Which AWS service or tool should the company use to meet these requirements?

A. AWS Systems Manager
B. Cost Explorer
C. AWS Trusted Advisor
D. AWS Organizations

Answer: D

Explanation:
AWS Organizations is a service that enables you to consolidate multiple AWS accounts into an organization that you create and centrally manage. With AWS Organizations, you can apply service control policies (SCPs) across multiple AWS accounts to restrict what services and actions users and roles can access. You can also use AWS Organizations to enable features such as consolidated billing, AWS Config rules

and conformance packs, and AWS CloudFormation StackSets across multiple accounts3. One of the benefits of using AWS Organizations is that you can share your Reserved Instances (RIs) with all of the accounts in your organization. This enables you to take advantage of the billing benefits of RIs without having to specify which account will use them4. AWS Systems Manager is a service that gives you visibility and control of your infrastructure on AWS. Cost Explorer is a tool that enables you to visualize, understand, and manage your AWS costs and usage over time. AWS Trusted Advisor is a service that provides real-time guidance to help you provision your resources following AWS best practices. None of these services or tools can help you manage the AWS accounts for all the departments so that the departments can share the Reserved Instances.

Question: 242

A manufacturing company has a critical application that runs at a remote site that has a slow internet connection. The company wants to migrate the workload to AWS. The application is sensitive to latency and interruptions in connectivity. The company wants a solution that can host this application with minimum latency.
Which AWS service or feature should the company use to meet these requirements?

A. Availability Zones
B. AWS Local Zones
C. AWS Wavelength
D. AWS Outposts

Answer: D

Explanation:
AWS Outposts is a service that offers fully managed and configurable compute and storage racks built with AWS-designed hardware that allow you to run your workloads on premises and seamlessly connect to AWS services in the cloud. AWS Outposts is ideal for workloads that require low latency, local data processing, or local data storage. With AWS Outposts, you can use the same AWS APIs, tools, and infrastructure across on premises and the cloud to deliver a truly consistent hybrid experience5. Availability Zones are isolated locations within each AWS Region that are engineered to be fault-tolerant and provide high availability. AWS Local Zones are extensions of AWS Regions that are placed closer to large population, industry, and IT centers where no AWS Region exists today. AWS Wavelength is a service that enables developers to build applications that deliver ultra-low latency to mobile devices and users by deploying AWS compute and storage at the edge of the 5G network. None of these services or features can help you host a critical application with minimum latency at a remote site that has a slow internet connection.

Question: 243

Which AWS services can a company use to host and run a MySQL database? (Select TWO.)

A. Amazon RDS
B. Amazon DynamoDB
C. Amazon S3
D. Amazon EC2
E. Amazon MQ

Answer: A,D

Explanation:
Amazon RDS and Amazon EC2 are two AWS services that you can use to host and run a MySQL database. Amazon RDS is a service that makes it easy to set up, operate, and scale a relational database in the cloud. You can use Amazon RDS to launch a MySQL database instance and let Amazon RDS manage common database tasks such as backups, patching, scaling, and replication6. Amazon EC2 is a service that provides secure, resizable compute capacity in the cloud. You can use Amazon EC2 to launch a virtual server and install MySQL software on it. You have complete control over your database configuration, but you are responsible for managing and maintaining the database software and the underlying infrastructure7. Amazon DynamoDB is a key-value and document database that delivers single-digit millisecond performance at any scale. Amazon S3 is an object storage service that offers industry-leading scalability, data availability, security, and performance. Amazon MQ is a managed message broker service for Apache ActiveMQ. None of these services can help you host and run a MySQL database.

Question: 244

A company wants its workload to perform consistently and correctly. Which benefit of AWS Cloud computing does this goal represent?

A. Security
B. Elasticity
C. Pay-as-you-go pricing
D. Reliability

Answer: D

Explanation:
Reliability is the benefit of AWS Cloud computing that ensures the workload performs consistently and correctly. According to the AWS Cloud Practitioner Essentials course, reliability means "theability of a system to recover from infrastructure or service disruptions, dynamically acquire computing resources to meet demand, and mitigate disruptions such as misconfigurations or transient network issues."1 Elasticity, security, and pay-as-you-go pricing are also benefits of AWS Cloud computing, but they do not directly relate to the goal of consistent and correct performance.

Question: 245

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

Explanation:
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.

Question: 246

Which design principle is included in the operational excellence pillar of the AWS Well-Architected

Framework?

A. Create annotated documentation.
B. Anticipate failure.
C. Ensure performance efficiency.
D. Optimize costs.

Answer: A

Explanation:
Create annotated documentation is the design principle that is included in the operational excellence pillar of the AWS Well-Architected Framework. According to the AWS Well-Architected Framework whitepaper, creating annotated documentation means "documenting your workload so that the team understands the architecture, how to operate the workload, and how the workload delivers value to customers."3 Anticipate failure, ensure performance efficiency, and optimize costs are design principles that belong to other pillars of the AWS Well-Architected Framework, such as reliability, performance efficiency, and cost optimization.

Question: 247

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

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

Answer: D

Explanation:
Deploying the application by using multiple Availability Zones is the best way to increase resilience for the application. According to the Amazon RDS User Guide, "Amazon RDS provides high availability and failover support for DB instances using Multi-AZ deployments. In a

Multi-AZ deployment, Amazon RDS automatically provisions and maintains a synchronous standby replica in a different Availability Zone. The primary DB instance is synchronously replicated across Availability Zones to a standby replica to provide data redundancy, eliminate I/O freezes, and minimize latency spikes during system backups."4 Deploying a copy of the application in another AWS account, using multiple VPCs, or using multiple subnets do not provide the same level of resilience as using multiple Availability Zones.

Question: 248

Which AWS services or tools are designed to protect a workload from SQL injections, cross-site scripting, and DDoS attacks? (Select TWO.)

A. VPC endpoint
B. Virtual private gatewayQ C. AWS Shield Standard
C. AWS Config
D. AWS WAF

Answer: C

Explanation:
AWS Shield Standard and AWS WAF are the AWS services or tools that are designed to protect a workload from SQL injections, cross-site scripting, and DDoS attacks. According to the AWS Shield Developer Guide, "AWS Shield is a managed Distributed Denial of Service (DDoS) protection service that safeguards applications running on AWS. AWS Shield provides always-on detection and automatic inline mitigations that minimize application downtime and latency, so there is no need to engage AWS Support to benefit from DDoS protection."5 According to the AWS WAF Developer Guide, “AWS WAF is a web application firewall that helps protect your web applications or APIs against common web exploits that may affect availability, compromise security, or consume excessive resources. AWS WAF gives you control over how traffic reaches your applications by enabling you to create security rules that block common attack patterns, such as SQL injection or cross-site scripting, and rules that filter out specific traffic patterns you define.” VPC endpoint, virtual private gateway, and AWS Config are not designed to protect a workload from these types of attacks.

Question: 249

A company wants guidance to optimize the cost and performance of its current AWS environment. Which AWS service or tool should the company use to identify areas for optimization?

A. Amazon QuickSight
B. AWS Trusted Advisor
C. AWS Organizations
D. AWS Budgets

Answer: B

Explanation:
AWS Trusted Advisor is the AWS service or tool that the company should use to identify areas for optimization. According to the AWS Trusted Advisor User Guide, “AWS Trusted Advisor is an online tool that provides you real time guidance to help you provision your resources following AWS best practices. AWS Trusted Advisor checks help optimize your AWS infrastructure, increase security and performance, reduce your overall costs, and monitor service limits.” Amazon QuickSight, AWS Organizations, and AWS Budgets are not

designed to provide optimization recommendations for the current AWS environment.

Question: 250 A new AWS user who has little cloud experience wants to build an application by using AWS services. The user wants to learn how to implement specific AWS services from other customer examples. The user also wants to ask questions to AWS experts.

Which AWS service or resource will meet these requirements?

A. AWS Online Tech Talks
B. AWS documentation
C. AWS Marketplace
D. AWS Health Dashboard

Answer: A

Explanation:
AWS Online Tech Talks are online presentations that cover a broad range of topics at varying technical levels and provide a live Q&A session with AWS experts. They are a great resource for new AWS users who want to learn how to implement specific AWS services from other customer examples and ask questions to AWS experts. AWS documentation, AWS Marketplace, and AWS Health Dashboard do not offer the same level of interactivity and guidance as AWS Online Tech Talks. Source: AWS Online Tech Talks

Topic 2, Exam Pool B

Question: 251

A user discovered that an Amazon EC2 instance is missing an Amazon Elastic Block Store (Amazon EBS) data volume. The user wants to determine when the EBS volume was removed.
Which AWS service will provide this information?

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

Answer: A

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

Question: 252

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

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

Answer: B

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

Question: 253

How should the company deploy the application to meet these requirements?

A. Ina single Availability Zone
B. On AWS Direct Connect
C. On Reserved Instances
D. In multiple Availability Zones

Answer: D

Explanation:
Deploying the application in multiple Availability Zones is the best way to ensure high availability for the application. Availability Zones are isolated locations within an AWS Region that are engineered to be fault-tolerant from failures in other Availability Zones. By deploying the application in multiple Availability Zones, the company can reduce the impact of outages and increase the resilience of the application. Deploying the application in a single Availability Zone, on AWS Direct Connect, or on Reserved Instances does not provide the same level of high availability as deploying the application in multiple Availability Zones. Source: Availability Zones

Question: 254

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

Which AWS service meets these requirements?

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

Answer: B

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

Question: 255

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 is a service that provides real-time guidance to help you provision your resources following AWS best practices. AWS Trusted Advisor provides recommended actions in five categories: cost optimization, performance, security, fault tolerance, and service quotas. Cost optimization helps you reduce your overall AWS costs by identifying idle and underutilized resources. Service quotas helps you monitor and manage your usage of AWS service quotas and request quota increases. Operating system patches, repetitive tasks, and account activity records are not categories that AWS Trusted Advisor provides recommended actions for. Source: [AWS Trusted Advisor]

Question: 256

Which actions are examples of a company's effort to right size its AWS resources to control cloud costs? (Select TWO.)

A. Switch from Amazon RDS to Amazon DynamoDB to accommodate NoSQL datasets.Q B. Base the selection of Amazon EC2 instance types on past utilization patterns.
B. Use Amazon S3 Lifecycle policies to move objects that users access infrequently to lower-cost storage tiers.
C. Use Multi-AZ deployments for Amazon RDS.
D. Replace existing Amazon EC2 instances with AWS Elastic Beanstalk.

Answer: B,C

Explanation:
Basing the selection of Amazon EC2 instance types on past utilization patterns is a way to right size the AWS resources and optimize the performance and cost. Using Amazon S3 Lifecycle policies to move objects that users access infrequently to lower-cost storage tiers is

another way to reduce the storage costs and align them with the business value of the data. These two actions are recommended by the AWS Cost Optimization Pillar1. Switching from Amazon RDS to Amazon DynamoDB is not necessarily a cost-saving action, as it depends on the use case and the data model. Using Multi-AZ deployments for Amazon RDS is a way to improve the availability and durability of the database, but it also increases the cost. Replacing existing Amazon EC2 instanceswith AWS Elastic Beanstalk is a way to simplify the deployment and management of the application, but it does not affect the cost of the underlying EC2 instances.

Question: 257

A company has a single Amazon EC2 instance. The company wants to adopt a highly available architecture.
What can the company do to meet this requirement?

A. Scale vertically to a larger EC2 instance size.
B. Scale horizontally across multiple Availability Zones.
C. Purchase an EC2 Dedicated Instance.
D. Change the EC2 instance family to a compute optimized instance.

Answer: B

Explanation:
Scaling horizontally across multiple Availability Zones is a way to adopt a highly available architecture, as it increases the fault tolerance and resilience of the application. Scaling vertically to a larger EC2 instance size is a way to improve the performance of the application, but it does not improve the availability. Purchasing an EC2 Dedicated Instance is a way to isolate the instance from other AWS customers, but it does not improve the availability. Changing the EC2 instance family to a compute optimized instance is a way to optimize the instance type for the workload, but it does not improve the availability. These concepts are explained in the AWS Well-Architected Framework2.

Question: 258 A company is running an application that is hosted on Amazon EC2 instances. The usage of the EC2 instances is higher during daytime hours than nighttime hours. The company wants to optimize the number of EC2 instances based on this usage pattern.

Which AWS service or instance purchasing option should the company use to meet these requirements?

A. Spot Instances
B. Reserved Instances
C. AWS CloudFormation
D. AWS Auto Scaling

Answer: D

Explanation:
AWS Auto Scaling is the AWS service that allows users to optimize the number of EC2 instances based on the usage pattern, as it automatically adjusts the capacity to maintain steady and predictable performance at the lowest possible cost. Spot Instances are a way to reduce the cost ofEC2 instances by bidding on unused EC2 capacity, but they are not suitable for applications that require steady and reliable performance. Reserved Instances are a way to reduce the cost of EC2 instances by committing to a certain amount of usage for a period of time, but they are not flexible to adjust to the usage pattern. AWS CloudFormation is a way to automate the creation and management of AWS resources, but it does not optimize the number of EC2 instances based on the usage pattern. These concepts are explained in the AWS Cloud Practitioner Essentials course3.

Question: 259

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

Question: 260

Which AWS service or tool provides on-demand access to AWS security and compliance reports and AWS online agreements?

A. AWS Artifact
B. AWS Trusted Advisor
C. Amazon Inspector
D. AWS Billing console

Answer: A

Explanation:
AWS Artifact is the AWS service or tool that provides on-demand access to AWS security and compliance reports and AWS online agreements. AWS Trusted Advisor is a tool that provides realtime guidance to help users provision their resources following AWS best practices. Amazon Inspector is a service that helps users improve the security and compliance of their applications. AWS Billing console is a tool that helps users manage their AWS costs and usage. These concepts are explained in the AWS Cloud Practitioner Essentials course3.

Question: 261

A company wants to move its iOS application development and build activities to AWS. Which AWS service or resource should the

company use for these activities?

A. AWS CodeCommit
B. Amazon EC2 M1 Mac instances
C. AWS Amplify
D. AWS App Runner

Answer: B

Explanation:
Amazon EC2 M1 Mac instances are the AWS service or resource that the company should use for its iOS application development and build activities, as they enable users to run macOS on AWS and access a broad and growing set of AWS services. AWS CodeCommit is a service that provides a fully managed source control service that hosts secure Git-based repositories. AWS Amplify is a set of tools and services that enable developers to build full-stack web and mobile applications using AWS. AWS App Runner is a service that makes it easy for developers to quickly deploy containerized web applications and APIs. These concepts are explained in the AWS Developer Tools page4.

Question: 262

Which statements explain the business value of migration to the AWS Cloud? (Select TWO.)

A. The migration of enterprise applications to the AWS Cloud makes these applications automatically available on mobile devices.S B. AWS availability and security provide the ability to improve service level agreements (SLAs) while reducing risk and unplanned downtime.
B. Companies that migrate to the AWS Cloud eliminate the need to plan for high availability and disaster recovery.
C. Companies that migrate to the AWS Cloud reduce IT costs related to infrastructure, freeing budget for reinvestment in otherareas.
D. Applications are modernized because migration to the AWS Cloud requires companies to rearchitect and rewrite allenterprise applications.

Answer: B,D

Explanation:
B and D are correct because AWS availability and security enable customers to improve their SLAs while reducing risk and unplanned downtime1, and AWS reduces IT costs related to infrastructure, allowing customers to reinvest in other areas2. A is incorrect because migrating to the AWS Cloud does not automatically make applications available on mobile devices, as it depends on the application design and compatibility. C is incorrect because companies that migrate to the AWS Cloud still need to plan for high availability and disaster recovery, as AWS is a shared responsibility model3. E is incorrect because migrating to the AWS Cloud does not require companies to rearchitect and rewrite all enterprise applications, as AWS offers different migration strategies depending on the application complexity and business objectives4.

Question: 263

Which AWS service is designed to help users build conversational interfaces into applications using voice and text?

A. Amazon Lex

B. Amazon Transcribe
C. Amazon Comprehend
D. Amazon Timestream

Answer: A

Explanation:
A is correct because Amazon Lex is the AWS service that helps users build conversational interfaces into applications using voice and text. B is incorrect because Amazon Transcribe is the AWS service that helps users convert speech to text. C is incorrect because Amazon Comprehend is the AWS service that helps users analyze text using natural language processing. D is incorrect because Amazon Timestream is the AWS service that helps users collect, store, and process time series data.

Question: 264 A company wants to develop a shopping application that records customer orders. The application needs to use an AWS managed database service to store data.

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

A. Amazon RDS
B. Amazon Redshift
C. Amazon ElastiCache
D. Amazon Neptune

Answer: A

Explanation:
A is correct because Amazon RDS is the AWS service that provides a managed relational database service that supports various database engines, such as MySQL, PostgreSQL, Oracle, and SQL Server. B is incorrect because Amazon Redshift is the AWS service that provides a managed data warehouse service that is optimized for analytical queries. C is incorrect because Amazon ElastiCache is the AWS service that provides a managed in-memory data store service that supports Redis and Memcached. D is incorrect because Amazon Neptune is the AWS service that provides a managed graph database service that supports property graph and RDF models.

Question: 265

A company wants to use Amazon EC2 instances for a stable production workload that will run for 1 year.
Which instance purchasing option meets these requirements MOST cost-effectively?

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

Answer: B

Explanation:
B is correct because Reserved Instances are the instance purchasing option that offers the most costeffective way to use Amazon EC2

instances for a stable production workload that will run for 1 year, as they provide significant discounts compared to On-Demand Instances in exchange for a commitment to use a specific amount of computing power for a period of time. A is incorrect because Dedicated Hosts are the instance purchasing option that allows customers to use physical servers that are fully dedicated to their use, which is more expensive and less flexible than Reserved Instances. C is incorrect because On-Demand Instances are the instance purchasing option that allows customers to pay for compute capacity by the hour or second with no long-term commitments, which is more suitable for short-term, variable, and unpredictable workloads. D is incorrect because Spot Instances are the instance purchasing option that allows customers to bid on spare Amazon EC2 computing capacity, which is more suitable for flexible, scalable, and fault- tolerant workloads that can tolerate interruptions.

Question: 266

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:
A and B are correct because AWS CodeCommit is the AWS service that provides a fully managed source control service that hosts secure Git-based repositories1, and AWS CodeDeploy is the AWS service that automates code deployments to any instance, including Amazon EC2 instances and servers running on-premises2. These two services can be used together to store source code and update the running software when the code changes. C is incorrect because Amazon DynamoDB is the AWS service that provides a fully managed NoSQL database service that supports key-value and document data models3. It is not related to storing source code or updating software. D is incorrect because Amazon S3 is the AWS service that provides object storage through a web service interface4. It can be used to store source code, but it does not provide source control features or update software. E is incorrect because Amazon Elastic Container Service (Amazon ECS) is the AWS service that allows users to run, scale, and secure Docker container applications. It can be used to deploy containerized software, but it does not store source code or update software.

Question: 267

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

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

Answer: C

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

Question: 268

A company wants to run its production workloads on AWS. The company needs concierge service, a designated AWS technical account manager (TAM), and technical support that is available 24 hours a day, 7 days a week.
Which AWS Support plan will meet these requirements?

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

Answer: B

Explanation:
B is correct because AWS Enterprise Support is the AWS Support plan that provides concierge service, a designated AWS technical account manager (TAM), and technical support that is available 24 hours a day, 7 days a week. This plan is designed for customers who run mission-critical workloads on AWS and need the highest level of support. A is incorrect because AWS Basic Support is the AWS Support plan that provides customer service and support for billing and account issues, service limit increases, and technical support for a limited set of AWS services. It does not provide concierge service, a designated TAM, or 24/7 technical support. C is incorrect because AWS Business Support is the AWS Support plan that provides customer service and support for billing and account issues, service limit increases, and technical support for all AWS services, as well as access to AWS Trusted Advisor and AWS Support API. It does not provide concierge service or a designated TAM. D is incorrect because AWS Developer Support is the AWS Support plan that provides customer service and support for billing and account issues, service limit increases, and technicalsupport for all AWS services, as well as access to AWS Trusted Advisor. It does not provide concierge service, a designated TAM, or 24/7 technical support.

Question: 269

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

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

D. Security groups

Answer: D

Explanation:

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

Question: 270

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

Which type of migration is this?

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

Answer: C

Explanation:

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

Question: 271

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

A. AWS Direct Connect
B. AWS Outposts
C. AWS Systems Manager hybrid activations

D. AWS Storage Gateway

Answer: B

Explanation:
AWS Outposts is a fully managed service that extends AWS infrastructure, AWS services, APIs, and tools to virtually any datacenter, co-location space, or on-premises facility for a truly consistent hybrid experience. AWS Outposts enables you to run AWS services in your on-premises data center1.

Question: 272

A company provides a web-based ecommerce service that runs in two Availability Zones within a single AWS Region. The web service distributes content that is stored in the Amazon S3 Standard storage class. The company wants to improve the web service's performance globally.
What should the company do to meet this requirement?

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

Answer: B

Explanation:
Amazon CloudFront is a fast content delivery network (CDN) service that securely delivers data, videos, applications, and APIs to customers globally with low latency, high transfer speeds, all within a developer-friendly environment. CloudFront can cache web server content in edge locations, which are located closer to the end users, toimprove the web service's performance globally2.

Question: 273

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

Explanation:
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.

Question: 274

Which AWS service is always available free of charge to users?

A. Amazon Athena
B. AWS Identity and Access Management (IAM)
C. AWS Secrets Manager
D. Amazon ElastiCacheA company has only basic knowledge of AWS technologies.

Answer: B

Explanation:

AWS Identity and Access Management (IAM) is a web service that helps you securely control access to AWS resources for your users. You use IAM to control who can use your AWS resources (authentication) and what resources they can use and in what ways (authorization).

IAM is always available free of charge to users4.

Question: 275

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

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

Answer: D

Explanation:

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

Lightsail provides the simplest way for the company to establish a website on AWS.

Question: 276

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:
Pay-as-you-go pricing is one of the main benefits of AWS. With pay-as-you-go pricing, you pay only for what you use, when you use it. There are no long-term contracts, termination fees, or complex licensing. You replace upfront expenses with lower variable costs and pay only for the resources you consume.

Question: 277

A company manages factory machines in real time. The company wants to use AWS technology to deploy its monitoring applications as close to the factory machines as possible.
Which AWS solution will meet these requirements with the LEAST latency?

A. AWS Outposts
B. Amazon EC2
C. AWS App Runner
D. AWS Batch

Answer: A

Explanation:
AWS Outposts is a fully managed service that extends AWS infrastructure, AWS services, APIs, and tools to virtually any datacenter, co-location space, or on-premises facility for a truly consistent hybrid experience. AWS Outposts enables you to run AWS services in your on-premises data center1.

Question: 278

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

A. Patch management
B. Cost optimization
C. Business technology strategy
D. Physical and environmental controls

Answer: B

Explanation:
The AWS Well-Architected Framework helps you understand the pros and cons of decisions you make while building systems on AWS. By using the Framework, you will learn architectural best practices for designing and operating reliable, secure, efficient, and cost-effective systems in the cloud. The Framework consists of five pillars: operational excellence, security, reliability, performance efficiency, and cost optimization2.

Question: 279

A company is collecting user behavior patterns to identify how to meet goals for sustainability impact.
Which guidelines are best practices for the company to implement to meet these goals? (Select TWO.)

A. Scale infrastructure with user load.
B. Maximize the geographic distance between workloads and user locations.
C. Eliminate creation and maintenance of unused assets.
D. Scale resources with excess capacity and remove auto scaling.
E. Scale infrastructure based on the number of users.

Answer: A,C

Explanation:
To meet the goals for sustainability impact, the company should follow the best practices of scaling infrastructure with user load and eliminating creation and maintenance of unused assets.
Scalinginfrastructure with user load means adjusting the capacity of the infrastructure to match the demand of the users, which can reduce the energy consumption and carbon footprint of the system. Eliminating creation and maintenance of unused assets means avoiding the waste of resources and money on assets that are not needed or used, which can also improve the environmental and economic efficiency of the system3.

Question: 280

A company is running an application on AWS. The company wants to identify and prevent the accidental
Which AWS service or feature will meet these requirements?

A. Amazon GuardDuty
B. Network ACL
C. AWS WAF
D. AWS Network Firewall

Answer: A

Explanation:
Amazon GuardDuty is a threat detection service that continuously monitors for malicious activity and unauthorized behavior to protect your AWS accounts, workloads, and data stored in Amazon S3. With the cloud, the collection and aggregation of account and network activities is simplified, but it can be time consuming for security teams to continuously analyze event log data for potential threats. With GuardDuty, you can automate anomaly detection and get actionable findings to help you protect your AWS resources4.

Question: 281

A company has an Amazon S3 bucket containing images of scanned financial invoices. The company is building an artificial intelligence (AI)-based application on AWS. The company wants the application to identify and read total balance amounts on the invoices.

Which AWS service will meet these requirements?

A. Amazon Forecast
B. Amazon Textract
C. Amazon Rekognition
D. Amazon Lex

Answer: B

Explanation:
Amazon Textract is a service that automatically extracts text and data from scanned documents. Amazon Textract goes beyond simple optical character recognition (OCR) to also identify thecontents of fields in forms and information stored in tables. Amazon Textract can analyze images of scanned financial invoices and extract the total balance amounts, as well as other relevant information, such as invoice number, date, vendor name, etc5.

Question: 282

A company migrated its core application onto multiple workloads in the AWS Cloud. The company wants to improve the application's reliability.
Which cloud design principle should the company implement to achieve this goal?

A. Maximize utilization.
B. Decouple the components.
C. Rightsize the resources.
D. Adopt a consumption model.

Answer: B

Explanation:
Decoupling the components of an application means reducing the dependencies and interactions between them, which can improve the application's reliability, scalability, and performance. Decoupling can be achieved by using services such as Amazon Simple Queue Service (Amazon SQS), Amazon Simple Notification Service (Amazon SNS), and AWS Lambda1

Question: 283

A company is using AWS Organizations to configure AWS accounts.
A company is planning its migration to the AWS Cloud. The company is identifying its capability gaps by using the AWS Cloud Adoption Framework (AWS CAF) perspectives.
Which phase of the cloud transformation journey includes these identification activities?

A. Envision
B. Align
C. Scale
D. Launch

Answer: A

Explanation:
The Envision phase of the cloud transformation journey is where the company defines its vision, business drivers, and desired outcomes for the cloud adoption. The company also identifies its capability gaps by using the AWS Cloud Adoption Framework (AWS CAF) perspectives, which are business, people, governance, platform, security, and operations2.

Question: 284

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

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

Answer: A

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

Question: 285

A company wants to migrate its on-premises application to the AWS Cloud. The company is legally obligated to retain certain data in its onpremises data center.
Which AWS service or feature will support this requirement?

A. AWS Wavelength
B. AWS Local Zones
C. VMware Cloud on AWS
D. AWS Outposts

Answer: D

Explanation:
AWS Outposts is a fully managed service that extends AWS infrastructure, AWS services, APIs, and tools to virtually any datacenter, co-location space, or on-premises facility for a truly consistent hybrid experience. AWS Outposts enables you to run AWS services in your on-premises data center, which can support the requirement of retaining certain data on-premises due to legal obligations5.

Question: 286

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:
To make a subnet public, the company should use an Amazon VPC internet gateway and an Amazon VPC route table. An internet gateway is a horizontally scaled, redundant, and highly available VPC component that allows communication between your VPC and the internet. A route table contains a set of rules, called routes, that are used to determine where network traffic from your subnet or gateway is directed. To enable internet access for a subnet, you need to attach an internet gateway to your VPC and add a route to the internet gateway in the route table associated with the subnet.

Question: 287

A company has a compliance requirement to record and evaluate configuration changes, as well as perform remediation actions on AWS resources.
Which AWS service should the company use?

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

Answer: A

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

Question: 288

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

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

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

Answer: B

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

Question: 289 A company has an application workload that is stateless by design and can sustain occasional downtime. The application performs massively parallel computations.
Which Amazon EC2 pricing model should the company choose for its application to reduce cost?

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

Answer: B

Explanation:
Amazon EC2 Spot Instances let you take advantage of unused EC2 capacity in the AWS cloud. Spot Instances are available at up to a 90% discount compared to On-Demand prices. You can use Spot Instances for various stateless, fault-tolerant, or flexible applications such as big data, containerized workloads, CI/CD, web servers, high-performance computing (HPC), and other test & development workloads. Spot Instances are well-suited for massively parallel computations, as they can provide large amounts of compute capacity at a low cost, and can be interrupted with a two-minute notice3

Question: 290

A company wants to store data with high availability, encrypt the data at rest, and have direct access to the data over the internet.
Which AWS service will meet these requirements MOST cost-effectively?

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

Answer: C

Explanation:

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

Amazon EFS offers two storage classes: the Standard storage class, and the Infrequent Access storage class (EFS IA). EFS IA provides price/performance that is cost-optimized for files not accessed every day. Amazon EFS encrypts data at rest and in transit, and supports direct access over the internet4.

Question: 291

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

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

Answer: B

Explanation:

Server-side encryption is an encryption option that Amazon S3 provides to encrypt data at rest in Amazon S3. With server-side encryption, Amazon S3 encrypts an object before saving it to disk in its data centers and decrypts it when you download the objects. You have three server-side encryption options to choose from: SSE-S3, SSE-C, and SSE-KMS. SSE-S3 uses keys that are managed by Amazon S3. SSE-C allows you to manage your own encryption keys. SSE-KMS uses keys that are managed by AWS Key Management Service (AWS KMS)5.

Question: 292

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

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

Answer: D

Explanation:

AWS Artifact is your go-to, central resource for compliance-related information that matters to you. It provides on-demand access to AWS' security and compliance reports and select online agreements. Reports available in AWS Artifact include our Service Organization Control (SOC) reports, Payment Card Industry (PCI) reports, and certifications from accreditation bodies across geographies and compliance

verticals that validate the implementation and operating effectiveness of AWS security controls. Agreements available in AWS Artifact include the Business Associate Addendum (BAA) and the Nondisclosure Agreement (NDA). You can use AWS Artifact to download the applicable report for AWS security controls and provide it to the auditor.

Question: 293

Which AWS service provides encryption at rest for Amazon RDS and for Amazon Elastic Block Store (Amazon EBS) volumes?

A. AWS Lambda
B. AWS Key Management Service (AWS KMS)
C. AWSWAF
D. Amazon Rekognition

Answer: B

Explanation:
AWS Key Management Service (AWS KMS) is a managed service that enables you to easily encrypt your data. AWS KMS provides you with centralized control of the encryption keys used to protect your data. You can use AWS KMS to encrypt data in Amazon RDS and Amazon EBS volumes12

Question: 294

Which task can only an AWS account root user perform?

A. Changing the AWS Support plan
B. Deleting AWS resources
C. Creating an Amazon EC2 instance key pair
D. Configuring AWS WAF

Answer: A

Explanation:
The AWS account root user is the email address that you use to sign up for AWS. The root user has complete access to all AWS services and resources in the account. The root user can perform tasks that only the root user can do, such as changing the AWS Support plan, closing the account, and restoring IAM user permissions34

Question: 295

A company is considering migration to the AWS Cloud. The company wants a fully managed service or feature that can transfer streaming data from multiple sources to an Amazon S3 bucket.
Which AWS service or feature should the company use to meet these requirements?

A. AWS DataSync
B. Amazon Kinesis Data Firehose
C. S3 Select

D. AWS Transfer Family

Answer: B

Explanation:
Amazon Kinesis Data Firehose is a fully managed service that delivers real-time streaming data to destinations such as Amazon S3, Amazon Redshift, Amazon Elasticsearch Service, and Splunk. You can use Amazon Kinesis Data Firehose to capture, transform, and load streaming data from multiple sources, such as web applications, mobile devices, IoT sensors, and social media.

Question: 296

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

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

Answer: A

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

Question: 297

A company is launching a mobile app. The company wants customers to be able to use the app without upgrading their mobile devices.
Which pillar of the AWS Well-Architected Framework does this goal represent?

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

Answer: C

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

Question: 298

Which AWS service can a company use to find security and compliance reports, including International Organization for Standardization (ISO) reports?

A. AWS Artifact
B. Amazon CloudWatch
C. AWS Config
D. AWS Audit Manager

Answer: A

Explanation:
AWS Artifact is a self-service portal that provides on-demand access to AWS security and compliance reports and select online agreements. You can use AWS Artifact to download AWS service audit reports, such as ISO, PCI, and SOC, and to accept and manage agreements with AWS, such as the Business Associate Addendum (BAA).

Question: 299

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

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

Answer: C,D

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

Question: 300

A company's headquarters is located on a different continent from where the majority of the company's customers live. The company wants an AWS Cloud environment setup that will provide the lowest latency to the customers.

A company wants to automate the creation of new AWS accounts and automatically prevent all users from creating Amazon EC2 instances.

Which AWS service provides this functionality?

A. AWS Service Catalog
B. AWS Organizations

C. EC2 Image Builder
D. AWS Systems Manager

Answer: B

Explanation:
AWS Organizations is a service that enables you to create and manage multiple AWS accounts centrally. You can use AWS Organizations to automate account creation, apply policies to control access and permissions, and consolidate billing across your accounts. You can also use AWS Organizations to prevent users from creating Amazon EC2 instances in certain regions or with certain configurations2

Question: 301

A company needs to set up user authentication for a new application. Users must be able to sign in directly with a user name and password, or through a third-party provider.
Which AWS service should the company use to meet these requirements?

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

Answer: C

Explanation:
Amazon Cognito is a service that provides user authentication and authorization for web and mobile applications. You can use Amazon Cognito to enable users to sign in directly with a user name and password, or through a third-party provider, such as Facebook, Google, or Amazon. You can also use Amazon Cognito to manage user profiles, preferences, and security settings3

Question: 302

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: A,B

Explanation:

The AWS Cloud offers many benefits, such as:
Trade variable expenses for capital expenses: You can pay only for the resources you use, instead of investing in fixed costs upfront. This

reduces the risk and complexity of planning and managing your IT infrastructure4

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

Question: 303

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:

AWS Marketplace is a digital catalog that offers thousands of software products and solutions from independent software vendors (ISVs) and AWS partners. Customers can use AWS Marketplace to find, buy, and deploy software on AWS. Some of the benefits of using AWS Marketplace are: Speed of business: You can quickly and easily discover and deploy software that meets your business needs, without having to go through lengthy procurement processes. You can also use AWS Marketplace to test and compare different solutions before making a purchase decision.

Fewer legal objections: You can benefit from standardized contract terms and conditions that are prenegotiated between AWS and the ISVs. This reduces the time and effort required to review and approve legal agreements.

Question: 304

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

Which AWS offering will meet these requirements?

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

Answer: C

Explanation:

AWS Budgets is a service that enables you to plan your service usage, service costs, and instance reservations. You can use AWS Budgets to create custom budgets that alert you when your costs or usage exceed (or are forecasted to exceed) your budgeted amount. You can also use AWS Budgets to monitor how close your usage and costs are to meeting your reservation purchases1

Question: 305

According to the AWS shared responsibility model, which task is the customer's responsibility?

A. Maintaining the infrastructure needed to run AWS Lambda
B. Updating the operating system of Amazon DynamoDB instances
C. Maintaining Amazon S3 infrastructure
D. Updating the guest operating system on Amazon EC2 instances

Answer: D

Explanation:
The AWS shared responsibility model describes 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 hardware, software, networking, and facilities that run AWS services. The customer is responsible for security in the cloud, which includes the customer data, applications, operating systems, and network and firewall configurations. Therefore, updating the guest operating system on Amazon EC2 instances is the customer's responsibility2

Question: 306

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) is a service that enables you to manage access to AWS services and resources securely. You can use IAM to perform the following actions: Control access to AWS service APIs and to other specific resources: You can create users, groups, roles, and policies that define who can access which AWS resources and how. You can also use IAM to grant temporary access to users or applications that need to perform certain tasks on your behalf3 Protect the AWS environment using multi-factor authentication (MFA): You can enable MFA for your IAM users and root user to add an extra layer of security to your AWS account. MFA requires users to provide a unique authentication code from an approved device or SMS text message, in addition to their user name and password, when they sign in to AWS4

Question: 307

A company needs to securely store important credentials that an application uses to connect users to a database.
Which AWS service can meet this requirement with the MINIMAL amount of operational overhead?

A. AWS Key Management Service (AWS KMS)
B. AWS Config
C. AWS Secrets Manager
D. Amazon GuardDuty

Answer: C

Explanation:
AWS Secrets Manager is a service that helps you protect secrets needed to access your applications, services, and IT resources. You can use AWS Secrets Manager to store, rotate, and retrieve database credentials, API keys, and other secrets throughout their lifecycle. AWS Secrets Manager eliminates the need to hardcode sensitive information in plain text, and reduces the risk of unauthorized access or leakage. AWS Secrets Manager also integrates with other AWS services, such as AWS Lambda, Amazon RDS, and AWS CloudFormation, to simplify the management of secrets across your environment5

Question: 308

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

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

Answer: B

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

Question: 309

Which task does AWS perform automatically?

A. Encrypt data that is stored in Amazon DynamoDB.
B. Patch Amazon EC2 instances.
C. Encrypt user network traffic.
D. Create TLS certificates for users' websites.

Answer: B

Explanation:
AWS performs some tasks automatically to help you manage and secure your AWS resources. One of these tasks is patching Amazon EC2 instances. AWS provides two options for patching your EC2 instances: managed instances and patch baselines. Managed instances are a

group of EC2 instances or on-premises servers that you can manage using AWS Systems Manager. Patch baselines define the patches that AWS Systems Manager applies to your instances. You can use AWS Systems Manager to automate the process of patching your instances based on a schedule or a maintenance window.

Question: 310

A company is migrating its data center to AWS. The company needs an AWS Support plan that provides chat access to a cloud sup engineer 24 hours a day, 7 days a week. The company does not require access to infrastructure event management.
What is the MOST cost-effective AWS Support plan that meets these requirements?

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

Answer: B

Explanation:
AWS Business Support is the most cost-effective AWS Support plan that provides chat access to a cloud support engineer 24/7. AWS Business Support also offers phone and email support, as well as a response time of less than one hour for urgent issues. AWS Business Support does not include access to infrastructure event management, which is a feature of AWS Enterprise Support. AWS Enterprise Support is more expensive and provides additional benefits, such as a technical account manager, a support concierge, and a response time of less than 15 minutes for critical issues. AWS Developer Support and AWS Basic Support do not provide chat access to a cloud support engineer. AWS Developer Support provides email support and a response time of less than 12 hours for general guidance issues. AWS Basic Support provides customer service and account support, as well as access to forums and documentation1

Question: 311

In the AWS shared responsibility model, which tasks are the responsibility of AWS? (Select TWO.)

A. Patch an Amazon EC2 instance operating system.
B. Configure a security group.
C. Monitor the health of an Availability Zone.
D. Protect the infrastructure that runs Amazon EC2 instances.
E. Manage access to the data in an Amazon S3 bucket

Answer: C,D

Explanation:
According to the AWS shared responsibility model, AWS is responsible for the security of the cloud, which includes the tasks of monitoring the health of an Availability Zone and protecting the infrastructure that runs Amazon EC2 instances. An Availability Zone is a physically isolated location within an AWS Region that has its own power, cooling, and network connectivity. AWS monitors the health and performance of each Availability Zone and notifies customers of any issues or disruptions. AWS also protects the infrastructure that runs AWS services, such as Amazon EC2, by implementing physical, environmental, and operational security measures. AWS is not responsible for patching an Amazon EC2 instance operating system, configuring a security group, or managing access to the data in an Amazon S3 bucket. These are the customer's responsibilities for security in the cloud. The customer must ensure that the operating system and applications on their EC2 instances are up to date and secure. The customer must also configure the security group rules that control the

inbound and outbound traffic for their EC2 instances. The customer must also manage the access permissions and encryption settings for their S3 buckets and objects2

Question: 312

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

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

Answer: B

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

Question: 313

A company needs to store infrequently used data for data archives and long-term backups.

A company needs a history report about how its Amazon EC2 instances were modified last month. Which AWS service can be used to meet this requirement?

A. AWS Service Catalog
B. AWS Config
C. Amazon CloudWatch
D. AWS Artifact

Answer: B

Explanation:
AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. AWS Config continuously monitors and records your AWS resource configurations and allows you to automate the evaluation of recorded configurations against desired configurations. AWS Config can also track changes to your EC2 instances over time and provide a history report of the

modifications. AWS Service Catalog, Amazon CloudWatch, and AWS Artifact are not the best services to meet this requirement. AWS Service Catalog is a service that allows you to create and manage catalogs of IT services that are approved for use on AWS. Amazon CloudWatch is a service that monitors your AWS resources and applications and provides metrics, alarms, dashboards, and logs. AWS Artifact is a service that provides on-demand access to AWS security and compliance reports and online agreements

Question: 314

A company wants to use the latest technologies and wants to minimize its capital investment. Instead of upgrading on-premises infrastructure, the company wants to move to the AWS Cloud. Which AWS Cloud benefit does this scenario describe?

A. Increased speed to market
B. The trade of infrastructure expenses for operating expenses
C. Massive economies of scale
D. The ability to go global in minutes

Answer: B

Explanation:
The trade of infrastructure expenses for operating expenses is one of the benefits of the AWS Cloud. By moving to the AWS Cloud, the company can avoid the upfront costs of purchasing and maintaining on-premises infrastructure, such as servers, storage, network, and software. Instead, the company can pay only for the AWS resources and services that they use, as they use them. This reduces the risk and complexity of planning and managing IT infrastructure, and allows the company to focus on innovation and growth. Increased speed to market, massive economies of scale, and the ability to go global in minutes are also benefits of the AWS Cloud, but they are not the best ones to describe this scenario. Increased speed to market means that the company can launch new products and services faster by using AWS services and tools. Massive economies of scale means that the company can benefit from the lower costs and higher performance that AWS achieves by operating at a large scale. The ability to go global in minutes means that the company can deploy their applications and data in multiple regions and availability zones around the world to reach their customers faster and improve performance and reliability5

Question: 315

Which AWS service provides threat detection by monitoring for malicious activities and unauthorized actions to protect AWS accounts, workloads, and data that is stored in Amazon S3?

A. AWS Shield
B. AWS Firewall Manager
C. Amazon GuardDuty
D. Amazon Inspector

Answer: C

Explanation:
Amazon GuardDuty is a service that provides intelligent threat detection and continuous monitoring for your AWS accounts, workloads, and data. Amazon GuardDuty analyzes and processes data sources, such as VPC Flow Logs, AWS CloudTrail event logs, and DNS logs, to identify malicious activities and unauthorized actions, such as reconnaissance, instance compromise, account compromise, and data

exfiltration. Amazon GuardDuty can also detect threats to your data stored in Amazon S3, such as API calls from unusual locations or disabling of preventative controls. Amazon GuardDuty generates findings that summarize the details of the detected threats and provides recommendations for remediation. AWS Shield, AWS Firewall Manager, and Amazon Inspector are not the best services to meet this requirement. AWS Shield is a service that provides protection against distributed denial of service (DDoS) attacks. AWS Firewall Manager is a service that allows you to centrally configure and manage firewall rules across your accounts and resources. Amazon Inspector is a service that assesses the security and compliance of your applications running on EC2 instances.

Question: 316

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

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

Answer: B

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

Question: 317

A company wants to migrate its server-based applications to the AWS Cloud. The company wants to determine the total cost of ownership for its compute resources that will be hosted on the AWS Cloud.
Which combination of AWS services or tools will meet these requirements?

A. AWS Pricing Calculator
B. Migration Evaluator
C. AWS Support Center
D. AWS Application Discovery Service
E. AWS Database Migration Service (AWS DMS)

Answer: A,D

Explanation:
AWS Pricing Calculator and AWS Application Discovery Service are the best combination of AWS services or tools to meet the requirements of determining the total cost of ownership for compute resources that will be hosted on the AWS Cloud. AWS Pricing Calculator is a tool that enables you to estimate the cost of using AWS services based on your usage scenarios and requirements. You can use AWS Pricing

Calculator to compare the costs of running your applications on-premises or on AWS, and to optimize your AWS spending. AWS Application Discovery Service is a service that helps you plan your migration to the AWS Cloud by collecting and analyzing information about your on-premises servers, applications, and dependencies. You can use AWS Application Discovery Service to identify the inventory of your on-premises infrastructure, group servers by applications, and estimate the performance and resource utilization of your applications45

Question: 318

A company is planning to migrate to the AWS Cloud and wants to become more responsive to customer inquiries and feedback. The company wants to focus on organizational transformation. A company wants to give its customers the ability to view specific data that is hosted in Amazon S3 buckets. The company wants to keep control over the full datasets that the company shares with the customers. Which S3 feature will meet these requirements?

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

Answer: D

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

Question: 319

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

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

Answer: C,D

Explanation:
AWS CloudFormation and AWS Cloud Development Kit (AWS CDK) are AWS services that can limit manual errors by consistently provisioning AWS resources in multiple environments. AWS

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

Question: 320

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

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

Answer: A

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

Question: 321

Which tasks are customer responsibilities, according to the AWS shared responsibility model? (Select 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.
F. AWS Config

Answer: A,B

Explanation:
According to the AWS shared responsibility model, the customer is responsible for security in the cloud, which includes the tasks of configuring the AWS provided security group firewall and classifying company assets in the AWS Cloud. A security group is a virtual firewall that controls the inbound and outbound traffic for one or more EC2 instances. The customer must configure the security group rules to

allow or deny traffic based on protocol, port, or source and destination IP address2 Classifying company assets in the AWS Cloud means identifying the types, categories, and sensitivity levels of the data and resources that the customer stores and processes on AWS. The customer must also determine the applicable compliance requirements and regulations that apply to their assets, and implement the appropriate security controls and measures to protect them

Question: 322

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

An ecommerce company is using Amazon EC2 Auto Scaling groups to manage a fleet of web servers running on Amazon EC2. This architecture follows which AWS Well-Architected Framework best practice?

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

Answer: C

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

Question: 323

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

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

Answer: C,E

Explanation:
According to the AWS shared responsibility model, the customer is responsible for security in the cloud, which includes the tasks of managing data encryption and granting least privilege access to IAM users. Data encryption is the process of transforming data into an unreadable format that canonly be accessed with a key or a password. The customer must decide whether to encrypt their data at rest (when it is stored on AWS) or in transit (when it is moving between AWS and the customer or between AWS services). The

customer must also choose the encryption method, algorithm, and key management solution that best suit their needs. AWS provides various services and features that support data encryption, such as AWS Key Management Service (AWS KMS), AWS Certificate Manager (ACM), and AWS Encryption SDK5 IAM users are entities that represent the people or applications that interact with AWS resources and services. The customer must grant the IAM users the minimum permissions that they need to perform their tasks, and avoid giving them unnecessary or excessive access. This is known as the principle of least privilege, and it helps reduce the risk of unauthorized or malicious actions. The customer can use IAM policies, roles, groups, and permissions boundaries to manage the access of IAM users.

Question: 324

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

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

Answer: A

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

Question: 325 A company is migrating to the AWS Cloud to meet storage needs. The company wants to optimize costs based on the amount of storage that the company uses.
Which AWS offering or benefit will meet these requirements MOST cost-effectively?

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

Answer: D

Explanation:
Volume-based discounts are an AWS offering or benefit that can help the company optimize costs based on the amount of storage that the company uses. Volume-based discounts are discounts that AWS provides for some storage services, such as Amazon S3 and Amazon EBS, when the company stores a large amount of data. The more data the company stores, the lower the price per GB. For example, Amazon S3 offers six storage classes, each with a different price per GB. The price per GB decreases as the amount of data stored in each storage class increases

Question: 326

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

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

Answer: A

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

Question: 327

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

Which AWS pricing model or offering will meet these requirements?

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

Answer: B

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

Question: 328

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

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

Answer: A

Explanation:

Network ACLs (network access control lists) are an AWS service or feature that provides the functionality of applying security rules to a subnet for EC2 instances. A subnet is a logical partition of an IP network within a VPC (virtual private cloud). A VPC is a logically isolated section of the AWS Cloud where the company can launch AWS resources in a virtual network that they define. A network ACL is a virtual firewall that controls the inbound and outbound traffic for one or more subnets. The company can use network ACLs to allow or deny traffic based on protocol, port, or source and destination IP address. Network ACLs are stateless, meaning that they do not track the traffic that flows through them. Therefore, the company must create rules for both inbound and outbound traffic4

Question: 329

A company wants to migrate its high-performance computing (HPC) application to Amazon EC2 instances. The application has multiple components. The application must have fault tolerance and must have the ability to fail over automatically.
Which AWS infrastructure solution will meet these requirements with the LEAST latency between components?

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

Answer: C

Explanation:

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

Question: 330

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
D. AWS Budgets

Answer: A

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

Question: 331

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 an easy-to-use cloud platform that offers everything you need to build an application or website, plus a cost-effective, monthly plan1. It is designed for developers who have little or no prior cloud experience and want to launch and manage applications on AWS with minimal complexity2. Amazon SageMaker is a service for building, training, and deploying machine learning models3. AWS Lambda is a service that lets you run code without provisioning or managing servers4. Amazon Elastic Container Service (Amazon ECS) is a fully managed container orchestration service.

Question: 332

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

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

Answer: A

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

Question: 333

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

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

Answer: B

Explanation:
IAM access keys are long-term credentials that consist of an access key ID and a secret access key. You use access keys to sign programmatic requests that you make to AWS. If you need to access AWS services from an on-premises application, you can use IAM access keys to authenticate your requests. AWS account user name and password are used to sign in to the AWS Management Console. Amazon EC2 key pairs are used to connect to your EC2 instances using SSH. AWS Key Management Service (AWS KMS) keys are used to encrypt and decrypt your data using the AWS Encryption SDK or the AWS CLI.

Question: 334

A company wants to launch its web application in a second AWS Region. The company needs to determine which services must be regionally configured for this launch.
Which AWS services can be configured at the Region level? (Select TWO.)

A. Amazon EC2
B. Amazon Route 53
C. Amazon CloudFront
D. AWS WAF

E. Amazon DynamoDB

Answer: B,D

Explanation:

Amazon Route 53 and AWS WAF are AWS services that can be configured at the Region level. Amazon Route 53 is a highly available and scalable cloud Domain Name System (DNS) web service that lets you register domain names, route traffic to resources, and check the health of your resources. AWS WAF is a web application firewall that helps protect your web applications or APIs against common web exploits that may affect availability, compromise security, or consume excessive resources. Amazon EC2, Amazon CloudFront, and Amazon DynamoDB are AWS services that can be configured at the global level or the Availability Zone level .

Question: 335

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 a service that enables governance, compliance, operational auditing, and risk auditing of your AWS account. With CloudTrail, you can log, continuously monitor, and retain account activity related to actions across your AWS infrastructure. You can use CloudTrail to identify who accessed an AWS service and what action was performed for a given time period. Amazon CloudWatch, AWS Security Hub, and Amazon Inspector are AWS services that provide different types of monitoring and security capabilities.

Question: 336

A company is running its application in the AWS Cloud and wants to protect against a DDoS attack. The company's security team wants near real-time visibility into DDoS attacks.

Which AWS service or traffic filter will meet these requirements with the MOST features for DDoS protection?

A. AWS Shield Advanced
B. AWS Shield
C. Amazon GuardDuty
D. Network ACLs

Answer: A

Explanation:

AWS Shield Advanced is a managed Distributed Denial of Service (DDoS) protection service that safeguards applications running on AWS. AWS Shield Advanced provides you with 24x7 access to the AWS DDoS Response Team (DRT) and protection against DDoS attacks of any size or duration. AWS Shield Advanced also provides near real-time visibility into attacks, advanced attack mitigation capabilities, and integration with AWS WAF and AWS Firewall Manager1. AWS Shield is a standard service that provides always-on detection and automatic inline mitigations to minimize application downtime and latency, but it does not offer the same level of features and support as AWS Shield

Advanced2. Amazon GuardDuty is a threat detection service that continuously monitors for malicious activity and unauthorized behavior, but it does not provide DDoS protection3. Network ACLs are stateless filters that can be associated with a subnet to control the traffic to and from the subnet, but they are not designed to protect against DDoS attacks

Question: 337

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

Explanation:

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.

Question: 338

Which task must a user perform by using the AWS account root user credentials?

A. Make changes to AWS production resources.
B. Change AWS Support plans.
C. Access AWS Cost and Usage Reports.
D. Grant auditors' access to an AWS account for a compliance audit.

Answer: B

Explanation:

Changing AWS Support plans is a task that must be performed by using the AWS account root user credentials. The root user is the email address that you used to sign up for AWS. It has complete access to all AWS services and resources in the account. You should use the root

user only to perform a few account and service management tasks, such as changing AWS Support plans, closing the account, or changing the account name or email address. Making changes to AWS production resources, accessing AWS Cost and Usage Reports, and granting auditors access to an AWS account for a compliance audit are tasks that can be performed by using IAM users or roles, which are entities that you create in AWS to delegate permissions to access AWS services and resources.

Question: 339

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

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

Answer: C

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

Question: 340

A company needs to control inbound and outbound traffic for an Amazon EC2 instance.
Which AWS service or feature can the company associate with the EC2 instance to meet this requirement?

A. Network ACL
B. Security group
C. AWS WAF
D. VPC route tables

Answer: B

Explanation:
A security group is a virtual firewall that can be associated with an Amazon EC2 instance to control the inbound and outbound traffic for the instance. You can specify which protocols, ports, and source or destination IP ranges are allowed or denied by the security group. A network ACL is a stateless filter that can be associated with a subnet to control the traffic to and from the subnet, but it is not associated with an EC2 instance4. AWS WAF is a web application firewall that helps protect your web applications or APIs against common web exploits that may affect availability, compromise security, or consume excessive resources. VPC route tables are used to determine where network traffic is directed within a VPC or to an internet gateway, virtual private gateway, NAT device, VPC peering connection, or VPC

endpoint.

Question: 341

A company is expecting a short-term spike in internet traffic for its application. During the traffic increase, the application cannot be interrupted. The company also needs to minimize cost and maximize flexibility.

A company needs to use a serverless interactive query service to analyze data in Amazon S3. The query service must support standard SQL.

Which AWS service will meet these requirements?

A. Amazon Redshift
B. AWS Glue
C. Amazon Athena
D. Amazon Kinesis Data Streams

Answer: C

Explanation:

Amazon Athena is a serverless interactive query service that makes it easy to analyze data in Amazon S3 using standard SQL. Athena is ideal for quick, ad-hoc querying but it can also handle complex analysis, including large joins, window functions, and arrays. Athena scales automatically—executing queries in parallel—so results are fast, even with large datasets and complex queries. Amazon Redshift is a fully managed, petabyte-scale data warehouse service that can run complex analytic queries against structured and semi-structured data using standard SQL. However, it is not a serverless service and requires provisioning and managing clusters of nodes. AWS Glue is a fully managed extract, transform, and load (ETL) service that makes it easy to prepare and load your data for analytics. However, it is not a query service and does not support standard SQL. Amazon Kinesis Data Streams is a service that enables you to build custom applications that process or analyze streaming data for specialized needs. However, it is not a query service and does not support standard SQL.

Question: 342

A company needs to run a workload for several batch image rendering applications. It is acceptable for the workload to experience downtime.

Which Amazon EC2 pricing model would be MOST cost-effective in this situation?

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

Answer: D

Explanation:

Amazon EC2 Spot Instances are instances that use spare EC2 capacity that is available at up to a 90% discount compared to On-Demand prices. You can use Spot Instances for various stateless, fault- tolerant, or flexible applications such as big data, containerized workloads, high-performance computing (HPC), and test & development workloads. Spot Instances are ideal for workloads that can be interrupted, such as batch image rendering applications1. On-Demand Instances are instances that let you pay for compute capacity by the hour or

second (minimum of 60 seconds) with no longterm commitments. This frees you from the costs and complexities of planning, purchasing, and maintaining hardware and transforms what are commonly large fixed costs into much smaller variable costs2. Reserved Instances are instances that provide you with a significant discount (up to 75%) compared to On-Demand Instance pricing. In exchange, you select a term and make an upfront payment to reserve a certain amount of compute capacity for that term3. Dedicated Instances are instances that run in a VPC on hardware that's dedicated to a single customer. Your Dedicated Instances are physically isolated at the host hardware level from instances that belong to other AWS accounts4.

Question: 343

A company has an application that runs periodically in an on-premises environment. The application runs for a few hours most days, but runs for 8 hours a day for a week at the end of each month. Which AWS service or feature should be used to host the application in the AWS Cloud?

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

Answer: B

Explanation:
Amazon EC2 On-Demand Instances are instances that let you pay for compute capacity by the hour or second (minimum of 60 seconds) with no long-term commitments. This frees you from the costs and complexities of planning, purchasing, and maintaining hardware and transforms what are commonly large fixed costs into much smaller variable costs. On-Demand Instances are suitable for applications with short-term, irregular, or unpredictable workloads that cannot be interrupted, such as periodic applications that run for a few hours most days, but run for 8 hours a day for a week at the end of each month2. Amazon EC2 Standard Reserved Instances are instances that provide you with a significant discount (up to 75%) compared to On-Demand Instance pricing. In exchange, you select a term and make an upfront payment to reserve a certain amount of compute capacity for that term. Reserved Instances are suitable for applications with steady state or predictable usage that require reserved capacity3. AWS Wavelength is a service that enables developers to build applications that deliver ultra-low latency to mobile devices and users by deploying AWS compute and storage at the edge of the 5G network. Wavelength is suitable for applications that requiresingle- digit millisecond latencies, such as game and live video streaming, machine learning inference at the edge, and augmented and virtual reality (AR/VR). Application Load Balancer is a service that operates at the request level (layer 7) and distributes incoming application traffic across multiple targets, such as EC2 instances, containers, Lambda functions, and IP addresses. Application Load Balancer is suitable for applications that need advanced routing capabilities, such as microservices or container-based architectures.

Question: 344

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: E

Explanation:
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.

Question: 345

A company is building an application on AWS. The application needs to comply with credit card regulatory requirements. The company needs proof that the AWS services and deployment are in compliance.
Which actions should the company take to meet these requirements? (Select TWO.)

A. Use Amazon Inspector to submit the application for certification.
B. Ensure that the application's underlying hardware components comply with requirements.
C. Use AWS Artifact to access AWS documents about the compliance of the services.
D. Get the compliance of the application certified by a company assessor.
E. Use AWS Security Hub to certify the compliance of the application.

Answer: C,D

Explanation:
Using AWS Artifact to access AWS documents about the compliance of the services, and getting the compliance of the application certified by a company assessor are actions that the company should take to meet the requirements of complying with credit card regulatory requirements. AWS Artifact is a service that provides on-demand access to AWS security and compliance reports and select online agreements. Reports available in AWS Artifact include our Service Organization Control (SOC) reports, Payment Card Industry (PCI) reports, and certifications from accreditation bodies across geographies and compliance verticals that validate the implementation and operating effectiveness of AWS security controls. AWS Artifact can help you demonstrate compliance with credit card regulatory requirements by providing you with proof that the AWS services and deployment are in compliance. Getting the compliance of the application certified by a company assessor is an action that the company should take to ensure that the application meets the specific requirements of the credit card industry. A company assessor is an independent third-party entity that is qualified to assess the compliance of the application with the relevant standards and regulations. Using Amazon Inspector to submit the application for certification is not an action that the company

should take, because Amazon Inspector is a service that helps you improve the security and compliance of your applications deployed on AWS by automatically assessing them for vulnerabilities and deviations from best practices, but it does not provide certification for the applications. Ensuring that the application's underlying hardware components comply with requirements is not an action that the company should take, because the application is deployed on AWS, and AWS is responsible for the security and compliance of the underlying hardware components. This is part of the shared responsibility model, where AWS is responsible for security of the cloud, and customers are responsible for security in the cloud. Using AWS Security Hub to certify the compliance of the application is not an action that the company should take, because AWS Security Hub is a service that gives you a comprehensive view of your security posture acrossyour AWS accounts and helps you check your environment against security industry standards and best practices, but it does not provide certification for the applications.

Question: 346

A company has set up a VPC on AWS. The company needs a dedicated connection between the VPC and the company's on-premises network.
Which action should the company take to meet this requirement?

A. Establish a VPN connection between the VPC and the company's on-premises network.
B. Establish an AWS Direct Connect connection between the VPC and the company's on- premisesnetwork.
C. Attach an internet gateway to the VPC. Use the AWS public endpoints for connectivity.
D. Configure Amazon Connect to provide connectivity between the VPC and the company's on- premisesnetwork.

Answer: B

Explanation:
Establishing an AWS Direct Connect connection between the VPC and the company's on-premises network is the action that the company should take to meet the requirement of having a dedicated connection between the VPC and the company's on-premises network. AWS Direct Connect is a service that lets you establish a dedicated network connection between your network and one of the AWS Direct Connect locations. Using AWS Direct Connect, you can create a private connection between AWS and your datacenter, office, or colocation environment, which can reduce your network costs, increase bandwidth throughput, and provide a more consistent network experience than internet-based connections. Establishing a VPN connection between the VPC and the company's on-premises network is an action that the company can take to create a secure and encrypted connection between the VPC and the company's on-premises network, but it is not a dedicated connection, as it uses the public internet as the transport mechanism. Attaching an internet gateway to the VPC and using the AWS public endpoints for connectivity is an action that the company can take to enable communication between the VPC and the internet, but it is not a dedicated connection, as it also uses the public internet as the transport mechanism. Configuring Amazon Connect to provide connectivity between the VPC and the company's on-premises network is not an action that the company can take, because Amazon Connect is a service that lets you set up and manage a contact center in the cloud, but it does not provide network connectivity between the VPC and the company's on-premises network.

Question: 347

A company has deployed an application in the AWS Cloud. The company wants to ensure that the application is highly resilient.
Which component of AWS infrastructure can the company use to meet this requirement?

A. Content delivery network (CDN)
B. Edge locations
C. Wavelength Zones

D. Availability Zones

Answer: D

Explanation:
Availability Zones are components of AWS infrastructure that can help the company ensure that the application is highly resilient. Availability Zones are multiple, isolated locations within each AWS Region. Each Availability Zone has independent power, cooling, and physical security, and is connected to the other Availability Zones in the same Region via low-latency, high-throughput, and highly redundant networking. Availability Zones allow you to operate production applications and databases that are more highly available, fault tolerant, and scalable than would be possible from a single data center.

Question: 348

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

Explanation:
AWS Site-to-Site VPN and AWS Direct Connect are AWS services that are connectivity services for a VPC. AWS Site-to-Site VPN is a service that enables you to securely connect your on-premises network or branch office site to your Amazon Virtual Private Cloud (Amazon VPC). You can establish VPN connections over the internet or over AWS Direct Connect1. AWS Direct Connect is a service that lets you establish a dedicated network connection between your network and one of the AWS Direct Connect locations. Using AWS Direct Connect, you can create a private connection between AWS and your datacenter, office, or colocation environment, which can reduce your network costs,increase bandwidth throughput, and provide a more consistent network experience than internet-based connections2. Amazon Connect is a service that lets you set up and manage a contact center in the cloud, but it does not provide network connectivity between the VPC and your onpremises network. AWS Key Management Service (AWS KMS) is a service that makes it easy for you to create and manage cryptographic keys and control their use across a wide range of AWS services and inyour applications, but it does not provide network connectivity between the VPC and your onpremises network. AWS Identity and Access Management (IAM) is a service that enables you to manage access to AWS services and resources securely, but it does not provide network connectivity between the VPC and your on-premises network.

Question: 349

A company wants a key-value NoSQL database that is fully managed and serverless. Which AWS service will meet these requirements?

A. Amazon DynamoDB
B. Amazon RDS
C. Amazon Aurora
D. Amazon Memory DB for Redis

Answer: A

Explanation:

Amazon DynamoDB is a key-value and document database that delivers single-digit millisecond performance at any scale. It is a fully managed, serverless database that does not require provisioning, patching, or backup. It offers built-in security, backup and restore, and in-memory caching3. Amazon RDS is a relational database service that makes it easy to set up, operate, and scale a relational database in the cloud. It provides cost-efficient and resizable capacity while automating time-consuming administration tasks such as hardware provisioning, database setup, patching, and backups. However, it is not a key-value NoSQL database, and it is not serverless, as it requires you to choose an instance type and size4. Amazon Aurora is a MySQL and PostgreSQL- compatible relational database built for the cloud, that combines the performance and availability of traditional enterprise databases with the simplicity and cost-effectiveness of open source databases. However, it is also not a key-value NoSQL database, and it is not serverless, as it requires you to choose an instance type and size. Amazon MemoryDB for Redis is a Redis-compatible, durable, inmemory database service that delivers ultra-fast performance and multi-AZ reliability for the most demanding applications. However, it is also not a key-value NoSQL database, and it is not serverless, as it requires you to choose a node type and size.

Question: 350

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

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

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

Answer: D

Explanation:

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

Question: 351

A software engineer wants to launch a virtual machine (VM) and MySQL database on AWS. Which AWS service will meet these

requirements with the LEAST operational effort?

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

Answer: B

Explanation:
AWS Elastic Beanstalk is a service that enables you to quickly deploy and manage applications in the AWS Cloud without worrying about the infrastructure that runs those applications. You simply upload your application, and Elastic Beanstalk automatically handles the details of capacity provisioning, load balancing, scaling, and application health monitoring. Elastic Beanstalk supports several platform configurations for Java, .NET, PHP, Node.js, Python, Ruby, Go, and Docker web applications that can run on familiar servers such as Apache, Nginx, Passenger, and IIS. You can also use Elastic Beanstalk to launch a virtual machine (VM) and MySQL database on AWS with theleast operational effort. Amazon Elastic Container Service (Amazon ECS) is a fully managed container orchestration service that enables you to easily run, scale, and secure Docker containerized applications on AWS. However, it requires more operational effort than Elastic Beanstalk, as you need to define your application architecture and the specifications of the containers that run it. Amazon Lightsail is an easy-to-use cloud platform that offers everything you need to build an application or website, plus a cost-effective, monthly plan. It is designed for developers who have little or no prior cloud experience and want to launch and manage applications on AWS with minimal complexity. However, it does not support MySQL databases, and it requires more operational effort than Elastic Beanstalk, as you need to configure your VM and database settings. Amazon EC2 is a web service that provides secure, resizable compute capacity in the cloud. It allows you to launch a virtual machine (VM) and MySQL database on AWS, but it requires the most operational effort, as you need to provision, monitor, and manage your EC2 instances and database.

Question: 352

A company runs business applications in an on-premises data center and in the AWS Cloud. The company needs a shared file system that can be available to both environments. Which AWS service meets these requirements?

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

Answer: D

Explanation:
Amazon Elastic File System (Amazon EFS) is a service that provides a simple, scalable, fully managed elastic NFS file system for use with AWS Cloud services and on-premises resources. It is built to scale on demand to petabytes without disrupting applications, growing and shrinking automatically as you add and remove files, eliminating the need to provision and manage capacity to accommodate growth. You can use Amazon EFS to create a shared file system that can be available to both your onpremises data center and your AWS Cloud environment. Amazon Elastic Block Store (Amazon EBS) is a service that provides persistent block storage volumes for use with Amazon EC2 instances in the AWS Cloud. Each Amazon EBS volume is automatically replicated within its Availability Zone to protect you from component failure, offering high availability and durability. However, Amazon EBS volumes are not shared file systems, and they cannot be available to both your on-premises data center and your AWS Cloud environment. Amazon S3 is a service that provides object storage

through a web services interface. You can use Amazon S3 to store and protect any amount of data for a range of use cases, such as data lakes, websites, mobile applications, backup and restore, archive, enterprise applications, IoT devices, and big dataanalytics. However, Amazon S3 is not a shared file system, and it cannot be available to both your on-premises data center and your AWS Cloud environment without additional configuration. Amazon ElastiCache is a service that enables you to seamlessly set up, run, and scale popular open-source compatible in-memory data stores in the cloud. You can use Amazon ElastiCache to improve the performance of your applications by allowing you to retrieve information from fast, managed, in-memory data stores, instead of relying entirely on slower disk-based databases. However, Amazon ElastiCache is not a shared file system, and it cannot be available to both your on-premises data center and your AWS Cloud environment.

Question: 353

Which option is AWS responsible for under the AWS shared responsibility model?

A. Network and firewall configuration
B. Client-side data encryption
C. Management of user permissions
D. Hardware and infrastructure

Answer: D

Explanation:
Hardware and infrastructure is the option that AWS is responsible for under the AWS shared responsibility model. The AWS shared responsibility model describes how AWS and customers share responsibilities for security and compliance in the cloud. AWS is responsible for security of the cloud, which means protecting the infrastructure that runs all the services offered in the AWS Cloud. This infrastructure is composed of the hardware, software, networking, and facilities that run AWS Cloud services. Customers are responsible for security in the cloud, which means taking care of the security of their own applications, data, and operating systems. This includes network and firewall configuration, client-side data encryption, management of user permissions, and more.

Question: 354

A company needs to run some of its workloads on premises to comply with regulatory guidelines. The company wants to use the AWS Cloud to run workloads that are not required to be on premises. The company also wants to be able to use the same API calls for the on-premises workloads and the cloud workloads.
Which AWS service or feature should the company use to meet these requirements?

A. Dedicated Hosts
B. AWS Outposts
C. Availability Zones
D. AWS Wavelength

Answer: B

Explanation:
AWS Outposts is a fully managed service that extends AWS infrastructure, AWS services, APIs, and tools to virtually any datacenter, co-location space, or on-premises facility for a truly consistent hybrid experience1. AWS Outposts enables customers to run workloads on

premises using the same AWS APIs, tools, and services that they use in the cloud2. Dedicated Hosts are physical servers with EC2 instance capacity fully dedicated to a customer's use3. Availability Zones are one or more discrete data centers, each with redundant power, networking, and connectivity, housed in separate facilities within an AWS Region4. AWS Wavelength is an AWS Infrastructure offering optimized for mobile edge computing applications.

Question: 355

A company wants to set up a high-speed connection between its data center and its applications that run on AWS. The company must not transfer data over the internet.
Which action should the company take to meet these requirements?

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

Answer: D

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

Question: 356

A company is using a central data platform to manage multiple types of data for its customers. The company wants to use AWS services to discover, transform, and visualize the data.
Which combination of AWS services should the company use to meet these requirements? (Select TWO.)

A. AWS Glue
B. Amazon Elastic File System (Amazon EFS)
C. Amazon Redshift
D. Amazon QuickSight
E. Amazon Quantum Ledger Database (Amazon QLDB)

Answer: A,C

Explanation:
AWS Glue is a fully managed extract, transform, and load (ETL) service that makes it easy to prepare and load data for analytics. AWS Glue can discover data sources, transform data, and make it available for analysis by using data catalogs and workflows. Amazon Redshift is a fully managed, petabyte-scale data warehouse service in the cloud that enables customers to analyze data using standard SQL and existing business intelligence tools. Amazon Redshift can also integrate with other AWS services to visualize and transform data. Amazon Elastic File System (Amazon EFS) provides a simple, scalable, fully managed elastic NFS file system for use with AWS Cloud services and onpremises

resources. Amazon QuickSight is a fast, cloud-powered business intelligence service that makes it easy to deliver insights to everyone in an organization. Amazon Quantum Ledger Database (Amazon QLDB) is a fully managed ledger database that provides a transparent, immutable, and cryptographically verifiable transaction log owned by a central trusted authority.

Question: 357

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

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

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

Answer: C

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

Question: 358

A company has all of its servers in the us-east-1 Region. The company is considering the deployment of additional servers different Region.

Which AWS tool should the company use to find pricing information for other Regions?

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

Answer: D

Explanation:
AWS Pricing Calculator lets customers explore AWS services, and create an estimate for the cost of their use cases on AWS. AWS Pricing Calculator can also compare the costs of different AWS Regions and configurations. Cost Explorer is a tool that enables customers to visualize, understand, and manage their AWS costs and usage over time. AWS Budgets gives customers the ability to set custom budgets that alert them when their costs or usage exceed (or are forecasted to exceed) their budgeted amount. AWS Purchase Order Management is a feature that allows customers to pay for their AWS invoices using purchase orders.

Question: 359

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

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

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

Answer: D

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

Question: 360

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

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

Answer: B

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

Question: 361

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

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

Answer: A,D

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

Question: 362

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 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].

Question: 363

AWS has the ability to achieve lower pay-as-you-go pricing by aggregating usage across hundreds of thousands of users.
This describes which advantage of the AWS Cloud?

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

Answer: C

Explanation:

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

Question: 364

A company wants to use guidelines from the AWS Well-Architected Framework to limit human error and facilitate consistent responses to events.
Which of the following is a Well-Architected design principle that will meet these requirements?

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

Answer: B

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

Question: 365

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

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

Answer: A

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

Question: 366

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

A. Operational excellence
B. Cost optimization

C. Performance efficiency

D. Security

Answer: A

Explanation:

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

Question: 367

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

A. Consolidate key components into monolithic architectures.

B. Test systems at production scale.

C. Provision more capacity than a workload is expected to need.

D. Drive architecture design based on data collected about the workload behavior and requirements.

E. Make AWS Cloud architectural decisions static, one-time events.

Answer: B,D

Explanation:

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

Question: 368

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:

These are two scenarios that represent the concept of elasticity on AWS. Elasticity means the ability to adjust the resources and capacity of the system in response to changes in demand or environment. Scaling the number of Amazon EC2 instances based on traffic means using services such as AWS Auto Scaling or Elastic Load Balancing to add or remove instances as the traffic increases or decreases. Resizing Amazon RDS instances as business needs change means using the Amazon RDS console or API to modify the instance type, storage type, or storage size of the database as the workload grows or shrinks. You can learn more about the concept of elasticity on AWS from [this webpage] or [this digital course].

Question: 369

An ecommerce company wants to distribute traffic between the Amazon EC2 instances that host its website.

Which AWS service or resource will meet these requirements?

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

Answer: A

Explanation:

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

Question: 370

Which AWS service will allow a user to set custom cost and usage limits, and will alert when the thresholds are exceeded?

A. AWS Organizations
B. AWS Budgets
C. Cost Explorer
D. AWS Trusted Advisor

Answer: B

Explanation:

AWS Budgets allows you to set custom budgets that alert you when your costs or usage exceed (or are forecasted to exceed) your budgeted amount. You can also use AWS Budgets to set reservation utilization or coverage targets and receive alerts when your utilization drops below the threshold you define. AWS Budgets provides you with a comprehensive view of your cost and usage, as well as your reservation utilization and coverage1.

Question: 371

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

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

Answer: B

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

Question: 372

A company must archive Amazon S3 data that the company's business units no longer need to access.
Which S3 storage class will meet this requirement MOST cost-effectively?

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

Answer: C

Explanation:
S3 Glacier Deep Archive is Amazon S3's lowest-cost storage class and supports long-term retention and digital preservation for data that may be accessed once or twice in a year. It is designed forcustomers — particularly those in highly-regulated industries, such as the Financial Services, Healthcare, and Public Sectors — that retain data sets for 7-10 years or longer to meet regulatory compliance requirements. Customers can store large amounts of data at a very low cost, and reliably access it with a wait time of 12 hours3.

Question: 373

A company wants to build a new web application by using AWS services. The application must meet the on-demand load for periods of heavy activity.
Which AWS services or resources provide the necessary workload adjustments to meet these requirements? (Select TWO.)

A. Amazon Machine Image (AMI)
B. Amazon EC2 Auto Scaling
C. Amazon EC2 instance

D. AWS Lambda
E. EC2 Image Builder

Answer: B,D

Explanation:
Amazon EC2 Auto Scaling helps you ensure that you have the correct number of Amazon EC2 instances available to handle the load for your application. You create collections of EC2 instances, called Auto Scaling groups. You can specify the minimum number of instances in each Auto Scaling group, and Amazon EC2 Auto Scaling ensures that your group never goes below this size. You can specify the maximum number of instances in each Auto Scaling group, and Amazon EC2 Auto Scaling ensures that your group never goes above this size4. AWS Lambda lets you run code without provisioning or managing servers. You pay only for the compute time you consume. With Lambda, you can run code for virtually any type of application or backend service - all with zero administration. Just upload your code and Lambda takes care of everything required to run and scale your code with high availability. You can set up your code to automatically trigger from other AWS services or call it directly from any web or mobile app.

Question: 374

Which AWS service or feature is an example of a relational database management system?

A. Amazon Athena

B. Amazon Redshift
C. Amazon S3 Select
D. Amazon Kinesis Data Streams

Answer: B

Explanation:
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 a relational database management system (RDBMS), so it is compatible with other RDBMS applications. You can use standard SQL to query the data.

Question: 375

A company needs to apply security rules to specific Amazon EC2 instances. Which AWS service or feature provides this functionality?

A. AWS Shield
B. Network ACLs
C. Security groups
D. AWS Firewall Manager

Answer: C

Explanation:

Security groups act as a firewall for associated Amazon EC2 instances, controlling both inbound and outbound traffic at the instance level. You can use security groups to set rules that allow or deny traffic to or from your instances. You can modify the rules for a security group at any time; the new rules are automatically applied to all instances that are associated with the security group.

Question: 376

Which AWS service is deployed to VPCs and provides protection from common network threats?

A. AWSShield
B. AWSWAF
C. AWS Network Firewall
D. AWS FirewallManager

Answer: C

Explanation:

AWS Network Firewall is a managed service that makes it easy to deploy essential network protections for all of your Amazon Virtual Private Clouds (VPCs). The service can be set up with just a few clicks from the AWS console or using APIs. AWS Network Firewall automatically scales with your network traffic, so you don't have to worry about deploying and managing any infrastructure. AWS Network Firewall provides protection from common network threats such as SQL injection, cross-site scripting, and DDoS attacks1.

Question: 377

Which option is a perspective that includes foundational capabilities of the AWS Cloud Adoption Framework (AWS CAF)?

A. Sustainability
B. Security
C. Performance efficiency
D. Reliability

Answer: B

Explanation:

The AWS Cloud Adoption Framework (AWS CAF) helps organizations understand how cloud adoption transforms the way they work, and it provides structure to identify and address gaps in skills and processes. The AWS CAF organizes guidance into six areas of focus, called perspectives. Each perspective reflects a different stakeholder viewpoint with its own distinct responsibilities, skills, and attributes. The Security Perspective helps you structure the selection and implementation of security controls that meet your organization's needs2.

Question: 378

A company needs to store data from a recommendation engine in a database. Which AWS service provides this functionality with the LEAST operational overhead?

A. Amazon RDS for PostgreSQL
B. Amazon DynamoDB
C. Amazon Neptune
D. Amazon Aurora

Answer: B

Explanation:
Amazon DynamoDB is a key-value and document database that delivers single-digit millisecond performance at any scale. It's a fully managed, multi-region, multi-active, durable database with built-in security, backup and restore, and in-memory caching for internet-scale applications. DynamoDB can handle more than 10 trillion requests per day and can support peaks of more than 20 million requests per second. DynamoDB provides the least operational overhead for storing data from a recommendation engine, as it does not require any server provisioning, patching, or

maintenance3

Question: 379

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

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

Answer: C

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

Question: 380

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

A. Amazon Aurora

B. Amazon RDS
C. Amazon DynamoDB
D. Amazon ElastiCache

Answer: D

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

Question: 381

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

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

Answer: A

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

Question: 382

A company uses AWS Organizations. The company wants to apply security best practices from the AWS Well-Architected Framework to all of its AWS accounts.
Which AWS service will meet these requirements?

A. Amazon Macie
B. Amazon Detective
C. AWS Control Tower
D. AWS Secrets Manager

Answer: A

Explanation:
AWS Control Tower is the easiest way to set up and govern a secure, multi-account AWS environment based on best practices established through AWS's experience working with thousands of enterprises as they move to the cloud. With AWS Control Tower, builders can provision new AWS accounts in a few clicks, while you have peace of mind knowing your accounts conform to your organization's policies. AWS Control Tower automates the setup of a baseline environment, or landing zone, that is a secure, well-architected multi-account AWS environment1. AWS Control Tower helps you apply security best practices from the AWS Well-Architected Framework to all of your AWS accounts2.

Question: 383

A company uses AWS for its web application. The company wants to minimize latency and perform compute operations for the application as close to end users as possible.
Which AWS service or infrastructure component will provide this functionality?

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

Answer: C

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

Question: 384

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

A. AWS Compute Optimizer
B. AWS Elastic Beanstalk
C. AWS AppSync
D. AWS Systems Manager

Answer: D

Explanation:
AWS Systems Manager gives you visibility and control of your infrastructure on AWS. Systems Manager provides a unified user interface so you can view operational data from multiple AWS services and allows you to automate operational tasks across your AWS resources. You

can use Systems Manager to apply OS patches, create system images, configure Windows and Linux operating systems, and execute PowerShell commands5. Systems Manager can help you ensure that all of your Amazon EC2 instances have compliant operating system patches by using the Patch Manager feature.

Question: 385

A company wants to integrate natural language processing (NLP) into business intelligence (BI) dashboards. The company wants to ask questions and receive answers with relevant visualizations. Which AWS service or tool will meet these requirements?

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

Answer: C

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

Question: 386

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

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

Answer: B

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

Question: 387

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: C

Explanation:
Identity and access management is one of the foundational capabilities for the operations perspective of the AWS Cloud Adoption Framework (AWS CAF). It involves managing the identities, roles, permissions, and credentials of users and systems that interact with AWS resources. Performance and capacity management is a capability for the platform perspective. Application portfolio management is a capability for the business perspective. Product management is a capability for the governance perspective.

Question: 388

A company needs to implement identity management for a fleet of mobile apps that are running in the AWS Cloud.
Which AWS service will meet this requirement?

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

Answer: A

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

Question: 389

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

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

A. Amazon Comprehend
B. Amazon Forecast

C. Amazon Personalize
D. Amazon SageMaker Studio

Answer: C

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

Question: 390

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

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

Answer: A

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

Question: 391

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

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

Answer: C

Explanation:
A network access control list (network ACL) is a feature that acts as a firewall for controlling traffic in and out of one or more subnets in a virtual private cloud (VPC). Network ACLs can be configured with rules that allow or deny traffic based on the source and destination IP

addresses, ports, and protocols1. AWS Security Hub is a service that provides a comprehensive view of the security posture of AWS accounts and resources2. Security groups are features that act as firewalls for controlling traffic at the instance level3. AWS WAF is a web application firewall that helps protect web applications from common web exploits4.

Question: 392

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

Explanation:
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.

Question: 393

A company wants to migrate its workloads to AWS, but it lacks expertise in AWS Cloud computing. Which AWS service or feature will help the company with its migration?

A. AWS Trusted Advisor
B. AWS Consulting Partners
C. AWS Artifacts
D. AWS Managed Services

Answer: D

Explanation:
AWS Managed Services is a service that provides operational management for AWS infrastructure and applications. It helps users migrate their workloads to AWS and provides ongoing support, security, compliance, and automation. AWS Trusted Advisor is a service that provides best practices and recommendations for cost optimization, performance, security, and fault tolerance. AWS Consulting Partners are professional services firms that help customers design, architect, build, migrate, and manage their workloads and applications on AWS. AWS Artifacts is a service that provides on-demand access to AWS compliance reports and select online agreements.

Question: 394

A company deployed an application on an Amazon EC2 instance. The application ran as expected for 6 months. In the past week, users have reported latency issues. A system administrator found that the CPU utilization was at 100% during business hours. The company wants a scalable solution to meet demand.

Which AWS service or feature should the company use to handle the load for its application during periods of high demand?

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

Answer: A

Explanation:

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

Question: 395

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

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

Answer: B

Explanation:

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

Question: 396

For which AWS service is the customer responsible for maintaining the underlying operating system?

A. Amazon DynamoDB
B. Amazon S3
C. Amazon EC2
D. AWS Lambda

Answer: C

Explanation:
Amazon EC2 is a service that provides resizable compute capacity in the cloud. Users can launch and manage virtual servers, known as instances, that run on the AWS infrastructure. Users are responsible for maintaining the underlying operating system of the instances, as well as any applications or software that run on them. Amazon DynamoDB is a service that provides a fully managed NoSQL database that delivers fast and consistent performance at any scale. Users do not need to manage the underlying operating system or the database software. Amazon S3 is a service that provides scalable and durable object storage in the cloud. Users do not need to manage the underlying operating system or the storage infrastructure. AWS Lambda is a service that allows users to run code without provisioning or managing servers. Users only need to upload their code and configure the triggers and parameters. AWS Lambda takes care of the underlying operating system and the execution environment.

Question: 397

According to the AWS shared responsibility model, who is responsible for the virtualization layer down to the physical security of the facilities in which AWS services operate?

A. It is the sole responsibility of the customer.
B. It is the sole responsibility of AWS.
C. It is a shared responsibility between AWS and the customer.
D. The customer's AWS Support plan tier determines who manages the configuration.

Answer: B

Explanation:
According to the AWS shared responsibility model, AWS is responsible for the security of the cloud, which includes the virtualization layer down to the physical security of the facilities in which AWSservices operate1. The customer is responsible for the security in the cloud, which includes the configuration and management of the AWS resources and applications that they use1.

Question: 398

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

A. Access to a technical project manager
B. Access to a technical account manager (TAM)

C. Access to a cloud support engineer

D. Access to a solutions architectA company wants to automatically set up and govern a multiaccount AWS environment.

Answer: B

Explanation:

AWS Enterprise Support plan is the highest level of support that AWS offers to its customers. One of the exclusive benefits of this plan is the access to a technical account manager (TAM), who is adedicated point of contact for guidance, advocacy, and support2. A technical project manager, a cloud support engineer, and a solutions architect are not exclusive benefits of the AWS Enterprise Support plan, as they are also available to customers with lower-tier support plans or through other AWS services or programs345.

Question: 399

Which AWS service provides this functionality?

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

Answer: D

Explanation:

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

Question: 400

A company wants its AWS usage to be more sustainable. The company wants to track, measure, review, and forecast polluting emissions that result from its AWS applications. Which AWS service or tool can the company use to meet these requirements?

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

Answer: B

Explanation:

AWS customer carbon footprint tool is a tool that helps customers measure and manage their carbon emissions from their AWS usage. It provides data on the carbon intensity, energy consumption, and estimated emissions of AWS services across regions and time periods. It

alsoenables customers to review and forecast their emissions, and compare them with industry benchmarks. AWS Health Dashboard is a service that provides personalized information about the health and performance of AWS services and resources. AWS Support Center is a service that provides access to AWS support resources, such as cases, forums, and documentation. Amazon QuickSight is a service that provides business intelligence and analytics for AWS data sources.

Question: 401

A company has a large number of Linux Amazon EC2 instances across several Availability Zones in an AWS Region. Applications that run on the EC2 instances need access to a common set of files. Which AWS service or device should the company use to meet this requirement?

A. AWS Backup
B. Amazon Elastic File System (Amazon EFS)
C. Amazon Elastic Block Store (Amazon EBS)
D. AWS Snowball Edge Storage Optimized

Answer: B

Explanation:
Amazon Elastic File System (Amazon EFS) is a service that provides a scalable and elastic file system for Linux-based workloads. It can be mounted on multiple Amazon EC2 instances across different Availability Zones within a region, allowing applications to access a common set of files1. AWS Backup is a service that provides a centralized and automated way to back up data across AWS services. Amazon Elastic Block Store (Amazon EBS) is a service that provides persistent block storage volumes for Amazon EC2 instances. AWS Snowball Edge Storage Optimized is a device that provides a petabyte-scale data transport and edge computing solution.

Question: 402

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

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

Answer: B

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

Question: 403

Which of the following is a benefit of operating in the AWS Cloud?

A. The ability to migrate on-premises network devices to the AWS Cloud
B. The ability to expand compute, storage, and memory when needed
C. The ability to host custom hardware in the AWS Cloud
D. The ability to customize the underlying hypervisor layer for Amazon EC2

Answer: B

Explanation:
One of the benefits of operating in the AWS Cloud is the ability to expand compute, storage, and memory when needed, which enables users to scale their applications and resources up or down based on demand. This also helps users optimize their costs and performance. The ability to migrate on-premises network devices to the AWS Cloud, the ability to host custom hardware in the AWS Cloud, and the ability to customize the underlying hypervisor layer for Amazon EC2 are not benefits of operating in the AWS Cloud, as they are either not possible or not recommended by AWS .

Question: 404

A company is operating several factories where it builds products. The company needs the ability to process data, store data, and run applications with local system interdependencies that require low latency.
Which AWS service should the company use to meet these requirements?

A. AWS IoT Greengrass
B. AWS Lambda
C. AWS Outposts
D. AWS Snowball Edge

Answer: C

Explanation:
AWS Outposts is a service that provides fully managed AWS infrastructure and services on premises. It allows users to run applications that require low latency and local data processing, while seamlessly connecting to the AWS Cloud for a consistent hybrid experience. AWS IoT Greengrass isa service that provides local compute, messaging, data caching, sync, and ML inference capabilities for connected devices. AWS Lambda is a service that allows users to run code without provisioning or managing servers. AWS Snowball Edge is a device that provides a petabyte-scale data transport and edge computing solution.

Question: 405

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

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

Answer: B

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

Question: 406

Which AWS service helps developers use loose coupling and reliable messaging between microservices?

A. Elastic Load Balancing
B. Amazon Simple Notification Service (Amazon SNS)
C. Amazon CloudFront
D. Amazon Simple Queue Service (Amazon SQS)

Answer: D

Explanation:
Amazon Simple Queue Service (Amazon SQS) is a service that provides fully managed message queues for asynchronous communication between microservices. It helps developers use loose coupling and reliable messaging by allowing them to send, store, and receive messages between distributed components without losing them or requiring each component to be always available1.Elastic Load Balancing is a service that distributes incoming traffic across multiple targets, such as Amazon EC2 instances, containers, and IP addresses. Amazon Simple Notification Service (Amazon SNS) is a service that provides fully managed pub/sub messaging for event-driven and push-based communication between microservices. Amazon CloudFront is a service that provides a fast and secure content delivery network (CDN) for web applications.

Question: 407

A company is building a mobile app to provide shopping recommendations to its customers. The company wants to use a graph database as part of the shopping recommendation engine. Which AWS database service should the company choose?

A. Amazon DynamoDB
B. Amazon Aurora
C. Amazon Neptune
D. Amazon DocumentDB (with MongoDB compatibility)

Answer: C

Explanation:
Amazon Neptune is a service that provides a fully managed graph database that supports property graphs and RDF graphs. It can be used to build applications that work with highly connected datasets, such as shopping recommendations, social networks, fraud detection, and knowledge graphs2. Amazon DynamoDB is a service that provides a fully managed NoSQL database that delivers fast and consistent performance at any scale. Amazon Aurora is a service that provides a fully managed relational database that is compatible with MySQL and PostgreSQL. Amazon DocumentDB (with MongoDB compatibility) is a service that provides a fully managed document database that is compatible with MongoDB.

Question: 408

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

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

Answer: A

Explanation:
On-Demand Instances are the default pricing model for Amazon EC2 instances. They allow users to pay for compute capacity by the second, with no long-term commitments or upfront payments. They are suitable for applications with short-term, irregular, or unpredictable workloads that cannot beinterrupted3. Savings Plans are a pricing model that offer significant savings on Amazon EC2 and AWS Fargate usage, in exchange for a commitment to a consistent amount of usage (measured in $/hour) for a 1-year or 3-year term. Spot Instances are a pricing model that offer spare Amazon EC2 compute capacity at up to 90% discount compared to On-Demand prices, but they can be interrupted by AWS with a two-minute notice when the demand exceeds the supply. Reserved Instances are a pricing model that offer up to 75% discount compared to On-Demand prices, in exchange for a commitment to use a specific instance type and size in a specific region for a 1-year or 3-year term.

Question: 409

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

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

Answer: C

Explanation:
AWS Direct Connect is a service that provides a dedicated network connection from on premises to the AWS Cloud. It can reduce network

costs, increase bandwidth throughput, and provide a more consistent network experience than internet-based connections. It can also provide low latency for applications that require real-time data transfer4. Amazon VPC is a service that provides a logically isolated section of the AWS Cloud where users can launch AWS resources in a virtual network that they define. Amazon Kinesis Data Streams is a service that provides a scalable and durable stream of data records for real-time data processing. Amazon OpenSearch Service is a service that provides a fully managed, scalable, and secure search and analytics solution that is compatible with Elasticsearch.

Question: 410

A company simulates workflows to review and validate that all processes are effective and that staff are familiar with the processes. Which design principle of the AWS Well-Architected Framework is the company following with this practice?

A. Perform operations as code.
B. Refine operation procedures frequently.
C. Make frequent, small, reversible changes.
D. Structure the company to support business outcomes.

Answer: B

Explanation:

Refine operation procedures frequently is one of the design principles of the operational excellence pillar of the AWS Well-Architected Framework. It means that users should continuously review and validate their operational processes to ensure that they are effective and that staff are familiar with them. It also means that users should identify and address any gaps or issues in their processes, and incorporate feedback and lessons learned from operational events5. Perform operations as code is another design principle of the operational excellence pillar, which means that users should automate and script their operational tasks to reduce human error and enable consistent and repeatable execution. Make frequent, small, reversible changes is a design principle of the reliability pillar, which means that users should deploy changes in small increments that can be easily tested and rolled back if necessary. Structure the company to support business outcomes is a design principle of the performance efficiency pillar, which means that users should align their organizational structure and culture with their business goals and cloud strategy.

Question: 411

A company has designed its AWS Cloud infrastructure to run its workloads effectively. The company also has protocols in place to continuously improve supporting processes.
Which pillar of the AWS Well-Architected Framework does this scenario represent?

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

Answer: D

Explanation:

The scenario represents the operational excellence pillar of the AWS Well-Architected Framework, which focuses on running and

monitoring systems to deliver business value and continually improve supporting processes and procedures1. Security, performance efficiency, cost optimization, and reliability are the other four pillars of the framework1.

Question: 412

Which AWS service is a continuous delivery and deployment solution?

A. AWSAppSync
B. AWS CodePipeline
C. AWS Cloud9
D. AWS CodeCommit

Answer: B

Explanation:
AWS CodePipeline is a continuous delivery and deployment service that automates the release process of software applications across different stages, such as source code, build, test, and deploy2. AWSAppSync, AWS Cloud9, and AWS CodeCommit are other AWS services related to application development, but they do not provide continuous delivery and deployment solutions34 .

Question: 413

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

Explanation:
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 .

Question: 414

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

Question: 415

A company's application has high customer usage during certain times of the day. The company wants to reduce the number of Amazon EC2 instances that run when application usage is low.
Which AWS service or instance purchasing option should the company use to meet this requirement?

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

Answer: D

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

Question: 416

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: A

Explanation:

The AWS best practice that ensures the most cost-effective architecture for the workload is rightsizing. Rightsizing means selecting the most appropriate instance type or resourceconfiguration that matches the needs of the workload. Rightsizing can help optimize performance and reduce costs by avoiding over-provisioning or under-provisioning of resources1. Loose coupling, caching, and redundancy are other AWS best practices that can improve the scalability, availability, and performance of the workload, but they do not necessarily ensure the most cost-effective architecture.

Question: 417

A company is looking for a managed machine learning (ML) service that can recommend products based on a customer's previous behaviors.
Which AWS service meets this requirement?

A. Amazon Personalize
B. Amazon SageMaker
C. Amazon Pinpoint
D. Amazon Comprehend

Answer: A

Explanation:
The AWS service that meets the requirement of providing a managed machine learning (ML) service that can recommend products based on a customer's previous behaviors is Amazon Personalize. Amazon Personalize is a fully managed service that enables developers to create personalized recommendations for customers using their own data. Amazon Personalize can automatically process and examine the data, identify what is meaningful, select the right algorithms, and train andoptimize a personalized recommendation model2. Amazon SageMaker, Amazon Pinpoint, and Amazon Comprehend are other AWS services related to machine learning, but they do not provide the specific functionality of product recommendation.

Question: 418

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

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

Answer: C

Explanation:
The solution that achieves the goal of having Amazon EC2 instances share the same geographic area but use multiple independent underlying power sources is to use EC2 instances in multiple Availability Zones in the same AWS Region. An Availability Zone is a physically isolated location within an AWS Region that has its own power, cooling, and network connectivity. An AWS Region is a geographical area that consists of two or more Availability Zones. By using multiple Availability Zones, users can increase the fault tolerance and resilience of their applications, as well as reduce latency for end users3. Using EC2 instances in a single Availability Zone, multiple AWS Regions, or the

same edge location and the same AWS Region would not meet the requirement of having multiple independent power sources.

Question: 419

Which AWS service should be used when a company needs to provide its remote employees with virtual desktops?

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

Answer: D

Explanation:
The AWS service that should be used when a company needs to provide its remote employees with virtual desktops is Amazon WorkSpaces. Amazon WorkSpaces is a fully managed, secure desktop-asa-service (DaaS) solution that runs on AWS. Amazon WorkSpaces allows users to provision cloudbased virtual desktops and provide their end users access to the documents, applications, and resources they need from any supported device, including Windows and Mac computers, Chromebooks, iPads, Fire tablets, and Android tablets4. Amazon Identity and Access Management (IAM), AWS Directory Service, and AWS IAM Identity Center (AWS Single Sign-On) are other AWS services related to identity and access management, but they do not provide virtual desktops.

Question: 420

A company needs a graph database service that is scalable and highly available. Which AWS service meets these requirements?

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

Answer: D

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

Question: 421

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

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

Answer: B

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

Question: 422

A company wants to design a reliable web application that is hosted on Amazon EC2. Which approach will achieve this goal?

A. Launch large EC2 instances in the same Availability Zone.
B. Spread EC2 instances across more than one security group.
C. Spread EC2 instances across more than one Availability Zone.
D. Use an Amazon Machine Image (AMI) from AWS Marketplace.

Answer: C

Explanation:
The approach that will achieve the goal of designing a reliable web application that is hosted on Amazon EC2 is to spread EC2 instances across more than one Availability Zone. An Availability Zone is a physically isolated location within an AWS Region that has its own power, cooling, and network connectivity. By spreading EC2 instances across multiple Availability Zones, users can increase the fault tolerance and availability of their web applications, as well as reduce latency for end users2. Launching large EC2 instances in the same Availability Zone, spreading EC2 instances across more than one security group, or using an Amazon Machine Image (AMI) from AWS Marketplace are not sufficient to ensure reliability, as they do not provide redundancy or resilience in case of an outage in one Availability Zone.

Question: 423

A company has a MySQL database running on a single Amazon EC2 instance. The company now requires higher availability in the event of an outage.
Which set of tasks would meet this requirement?

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

Answer: C

Explanation:

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

Question: 424

A company wants to verify if multi-factor authentication (MFA) is enabled for all users within its AWS accounts.
Which AWS service or resource will meet this requirement?

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

Answer: B

Explanation:

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

Question: 425

A company has migrated its workloads to AWS. The company wants to adopt AWS at scale and operate more efficiently and securely.

Which AWS service or framework should the company use for operational support?

A. AWS Support
B. AWS Cloud Adoption Framework (AWS CAF)
C. AWS Managed Services (AMS)
D. AWS Well-Architected Framework

Answer: D

Explanation:

The AWS Well-Architected Framework is a set of best practices and guidelines for designing and operating workloads on AWS. It helps customers achieve operational excellence, security, reliability, performance efficiency, cost optimization, and sustainability. The framework is based on six pillars, each with its own design principles, best practices, and questions. Customers can use the framework to assess their current state, identify gaps, and implement improvements12.

AWS Support is a service that provides technical assistance, guidance, and resources for AWS customers. It offers different plans with varying levels of access to AWS experts, response times, and features3. AWS Support does not provide a comprehensive framework for operational support.

AWS Cloud Adoption Framework (AWS CAF) is a guidance tool that helps customers plan and execute their cloud migration journey. It provides a set of perspectives, capabilities, and best practices to align the business and technical aspects of cloud adoption4. AWS CAF does not focus on operational support for existing workloads on AWS.

AWS Managed Services (AMS) is a service that operates AWS infrastructure on behalf of customers. It provides a secure and compliant environment, automates commonactivities, and applies best practices for provisioning, patching, backup, recovery, and monitoring5. AMS does not provide a framework for customers to operate their own workloads on AWS.

Question: 426

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

Which AWS service will meet these requirements?

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

Answer: D

Explanation:

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

AWS Key Management Service (AWS KMS) is a service that provides encryption and decryption services for data and keys. It does not provide temporary security credentials2.

AWS CloudHSM is a service that provides hardware security modules (HSMs) for cryptographic operations and key management. It does not provide temporary security credentials3.

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

Question: 427

Which AWS service offers object storage?

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

Answer: C

Explanation:
Amazon S3 is the AWS service that offers object storage. Object storage is a technology that stores and manages data in an unstructured format called objects. Each object consists of the data, metadata, and a unique identifier. Object storage is ideal for storing large amounts of unstructured data, such as photos, videos, email, web pages, sensor data, and audio files1. Amazon S3 provides

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

Question: 428

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

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

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

Answer: B

Explanation:
AWS Elastic Beanstalk is the AWS service that allows customers to deploy applications in the AWS Cloud as quickly as possible. AWS Elastic Beanstalk automatically handles the deployment, from capacity provisioning, load balancing, and auto-scaling to application health monitoring. Customers can upload their code and Elastic Beanstalk will take care of the rest1. AWS Elastic Beanstalk also minimizes the complexity that is related to the management of AWS resources. Customers can retain full control of the underlying AWS resources powering their applications and adjust the settings to suittheir needs1. Customers can also use the AWS Management Console, the AWS Command Line Interface (AWS CLI), or APIs to manage their applications1.
AWS Config is the AWS service that enables customers to assess, audit, and evaluate the configurations of their AWS resources. AWS Config continuously monitors and records the configuration changes of the resources and evaluates them against desired configurations or best practices2. AWS Config does not help customers deploy applications in the AWS Cloud as quickly as possible or minimize the complexity that is related to the management of AWS resources.

Amazon EC2 is the AWS service that provides secure, resizable compute capacity in the cloud. Customers can launch virtual servers called instances and choose from various configurations of CPU, memory, storage, and networking resources3. Amazon EC2 does not automatically handle the deployment or management of AWS resources for customers. Customers have to manually provision, configure, monitor, and scale their instances and other related resources.
Amazon Personalize is the AWS service that enables customers to create personalized recommendations for their users based on their behavior and preferences. Amazon Personalize uses machine learning to analyze data and deliver real-time recommendations4. Amazon Personalize does not help customers deploy applications in the AWS Cloud as quickly as possible or minimize the complexity that is related to the management of AWS resources.

Question: 429

Which AWS service or feature can a company use to apply security rules to specific Amazon EC2 instances?

A. Network ACLs
B. Security groups
C. AWS Trusted Advisor
D. AWS WAF

Answer: B

Explanation:
Security groups are the AWS service or feature that can be used to apply security rules to specific Amazon EC2 instances. Security groups are virtual firewalls that control the inbound and outbound traffic for one or more instances. Customers can create security groups and add rules that reflect the role of the instance that is associated with the security group. For example, a web server instance needs security group rules that allow inbound HTTP and HTTPS access, while a database instance needs rules that allow access for the type of database12. Security groups are stateful, meaning that the responses to allowed inbound traffic are alsoallowed, regardless of the outbound rules1. Customers can assign multiple security groups to an instance, and the rules from each security group are effectively aggregated to create one set of rules1.
Network ACLs are another AWS service or feature that can be used to control the traffic for a subnet. Network ACLs are stateless, meaning that they do not track the traffic that they allow. Therefore, customers must add rules for both inbound and outbound traffic3. Network ACLs are applied at the subnet level, not at the instance level.

AWS Trusted Advisor is an AWS service that provides best practice recommendations for security, performance, cost optimization, and fault tolerance. AWS Trusted Advisor does not apply security rules to specific Amazon EC2 instances, but it can help customers identify security gaps and improve their security posture4.

AWS WAF is an AWS service that helps protect web applications from common web exploits, such as SQL injection, cross-site scripting, and bot attacks. AWS WAF does not apply security rules to specific Amazon EC2 instances, but it can be integrated with other AWS services, such as Amazon CloudFront, Amazon API Gateway, and Application Load Balancer.

Question: 430

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

Answer: C,D

Explanation:

The AWS account root user is the identity that has complete access to all AWS services and resources in the account. It is accessed by signing in with the email address and password that were used to create the account1. The root user should be protected and used only for a few account and service management tasks that require it1. Therefore, the following actions are best practices for an AWS account root user:

Enable multi-factor authentication (MFA) on the root user. MFA is a security feature that requires users to provide two or more pieces of information to authenticate themselves, such as a password and a code from a device. MFA adds an extra layer of protection for the root user credentials, which can access sensitive information and perform critical operations in the account2.

Create an IAM user with administrator privileges for daily administrative tasks, instead of using the root user. IAM is a service that helps customers manage access to AWS resources for users and groups. Customers can create IAM users and assign them permissions to perform specific tasks on specific resources. Customers can also create IAM roles and policies to delegate access to other AWS services or external entities3. By creating an IAM user with administrator privileges, customers can avoid using the root user for everyday tasks and reduce the risk of accidental or malicious changes to the account1.

Question: 431

A company wants an automated process to continuously scan its Amazon EC2 instances for software vulnerabilities.

Which AWS service will meet these requirements?

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

Answer: B

Explanation:

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

Question: 432

Which AWS service can identify when an Amazon EC2 instance was terminated?

A. AWS Identity and Access Management (IAM)
B. AWS CloudTrail
C. AWS Compute Optimizer
D. Amazon EventBridge

Answer: B

Explanation:

AWS CloudTrail is the AWS service that can identify when an Amazon EC2 instance was terminated. AWS CloudTrail is a service that records API calls and events for AWS accounts and resources. AWS CloudTrail can capture the TerminateInstances event, which is triggered when an EC2 instance is terminated by a user or an AWS service. The event contains information such as the instance ID, the user identity, the source IP address, the time, and the reason for the termination12. Customers can use the CloudTrail console, the AWS CLI, or the AWS SDKs to viewand search for the TerminateInstances events in their event history or in their S3 buckets where they store their CloudTrail logs13.

Question: 433

A company needs to categorize and track AWS usage cost based on business categories.

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

A. Cost allocation tags
B. AWS Organizations
C. AWS Security Hub
D. AWS Cost and Usage Report

Answer: A

Explanation:

The AWS service or feature that the company should use to categorize and track AWS usage cost based on business categories is cost

allocation tags. Cost allocation tags are key-value pairs that users can attach to AWS resources to organize and track their AWS costs. Users can use cost allocation tags to filter and group their AWS costs by categories such as project, department, environment, or application. Users can also use cost allocation tags to generate detailed billing reports that show the costs associated with each tag3. AWS Organizations, AWS Security Hub, and AWS Cost and Usage Report are other AWS services or features that can help users with different aspects of their AWS usage, such as managing multiple accounts, monitoring security issues, or analyzing billing data, but they do not enable users to categorize and track AWS costs based on business categories.

Question: 434

Which options are AWS Cloud Adoption Framework (AWS CAF) cloud transformation journey recommendations? (Select 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) cloud transformation journey is a four-phase process that helps customers plan and execute their cloud migration and digital transformation. The four phases are:
Envision phase: This phase focuses on demonstrating how cloud will help accelerate the business outcomes of the customer. It involves identifying and prioritizing transformation opportunities across four domains: business, people, governance, and platform. It also involves associating the transformation initiatives with key stakeholders and measurable business outcomes1.
Align phase: This phase focuses on identifying capability gaps across six perspectives: business, people, governance, platform, security, and operations. It also involves identifying cross- organizational dependencies and surfacing stakeholder concerns and challenges. The goal of this phase is to create strategies for improving the cloud readiness, ensure stakeholder alignment, and facilitate relevant organizational change management activities1.
Launch phase: This phase focuses on delivering pilot initiatives in production and demonstrating incremental business value. Pilots should be highly impactful and influence future direction. The customer should learn from the pilots and adjust their approach before scaling to full production1.

Scale phase: This phase focuses on expanding production pilots and business value to the desired scale and ensuring that the business benefits associated with the cloud investments are realized and sustained1.

Question: 435

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

A. Amazon DynamoDB
B. AWS Lambda
C. AWS Fargate

D. Amazon EC2

Answer: D

Explanation:

Amazon EC2 is the AWS service that requires the customer to be fully responsible for applying operating system patches. Amazon EC2 is a service that provides secure, resizable compute capacity in the cloud. Customers can launch virtual servers called instances and choose from various configurations of CPU, memory, storage, and networking resources1. Customers have full control and access to their instances, which means they are also responsible for managing and maintaining them, including applying operating system patches2. Customers can use AWS Systems Manager Patch Manager, a feature of AWS Systems Manager, to automate the process of patching their EC2 instances with both security-related updates and other types of updates3.

A company wants an AWS service to collect and process 10 TB of data locally and transfer the data to AWS. The company has intermittent connectivity.

Which AWS service will meet these requirements?

AWS Database Migration Service (AWS DMS)
AWS DataSync
AWS Backup
AWS Snowball Edge

Answer: D

The correct answer is D. AWS Snowball Edge.
AWS Snowball Edge is a physical device that can be used to collect and process data locally and then transfer it to AWS. It is designed for situations where there is limited or intermittent network connectivity, or where bandwidth costs are high.AWS Snowball Edge can store up to 80 TB of data and has compute and storage capabilities to run applications on the device1.
AWS Database Migration Service (AWS DMS) is a service that helps migrate databases to AWS.It does not collect or process data locally, nor does it work ofline2.

AWS DataSync is a service that helps transfer data between on-premises storage systems and AWS storage services.It does not collect or process data locally, and it requires a network connection to work3.
AWS Backup is a service that helps automate and manage backups across AWS services. It does not collect or process data locally, nor does it transfer data to AWS.It only backs up data that is already in AWS4.
Reference:
1:AWS Snowball Edge2:AWS Database Migration Service (AWS DMS)3:AWS DataSync4:AWS Backup

Question: 436

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

Which AWS CAF governance perspective capability will meet these requirements?

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

Answer: A

Explanation:
The correct answer is A. Benefits management.
Benefits management is the AWS CAF governance perspective capability that helps you define and track business outcomes as part of your cloud transformation journey.Benefits management helps you align your cloud initiatives with your business objectives, measure the value and impact of your cloud investments, and communicate the benefits of cloud adoption to your stakeholders12. Risk management is the AWS CAF governance perspective capability that helps you identify and mitigate the potential risks associated with cloud adoption, such as security, compliance, legal, and operational risks12.
Application portfolio management is the AWS CAF governance perspective capability that helps you assess and optimize your existing application portfolio for cloud migration or modernization.Application portfolio management helps you categorize your applications based on their business value and technical fit, prioritize them for cloud adoption, and select the best migration or modernization strategy for each application12.
Cloud financial management is the AWS CAF governance perspective capability that helps you manage and optimize the costs and value of your cloud resources.Cloud financial management helps you plan and budget for cloud adoption, track and allocate cloud costs, implement cost optimization strategies, and report on cloud financial performance12.
Reference:
1:AWS Cloud Adoption Framework: Governance Perspective2:All you need to know about AWS Cloud Adoption Framework — Governance Perspective

Question: 437

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

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

Answer: C

Explanation:

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

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

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

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

Reference:

1:Foundational capabilities - An Overview of the AWS Cloud Adoption Framework2: [AWS Cloud Adoption Framework: Platform Perspective]3: [AWS Cloud Adoption Framework: Security Perspective]4: [AWS Cloud Adoption Framework: Governance Perspective] : [AWS Cloud Adoption Framework: Operations Perspective]

Question: 438

A company needs an automated vulnerability management service that continually scans AWS workloads for software vulnerabilities.

Which AWS service will meet these requirements?

A. Amazon GuardDuty

B. Amazon Inspector

C. AWS Security Hub

D. AWS Shield

Answer: B

Explanation:

The correct answer is B. Amazon Inspector.

Amazon Inspector is an automated vulnerability management service that continually scans AWS workloads for software vulnerabilities and unintended network exposure.Amazon Inspector automatically discovers workloads, such as Amazon EC2 instances, containers, and Lambda functions, and scans them for software vulnerabilities and unintended network exposure12.

Amazon GuardDuty is a threat detection service that monitors your AWS accounts and workloads for malicious or unauthorized activity.Amazon GuardDuty does not scan for software vulnerabilities, but rather analyzes AWS CloudTrail, Amazon VPC Flow Logs, and DNS logs to detect threats such as compromised credentials, backdoors, or crypto mining3.

AWS Security Hub is a security and compliance service that aggregates and prioritizes security findings from multiple AWS services and partner solutions.AWSSecurity Hub does not scan for software vulnerabilities, but rather provides a comprehensive view of your security posture across your AWS accounts4.

AWS Shield is a managed service that protects your web applications and network resources from distributed denial-of-service (DDoS) attacks.AWS Shield does not scan for software vulnerabilities, but rather provides detection and mitigation of DDoS attacks at the network and application layers5. Reference:

1:Automated Software Vulnerability Management - Amazon Inspector - AWS3: [Amazon GuardDuty - Intelligent Threat Detection Made Easy]2:AWS Re-LaunchesAmazon Inspector with New Architecture and Features - InfoQ4: [AWS Security Hub - Unified Security and

Compliance Center]5: [AWS Shield - Managed DDoS Protection]

Question: 439

A developer has been hired by a large company and needs AWS credentials.

Which are security best practices that should be followed? (Select TWO.)

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

Answer: A,E

Explanation:

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

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

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

B . Share the AWS account root user credentials with the developer. This is a bad practice that should be avoided. The root user has full access to all AWS resources and services, and can perform sensitive actions such as changing billing information, closing the account, or deleting all resources. Sharing the root user credentials exposes your account to potential compromise or misuse.You should never share your root user credentials with anyone, and use them only for account administration tasks5. C . Add the developer to the administrator's group in IAM. This is also a bad practice that should be avoided. The administrator's group has full access to all AWS resources and services, which is more than what a developer needs to perform their job. Adding the developer to the administrator's group violates the principle of least privilege and increases the risk of unauthorized access, data leakage, or accidental damage to AWS resources.You should create a custom group for the developer that grants only the necessary permissions for their role12.

D . Configure a password policy that ensures the developer's password cannot be changed. This is another bad practice that should be avoided. Preventing the developer from changing their password reduces their ability to protect their credentials and comply with security policies. For example, if the developer's password is compromised, they cannot change it to prevent further unauthorized access. Or if the company requires periodic password rotation, they cannot update their password to meet this requirement.You should allow the developer to change their password as needed, and enforce a password policy that sets reasonable rules for password management34.

Question: 440

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

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

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

Answer: A

Explanation: The AWS service that the company should use to meet these requirements is A. AWS Snowmobile. AWS Snowmobile is a service that allows you to migrate large amounts of data to AWS using a 45- foot long ruggedized shipping container that can store up to 100 petabytes of data. AWS Snowmobile is designed for situations where you need to move massive amounts of data to the cloud in a fast, secure, and cost-effective way.AWS Snowmobile has the least possible operational overhead because it eliminates the need to buy, configure, or manage hundreds or thousands of storage devices12. AWS Snowball Edge is a service that allows you to migrate data to AWS using a physical device that can store up to 80 terabytes of data and has compute and storage capabilities to run applications on the device. AWS Snowball Edge is suitable for situations where you have limited or intermittent network connectivity, or where bandwidth costs are high.However, AWS Snowball Edge has more operational overhead than AWS Snowmobile because you need to request multiple devices and transfer your data onto them using the client3.

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

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

Reference:

1:AWS Snowmobile - Move Exabytes of Data to the Cloud in Weeks2:AWS Snowmobile - Amazon Web Services3:Automated Software Vulnerability Management - Amazon Inspector - AWS4:AWS Data Exchange - Find, subscribe to, and use third-party data in ...5:AWS Database Migration Service - Amazon Web Services

Question: 441

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

Which AWS service will meet this requirement?

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

Answer: C

Explanation:

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

AWS Backup is a service that allows you to define a central data protection policy that works across AWS services for compute, storage, and database resources. You can use AWS Backup to create backup plans that specify the frequency, retention, and lifecycle of your backups, and apply them to your AWS resources using tags or resource IDs.AWS Backup supports various AWS services, such as Amazon

EC2,Amazon EBS, Amazon RDS, Amazon DynamoDB, Amazon EFS, Amazon FSx, and AWS Storage Gateway12.

AWS Batch is a service that allows you to run batch computing workloads on AWS.AWS Batch does not provide a central data protection policy, but rather enables you to optimize the allocation and utilization of your compute resources3.

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

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

Reference:

1:AWS Backup - Centralized backup across AWS services3:AWS Batch - Run Batch Computing Jobs on AWS2:Data Protection Reference Architectures with AWSBackup4:AWS Elastic Disaster Recovery - Prepare for and recover from disastersusing AWS5:Amazon FSx - Fully managed file storage for Windows and Linux applications

Question: 442

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

Which AWS service or resource will meet these requirements?

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

Answer: D

Explanation:

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

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

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

AWS Service Catalog is a service that allows customers to create and manage catalogs of IT services that are approved for use on AWS. AWS Service Catalogenables customers to control the configuration, deployment, and governance of their IT services.AWS Service Catalog does not directly engage third-party consultants, but rather helps customers standardize and streamline their IT service delivery5.

Reference:

1:AWS Partner Network (APN) - Amazon Web Services (AWS)2:Find an APN Partner - Amazon Web Services (AWS)3:AWS Support - Amazon Web Services4:AWS Organizations - Amazon Web Services5:AWS Service Catalog - Amazon Web Services

Question: 443

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

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

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

Answer: A

Explanation:

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

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

A single-region deployment model is a cloud computing architecture that runs an application and its data within a single geographic region. A single-region deployment model is simpler and cheaper than a multi-region deployment model, but it has limited scalability, availability, and performance.A single-region deployment model may notbe suitable for a company that wants to deploy an application globally, as it may face challenges such as network latency, regional outages, or regulatory compliance12.

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

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

Reference:

1:AWS Cloud Computing - W3Schools2:Understand the Different Cloud Computing Deployment Models Unit - Trailhead3:Regions and Availability Zones - Amazon Elastic Compute Cloud4:AWS Reference Architecture Diagrams

Question: 444

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

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

Answer: B

Explanation:

The option that is a customer responsibility under the AWS shared responsibility model is B. Application data security.
According to the AWS shared responsibility model, AWS is responsible for the security of the cloud, while the customer is responsible for the security in the cloud.This means that AWS manages the security of the underlying infrastructure, such as the hardware, software, networking, and facilities that run the AWS services, while the customer manages the security of their applications, data, and resources that they use on top of AWS12.
Application data security is one of the customer responsibilities under the AWS shared responsibility model. This means that the customer is responsible for protecting their application data from unauthorized access, modification, deletion, or leakage.The customer can use various AWS services and features to help withapplication 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.
Reference:

1:Shared Responsibility Model - Amazon Web Services (AWS)2:AWS Cloud Computing - W3Schools3: [Amazon EC2 FAQs - Amazon Web Services]4: [AWS Security - Amazon Web Services]5: [Amazon Virtual Private Cloud (VPC) - Amazon Web Services]

Question: 445

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

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

Answer: A

Explanation:

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

reduce costs and improve performance. You can view the recommendations on the AWS Compute Optimizer console or the Amazon EC2 console12.

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

Question: 446

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

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

Answer: B,C

Explanation:

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

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

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

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

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

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

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

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

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

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

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

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

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

Question: 447

How does the AWS Enterprise Support Concierge team help users?

A. Supporting application development
B. Providing architecture guidance
C. Answering billing and account inquiries
D. Answering questions regarding technical support cases

Answer: C

Explanation:

The AWS Enterprise Support Concierge team is a group of billing and account experts who specialize in working with enterprise customers. They can help customers with questions about billing, account management, cost optimization, and other non-technical issues. They can also assist customers with navigating and optimizing their AWS environment, such as setting up consolidated billing, applying for service limit increases, or requesting refunds.

Reference:

AWS Support Plan Comparison

AWS Enterprise Support Plan

Answer Explained: Which AWS Support plan provides access to AWS Concierge Support team for account assistance?

Question: 448

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

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

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

Answer: B

Explanation:

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

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

Amazon Kinesis is a platform for streaming data on AWS, enabling users to collect, process, and analyze real-time data. It is not designed

for querying data stored in Amazon S3. Amazon Kinesis consists of four services: Kinesis Data Streams, Kinesis Data Firehose, Kinesis Data Analytics, and Kinesis Video Streams3.

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

Reference:

Interactive SQL - Serverless Query Service - Amazon Athena - AWS

[Amazon Redshift - Data Warehouse Solution - AWS]

[Amazon Kinesis - Streaming Data Platform - AWS]

[Amazon Relational Database Service (RDS) - AWS]

Question: 449

A company needs to search for text in documents that are stored in Amazon S3.

Which AWS service will meet these requirements?

A. Amazon Kendra
B. Amazon Rekognition
C. Amazon Polly
D. Amazon Lex

Answer: A

Explanation:

Amazon Kendra is a highly accurate and easy to use intelligent search service powered by machine learning. It enables users to easily find the content they are looking for, even when it is scattered across multiple locations and content repositories within their organization. Amazon Kendra supports natural language queries, and can search fortext in documents stored in Amazon S3, as well as other sources such as SharePoint, OneDrive, Salesforce, ServiceNow, and more1.

Amazon Rekognition is a computer vision service that makes it easy to add image and video analysis to applications. It can detect objects, faces, text, scenes, activities, and emotions in images and videos. However, it is not designed for searching for text in documents stored in Amazon S32. Amazon Polly is a text-to-speech service that turns text into lifelike speech. It can create audio versions of books, articles, podcasts, and more. However, it is not designed for searching for text in documents stored in Amazon S33.

Amazon Lex is a service for building conversational interfaces using voice and text. It can create chatbots that can interact with users using natural language. However, it is not designed for searching for text in documents stored in Amazon S34.

Reference:

Amazon Kendra - Intelligent Search Service Powered by Machine Learning

Amazon Rekognition - Video and Image - AWS

Amazon Polly - Text-to-Speech Service - AWS

Amazon Lex - Build Conversation Bots - AWS

Question: 450

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

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

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

Answer: B

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

Reference:
What is Amazon Relational Database Service (Amazon RDS)?
Perform common DBA tasks for Amazon RDS DB instances

Question: 451

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

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

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

Answer: C

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

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

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

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

Question: 452

Which AWS service converts text to lifelike voices?

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

Answer: C

Explanation:

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

Amazon Transcribe is a service that converts speech to text, enabling you to create text transcripts from audio or video files. It can recognize multiple speakers, different languages, accents, dialects, and background noises. It can also add punctuation and formatting to the transcripts. Amazon Transcribe is useful for applications such as subtitling, captioning, transcription, and voice search. Amazon Rekognition is a service that provides image and video analysis using computer vision and deep learning. It can detect objects, faces, text, scenes, activities, and emotions in images and videos. It can also perform face recognition, face comparison, face search, celebrity recognition, and facial analysis. Amazon Rekognition is useful for applications such as security, social media, ecommerce, and media and entertainment.

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

Reference:

Text to Speech Software - Amazon Polly - Amazon Web Services

What is Text to Speech - Amazon Web Services (AWS)

AWS Amazon Polly - Text to Speech Converter - CodeCanyon

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

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

[Automatic Speech Recognition - Amazon Transcribe - AWS]
[Amazon Rekognition - Video and Image - AWS]
[Extract Text & Data - OCR - Amazon Textract - AWS]

Question: 453

A company wants to monitor its workload performance. The company wants to ensure that the cloud services are delivered at a level that meets its business needs.

Which AWS Cloud Adoption Framework (AWS CAF) perspective will meet these requirements?

A. Business
B. Governance
C. Platform
D. Operations

Answer: D

Explanation:
The Operations perspective helps you monitor and manage your cloud workloads to ensure that they are delivered at a level that meets your business needs. Common stakeholders include chief operations officer (COO), cloud director, cloud operations manager, and cloud operations engineers1. The Operations perspective covers capabilities such as workload health monitoring, incident management, change management, release management, configuration management, and disaster recovery2.
The Business perspective helps ensure that your cloud investments accelerate your digital transformation ambitions and business outcomes. Common stakeholders include chief executive officer (CEO), chief financial officer (CFO), chief information officer (CIO), and chief technology officer (CTO). The Business perspective covers capabilities such as business case development, value realization, portfolio management, and stakeholder management3.
The Governance perspective helps you orchestrate your cloud initiatives while maximizing organizational benefits and minimizing transformation-related risks. Common stakeholders include chief transformation officer, CIO, CTO, CFO, chief data officer (CDO), and chief risk officer (CRO). The Governance perspective covers capabilities such as governance framework, budget and cost management, compliance management, and data governance4.
The Platform perspective helps you build an enterprise-grade, scalable, hybrid cloud platform, modernize existing workloads, and implement new cloud-native solutions. Common stakeholders include CTO, technology leaders, architects, and engineers. The Platform perspective covers capabilities such as platform design and implementation, workload migration and modernization, cloud-native development, and DevOps5.
Reference:
AWS Cloud Adoption Framework: Operations Perspective
AWS Cloud Adoption Framework - Operations Perspective
AWS Cloud Adoption Framework: Business Perspective
AWS Cloud Adoption Framework: Governance Perspective
AWS Cloud Adoption Framework: Platform Perspective

Question: 454

A company wants a list of all users in its AWS account, the status of all of the users' access keys, and if multi-factor authentication (MFA) has been configured.

Which AWS service or feature will meet these requirements?

A. AWS Key Management Service (AWS KMS)
B. IAM Access Analyzer
C. IAM credential report
D. Amazon CloudWatch

Answer: C

Explanation:
IAM credential report is a feature that allows you to generate and download a report that lists all IAM users in your AWS account and the status of their various credentials, including access keys and MFA devices. You can use this report to audit the security status of your IAM users and ensure that they follow the best practices for using AWS1.
AWS Key Management Service (AWS KMS) is a service that allows you to create and manage encryption keys to protect your data. It does not provide information about IAM users or their credentials2.
IAM Access Analyzer is a feature that helps you identify the resources in your AWS account, such as S3 buckets or IAM roles, that are shared with an external entity. It does not provide information about IAM users or their credentials3.
Amazon CloudWatch is a service that monitors and collects metrics, logs, and events from your AWS resources and applications. It does not provide information about IAM users or their credentials4.

Reference:
Getting credential reports for your AWS account - AWS Identity and Access Management
AWS Key Management Service - Amazon Web Services
IAM Access Analyzer - AWS Identity and Access Management
Amazon CloudWatch - Amazon Web Services

Question: 455

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

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

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

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

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

Question: 456

A company wants an AWS service to provide product recommendations based on its customer data.

Which AWS service will meet this requirement?

A. Amazon Polly
B. Amazon Personalize
C. Amazon Comprehend
D. Amazon Rekognition

Answer: B

Explanation:

Amazon Personalize is an AWS service that helps developers quickly build and deploy a custom recommendation engine with real-time personalization and user segmentation1. It uses machine learning (ML) to analyze customer data and provide relevant recommendations based on their preferences, behavior, and context.Amazon Personalize can be used for various use cases such as optimizing recommendations, targeting customers more accurately, maximizing the value of unstructured text, and promoting items using

business rules1.

The other options are not suitable for providing product recommendations based on customer data. Amazon Polly is a service that converts text into lifelike speech. Amazon Comprehend is a service that uses natural language processing (NLP) toextract insights from text and documents. Amazon Rekognition is a service that uses computer vision (CV) to analyze images and videos for faces, objects, scenes, and activities.

1: Cloud Products - Amazon Web Services (AWS)
2: Recommender System - Amazon Personalize - Amazon Web Services
3: Top 25 AWS Services List 2023 - GeeksforGeeks
4: AWS to Azure services comparison - Azure Architecture Center
5: The 25+ Best AWS Cost Optimization Tools (Updated 2023) - CloudZero
6: Amazon Polly - Text-to-Speech Service - AWS
7: Natural Language Processing - Amazon Comprehend - AWS
8: Image and Video Analysis - Amazon Rekognition - AWS

Question: 457

A company wants to launch multiple workloads on AWS. Each workload is related to a different business unit. The company wants to separate and track costs for each business unit.

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

A. Use AWS Organizations and create one account for each business unit.
B. Use a spreadsheet to control the owners and cost of each resource.
C. Use an Amazon DynamoDB table to record costs for each business unit.

D. Use the AWS Billing console to assign owners to resources and track costs.

Answer: A

Explanation:

AWS Organizations is a service that helps you centrally manage and govern your AWS environment. You can use AWS Organizations to create multiple accounts for different business units, and group them into organizational units (OUs) that reflect your organizational structure1. By doing so, you can separate and track costs for each business unit using the account ID as a cost allocation tag2. You can also use AWSOrganizations to apply policies and controls to your accounts, such as service control policies (SCPs) and tag policies1.

The other options are not suitable for meeting the requirements with the least operational overhead. Using a spreadsheet or a DynamoDB table to control and record costs for each business unit would require manual data entry and maintenance, which is prone to errors and inconsistencies. Using the AWS Billing console to assign owners to resources and track costs would also require manual tagging of each resource, which is time-consuming and inefficient.

1: What Is AWS Organizations? - AWS Organizations
2: Cost Tagging and Reporting with AWS Organizations | AWS Cloud Financial Management

Question: 458

A company wants a time-series database service that makes it easier to store and analyze trillions of events each day.

Which AWS service will meet this requirement?

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

Answer: B

Explanation:

Amazon Timestream is a fast, scalable, and serverless time-series database service for IoT and other operational applications that makes it easy to store and analyze trillions of events per day up to 1,000 times faster and at as little as 1/10th thecost of relational databases1. Amazon Timestream saves you time and cost in managing the lifecycle of time series data, and its purpose-built query engine lets you access and analyze recent and historical data together with a single query1. Amazon Timestream has built-in time series analytics functions, helping you identify trends and patterns in near real time1.

The other options are not suitable for storing and analyzing trillions of events per day. Amazon

Neptune is a graph database service that supports highly connected data sets. Amazon Forecast is a machine learning service that generates accurate forecasts based on historical data. Amazon DocumentDB (with MongoDB compatibility) is a document database service that supports MongoDB workloads.

1: Time Series Database - Amazon Timestream - Amazon Web Services

Question: 459

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

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

Answer: A,C

Explanation:

The AWS services that are supported by Savings Plans are:

Amazon EC2: Amazon EC2 is a service that provides scalable computing capacity in the AWS cloud. You can use Amazon EC2 to launch virtual servers, configure security and networking, and manage storage. Amazon EC2 is eligible for both Compute Savings Plans and EC2 Instance Savings Plans12. Amazon SageMaker: Amazon SageMaker is a service that helps you build and deploy machine learning models. You can use Amazon SageMaker to access Jupyter notebooks, use common machine learning algorithms, train and tune models, and deploy them to a hosted environment. Amazon SageMaker is eligible for SageMaker Savings Plans13.

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

Question: 460

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

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

Answer: D

Explanation:

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

The other options are not correct for the following reasons:

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

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

AWS Service Catalog is a service that allows you to create and manage catalogs of IT services that are approved for use on AWS4. It does not track the progress of application migrations, but rather helps you manage the provisioning and governance of your IT services.

1: What Is AWS Migration Hub? - AWS Migration Hub

2: What Is AWS Application Discovery Service? - AWS Application Discovery Service

3: App Migration Tool - AWS Application Migration Service - AWS

4: What Is AWS Service Catalog? - AWS Service Catalog

Question: 461

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

Which AWS service will meet the requirement?

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

Answer: B

Explanation:

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

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

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

Question: 462

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

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

Answer: A,C

Explanation:

The AWS Cloud Adoption Framework (AWS CAF) people perspective capabilities are the organizational skills and processes that enable effective cloud adoption.According to the AWS CAF people perspective whitepaper1, there are seven capabilities in this perspective, two of which are: Organizational alignment: This capability helps you align your organizational structure, roles, and responsibilities to support your cloud transformation goals and objectives.It involves assessing your current and desired state ofalignment, identifying gaps and

misalignments, and designing and implementing changes to optimize your cloud performance1.

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

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

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

2: AWS Cloud Adoption Framework - AWS Cloud Adoption Framework

Question: 463

Which Amazon EC2 instan ce pricing model can provide discounts of up to 90%?

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

Answer: D

Explanation:

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 arerecommended for applications that are flexible, fault- tolerant, or have low priority, such as batch processing, data analysis, or testing and development.

Question: 464

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

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

Answer: D

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

Question: 465

A company has teams that have different job roles and responsibilities. The company's employees often change teams. The company needs to manage permissions for the employees so that the permissions are appropriate for the job responsibilities.
Which IAM resource should the company use to meet this requirement with the LEAST operational overhead?

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

Answer: B

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

Question: 466

A company is storing sensitive customer data in an Amazon S3 bucket. The company wants to protect the data from accidental deletion or overwriting.
Which S3 feature should the company use to meet these requirements?

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

Answer: B

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

Question: 467

A company plans to migrate to the AWS Cloud. The company is gathering information about its onpremises infrastructure and requires information such as the hostname, IP address, and MAC address.
Which AWS service will meet these requirements?

A. AWS DataSync
B. AWS Application Migration Service
C. AWS Application Discovery Service
D. AWS Database Migration Service (AWS DMS)

Answer: C

Explanation:
AWS Application Discovery Service is a service that helps you plan your migration to the AWS Cloud by collecting usage and configuration data about your on-premises servers and databases. This data includes information such as the hostname, IP address, and MAC address of each server, as well as the performance metrics, network connections, and processes running on them. You can use AWS Application Discovery Service to discover your on-premises inventory, map the dependencies between servers and applications, and estimate the cost and effort of migrating to AWS. You can also export the data to other AWS services, such as AWS Migration Hub and AWS Database Migration Service, to support your migration tasks. AWS Application Discovery Service offers two ways of performing discovery: agentless discovery and agent-based discovery. Agentless discovery uses a virtual appliance that you deploy on your VMware vCenter to collect data from your virtual machines and hosts. Agent-based discovery uses an agent that you install on each of your physical or virtual servers to collect data. You can choose the method that best suits your environment and needs. AWS DataSync is a service that helps you transfer data between your on-premises storage and AWS storage services, such as Amazon S3, Amazon EFS, and Amazon FSx for Windows File Server. AWS DataSync does not collect information about your on-premises infrastructure, but rather focuses on optimizing the data transfer speed, security, and reliability. AWS Application Migration Service is a service that helps you migrate your applications from your on-premises or cloud environment to AWS without making any changes to the applications, their architecture, or the migrated servers. AWS Application Migration Service does not collect information about your on-premises infrastructure, but rather uses a lightweight agent to replicate your servers as Amazon Machine Images (AMIs) and launch them as EC2 instances on AWS. AWS Database Migration Service is a service that helps you migrate your databases from your on-premises or cloud environment to AWS, either as a one-time migration or as a continuous replication. AWS Database MigrationService does not collect information about your on-premises infrastructure, but rather uses a source and a target endpoint to connect to your databases and transfer the data. Reference: AWS Application Discovery Service, AWS DataSync, AWS Application Migration Service, [AWS Database Migration Service]

Question: 468

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

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

Answer: C

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

Question: 469

Which AWS Cloud Adoption Framework (AWS CAF) capability belongs to the people perspective?

A. Data architecture
B. Event management
C. Cloud fluency
D. Strategic partnership

Answer: C

Explanation:
Cloud fluency is a capability that belongs to the people perspective of the AWS Cloud Adoption Framework (AWS CAF). Cloud fluency is the ability of the workforce to understand the benefits, challenges, and best practices of cloud computing, and to apply them to their roles and responsibilities. Cloud fluency helps the organization to adopt a cloud mindset, culture, and skills, and to leverage the full potential of the cloud. Cloud fluency can be achieved through various methods, such as training, certification, mentoring, coaching, and hands-on experience. Cloud

fluency is one of the four capabilities of the people perspective, along with culture, organizational structure, and leadership. The other three capabilities belong to different perspectives of the AWS CAF. Data architecture is a capability of the platform perspective, which helps you design and implement data solutions that meet your business and technical requirements. Event management is a capability of the operations perspective, which helps you monitor and respond to events that affect the availability, performance, and security of your cloud resources. Strategic partnership is a capability of the business perspective, which helps you establish and maintain relationships with external stakeholders, such as customers, partners, suppliers, and regulators, to create value and achieve your business goals. Reference: AWS Cloud Adoption Framework: People Perspective, AWS CAF - Cloud Adoption Framework - W3Schools

Question: 470

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

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

Answer: A

Explanation:

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

Question: 471

A company has a centralized group of users with large file storage requirements that have exceeded the space available on premises. The company wants to extend its file storage capabilities for this group while retaining the performance benefit of sharing content locally.
What is the MOST operationally efficient AWS solution for this scenario?

A. Create an Amazon S3 bucket for each user. Mount each bucket by using an S3 file system mounting utility.
B. Configure and deploy an AWS Storage Gateway file gateway. Connect each user's workstation to the file gateway.
C. Move each user's working environment to Amazon Workspaces. Set up an Amazon WorkDocs account for each user.
D. Deploy an Amazon EC2 instance and attach an Amazon Elastic Block Store (Amazon EBS) Provisioned IOPS volume. Share the EBS volume directly with the users.

Answer: B

Explanation:
AWS Storage Gateway is a hybrid cloud storage service that allows you to extend your on-premises file storage capabilities to the AWS Cloud. AWS Storage Gateway file gateway enables you to store and access your files in Amazon S3 using industry-standard file protocols such as NFS and SMB. File gateway caches frequently accessed files locally, providing low-latency access to your data. File gateway also optimizes the transfer of data between your on-premises environment and AWS, minimizing the amount of bandwidth consumed. By using file gateway, you can retain the performance benefit of sharing content locally while leveraging the scalability, durability, and cost-effectiveness of Amazon S3. Reference: AWS Storage Gateway, File Gateway

Question: 472

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 a serverless compute engine for containers that works with both Amazon Elastic Container Service (Amazon ECS) and Amazon Elastic Kubernetes Service (Amazon EKS). AWS Fargate allows you to run containers without having to manage servers or clusters of Amazon EC2 instances. With AWS Fargate, you only pay for the compute resources you use to run your containers, and you don't need to worry about scaling, patching, securing, or maintaining the underlying infrastructure.
AWS Fargate simplifies the deployment and management of containerized applications, and enables you to focus on building and running your applications instead of managing the infrastructure. Reference: AWS Fargate, What is AWS Fargate?

Question: 473

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

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

Answer: A,C

Explanation:
Agility in AWS Cloud computing means the ability to rapidly provision and deprovision AWS resources as needed, and the ability to experiment quickly with new ideas and solutions. Agility helps businesses to respond to changing customer demands, market opportunities, and competitive threats, and to innovate faster and cheaper. Agility also reduces the risk of failure, as businesses can test and validate their assumptions before committing to large-scale deployments. Some of the benefits of agility in AWS Cloud computing are:
The speed at which AWS resources are implemented: AWS provides a variety of services and tools that allow you to create, configure, and launch AWS resources in minutes, using the AWS Management Console, the AWS Command Line Interface (AWS CLI), the AWS Software Development Kits (AWS SDKs), or the AWS CloudFormation templates. You can also use the AWS Cloud Development Kit (AWS CDK) to define your AWS resources as code using familiar programming languages, and synthesize them into AWS CloudFormation templates. You can also use the AWS Service Catalog to create and manage standardized portfolios of AWS resources that meet your organizational policies and best practices. AWS also offers on-demand, pay-as-you-go pricing models, soyou only pay for the resources you use,

and you can scale them up or down as your needs change12345

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

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

Question: 474

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

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

Answer: C

Explanation:

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

Question: 475

A company wants to integrate its online shopping website with social media login credentials. Which AWS service can the company use to make this integration?

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

Answer: C

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

Question: 476

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

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

Answer: D

Explanation:
Amazon Aurora is a fully managed MySQL-compatible database that combines the performance and availability of traditional enterprise databases with the simplicity and cost-effectiveness of opensource databases. Amazon Aurora is part of the Amazon Relational Database Service (Amazon RDS) family, which means it inherits the benefits of a fully managed service, such as automated backups, patches, scaling, monitoring, and security. Amazon Aurora also offers up to five times the throughput of standard MySQL, as well as high availability, durability, and fault tolerance with up to 15 read replicas, cross-Region replication, and self-healing storage. Amazon Aurora is compatible with the latest versions of MySQL, as well as PostgreSQL, and supports various features and integrations that enhance its functionality and usability123
: Amazon Aurora, Amazon RDS, AWS — Amazon Aurora Overview

Question: 477

A customer runs an On-Demand Amazon Linux EC2 instance for 3 hours, 5 minutes, and 6 seconds. For how much time will the customer be billed?

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

Answer: C

Explanation:
Amazon EC2 usage is calculated by either the hour or the second based on the size of the instance, operating system, and the AWS Region

where the instances are launched. Pricing is per instance-hour consumed for each instance, from the time an instance is launched until it's terminated or stopped. Each partial instance-hour consumed is billed per-second for Linuxinstances and as a full hour for all other instance types1. Therefore, the customer will be billed for 3 hours and 6 minutes for running an On-Demand Amazon Linux EC2 instance for 3 hours, 5 minutes, and 6 seconds. Reference: Understand Amazon EC2 instance-hours billing

Question: 478

Which AWS service supports a hybrid architecture that gives users the ability to extend AWS infrastructure, AWS services, APIs, and tools to data centers, co-location environments, or onpremises facilities?

A. AWS Snowmobile
B. AWS Local Zones
C. AWS Outposts
D. AWS Fargate

Answer: C

Explanation:
AWS Outposts is a service that delivers AWS infrastructure and services to virtually any on-premises or edge location for a truly consistent hybrid experience. AWS Outposts allows you to extend and run native AWS services on premises, and is available in a variety of form factors, from 1U and 2U Outposts servers to 42U Outposts racks, and multiple rack deployments. With AWS Outposts, you can run some AWS services locally and connect to a broad range of services available in the local AWS Region. Run applications and workloads on premises using familiarAWS services, tools, and APIs2. AWS Outposts is the only AWS service that supports a hybridarchitecture that gives users the ability to extend AWS infrastructure, AWS services, APIs, and tools to data centers, co-location environments, or on-premises facilities. Reference: On-Premises Infrastructure - AWS Outposts Family

Question: 479

A company website is experiencing DDoS attacks.
Which AWS service can help protect the company website against these attacks?

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

Answer: C

Explanation:

AWS Shield is a managed DDoS protection service that safeguards applications running on AWS from distributed denial of service (DDoS) attacks. DDoS attacks are malicious attempts to disrupt the normal functioning of a website or application by overwhelming it with a large volume of traffic from multiple sources. AWS Shield provides two tiers of protection: Standard and Advanced. AWS Shield Standard is automatically enabled for all AWS customers at no additional cost. It protects your AWS resources, such as Amazon CloudFront, AWS Global Accelerator, and Amazon Route 53, from the most common and frequently occurring network and transport layer DDoS attacks.

AWS Shield Advanced is an optional paid service that provides additional protection for your AWS resources and applications, such as Amazon Elastic Compute Cloud (Amazon EC2), Elastic Load Balancing (ELB), Amazon Simple Storage Service (Amazon S3), Amazon Relational Database Service (Amazon RDS), and AWS Elastic Beanstalk. AWSShield Advanced offers enhanced detection and mitigation capabilities, 24/7 access to the AWS DDoS Response Team (DRT), real-time visibility and reporting, and cost protection against DDoS-related spikes in your AWS bill12

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

Question: 480 A company wants a customized assessment of its current on-premises environment. The company wants to understand its projected running costs in the AWS Cloud.

Which AWS service or tool will meet these requirements?

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

Answer: D

Explanation:

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

Question: 481

A company that has multiple business units wants to centrally manage and govern its AWS Cloud environments. The company wants to automate the creation of AWS accounts, apply service control policies (SCPs), and simplify billing processes.

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

A. AWS Organizations
B. Cost Explorer
C. AWS Budgets
D. AWS Trusted Advisor

Answer: A

Explanation:

AWS Organizations is an AWS service that enables you to centrally manage and govern your AWS Cloud environments across multiple business units. AWS Organizations allows you tocreate an organization that consists of AWS accounts that you create or invite to join. You can group your accounts into organizational units (OUs) and apply service control policies (SCPs) to them. SCPs are a type of policy that specify the maximum permissions for the accounts in your organization, and can help you enforce compliance and security requirements.

AWS Organizations also simplifies billing processes by enabling you to consolidate and pay for all member accounts with a single payment method. You can also use AWS Organizations to automate the creation of AWS accounts by using APIs or AWS CloudFormation templates.
Reference: What is AWS Organizations?, Policy-Based Management - AWS Organizations

Question: 482

According to security best practices, how should an Amazon EC2 instance be given access to an Amazon S3 bucket?

A. Hard code an IAM user's secret key and access key directly in the application, and upload the file.
B. Store the IAM user's secret key and access key in a text file on the EC2 instance, read the keys, then upload the file.
C. Have the EC2 instance assume a role to obtain the privileges to upload the file.
D. Modify the S3 bucket policy so that any service can upload to it at any time.

Answer: C

Explanation:
According to security best practices, the best way to give an Amazon EC2 instance access to an Amazon S3 bucket is to have the EC2 instance assume a role to obtain the privileges to upload the file. A role is an AWS Identity and Access Management (IAM) entity that defines a set of permissions for making AWS service requests. You can use roles to delegate access to users, applications, or services that don't normally have access to your AWS resources. For example, you can create a role that allows EC2 instances to access S3 buckets, and then attach the role to the EC2 instance. This way, the EC2 instance can assume the role and obtain temporary security credentials to access the S3 bucket. This method is more secure and scalable than storing or hardcoding IAM user credentials on the EC2 instance, as it avoids the risk of exposing or compromising the credentials. It also allows you to manage the permissions centrally and dynamically, and to audit the access using AWS CloudTrail. For more information on how to create and use roles for EC2 instances, see Using an IAM role to grant permissions to applications running on Amazon EC2 instances1
The other options are not recommended for security reasons. Hardcoding or storing IAM user credentials on the EC2 instance is a bad practice, as it exposes the credentials to potential attackers or unauthorized users who can access the instance or the application code. It also makes it difficult to rotate or revoke the credentials, and to track the usage of the credentials. Modifying the S3 bucket policy to allow any serviceto upload to it at any time is also a bad practice, as it opens the bucket to potential data breaches, data loss, or data corruption. It also violates the principle of least privilege, which states that you should grant only the minimum permissions necessary for a task.
: Using an IAM role to grant permissions to applications running on Amazon EC2 instances

Question: 483

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

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

Answer: B

Explanation:

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

Question: 484

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

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

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

Answer: A,D

Explanation:

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

Question: 485

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

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

Answer: C

Explanation:

AWS Database Migration Service (AWS DMS) is a managed and automated service that helps you migrate your databases from your on-premises or cloud environment to AWS, either as a one-time migration or as a continuous replication. AWS DMS supports migration between 20-plus database and analytics engines, such as PostgreSQL, Oracle, MySQL, SQL Server, MongoDB, Amazon Aurora, Amazon RDS, Amazon Redshift, and Amazon S3. AWS DMS also provides schema conversion and validation tools, as well as monitoring and security features. AWS DMS is a cost-effective and reliable solution for database migration, as you only pay for the compute resources and additional log storage used during the migration process, and you can minimize the downtime and data loss with Multi-AZ and ongoing replication12

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

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

Create a source endpoint that points to your on-premises PostgreSQL database. You need to provide the connection details, such as the server name, port, database name, user name, and password. You also need to specify the engine name as postgres and the SSL mode as required4

Create a target endpoint that points to your Amazon RDS PostgreSQL DB instance. You need to provide the connection details, such as the server name, port, database name, user name, and password. You also need to specify the engine name as postgres and the SSL mode as verify-full. Create a migration task that defines the migration settings and options, such as the replication instance, the source and target endpoints, the migration type (full load, full load and change data capture, or change data capture only), the table mappings, the task settings, and the task monitoring role. You can also use the AWS Schema Conversion Tool (AWS SCT) to convert your source schema to the target schema and apply it to the target endpoint before or after creating the migration task. Start the migration task and monitor its progress and status using the AWS DMS console, the AWS CLI, or the AWS DMS API. You can also use AWS CloudFormation to automate the creation and execution of the migration task.

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

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

Tool], [AWS Migration Hub], [AWS Application Migration Service]

Question: 486

Which cloud concept is demonstrated by using AWS Compute Optimizer?

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

Answer: B

Explanation:

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

Question: 487

A company wants to migrate its on-premises relational databases to the AWS Cloud. The company wants to use infrastructure as close to its current geographical location as possible.
Which AWS service or resource should the company use to select its Amazon RDS deployment area?

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

Answer: C

Explanation:

AWS Regions are the AWS service or resource that the company should use to select its Amazon RDS deployment area. AWS Regions are separate geographic areas where AWS clusters its data centers. Each AWS Region consists of multiple, isolated, and physically separate Availability Zones within a geographic area. Each AWS Region is designed to be isolated from the other AWS Regions to achieve the highest possible fault tolerance and stability. AWS provides a more extensive global footprint than any other cloud provider, and to support its global footprint and ensure customers are served across the world, AWS opens new Regions rapidly. AWS maintains multiple geographic Regions, including Regions in North America, South America, Europe, China, Asia Pacific, South Africa, and the Middle East. Amazon RDS is available in several AWS Regions worldwide. To create or work with an Amazon RDS DB instance in a specific AWS Region, you must use

the corresponding regional service endpoint. You can choose the AWS Region that meets your latency or legal requirements. You can also use multiple AWS Regions to design a disaster recovery solution or to distribute your read workload. Reference: Global Infrastructure Regions & AZs - aws.amazon.com, Regions, Availability Zones, and Local Zones - Amazon Relational Database Service

Question: 488

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

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

Answer: B

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

Question: 489

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

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

Answer: A

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

that is the company's direct responsibility. The other costs are either included in the AWS service fees or covered by AWS.
: AWS License Manager Pricing, Software licensing: The blind spot in public cloud costs, Cost Optimization tips for SQL Server Licenses on AWS, Microsoft Licensing on AWS

Question: 490

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 and Amazon CloudWatch are two AWS services or tools that the company can use to receive a notification when a specific AWS cost threshold is reached. AWS Budgets allows users to set custom budgets to track their costs and usage, and respond quickly to alerts received from email or Amazon Simple Notification Service (Amazon SNS) notifications if they exceed their threshold. Users can create cost budgets with fixed or variable target amounts, and configure their notifications for actual or forecasted spend. Users can also set up custom actions to run automatically or through an approval process when a budget target is exceeded. For example, users could automatically apply a custom IAM policy that denies them the ability to provision additional resources within an account. Amazon CloudWatch is a service that monitors applications, responds to performance changes, optimizes resource use, and provides insights into operational health. Users can use CloudWatch to collect and track metrics, which are variables they can measure for their resources and applications. Users can create alarms that watch metrics and send notifications or automatically make changes to the resources they are monitoring when a threshold is breached. Users can use CloudWatch to monitor their AWS costs and usage by creating billing alarms that send notifications when their estimated charges exceed a specified threshold amount. Userscan also use CloudWatch to monitor their Reserved Instance (RI) or Savings Plans utilization and coverage, and receive notifications when they fall below a certain level.
: Cloud Cost And Usage Budgets - AWS Budgets, What is Amazon CloudWatch?, Creating a billing alarm - Amazon CloudWatch

Question: 491

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

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

Answer: B

Explanation:

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

Question: 492

To reduce costs, a company is planning to migrate a NoSQL database to AWS.

Which AWS service is fully managed and can automatically scale throughput capacity to meet database workload demands?

A. Amazon Redshift
B. Amazon Aurora
C. Amazon DynamoDB
D. Amazon RDS

Answer: C

Explanation:

Amazon DynamoDB is a fully managed, serverless, key-value NoSQL database service that can deliver consistent, single-digit millisecond performance at any scale. DynamoDB can automatically scale throughput capacity to meet the demands of the database workload, without requiring any manual intervention. DynamoDB is ideal for NoSQL applications that need high performance, availability, and scalability. DynamoDB also offers features such as encryption at rest, point-in-time recovery, global tables, and in-memory caching. Reference: What is NoSQL?, Amazon DynamoDB, [AWS Cloud Practitioner Essentials: Module 4 - Databases in the Cloud]

Question: 493

A company is running a monolithic on-premises application that does not scale and is difficult to maintain. The company has a plan to migrate the application to AWS and divide the application into microservices.

Which best practice of the AWS Well-Architected Framework is the company following with this plan?

A. Integrate functional testing as part of AWS deployment.
B. Use automation to deploy changes.
C. Deploy the application to multiple locations.
D. Implement loosely coupled dependencies.

Answer: D

Explanation:

The company is following the best practice of implementing loosely coupled dependencies by migrating the application to AWS and dividing the application into microservices. Loosely coupled dependencies are a design principle of the AWS Well-Architected Framework that helps to reduce the interdependencies between components and improve the scalability, reliability, and performance of the system. By breaking down the monolithic application into smaller, independent, and modular services, the company can reduce the complexity and maintenance costs, increase the agility and flexibility, and enable faster and more frequent deployments. AWS CloudFormation is an AWS service that provides the ability to manage infrastructure as code. Infrastructure as code is a process of defining and provisioning AWS resources using code or templates, rather than manual actions or scripts. AWS CloudFormation allows users to create and update stacks of AWS resources based on predefined templates that describe the desired state and configuration of the resources. AWS CloudFormation automates and simplifies the deployment and management of AWS resources, and ensures consistency and repeatability across different environments and regions. AWS CloudFormation also supports rollback, change sets, drift detection, and nested stacks features that help users to monitor andcontrol the changes to their infrastructure. Reference: Implementing Loosely Coupled Dependencies, What is AWS CloudFormation?

Question: 494

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

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

Answer: D

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

Question: 495

A company wants to grant users in one AWS account access to resources in another AWS account. The users do not currently have permission to access the resources.
Which AWS service will meet this requirement?

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

Answer: B

Explanation:

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

Question: 496

A company is planning to host its workloads on AWS.
Which AWS service requires the company to update and patch the guest operating system?

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

Answer: C

Explanation:

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

Question: 497

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

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

Answer: D

Explanation:

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

Question: 498

Which AWS service or feature will search for and identify AWS resources that are shared externally?

A. Amazon OpenSearch Service
B. AWS Control Tower
C. AWS IAM Access Analyzer
D. AWS Fargate

Answer: C

Explanation:
AWS IAM Access Analyzer is an AWS service that helps customers identify and review the resources in their AWS account that are shared with an external entity, such as another AWS account, a root user, an organization, or a public entity. AWS IAM Access Analyzer uses automated reasoning, a form of mathematical logic and inference, to analyze the resource-based policies in the account and generate comprehensive findings that show the access level, the source of the access, the affected resource, and the condition under which the access applies. Customers can use AWS IAM Access Analyzer to audit their shared resources, validate their access policies, and monitor any changes to the resource sharing status. Reference: AWSIAM Access Analyzer, Identify and review resources shared with external entities, How AWS IAM Access Analyzer works

Question: 499

Which AWS service or tool helps users visualize, understand, and manage spending and usage over time?

A. AWS Organizations
B. AWS Pricing Calculator
C. AWS Cost Explorer
D. AWS Service Catalog

Answer: C

Explanation:
AWS Cost Explorer is the AWS service or tool that helps users visualize, understand, and manage spending and usage over time. AWS Cost Explorer is a web-based interface that allows users to access interactive graphs and tables that display their AWS costs and usage data. Users can create custom reports that analyze cost and usage data by various dimensions, such as service, region, account, tag, and more. Users can also view historical data for up to the last 12 months, forecast future costs for up to the next 12 months, and get

recommendations for cost optimization. AWS Cost Explorer also provides preconfigured views that show common cost and usage scenarios, such as monthly spend by service, daily spend by linked account, and Reserved Instance utilization. Users can use AWS Cost Explorer to monitor their AWS spending and usage trends, identify cost drivers and anomalies, and optimize their resource allocation and budget planning. Reference: Cloud Cost Analysis - AWS Cost Explorer - AWS, Analyzing your costs with AWS Cost Explorer

Question: 500

A company is migrating its workloads to the AWS Cloud. The company must retain full control of patch management for the guest operating systems that host its applications.
Which AWS service should the company use to meet these requirements?

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

Answer: B

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

Topic 3, Exam Pool C

Question: 501

At what support level do users receive access to a support concierge?

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

Answer: D

Explanation:
Users receive access to a support concierge at the Enterprise Support level. A support concierge is a team of AWS billing and account

experts that specialize in working with enterprise accounts. They can help users with billing and account inquiries, cost optimization, FinOps support, cost analysis, and prioritized answers to billing questions. The support concierge is included as part of the Enterprise Support plan, which also provides access to a Technical Account Manager (TAM), Infrastructure Event Management, AWS Trusted Advisor, and 24/7 technical support. Reference: AWS Support Plan Comparison, AWS Enterprise Support Plan, AWS Support Concierge

Question: 502

Which AWS service can a company use to visually design and build serverless applications?

A. AWS Lambda
B. AWS Batch
C. AWS Application Composer
D. AWS App Runner

Answer: C

Explanation:
AWS Application Composer is a service that allows users to visually design and build serverless applications. Users can drag and drop components, such as AWS Lambda functions, Amazon API Gateway endpoints, Amazon DynamoDB tables, and Amazon S3 buckets, to create aserverless application architecture. Users can also configure the properties, permissions, and dependencies of each component, and deploy the application to their AWS account with a few clicks. AWS Application Composer simplifies the design and configuration of serverless applications, and reduces the need to write code or use AWS CloudFormation templates. Reference: AWS Application Composer, AWS releases Application Composer to make serverless 'easier' but initial scope is limited

Question: 503

A company wants to migrate to AWS and use the same security software it uses on premises. The security software vendor offers its security software as a service on AWS.

Where can the company purchase the security solution?

A. AWS Partner Solutions Finder
B. AWS Support Center
C. AWS Management Console
D. AWS Marketplace

Answer: D

Explanation:
AWS Marketplace is an online store that helps customers find, buy, and immediately start using the software and services that run on AWS. Customers can choose from a wide range of software products in popular categories such as security, networking, storage, machine learning, business intelligence, database, and DevOps. Customers can also use AWS Marketplace to purchase software as a service (SaaS) solutions that are integrated with AWS. Customers can benefit from simplified procurement, billing, and deployment processes, as well as flexible pricing options and free trials. Customers can also leverage AWS Marketplace to discover and subscribe to solutions offered by AWS Partners, such as the security software vendor mentioned in the question. Reference: AWS Marketplace, [AWS Marketplace: Software as a Service (SaaS)], [AWS Cloud Practitioner Essentials: Module 6 - AWS Pricing, Billing, and Support]

Question: 504

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

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

Answer: C

Explanation:
According to the AWS shared responsibility model, AWS is responsible for protecting the infrastructure that runs all ofthe services offered in the AWS Cloud, such as data centers, hardware, software, networking, and facilities1. This includes the configuration of infrastructure devices, such as routers, switches, firewalls, and load balancers2. Customers are responsible for managing their data, applications, operating systems, security groups, and other aspects of their AWS environment1. Therefore, options A, B, and D are customer responsibilities, not AWS responsibilities. Reference: 1: AWS Well-Architected Framework - Elasticity; 2: Reactive Systems on

AWS - Elastic

Question: 505

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 in the AWS Cloud refers to the ability to acquire resources as you need them and release resources when you no longer need them. In the cloud, you want to do this automatically1. This means that you can rightsized resources as demand shifts, and you can easily procure resources when they are needed. Elasticity is not related to how quickly an Amazon EC2 instance can be restarted, the maximum amount of RAM an Amazon EC2 instance can use, or the pay-as-you-go billing model. These are aspects of scalability, performance, and cost, respectively2.
For more information on elasticity, you can refer to the following sources:
Elasticity - AWS Well-Architected Framework
Elastic - Reactive Systems on AWS
What is the difference between scalability and elasticity?

Question: 506

A company wants to migrate its PostgreSQL database to AWS. The company does not use the database frequently.
Which AWS service or resource will meet these requirements with the LEAST management overhead?

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

Answer: D

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

Question: 507

A company is using Amazon DynamoDB for its application database.
Which tasks are the responsibility of AWS, according to the AWS shared responsibility model? (Select TWO.)

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

Answer: D,E

Explanation:
According to the AWS shared responsibility model, AWS is responsible for security of the cloud, while customers are responsible for security in the cloud. This means that AWS is responsible for protecting the infrastructure that runs AWS services, such as hardware, software, networking, and facilities. Customers are responsible for managing their data, classifying their assets, and using IAM tools to apply the appropriate permissions. For abstracted services, such as Amazon DynamoDB, AWS operates the infrastructure layer, the operating system, and platforms, and provides customers with public endpoints to store and retrieve data. Customers are responsible for classifying their data, managing their encryption options, and configuring their access permissions. Reference: Shared Responsibility Model, Security and compliance in Amazon DynamoDB, [AWS Cloud Practitioner Essentials: Module 2 - Security in the Cloud]

Question: 508

A company wants to create a globally accessible ecommerce platform for its customers. The company wants to use a highly available and scalable DNS web service to connect users to the platform.
Which AWS service will meet these requirements?

A. Amazon EC2
B. Amazon VPC
C. Amazon Route 53
D. Amazon RDS

Answer: C

Explanation:
Amazon Route 53 is a highly available and scalable Domain Name System (DNS) web service that can route internet traffic to the company's ecommerce platform1. Route 53 can also register domain names, check the health of resources, and provide global DNS features2. Route 53 can connect users to the platform by translating human-readable names like www.example.com into the numeric IP addresses that computers use to communicate witheach other2. Reference: 1: Amazon Route 53 | DNS Service | AWS; 2: What is Amazon Route 53? - Amazon Route 53

Question: 509

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

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

Answer: D

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

Question: 510

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

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

Answer: A

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

Question: 511

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

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

Answer: C,D

Explanation:

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

Question: 512

A company wants to generate a list of IAM users. The company also wants to view the status of various credentials that are associated with the users, such as password, access keys: and multifactor authentication (MFA) devices
Which AWS service or feature will meet these requirements?

A. IAM credential report

B. AWS IAM Identity Center (AWS Single Sign-On)

C. AWS Identity and Access Management Access Analyzer
D. AWS Cost and Usage Report

Answer: A

Explanation:
An IAM credential report is a feature of AWS Identity and Access Management (IAM) that allows you to view and download a report that lists all IAM users in your account and the status of their various credentials, such as passwords, access keys, and MFA devices. You can use this reportto audit the security status of your IAM users and ensure that they follow the best practices for credential management1.
Reference: 1: AWS Documentation - IAM User Guide - Getting credential reports for your AWS account

Question: 513

Which of the following is an AWS Well-Architected Framework design principle for operational excellence in the AWS Cloud?

A. Go global in minutes
B. Make frequent, small, reversible changes
C. Implement a strong foundation of identity and access management
D. Stop spending money on hardware infrastructure for data center operations

Answer: B

Explanation:
Making frequent, small, reversible changes is one of the design principles for operational excellence in the AWS Cloud, as defined by the AWS Well-Architected Framework. This principle means that you should design your workloads to allow for rapid and safe changes, such as deploying updates, rolling back failures, and experimenting with new features. By making small and reversible changes, you can reduce the risk of errors, minimize the impact of failures, and increase the speed of recovery2. Reference: 2: AWS Documentation - AWS Well-Architected Framework - Operational Excellence Pillar

Question: 514

Which type of AWS storage is ephemeral and is deleted when an Amazon EC2 instance is stopped or terminated?

A. Amazon Elastic Block Store (Amazon EBS)
B. Amazon EC2 instance store
C. Amazon Elastic File System (Amazon EFS)
D. Amazon S3

Answer: B

Explanation:
Amazon EC2 instance store provides temporary block-level storage for your EC2 instance. This storage is located on disks that are physically attached to the host computer. Instance store is ideal for temporary storage of information that changes frequently, such as buffers, caches, scratch data,and other temporary content. It can also be used to store temporary data that you replicate across a fleet of instances,

such as a load-balanced pool of web servers. An instance store consists of one or more instance store volumes exposed as block devices. The size of an instance store as well as the number of devices available varies by instance type and instance size. The virtual devices for instance store volumes are ephemeral[0-23]. Instance types that support one instance store volume have ephemeral0. Instance types that support two or more instance store volumes have ephemeral0, ephemeral1, and so on. Instance store pricing Instance store volumes are included as part of the instance's usage cost. The data on an instance store volume persists even if the instance is rebooted. However, the data does not persist if the instance is stopped, hibernated, or terminated. When the instance is stopped, hibernated, or terminated, every block of the instance store volume is cryptographically erased. Therefore, do not rely on instance store volumes for valuable, long-term data. If you need to retain the data stored on an instance store volume beyond the lifetime of the instance, you need to manually copy that data to more persistent storage, such as an Amazon EBS volume, an Amazon S3 bucket, or an Amazon EFS file system. There are some events that can result in your data not persisting throughout the lifetime of the instance. The following table indicates whether data on instance store volumes is persisted during specific events, for both virtualized and bare metal instances1. Reference: Amazon EC2 instance store - Amazon Elastic Compute Cloud

Question: 515

Which AWS Cloud deployment model uses AWS Outposts as part of the application deployment infrastructure?

A. On-premises
B. Serverless
C. Cloud-native
D. Hybrid

Answer: D

Explanation:

AWS Outposts is a fully managed service that extends AWS infrastructure, services, APIs, and tools to customer premises. By providing local access to AWS managed infrastructure, AWS Outposts enables customers to build and run applications on premises using the same programming interfaces as in AWS Regions, while using local compute and storage resources for lower latency and local data processing needs. An Outpost is a pool of AWS compute and storage capacity deployed at a customer site. AWS operates, monitors, and manages this capacity as part of an AWS Region. You can create subnets on your Outpost and specify them when you create AWS resources such as EC2 instances,

EBS volumes, ECS clusters, and RDS instances. Instances in Outpost subnets communicate with other instances in the AWS Region using private IP addresses, all within the same VPC. Outposts solutions allow you to extend and run native AWS services on premises, and isavailable in a variety of form factors, from 1U and 2U Outposts servers to 42U Outposts racks, and multiple rack deployments. With AWS Outposts, you can run some AWS services locally and connect to a broad range of services available in the local AWS Region2. AWS Outposts is a hybrid cloud deployment model that uses AWS Outposts as part of the application deployment infrastructure. Hybrid cloud is a cloud computing environment that uses a mix of on-premises, private cloud, and public cloud services with orchestration between the platforms. Hybrid cloud provides businesses with greater flexibility, more deployment options, and optimized costs. By using AWS Outposts, customers can benefit from the fully managed infrastructure, services, APIs, and tools of AWS on premises, while still having access to the full range of AWS services available in the Region for a truly consistent hybrid experience3. Reference: On-Premises Private Cloud - AWS Outposts Family - AWS, What is AWS Outposts? - AWS Outposts

Question: 516

A company wants to automatically add and remove Amazon EC2 instances. The company wants the EC2 instances to adjust to varying workloads dynamically.
Which service or feature will meet these requirements?

A. Amazon DynamoDB
B. Amazon EC2 Spot Instances
C. AWS Snow Family
D. Amazon EC2 Auto Scaling

Answer: D

Explanation:
Amazon EC2 Auto Scaling is a service that helps you maintain application availability and allows you to automatically add or remove EC2 instances according to definable conditions. You can create collections of EC2 instances, called Auto Scaling groups, and specify the minimum and maximum number of instances in each group. You can also define scaling policies that adjust the number of instances based on the demand on your application. Amazon EC2 Auto Scaling helps you improve the performance, reliability, and cost-efficiency of your EC2workloads123. Reference: 1: VDI Desktops - Amazon WorkSpaces Family - AWS, 2: What is Amazon EC2 Auto Scaling? - Amazon EC2 Auto Scaling, 3: Discover Amazon EC2 Auto Scaling Unit | Salesforce Trailhead

Question: 517

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

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

Answer: D

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

Question: 518

Which AWS service is a cloud security posture management (CSPM) service that aggregates alerts from various AWS services and

partner products in a standardized format?

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

Answer: A

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

Question: 519

Which AWS services can a company use to achieve a loosely coupled architecture? (Select TWO.)

A. Amazon Workspaces
B. Amazon Simple Queue Service (Amazon SQS)
C. Amazon Connect
D. AWS Trusted Advisor
E. AWS Step Functions

Answer: B,E

Explanation:
Amazon Simple Queue Service (Amazon SQS) and AWS Step Functions are AWS services that can be used to achieve a loosely coupled architecture. Amazon SQS is a fully managed message queuing service that enables you to decouple and scale microservices, distributed systems, and serverless applications. AWS Step Functions lets you coordinate multiple AWS services into serverless workflows so you can build and update apps quickly. Using Step Functions, you can design and run workflows that stitch together services such as AWS Lambda and Amazon SNS into feature-rich applications. Reference: Amazon SQS, AWS Step Functions

Question: 520

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

A. AWS Outposts
B. AWS Transfer Family

C. AWS Snow Family
D. AWS Migration Hub

Answer: C

Explanation:
AWS Snow Family is a group of devices that transport data in and out of AWS. AWS Snow Family devices are physical devices that can transfer up to exabytes of data. One exabyte is 1 000 000 000 000 megabytes. AWS Snow Family devices are designed for use in remote locations where internet connectivity is limited or unavailable. You can use these devices to collect and process data at the edge, and then ship them back to AWS for data upload. AWS Snow Family consists of three types of devices: AWS Snowcone, AWS Snowball, and AWS Snowmobile1234. Reference: 1: Edge Computing Devices, Secure Data Transfer - AWS Snow Family - AWS, 2: AWS Snow Family Documentation, 3: AWS Snow Family - W3Schools, 4: AWS Snow Family: Data Storage, Migration, and Computation

Question: 521

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

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

Answer: B,E

Explanation:
Consolidated billing is a feature of AWS Organizations that enables customers to consolidate billing and payment for multiple AWS accounts. With consolidated billing, customers can group multiple AWS accounts under one payer account, making it easier to manage billing and track costs across multiple accounts. Consolidated billing also offers benefits such as volume discounts, Reserved Instance discounts, and Savings Plans discounts. Consolidated billing is offered at no additional cost. Multiple AWS accounts is a feature of AWS Organizations that enables customers to create and manage multiple AWS accounts from a central location. With multiple AWS accounts, customers can isolate workloads for different departments, projects, or environments, and apply granular access controls and policies to each account. Multiple AWS accounts also helps customers improve security, compliance, and governance of their AWS resources56. Reference: 5: Consolidated billing for AWS Organizations - AWS Billing, 6: Understanding Consolidated Bills - AWS Billing, 7: AWS Consolidated Billing: Tutorial & Best Practices, 8: Simplifying Your Bills With Consolidated Billing on AWS - Aimably, 9: AWS Consolidated Billing - W3Schools

Question: 522

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

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

Answer: B

Explanation:
AWS Artifact is a service that provides on-demand access to security and compliance reports from AWS and Independent Software Vendors (ISVs) who sell their products on AWS Marketplace. You can use AWS Artifact to download auditor-issued reports, certifications, accreditations, and other third- party attestations of AWS compliance with various standards and regulations, such as PCI-DSS, HIPAA, FedRAMP, GDPR, and more1234. You can also use AWS Artifact to review, accept, and manage your agreements with AWS and apply them to current and future accounts within your organization2. Reference: 1: Cloud Compliance - Amazon Web Services (AWS), 2: Security Compliance Management - AWS Artifact - AWS, 3: AWS ComplianceContact Us - Amazon Web Services, 4: AWS SECURITY AND COMPLIANCE QUICK REFERENCE GUIDE

Question: 523

A company wants to manage its AWS Cloud resources through a web interface. Which AWS service will meet this requirement?

A. AWS Management Console
B. AWS CLI
C. AWS SDK
D. AWS Cloud

Answer: A

Explanation:
AWS Management Console is a web application that allows you to manage and monitor your AWS Cloud resources through a user-friendly interface. You can use the AWS Management Console to access and experiment with over 150 AWS services, view and modify your account and billing information, get in-console help from AWS Support, and customize your dashboard with widgets that display key metrics and information for your applications567. You can also use the AWS Management Console to launch and configure AWS resources using wizards andtemplates, without writing any code5. Reference: 5: Manage AWS Resources - AWS Management Console - AWS, 6: Getting Started with the AWS Management Console, 7: Manage AWS Resources - AWS Management Console Features - AWS

Question: 524

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

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

Answer: C

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

Question: 525

A company needs to block SQL injection attacks.
Which AWS service or feature can meet this requirement?

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

Answer: A

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

Question: 526

A company has a physical tape library to store data backups. The tape library is running out of space. The company needs to extend the tape library's capacity to the AWS Cloud.
Which AWS service should the company use to meet this requirement?

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

Answer: D

Explanation:
AWS Storage Gateway is a hybrid cloud storage service that provides on-premises access to virtually unlimited cloud storage. You can use AWS Storage Gateway to simplify storage management and reduce costs for key hybrid cloud storage use cases. One of these use cases is tape-based backup, which allows you to store data backups on virtual tapes in the AWS Cloud. You can use the Tape Gateway feature of

AWS Storage Gateway to extend your existing physical tape library to the AWS Cloud. Tape Gateway provides a virtual tape infrastructure that scales seamlessly with your backup needs and eliminates the operational burden of provisioning, scaling, and maintaining a physical tape infrastructure123. Reference: 1: CloudStorage Appliances, Hybrid Device - AWS Storage Gateway - AWS, 2: AWS Storage Gateway Documentation, 3: AWS Storage Gateway Features | Amazon Web Services

Question: 527

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

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

Answer: A

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

Question: 528

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

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

Answer: B,C

Explanation:
Amazon CloudFront and AWS Global Accelerator are two AWS services that make use of global edge locations. Edge locations are AWS sites that are deployed worldwide in major cities and places with a high population. Edge locations are used to cache data and reduce latency for end-user access1. Amazon CloudFront is a content delivery network (CDN) service that securely delivers data, videos, applications, and APIs to customers globally with low latency and high transfer speeds. Amazon CloudFront uses a global network of over 200 edge locations and 13 regional edge caches to cache your content closer to your viewers, improving performance and reducing costs23.

AWS Global Accelerator is a networking service that improves the availability and performance of your applications with local or global users. AWS Global Accelerator uses the AWS global network to route user traffic to the optimal endpoint based on health, performance, and policies. AWS Global Accelerator uses over 100 edge locations to bring your application endpoints closer to your users, reducing networkhops and improving user experience45. Reference: 1: AWS for the Edge - Amazon Web Services (AWS), 2: Content Delivery Network (CDN) - Amazon CloudFront - AWS, 3: Amazon CloudFront Documentation, 4: AWS Global Accelerator - Amazon Web Services, 5: AWS Global Accelerator Documentation

Question: 529

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

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

Answer: D

Explanation:
Amazon CloudWatch alarm is an AWS service or feature that can initiate an Amazon EC2 Auto Scaling action based on CPU utilization. Amazon CloudWatch is a monitoring and observability service that collects and tracks metrics, logs, events, and alarms for your AWS resources and applications. Amazon CloudWatch alarms are actions that you can configure to send notifications or automatically make changes to the resources you are monitoring based on rules that you define67. Amazon EC2 Auto Scaling is a service that helps you maintain application availability and allows you to automatically add or remove EC2 instances according to definable conditions. You can create dynamic scaling policies that track a specific CloudWatch metric, such as CPU utilization, and define what action to take when the associated CloudWatch alarm is in ALARM. When the policy is in effect, Amazon EC2 Auto Scaling adjusts the group's desired capacity up or down when the threshold of an alarm is breached89. Reference: 6: Cloud Monitoring - Amazon CloudWatch - AWS, 7: Amazon CloudWatch Documentation, 8: Dynamic scaling for Amazon EC2 Auto Scaling, 9: Amazon EC2 Auto Scaling Documentation

Question: 530

Which of the following services can be used to block network traffic to an instance? (Select TWO.)

A. Security groups
B. Amazon Virtual Private Cloud (Amazon VPC) flow logs
C. Network ACLs
D. Amazon CloudWatch
E. AWS CloudTrail

Answer: A,C

Explanation:
Security groups and network ACLs are two AWS services that can be used to block network traffic to an instance. Security groups are virtual firewalls that control the inbound and outbound traffic for your instances at the instance level. You can specify which protocols, ports, and

source or destination IP addresses are allowed or denied for each instance. Security groups are stateful, which means that they automatically allow return traffic for any allowed inbound or outbound traffic123. Network ACLs are virtual firewalls that control the inbound and outbound traffic for your subnets at the subnet level. You can create rules to allow or deny traffic based on protocols, ports, and source or destination IP addresses. Network ACLs are stateless, which means that you have to explicitly allow return traffic for any allowed inbound or outbound traffic456. Reference: 1: Security groups for your VPC - Amazon Virtual Private Cloud, 2: Security Groups for Your VPC - Amazon Elastic Compute Cloud, 3: AWS Security Groups: Everything You Need to Know, 4: Network ACLs - Amazon Virtual Private Cloud, 5: Control traffic to subnetsusing network ACLs - Amazon Virtual Private Cloud, 6: AWS Network ACLs: Everything You Need to Know

Question: 531

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

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

Answer: C,E

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

Question: 532

A company wants to quickly implement a continuous integration/continuous delivery (CI/CD) pipeline.
Which AWS service will meet this requirement?

A. AWS Config
B. Amazon Cognito
C. AWS DataSync
D. AWS CodeStar

Answer: D

Explanation:
AWS CodeStar is a service that enables you to quickly develop, build, and deploy applications on AWS. It provides a unified user interface for managing your application lifecycle, including code repositories, build pipelines, deployments, and project dashboards. AWS CodeStar also integrates with other AWS services, such as AWS CodeCommit, AWS CodeBuild, AWS CodeDeploy, and AWS CodePipeline, to create a complete CI/CD pipeline for your application12. Reference: AWS CodeStar
AWS Certified Cloud Practitioner Exam Guide

Question: 533

A development team wants to deploy multiple test environments for an application in a fast repeatable manner.
Which AWS service should the team use?

A. Amazon EC2
B. AWS CloudFormation
C. Amazon QuickSight
D. Amazon Elastic Container Service (Amazon ECS)

Answer: B

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

Question: 534

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

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

Answer: C

Explanation:
AWS Direct Connect is a cloud service solution that makes it easy to establish a dedicated network connection from your premises to AWS. Using AWS Direct Connect, you can establish private connectivity between AWS and your datacenter, office, or colocation environment, which in many cases can reduce your network costs, increase bandwidth throughput, and provide a more consistent network experience

than internet-based connections12. Reference: 1: Dedicated Network Connection - AWS Direct Connect - AWS, 2: What is AWS Direct Connect? - AWS Direct Connect

Question: 535

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

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

Answer: C

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

Question: 536

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

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

Answer: B,C

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

Question: 537

A company wants to migrate its database to a managed AWS service that is compatible with PostgreSQL.
Which AWS services will meet these requirements? (Select TWO)

A. Amazon Athena
B. Amazon RDS
C. Amazon EC2
D. Amazon DynamoDB
E. Amazon Aurora

Answer: B,E

Explanation:
Amazon RDS and Amazon Aurora are both managed AWS services that support the PostgreSQL database engine. Amazon RDS makes it easier to set up, operate, and scale PostgreSQL deployments on the cloud, while Amazon Aurora is a cloud-native database engine that is compatible with PostgreSQL and offers higher performance and availability. Amazon Athena is a serverless query service that does not support PostgreSQL, but can analyze data in Amazon S3 using standard SQL. Amazon EC2 is a compute service that allows users to launch virtual machines, but does not provide any database management features. Amazon DynamoDB is a NoSQL database service that is not compatible with PostgreSQL, but offers fast and consistent performance at any scale. Reference: Hosted PostgreSQL - Amazon RDS for PostgreSQL - AWS, Amazon RDS for PostgreSQL - Amazon Relational Database Service, AWS PostgreSQL: Managed or Self-Managed? - NetApp, AWS Announces Amazon Aurora Supports PostgreSQL 12 - InfoQ, Amazon Aurora vs PostgreSQL | What are the differences? - StackShare

Question: 538

Which AWS service is always provided at no charge?

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

Answer: C

Explanation:
AWS Identity and Access Management (IAM) is a web service that helps you securely control access to AWS resources. You can use IAM to create and manage AWS users and groups, and use permissions to allow and deny their access to AWS resources. IAM is always provided at nocharge12. Reference: 1: AWS Identity and Access Management (IAM) - Amazon Web Services (AWS), 2: Which aws service is always provided at no charge? - Brainly.in

Question: 539

Which characteristic of the AWS Cloud helps users eliminate underutilized CPU capacity'?

A. Agility
B. Elasticity
C. Reliability
D. Durability

Answer: B

Explanation:

Elasticity is a characteristic of the AWS Cloud that helps users eliminate underutilized CPU capacity. Elasticity refers to the ability to dynamically provision and de-provision computing resources as per demand, ensuring that the application or service always has the required resources to operate efficiently. Elasticity helps users optimize performance and costs, as theyonly pay for the resources they use and avoid wasting resources when the demand is low345. Reference: 3: Which characteristic of the aws cloud helps users eliminate ..., *4: AWS* Elastic Load Balancing and Application Load

Balancer, 5: Which characteristic of the AWS Cloud helps users eliminate ...

Question: 540

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

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

Answer: A

Explanation:

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

Question: 541

A company wants to run its workload on Amazon EC2 instances for more than 1 year. This workload will run continuously.
Which option offers a discounted hourly rate compared to the hourly rate of On-Demand Instances?

A. AWS Graviton processor
B. Dedicated Hosts
C. EC2 Instance Savings Plans
D. Amazon EC2 Auto Scaling instances

Answer: C

Explanation:
EC2 Instance Savings Plans are a flexible pricing model that offer discounted hourly rates on Amazon EC2 instance usage for a 1 or 3 year term. EC2 Instance Savings Plans provide savings up to 72% off On-Demand rates, in exchange for a commitment to a specific instance family in a chosen AWS Region (for example, M5 in Virginia). These plans automatically apply to usage regardless of size (for example, m5.xlarge, m5.2xlarge, etc.), OS (for example, Windows, Linux, etc.), and tenancy (Host, Dedicated, Default) within the specified family in a Region. With an EC2Instance Savings Plan, you can change your instance size within the instance family (for example, from c5.xlarge to c5.2xlarge) or the operating system (for example, from Windows to Linux), or move from Dedicated tenancy to Default and continue to receive the discounted rate provided by your EC2 Instance Savings
Plan4567. Reference: 4: Compute Savings Plans - Amazon Web Services, 5: What are Savings Plans? - Savings Plans, 6: How To Cut Your AWS Bill With Savings Plans (and avoid some common 7: AWS Savings Plans vs Reserved Instances - GorillaStack

Question: 542

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

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

Answer: D

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

Question: 543

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

A. Amazon Aurora
B. Amazon FSx

C. Amazon DynamoDB
D. Amazon Neptune

Answer: D

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

Question: 544

What is a benefit of using AWS serverless computing?

A. Application deployment and management are not required
B. Application security will be fully managed by AWS
C. Monitoring and logging are not needed
D. Management of infrastructure is ofloaded to AWS

Answer: A

Explanation:
AWS serverless computing is a way of building and running applications without thinking about servers. AWS manages the infrastructure for you, so you don't have to provision, scale, patch, or monitor servers. You only pay for the compute time you consume, and you can focus on your application logic instead of managing servers12. Reference: Serverless Computing - Amazon Web Services, AWS Serverless Computing, Benefits, Architecture and Use-cases - XenonStack

Question: 545

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

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

Answer: A

Explanation:

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

Question: 546

A company is migrating its applications from on-premises to the AWS Cloud. The company wants to ensure that the applications are assigned only the minimum permissions that are needed to perform all operations.
Which AWS service will meet these requirements'?

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

Answer: A

Explanation:
AWS Identity and Access Management (IAM) is a service that helps you securely control access to AWS resources for your users. You use IAM to control who can use your AWS resources (authentication) and what resources they can use and in what ways (authorization). IAM also enables you to follow the principle of least privilege, which means granting only the permissions that are necessary to perform a task1.
Reference: AWS Identity and Access Management (IAM) - AWS Documentation

Question: 547

A company wants to allow users to authenticate and authorize multiple AWS accounts by using a single set of credentials.
Which AWS service or resource will meet this requirement?

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

Answer: C

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

Documentation

Question: 548

A company is using Amazon DynamoDB.
Which task is the company's responsibility, according to the AWS shared responsibility model?

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

Answer: C

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

Question: 549

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

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

Answer: B

Explanation:
Amazon S3 is a service that provides durable storage for static content and infinitely scalable data storage infrastructure at the lowest cost. Amazon S3 is an object storage service that allows you to store and retrieve any amount of data from anywhere on the internet. Amazon S3 offers industryleading scalability, availability, and performance, as well as 99.999999999% (11 9s) of durability and multi-AZ resilience. Amazon S3 also provides various storage classes that offer different levels of performance and cost optimization, such as S3 Standard, S3 Intelligent-Tiering, S3 Standard-

Infrequent Access (S3 Standard-IA), S3 One Zone-Infrequent Access (S3One Zone-IA), and S3 Glacier456. Amazon S3 is ideal for storing static content, such as images, videos, documents, and web pages, as well as building data lakes, backup and archive solutions, big data analytics, and machine learning applications456. Reference: 4: Cloud Storage on AWS, 5: Object Storage - Amazon Simple Storage Service (S3) - AWS, 6: Amazon S3 Documentation

Question: 550

A company wants to run a NoSQL database on Amazon EC2 instances. Which task is the responsibility of AWS in this scenario"?

A. Update the guest operating system of the EC2 instances
B. Maintain high availability at the database layer
C. Patch the physical infrastructure that hosts the EC2 instances
D. Configure the security group firewall

Answer: C

Explanation:
When you run a NoSQL database on Amazon EC2 instances, you are responsible for managing the database layer and the guest operating system of the instances. This means that you need to perform tasks such as updating the operating system, maintaining high availability, and configuring the security group firewall. AWS is responsible for managing the physical infrastructure that hosts the EC2 instances. This means that AWS ensures that the hardware and firmware of the servers, routers, switches, and other devices are updated and secure. AWS also handles the power, cooling, networking, and security of the data centers12. Reference: CLF-C02: Which task is responsibility of AWS to run NoSQL database on Best Practices for Hosting NoSQL Databases on Amazon EC2

Question: 551

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

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

Answer: A

Explanation:
AWS CloudTrail is a service that provides a record of actions taken by a user, role, or an AWS service in your AWS account. CloudTrail captures all API calls for AWS services as events, including calls from the AWS Management Console, AWS SDKs, command line tools, and higher-level AWS services. You can use CloudTrail to monitor, audit, and troubleshoot your AWS accountactivity34. AWS Trusted Advisor is a service that provides best practices recommendations for cost optimization, performance, security, and fault tolerance in your AWS account5. Amazon Inspector is a service that helps you improve the security and compliance of your applications deployed on AWS by automatically assessing them for vulnerabilities and deviations from best practices6. AWS X-Ray is a service that helps you analyze and debug your applications by collecting data about the requests that your application serves, and providing tools to view, filter, and gain insights into that data7. Reference: Logging AWS Audit Manager API calls with CloudTrail, Logging AWS Account Management API calls using AWS CloudTrail, Review API calls in your AWS account using CloudTrail, Monitor the usage of AWS API calls using Amazon CloudWatch, Which service enables customers to audit API calls in their AWS .

Question: 552

A company needs a bridge between technology and business to help evolve to a culture of continuous growth and learning.
Which perspective in the AWS Cloud Adoption Framework (AWS CAF) serves as this bridge?

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

Answer: A

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

Question: 553

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

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

Answer: A

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

Question: 554

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

A. Amazon Macie
B. Amazon Detective
C. Amazon GuardDuty
D. AWS I AM Access Analyzer

Answer: A

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

Question: 555

Which AWS service or resource provides answers to the most frequently askedsecurity-related questions that AWS receives from its users'?

A. AWS Artifact
B. Amazon Connect
C. AWS Chatbot
D. AWS Knowledge Center

Answer: A

Explanation:
AWS Artifact is your go-to, central resource for compliance-related information that matters to you. It provides on-demand access to AWS's security and compliance reports and select online agreements. Reports available in AWS Artifact include our Service Organization Control (SOC) reports, Payment Card Industry (PCI) attestation of compliance, and certifications from accreditation bodies across geographies and compliance verticals that validate the implementation and operating effectiveness of AWS security controls. Agreements available in AWS Artifact include the Business Associate Addendum (BAA) and the Nondisclosure Agreement (NDA). AWS Artifact helps you answer the most frequently asked security and compliance questions that AWS receives from its users. Reference: Compliance FAQ, Compliance Solutions Guide

Question: 556

A company wants to move its on-premises databases to managed cloud database services by using a simplified migration process. Which AWS service or tool can help the company meet this requirement?

A. AWS Storage Gateway
B. AWS Application Migration Service
C. AWS DataSync
D. AWS Database Migration Service (AWS DMS)

Answer: D

Explanation:

AWS Database Migration Service (AWS DMS) is a cloud service that makes it possible to migrate relational databases, data warehouses, NoSQL databases, and other types of data stores. You can use AWS DMS to migrate your data into the AWS Cloud or between combinations of cloud and onpremises setups. With AWS DMS, you can discover your source data stores, convert your source schemas, and migrate your data. AWS DMS supports migration between 20-plus database and analytics engines, such as Oracle to Amazon Aurora MySQL-Compatible Edition, MySQL to Amazon Relational Database (RDS) for MySQL, Microsoft SQL Server to Amazon Aurora PostgreSQL-Compatible Edition, MongoDB to Amazon DocumentDB (with MongoDB compatibility), Oracle to Amazon Redshift, and Amazon Simple Storage Service (S3). You can perform one-time migrations or replicate ongoing changes to keep sources and targets in sync. AWS DMS automatically manages the deployment, management, and monitoring of all hardware and software needed for your migration. AWS DMS is a highly resilient, secure cloud service that provides database discovery, schema conversion, data migration, and ongoing replication to and from a wide range of databases and analytics systems12. Reference:
Database Migration - AWS Database Migration Service - AWS What is AWS Database Migration Service? - AWS Database Migration Service

Question: 557

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

Explanation:

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.
Reference: Instance purchasing options, Amazon EC2 Pricing, 4 Ways to Purchase Amazon EC2 Instances

Question: 558

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 a fully managed relational database service that supports several database engines, including PostgreSQL. Amazon RDS can run a managed PostgreSQL database that provides online transaction processing (OLTP), which is a type of database workload that handles frequent read and write operations on small amounts of data. Amazon RDS for PostgreSQL offers high performance, availability, scalability, security, and compatibility with the PostgreSQL community edition. Amazon RDS also provides automated backups, point-in-time recovery, encryption, monitoring, and maintenance for PostgreSQL databases. Reference:
Hosted PostgreSQL - Amazon RDS for PostgreSQL
OLTP Database, MySQL And PostgreSQL ManagedDatabase - Amazon Aurora
PostgreSQL options on AWS: Self- managed, managed, and serverless

Question: 559

Which complimentary AWS service or tool creates data-driven business cases for cloud planning?

A. Migration Evaluator
B. AWS Billing Conductor
C. AWS Billing Console
D. Amazon Forecast

Answer: A

Explanation:
Migration Evaluator is a cloud-based service that provides organizations with a comprehensive assessment of their current IT environment and estimates the cost savings and performance improvements that can be achieved by migrating to AWS. Migration Evaluator helps users build a data-driven business case for AWS by discovering over-provisioned on-premises instances,providing recommendations for cost-effective AWS alternatives, and analyzing existing licenses and cost comparisons of Bring Your Own License (BYOL) and License Included (LI) options

Question: 560

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:
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 is responsible for protecting the infrastructure that runs AWS services, such as DynamoDB,

while the customer is responsible for properly configuring the security of the provided service. For abstracted services, such as DynamoDB, the customer is primarily responsible for managing their data, classifying their assets, and using IAM tools to apply the appropriatepermissions12. Therefore, the customer is responsible for controlling the access to DynamoDB tables, such as by creating IAM policies, roles, and users, and using encryption and authentication mechanisms3. Reference:

Shared Responsibility Model - Amazon Web Services (AWS) Security and compliance in Amazon DynamoDB - Amazon DynamoDB What is Shared Responsibility Model? - Check Point Software

Question: 561

Using AWS Identity and Access Management (IAM) to grant access only to the resources needed to perform a task is a concept known as:

A. restricted access.
B. as-needed access.
C. least privilege access.
D. token access.

Answer: C

Explanation:

The concept of granting access only to the resources needed to perform a task is known as least privilege access. This is a security best practice in IAM that helps to reduce the risk of unauthorized or malicious actions. By applying least privilege access, you can limit the permissions of your IAM users, groups, and roles to the minimum required for their specific tasks. You can also use conditions, permissions boundaries, and IAM Access Analyzer to further restrict and verify access. Reference: Security best practices in IAM, Policies and permissions in IAM, Use IAM policies to grant the least privileges required to access Amazon RDS resources, How to Design a Least Privilege Architecture in AWS, 12 Azure & AWS IAM Security Best Practices

Question: 562

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

Explanation:

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. AWSre:Post is a great way to learn from the AWS community, share

your knowledge, and get inspired. Reference: AWS re:Post
Join the Conversation

Question: 563

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 a serverless data integration service that makes it easy to discover, prepare, move, and integrate data from multiple sources for analytics, machine learning, and application development. You can use various data integration engines, such as ETL, ELT, batch, and streaming, and manage your data in a centralized data catalog. AWS Glue is designed specifically for extract, transform, and load (ETL) data, whereas the other options are not.

Question: 564

Which AWS service or feature can be used to create a private connection between an on-premises workload and an AWS Cloud workload?

A. Amazon Route 53
B. Amazon Macie
C. AWS Direct Connect
D. AWS PrivaleLink

Answer: C

Explanation:
AWS Direct Connect is a service that establishes a dedicated network connection between your onpremises network and one or more AWS Regions. AWS Direct Connect can be used to create a private connection between an on-premises workload and an AWS Cloud workload, bypassing the public internet and reducing network costs, latency, and bandwidth issues. AWS Direct Connect can also provide increased security and reliability for your hybrid cloud applications and data transfers. Reference:
AWS Direct Connect
What is AWS Direct Connect?
AWS Direct Connect User Guide

Question: 565

What is the best resource for a user to find compliance-related information and reports about AWS?

A. AWS Artifact
B. AWS Marketplace
C. Amazon Inspector
D. Increase operational costs across data centers.

Answer: A

Explanation:
AWS Artifact is a self-service portal that provides on-demand access to AWS security and compliance reports and select online agreements. Users can download reports such as AWS ISO certifications, PCI reports, SOC reports, and GDPR DPA, and review and accept agreements such as BAA and NDA. AWS Artifact helps users to understand and meet compliance requirements for various standards and regulations that apply to AWS services and infrastructure. AWS Artifact is the best resource for a user to find compliance-related information and reports about AWS, whereas the other options are not

Question: 566

Which option is a benefit of the economies of scale based on the advantages of cloud computing?

A. The ability to trade variable expense for fixed expense
B. Increased speed and agility
C. Lower variable costs over fixed costs
D. Increased operational costs across data centers

Answer: B

Explanation:
Economies of scale are the cost advantages that result from increasing the scale of production or operation. In cloud computing, economies of scale are achieved by pooling resources and sharing them among multiple users, which reduces the unit cost of computing and storage. One of the benefits of economies of scale in cloud computing is increased speed and agility, which means the ability to deploy applications faster and respond to changing business needs more quickly. Cloud computing allows users to access computing resources on demand, without having to invest in expensive infrastructure or wait for lengthy provisioning processes. This enables users to scale up or down as needed, experiment with new ideas, and deliver value to customers faster123. Reference: Economics of Cloud Computing - GeeksforGeeks What is Cloud Economics? | VMware Glossary
ECONOMIES OF SCALE WITH CLOUD COMPUTING & SERVICES PRACTICE - IDC-Online

Question: 567

Which AWS service or feature improves network performance by sending traffic through the AWS worldwide network infrastructure?

A. Route table
B. AWS Transit Gateway

C. AWS Global Accelerator

D. Amazon VPC

Answer: C

Explanation:

AWS Global Accelerator is a service that improves network performance by sending traffic through the AWS worldwide network infrastructure. It uses the AWS global network to direct TCP or UDP traffic to a healthy application endpoint in the closest AWS Region to the client. This provides improvements in terms of latency, throughput, and jitter. Global Accelerator also introduces features such as TCP termination at the edge, jumbo frame support, and large receive side window and TCP buffers to optimize data transfer12. Route table, AWS Transit Gateway, and Amazon VPC are not services or features that improve network performance by sending traffic through the AWS worldwide network infrastructure. Route table is a resource that defines how traffic is routed within a VPC3. AWS Transit Gateway is a service that enables you to connect your VPCs and on-premises networks to a single gateway4. Amazon VPC is a service that lets you provision a logically isolated section of the AWS Cloud where you can launch AWS resources in a virtual network that you define5.
Reference: Achieve up to 60% better performance for internet traffic with AWS Global Accelerator, Improving Performance on AWS and Hybrid Networks, Route tables, AWS Transit Gateway, Amazon Virtual Private Cloud (VPC)

Question: 568

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

A. Establish the global infrastructure.

B. Perform client-side data encryption.

C. Configure 1AM credentials.

D. Secure edge locations.

E. Patch Amazon RDS DB instances.

Answer: B,C

Explanation:

According to the AWS shared responsibility model, AWS is responsible for the security of the cloud, while the customer is responsible for the security in the cloud. This means that AWS is responsible for protecting the infrastructure that runs all of the services offered in the AWS Cloud, such as the global network, the hardware, the software, and the facilities. The customer is responsible for properly configuring the security of the provided service, such as the guest operating system, the application software, the data, and the network traffic. For abstracted services, such as Amazon RDS, AWS operates the infrastructure layer, the operating system, and the database software, while the customer is responsible for managing their data, classifying their assets, and using IAM tools to apply the appropriate permissions12.
Therefore, the tasks that are the customer's responsibility are:
Perform client-side data encryption: The customer is responsible for encrypting their data before sending it to AWS, and decrypting it after receiving it from AWS. This ensures that the data is protected in transit and at rest. AWS provides various encryption options, such as AWS Key Management Service (AWS KMS), AWS CloudHSM, and AWS Certificate Manager (ACM)3.
Configure IAM credentials: The customer is responsible for creating and managing IAM users, groups, roles, and policies that control the access to AWS resources and services. IAM credentials include user names, passwords, access keys, and permissions4.
The tasks that are not the customer's responsibility are:

Establish the global infrastructure: AWS is responsible for building and maintaining the global network of regions, availability zones, and edge locations that provide low latency, high availability, and fault tolerance for the AWS Cloud5.
Secure edge locations: AWS is responsible for protecting the physical security of the edge locations, which are sites that deliver cached content to end users with improved performance6.
Patch Amazon RDS DB instances: AWS is responsible for applying patches and updates to the operating system and the database software of the Amazon RDS DB instances, which are managed relational database service for MySQL, PostgreSQL, Oracle, SQL Server, and Amazon Aurora. Reference:
Shared Responsibility Model - Amazon Web Services (AWS)
Shared responsibility model - Amazon Web Services: Risk and Compliance
Encryption - Amazon Web Services (AWS)
What Is IAM? - AWS Identity and Access Management
Global Infrastructure - Amazon Web Services (AWS)
Amazon CloudFront Features - Content Delivery Network (CDN)
[What Is Amazon Relational Database Service (Amazon RDS)? - Amazon Relational Database Service]

Question: 569

Which Amazon EC2 pricing model is the MOST cost efficient for an uninterruptible workload that runs once a year for 24 hours?

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

Answer: A

Explanation:

On-Demand Instances are the most cost-efficient pricing model for an uninterruptible workload that runs once a year for 24 hours. On-Demand Instances let you pay for compute capacity by the hour or second, depending on which instances you run. No long-term commitments or up-front payments are required. You can increase or decrease your compute capacity to meet the demands of your application and only pay the specified hourly rates for the instance you use1. This model is suitable for developing/testing applications with short-term or unpredictable workloads2. The other pricing models are not cost-efficient for this use case. Reserved Instances and Savings Plans require a commitment to a consistent amount of usage, in USD per hour, for a term of 1 or 3 years. They provide significant discounts compared to On-Demand Instances, but they are not flexible or scalable for workloads that run only once a year12. Spot Instances are the cheapest option, but they are not suitable for uninterruptible workloads, as they can be reclaimed by AWS at any time. They are recommended for applications that have flexible start and end times, or that are only feasible at very low compute prices12. 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 hardware12. Reference: Amazon EC2 - Secure and resizable compute capacity - AWS, Amazon EC2 - How AWS Pricing Works

Question: 570

Which AWS Cloud benefit gives a company the ability to quickly deploy cloud resources to access compute, storage, and database infrastructures in a matter of minutes?

A. Elasticity
B. Cost savings
C. Agility
D. Reliability

Answer: C

Explanation:
Agility is the AWS Cloud benefit that gives a company the ability to quickly deploy cloud resources to access compute, storage, and database infrastructures in a matter of minutes. Agility means that you can reduce the time to make IT resources available to your developers from weeks to just minutes, resulting in a dramatic increase in innovation and responsiveness1. AWS provides a range of services and tools that enable you to launch, scale, and manage your cloud applications with ease and speed, such as AWS CloudFormation, AWS Elastic Beanstalk, AWS CodeDeploy, and AWS Quick Starts2345. Reference:
Six advantages of cloud computing - Overview of Amazon Web Services
[AWS CloudFormation]
[AWS Elastic Beanstalk]
[AWS CodeDeploy]
AWS Quick Starts

Question: 571

Which pricing model will interrupt a running Amazon EC2 instance if capacity becomes temporarily unavailable?

A. On-Demand Instances
B. Standard Reserved Instances
C. Spot Instances
D. Convertible Reserved Instances

Answer: C

Explanation:
Spot Instances are a type of EC2 instance that let you bid on unused compute capacity, which AWS offers at a discount of up to 90% compared to On-Demand prices1. Spot Instances are suitable for fault-tolerant, stateless, or flexible applications that can handle interruptions2. Spot Instances can be interrupted with a two-minute warning when EC2 needs the capacity back3. The other options are not pricing models that will interrupt a running EC2 instance if capacity becomes temporarily unavailable

Question: 572

A company wants to provision and manage its AWS infrastructure by using the common programming languages TypeScript, Python, Java, and .NET. Which h AWS service will meet this requirement?

A. AWS CodeBuild
B. AWS CloudFormation
C. AWSCLI

D. AWS Cloud Development Kit (AWS CDK)

Answer: D

Explanation:

AWS Cloud Development Kit (AWS CDK) is an open source software development framework that allows you to model and provision your cloud infrastructure using familiar programming languages such as TypeScript, Python, Java, and .NET. AWS CDK enables you to use the expressive power of your favorite language to define your cloud resources, such as compute, storage, network, and application services. AWS CDK also provides a library of high-level constructs that represent AWS services and best practices. AWS CDK uses AWS CloudFormation in the background to deploy your resources in a safe and repeatable manner12. Reference:
AWS Cloud Development Kit (CDK) - TypeScript and Python are Now Generally Available
AWS Cloud Development Kit (AWS CDK) - Introduction to DevOps on AWS

Question: 573

A company has a set of ecommerce applications. The applications need to be able to send messages to each other. Which AWS service meets this requirement?

A. AWS Auto Scaling
B. Elastic Load Balancing
C. Amazon Simple Queue Service (Amazon SOS)
D. Amazon Kinesis Data Streams

Answer: C

Explanation:
Amazon Simple Queue Service (Amazon SQS) is a fully managed message queuing service that lets you send, store, and receive messages between software components at any volume, without losing messages or requiring other services to be available1. Amazon SQS is designed to provide a simpleand reliable way for customers to decouple and connect components (microservices) together using queues2. Queues are an important mechanism for providing fault tolerance and scalability in distributed systems, and help decouple different parts of your application3. The other options are not AWS services that are used specifically for sending messages between applications

Question: 574

A company wants to establish a schedule for rotating database user credentials.
Which AWS service will support this requirement with the LEAST amount of operational overhead?

A. AWS Systems Manager
B. AWS Secrets Manager
C. AWS License Manager
D. AWS Managed Services

Answer: B

Explanation:

AWS Secrets Manager is a service that helps you protect access to your applications, services, and IT resources. This service enables you to easily rotate, manage, and retrieve database credentials, API keys, and other secrets throughout their lifecycle. Users and applications retrieve secrets with a call to Secrets Manager APIs, eliminating the need to hardcode sensitive information in plain text. Secrets Manager offers secret rotation with built-in integration for Amazon RDS, Amazon Redshift, Amazon DocumentDB, and other AWS services1. You can also extend Secrets Manager to rotate other types of secrets, such as credentials for Oracle, SQL Server, or MongoDB databases, by using custom AWS Lambda functions2. Secrets Manager enables you to control access to secrets using fine-grained permissions and audit secret rotation centrally for resources in the AWS Cloud, third-party services, and on-premises3. Therefore, AWS Secrets Manager supports the requirement of rotating database user credentials with the least amount of operational overhead, compared to the other options.

Reference:

What Is AWS Secrets Manager? - AWS Secrets Manager

Rotating Your AWS Secrets Manager Secrets - AWS Secrets Manager

AWS Secrets Manager Features - AWS Secrets Manager

Question: 575

Which AWS service is used to provide encryption for Amazon EBS?

A. AWS Certificate Manager
B. AWS Systems Manager
C. AWS KMS
D. AWS Config

Answer: C

Explanation:

AWS KMS is the service that is used to provide encryption for Amazon EBS. AWS KMS is a managed service that enables you to easily create and control the encryption keys used to encrypt your data.Amazon EBS uses AWS KMS to encrypt and decrypt your EBS volumes and snapshots. You can choose to use either the default AWS managed CMK or your own customer managed CMK for encryption. AWS KMS also provides features such as key rotation, audit logging, and access control policies to help you manage your encryption keys and protect your data12. The other services are not used to provide encryption for Amazon EBS. AWS Certificate Manager is a service that lets you provision, manage, and deploy public and private SSL/TLScertificates for use with AWS services and your internal connected resources3. AWS Systems Manager is a service that provides a unified user interface to view and manage your AWS resources, automate common operational tasks, and apply compliance policies4. AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. Reference: Amazon EBS encryption, AWS Key Management Service, AWS Certificate Manager, AWS Systems Manager, [AWS Config]

Question: 576

A company operates a petabyte-scale data warehouse to analyze its dat

a. The company wants a solution that will not require manual hardware and software management. Which AWS service will meet these requirements?

A. Amazon DocumentDB (with MongoDB compatibility)

B. Amazon Redshift
C. Amazon Neptune
D. Amazon ElastiCache

Answer: B

Explanation:
Amazon Redshift is a fast, fully managed, petabyte-scale data warehouse service that makes it simple and cost-effective to analyze all your data using your existing business intelligence tools. You can start small with no commitments, and scale to petabytes for less than a tenth of the cost of traditional solutions. Amazon Redshift does not require manual hardware and software management, as AWS handles all the tasks such as provisioning, patching, backup, recovery, failure detection, and repair12. Amazon Redshift also offers serverless capabilities, which allow you to access and analyze data without any configurations or capacity planning. Amazon Redshift automatically scales the data warehouse capacity to deliver fast performance for even the most demanding and unpredictable workloads3. Therefore, Amazon Redshift meets the requirements of the company, compared to the other options.
The other options are not suitable for the company's requirements, because:
Amazon DocumentDB (with MongoDB compatibility) is a fast, scalable, highly available, and fully managed document database service that supports MongoDB workloads. It is not designed for petabyte-scale data warehousing or analytics4.
Amazon Neptune is a fast, reliable, and fully managed graph database service that makes it easy to build and run applications that work with highly connected datasets. It is not designed for petabytescale data warehousing or analytics5.
Amazon ElastiCache is a fully managed in-memory data store and cache service that supports Redis and Memcached. It is not designed for petabyte-scale data warehousing or analytics.

What is Amazon Redshift? - Amazon Redshift Amazon Redshift Features - Amazon Redshift Amazon Redshift Serverless - Amazon Redshift
What Is Amazon DocumentDB (with MongoDB compatibility)? - Amazon DocumentDB (with MongoDB compatibility)
What Is Amazon Neptune? - Amazon Neptune
[What Is Amazon ElastiCache for Redis? - Amazon ElastiCache for Redis]

Question: 577

A company needs to perform data processing once a week that typically takes about 5 hours to complete. Which AWS service should the company use for this workload?

A. AWS Lambda
B. Amazon EC2
C. AWS CodeDeploy
D. AWS Wavelength

Answer: B

Explanation:
Amazon EC2 is the most suitable AWS service for this workload. Amazon EC2 provides secure, resizable compute capacity in the cloud. You can launch virtual servers, called instances, and configure them according to your needs. You can choose from different instance types, sizes, and families, and pay only for the resources you use. Amazon EC2 also offers features such as auto scaling, load balancing, security groups, and placement groups to optimize your performance, availability, and security1. Amazon EC2 is ideal for workloads that require

consistent and reliable compute power, such as data processing, web hosting, gaming, and high-performance computing2. The other services are not suitable for this workload. AWS Lambda is a serverless compute service that lets you run code without provisioning or managing servers. You pay only for the compute time you consume. Lambda is best for short-lived, stateless, and event-driven workloads that can becompleted in under 15 minutes3. AWS CodeDeploy is a deployment service that automates application deployments to Amazon EC2 instances, on-premises instances, serverless Lambda functions, or Amazon ECS services. CodeDeploy is not a compute service, but a tool to help you update your applications with minimal downtime4. AWS Wavelength is a service that delivers ultralow latency applications for 5G devices. Wavelength embeds AWS compute and storage services at the edge of telecommunications providers' 5G networks. Wavelength is designed for mobile edge computing, such as interactive gaming, video streaming, and augmentedreality. Reference: Amazon EC2, Amazon EC2 Use Cases, AWS Lambda, AWS CodeDeploy, [AWS Wavelength]

Question: 578

A company wants to provide managed Windows virtual desktops and applications to its remote employees over secure network connections. Which AWS services can the company use to meet these requirements? (Select TWO.)

A. Amazon Connect
B. Amazon AppStream 2.0
C. Amazon Workspaces
D. AWS Site-to-Site VPN
E. Amazon Elastic Container Service (Amazon ECS)

Answer: B,C

Explanation:

Amazon AppStream 2.0 and Amazon WorkSpaces are AWS services that can be used to provide managed Windows virtual desktops and applications to remote employees over secure network connections. Amazon AppStream 2.0 is a fully managed application streaming service that allows users to access Windows desktop applications from any device, without installing or managing any software. Amazon AppStream 2.0 delivers applications over an encrypted connection and isolates them from the underlying infrastructure, ensuring security and compliance1. Amazon WorkSpaces is a fully managed desktop virtualization service that allows users to access Windows or Linux desktops from any device, with a consistent user experience. Amazon WorkSpaces provides persistent, cloudbased virtual desktops that can be customized and scaled according to the user's needs. Amazon WorkSpaces also offers encryption, backup, and monitoring features to ensure security and reliability2. Reference:

Amazon AppStream 2.0

Amazon WorkSpaces

Question: 579

Which AWS Cloud service can send alerts to customers if custom spending thresholds are exceeded?

A. AWS Budgets
B. AWS Cost Explorer
C. AWS Cost Allocation Tags
D. AWS Organizations

Answer: A

Explanation:

AWS Budgets is a service that allows you to set custom budgets for your AWS costs and usage, and receive alerts via email or Amazon SNS notifications if you exceed or are forecasted to exceed your budgeted amount1. You can create budgets based on different dimensions, such as service, linked account, tag, or purchase option, and define various types of alerts, such as actual, forecasted, or RI utilization alerts2. You can also configure custom actions to automatically execute remediation tasks or workflows when a budget threshold is breached3. AWS Budgets is the only service among the options that can send alerts to customers if custom spending thresholds are exceeded. The other options are not AWS services that provide this functionality.

Question: 580

Which AWS service or feature provides log information of the inbound and outbound traffic on network interfaces in a VPC?

A. Amazon CloudWatch Logs
B. AWS CloudTrail
C. VPC Flow Logs
D. AWS Identity and Access Management (IAM)

Answer: C

Explanation:

VPC Flow Logs is a feature that enables you to capture information about the IP traffic going to and from network interfaces in your VPC. Flow log data can be published to the following locations: Amazon CloudWatch Logs, Amazon S3, or Amazon Kinesis Data Firehose. You can use VPC Flow Logs to monitor network traffic, diagnose security issues, troubleshoot connectivity problems, and perform network forensics1. Reference:

Logging IP traffic using VPC Flow Logs - Amazon Virtual Private Cloud

Question: 581

Which tool should a developer use lo integrate AWS service features directly into an application?

A. AWS Software Development Kit
B. AWS CodeDeploy
C. AWS Lambda
D. AWS Batch

Answer: A

Explanation:

AWS Software Development Kit (SDK) is a set of platform-specific tools for developers that let them integrate AWS service features directly into their applications. AWS SDKs provide libraries, code samples, documentation, and other resources to help developers write code that interacts with AWS APIs. AWS SDKs support various programming languages, such as Java, Python, Ruby, .NET, Node.js, Go, and more. AWS SDKs make it easier for developers to access AWS services, such as Amazon S3, Amazon EC2, Amazon DynamoDB, AWS Lambda, and more,

from their applications. AWS SDKs also handle tasks such as authentication, error handling, retries, and data serialization, so developers can focus on their application logic.

Question: 582

Which AWS service enables companies to deploy an application dose to end users?

A. Amazon CloudFront
B. AWS Auto Scaling
C. AWS AppSync
D. Amazon Route S3

Answer: A

Explanation:

Amazon CloudFront is a fast content delivery network (CDN) service that securely delivers data, videos, applications, and APIs to customers globally with low latency, high transfer speeds, all within a developer-friendly environment. CloudFront enables companies to deploy an application close to end users by caching the application's content at edge locations that are geographically closer to the users. This reduces the network latency and improves the user experience. CloudFront also integrates with other AWS services, such as Amazon S3, Amazon EC2, AWS Lambda, AWS Shield, and AWS WAF, to provide a secure and scalable solution for delivering applications12. Reference:
What Is Amazon CloudFront? - Amazon CloudFront
Amazon CloudFront Features - Amazon CloudFront

Question: 583

A company needs to evaluate its AWS environment and provide best practice recommendations in five categories: cost, performance, service limits, fault tolerance, and security. Which AWS service can the company use to meet these requirements?

A. AWS Shield
B. AWS WAF
C. AWS Trusted Advisor
D. AWS Service Catalog

Answer: C

Explanation:
AWS Trusted Advisor is the service that can meet these requirements. AWS Trusted Advisor is a service that helps you optimize your AWS environment by providing recommendations based on AWS best practices. Trusted Advisor continuously evaluates your AWS resources and services across five categories: cost optimization, performance, service limits, fault tolerance, and security. You can view the recommendations on the Trusted Advisor console or access them programmatically using the Trusted Advisor API. You can also set up notifications and alerts for any changes in the status of your checks. Trusted Advisor can help you improve your AWS environment by reducing costs, enhancing performance, increasing security, and ensuring reliability12. The other services are not designed to provide best practice recommendations in five categories. AWS Shield is a service that protects your AWS resources from distributed denial-of-service (DDoS) attacks. AWS WAF is a service

that helps you protect your web applications from common web exploits. AWS Service Catalog is a service that enables you to create and manage catalogs of IT services that are approved for use on AWS34 . Reference: AWS Trusted Advisor, Achieve operational excellence with AWS Trusted Advisor, AWS Shield, AWS WAF, [AWS Service Catalog]

Question: 584

Which AWS service or feature allows users to create new AWS accounts, group multiple accounts to organize workflows, and apply policies to groups of accounts?

A. AWS Identity and Access Management (1AM)
B. AWS Trusted Advisor
C. AWS CloudFormation
D. AWS Organizations

Answer: D

Explanation:
AWS Organizations is the AWS service or feature that allows users to create new AWS accounts, group multiple accounts to organize workflows, and apply policies to groups of accounts. AWS Organizations enables users to centrally manage and govern their AWS environment across multiple accounts. Users can create organizational units (OUs) to group accounts based on their business needs, such as by function, project, or region. Users can also apply service control policies (SCPs) toOUs or individual accounts to define the permissions and restrictions for the AWS services and resources that they can access. AWS Organizations also offers features such as consolidated billing, account creation automation, and trusted access12. Reference:
AWS Organizations
What is AWS Organizations?

Question: 585

Which options are AWS Cloud Adoption Framework (AWS CAF) cloud transformation journey recommendations? (Select 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) is a tool that helps organizations plan and execute their cloud transformation journey. The AWS CAF defines four phases of the cloud transformation journey: Envision, Align, Launch, and Scale. Each phase has a specific purpose and outcome1: Envision: This phase helps you define your vision, goals, and expected outcomes for your cloud transformation. It also helps you identify and prioritize transformation opportunities across four domains: business, people, governance, and platform2. Align: This phase helps you identify capability gaps across six perspectives: business, people, governance, platform, security, and

operations. It also helps you create strategies for improving your cloud readiness, ensure stakeholder alignment, and facilitate relevant organizational change management activities3.

Launch: This phase helps you deliver pilot initiatives in production and demonstrate incremental business value. It also helps you learn from pilots and adjust your approach before scaling to full production4.

Scale: This phase helps you expand production pilots and business value to desired scale and ensure that the business benefits associated with your cloud investments are realized and sustained. The options A and B are the correct AWS CAF cloud transformation journey recommendations, as they are part of the four phases defined by the AWS CAF. The options C, D, and E are not AWS CAFcloud transformation journey recommendations, as they are not part of the four phases defined by the AWS CAF

Question: 586

Which responsibility belongs to AWS when a company hosts its databases on Amazon EC2 instances?

A. Database backups
B. Database software patches
C. Operating system patches
D. Operating system installations

Answer: C

Explanation:

When a company hosts its databases on Amazon EC2 instances, AWS and the customer share the responsibility for the security and management of the database environment. 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 is responsible for protecting the infrastructure that runs the EC2 instances, such as the hardware, software, networking, and facilities. The customer is responsible for properly configuring the security of the provided service, such as the guest operating system, the database software, the data, and the network traffic12. One of the tasks that belongs to AWS when a company hosts its databases on Amazon EC2 instances is operating system patches. AWS provides regular updates and patches to the operating system of the EC2 instances, which are applied automatically by default. The customer can also choose to manually apply the patches or schedule them for a specific time window3. Operating system patches are important for maintaining the security and performance of the EC2 instances and the databases running on them.

The other tasks that belong to AWS when a company hosts its databases on Amazon EC2 instances are:

Operating system installations: AWS provides a variety of operating system options for the EC2 instances, such as Linux, Windows, and Amazon Linux. The customer can choose the operating system that best suits their database needs and AWS will install it on the EC2 instances4. Server maintenance: AWS performs regular maintenance and repairs on the physical servers that host the EC2 instances, ensuring that they are in optimal condition and have adequate power, cooling, and network connectivity5.

Hardware lifecycle: AWS manages the lifecycle of the hardware that supports the EC2 instances, such as replacing faulty components, upgrading equipment, and decommissioning old servers.

The tasks that do not belong to AWS when a company hosts its databases on Amazon EC2 instances are:

Database backups: The customer is responsible for backing up their data and databases on the EC2 instances, using tools such as Amazon S3, Amazon EBS snapshots, or AWS Backup. Database backups are essential for data protection and recovery in case of failures or disasters.

Database software patches: The customer is responsible for applying patches and updates to the database software on the EC2 instances, such as MySQL, PostgreSQL, Oracle, or SQL Server. Database software patches are important for fixing bugs, improving features, and addressing security vulnerabilities.

Database software install: The customer is responsible for installing the database software on the EC2 instances, choosing the version and configuration that meets their requirements. AWS provides some preconfigured AMIs (Amazon Machine Images) that include common database software, or the customer can use their own custom AMIs.

Shared Responsibility Model - Amazon Web Services (AWS) Shared responsibility model - Amazon Web Services: Risk and Compliance
Patching Amazon EC2 instances - AWS Systems Manager Amazon EC2 FAQs - Amazon Web Services
Maintenance and Retirements - Amazon Elastic Compute Cloud
[Hardware Lifecycle - Amazon Web Services (AWS)]
[Backing Up Your Data - Amazon Web Services (AWS)] [Database Patching - Amazon Web Services (AWS)]
[Installing Database Software on Amazon EC2 Instances - Amazon Web Services (AWS)]

Question: 587

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

A. AWS CloudHSM
B. AWS CloudShell
C. Amazon Workspaces
D. AWS Cloud Map

Answer: B

Explanation:
AWS CloudShell is the service that provides command line access to AWS tools and resources directly from a web browser. AWS CloudShell is a browser-based shell that makes it easy to securely manage, explore, and interact with your AWS resources. It comes pre-authenticated with your console credentials and common development and administration tools are pre-installed, so no local installation or configuration is required. You can open AWS CloudShell from the AWS Management Console with a single click and start running commands and scripts using the AWS Command Line Interface (AWS CLI), Git, or SDKs. AWS CloudShell also provides persistent home directories with 1 GB of storage per AWS Region12. The other services do not provide command line access to AWS tools and resources directly from a web browser. AWS CloudHSM is a service that helps you meet corporate, contractual and regulatory compliance requirements for data security by using dedicated Hardware Security Module (HSM) appliances within the AWS Cloud3. Amazon WorkSpaces is aservice that provides a fully managed, secure Desktop-as-a-Service (DaaS) solution that runs on AWS4. AWS Cloud Map is a service that makes it easy for your applications to discover and connect to each other usinglogical names and attributes5. Reference: AWS CloudShell, AWS CloudShell - Command-Line Access to AWS Resources, AWS CloudHSM, Amazon WorkSpaces, AWS Cloud Map

Question: 588

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

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

Answer: D

Explanation:

AWS CloudFormation is a service that allows developers to model and provision their AWS infrastructure in a repeatable and declarative way, using code and templates. AWS CloudFormation enables developers to define the resources they need for their development and production environments, such as compute, storage, network, and application services, and automate their creation and configuration. AWS CloudFormation also provides features such as change sets, nested stacks, and rollback triggers to help developers manage and update their infrastructure safely and efficiently12. Reference:

AWS CloudFormation

What is AWS CloudFormation?

Question: 589

A company wants to migrate its applications to the AWS Cloud. The company plans to identity and prioritize any business transformation opportunities and evaluate its AWS Cloud readiness. Which AWS service or tool should the company use to meet these requirements?

A. AWS Cloud Adoption Framework (AWS CAF)
B. AWS Managed Services (AMS)
C. AWS Well-Architected Framework
D. AWS Migration Hub

Answer: A

Explanation:

AWS Cloud Adoption Framework (AWS CAF) is a set of best practices, tools, and guidance that helps organizations get started with cloud technologies. AWS CAF helps organizations identify and prioritize transformation opportunities, evaluate and improve their cloud readiness, and iteratively evolve their transformation roadmap. AWS CAF groups its capabilities in six perspectives: Business, People, Governance, Platform, Security, and Operations. Each perspective comprises a set of capabilities that functionally related stakeholders own or manage in the cloud transformation journey1

AWS Managed Services (AMS) is a service that operates AWS infrastructure on behalf of customers, providing a secure AWS Landing Zone, features that help meet various compliance program requirements, a proven enterprise operating model, on-going cost optimization, and day-to-day infrastructure management. AMS does not help customers identify and prioritize business transformation opportunities or evaluate their cloud readiness2

AWS Well-Architected Framework is a set of six pillars and lenses that help cloud architects design and run workloads in the cloud. It provides a consistent approach for customers and AWS Partners to evaluate and implement designs that scale with their needs. AWS Well-Architected Framework helps customers understand the pros and cons of decisions they make while building systems on AWS, but it does not help them identify and prioritize business transformation opportunities3

AWS Migration Hub is a tool that lets customers discover, plan, and track their existing servers and applications for migration to AWS. It offers journey templates, cross-team collaboration, application and server discovery, strategy recommendations, orchestration and simple dashboard. AWS Migration Hub simplifies the migration and modernization process, but it does not help customers identify and prioritize business transformation opportunities or evaluate their cloud readiness4 Reference: 1: AWS Cloud Adoption Framework 2: Cloud Management Services - AWS Managed Services - AWS 3: AWS Well-Architected - Build secure, efficient cloud applications 4: Cloud Inventory Management - AWS Migration Hub - AWS

Question: 590

A social media company wants to protect its web application from common web exploits such as SQL injections and cross-site scripting. Which AWS service will meet these requirements?

A. Amazon Inspector
B. AWS WAF
C. Amazon GuardDuty
D. Amazon CloudWatch

Answer: B

Explanation:
AWS WAF is a web application firewall service that helps protect web applications from common web exploits that could affect availability, compromise security, or consume excessive resources.

AWS WAF gives you control over which traffic to allow or block to your web applications by definingcustomizable web security rules. You can use AWS WAF to create rules that block common attack patterns, such as SQL injection or cross-site scripting, and rules that filter out specific traffic patterns you define1. AWS WAF also integrates with other AWS services, such as Amazon CloudFront, Amazon API Gateway, AWS AppSync, and AWS Load Balancer, to provide a comprehensive defense against web attacks2. Therefore, AWS WAF meets the requirements of the social media company, compared to the other options.
The other options are not suitable for the social media company's requirements, because: Amazon Inspector is an automated security assessment service that helps improve the security and compliance of applications deployed on AWS. Amazon Inspector automatically assesses applications for exposure, vulnerabilities, and deviations from best practices. However, Amazon Inspector does not provide a web application firewall service that can block malicious web requests3.
Amazon GuardDuty is a threat detection service that continuously monitors for malicious activity and unauthorized behavior to protect your AWS accounts, workloads, and data stored in Amazon S3. Amazon GuardDuty analyzes and processes the following data sources: VPC Flow Logs, AWS CloudTrail event logs, and DNS logs. However, Amazon GuardDuty does not provide a web application firewall service that can block malicious web requests4.
Amazon CloudWatch is a monitoring and observability service that provides data and actionable insights to monitor your applications, respond to system-wide performance changes, optimize resource utilization, and get a unified view of operational health. Amazon CloudWatch collects monitoring and operational data in the form of logs, metrics, and events, and visualizes it using automated dashboards, alarms, and notifications. However, Amazon CloudWatch does not provide a web application firewall service that can block malicious web requests.

What Is AWS WAF? - AWS WAF, AWS Firewall Manager, and AWS Shield Advanced AWS WAF Features - AWS WAF, AWS Firewall Manager, and AWS Shield Advanced What Is Amazon Inspector? - Amazon Inspector
What Is Amazon GuardDuty? - Amazon GuardDuty
[What Is Amazon CloudWatch? - Amazon CloudWatch]

Question: 591

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:

The correct answer is B and C. EC2 Amazon Machine Images (AMIs) and Amazon Elastic Block Store (Amazon EBS) snapshots are two AWS services that provide disaster recovery solutions for Amazon EC2 instances.

EC2 AMIs are preconfigured templates that contain the software configuration and data required to launch an EC2 instance. You can create AMIs from your running EC2 instances and use them to launch new instances in the same or different AWS Regions. This way, you can quickly recover your EC2 instances in case of a disaster that affects your primary Region or Availability Zone1.

Amazon EBS snapshots are incremental backups of your Amazon EBS volumes. You can create snapshots of your volumes and store them in Amazon S3, which is a highly durable and scalable storage service. You can use snapshots to restore your volumes to a previous point in time or to create new volumes from snapshots. Snapshots can also be copied across AWS Regions, enabling you to recover your data in another Region in case of a disaster2.

The other options are not directly related to disaster recovery for EC2 instances:

EC2 Reserved Instances are a pricing model that allows you to reserve EC2 capacity for a specific period of time and receive a discount on the hourly charge. Reserved Instances do not provide any disaster recovery benefits, as they are only a billing option3.

AWS Shield is a managed service that protects your AWS resources from distributed denial-of-service (DDoS) attacks. AWS Shield provides basic protection for all AWS customers at no additional charge, and advanced protection for customers who need higher levels of detection and mitigation. AWS Shield does not provide any disaster recovery benefits, as it is only a security service4. Amazon GuardDuty is a threat detection service that monitors your AWS account and workloads for malicious or unauthorized activity. Amazon GuardDuty analyzes various data sources, such as AWS CloudTrail, Amazon VPC Flow Logs, and DNS logs, to identify potential threats and alert you via Amazon CloudWatch Events or AWS Lambda. Amazon GuardDuty does not provide any disaster recovery benefits, as it is only a monitoring service5.

Question: 592

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

A. Maintain the security of the AWS Cloud.
B. Configure firewalls and networks.
C. Patch the operating system of Amazon RDS instances.
D. Implement physical and environmental controls.

Answer: B

Explanation:

According to the AWS shared responsibility model, the customer is responsible for the security in the cloud, which includes configuring firewalls and networks. AWS provides security groups and network access control lists (NACLs) as firewall features that customers can use to control the traffic to andfrom their AWS resources. Customers are also responsible for managing their own virtual private clouds (VPCs), subnets, route tables, internet gateways, and other network components. AWS is responsible for the security of the cloud, which includes the physical security of the facilities, thehost operating system and virtualization layer, and the AWS global network infrastructure12.

Reference: Shared Responsibility Model - Amazon Web Services (AWS)

Shared responsibility model - Amazon Web Services: Risk and Compliance

Question: 593 A network engineer needs to build a hybrid cloud architecture connecting on-premises networks to the AWS Cloud using AWS Direct Connect. The company has a few VPCs in a single AWS Region and expects to increase the number of VPCs to hundreds over time.

Which AWS service or feature should the engineer use to simplify and scale this connectivity as the VPCs increase in number?

A. VPC endpoints
B. AWS Transit Gateway
C. Amazon Route 53
D. AWS Secrets Manager

Answer: B

Explanation:
AWS Transit Gateway is a network transit hub that you can use to interconnect your VPCs and onpremises networks through a central gateway. AWS Transit Gateway simplifies and scales the connectivity between your on-premises networks and AWS, as you only need to create and manage a single connection from the central gateway to each on-premises network, rather than individual connections to each VPC. You can also use AWS Transit Gateway to connect to other AWS services, such as Amazon S3, Amazon DynamoDB, and AWS PrivateLink12. AWS Transit Gateway supports thousands of VPCs per gateway, and enables you to peer Transit Gateways across AWS Regions3. The other options are not AWS services or features that can simplify and scale the connectivity between on-premises networks and hundreds of VPCs using AWS Direct Connect. VPC endpoints enable private connectivity between your VPCs and supported AWS services, but do not support onpremises networks4. Amazon Route 53 is a DNS service that helps you route internet traffic to your resources, but does not provide network connectivity5. AWS Secrets Manager is a service that helps you securely store and manage secrets, such as database credentials and API keys, but does not relate to network connectivity

Question: 594

A systems administrator created a new 1AM user for a developer and assigned the user an access key instead of a user name and password. What is the access key used for?

A. To access the AWS account as the AWS account root user
B. To access the AWS account through the AWS Management Console
C. To access the AWS account through a CLI
D. To access all of a company's AWS accounts

Answer: C

Explanation:

An access key is a pair of long-term credentials that consists of an access key ID and a secret access key. An access key is used to sign programmatic requests to the AWS CLI or AWS API (directly or using the AWS SDK). An access key allows a user to access the AWS account through a CLI, which is a tool that enables users to interact with AWS services using commands in a terminal or a script12.
The other options are not correct, because:
To access the AWS account as the AWS account root user, a user needs the email address and password associated with the account. The root user has complete access to all AWS resources and services in the account. However, it is not recommended to use the root user for

everyday tasks3. To access the AWS account through the AWS Management Console, a user needs a user name and password. The console is a web-based interface that allows users to manage their AWS resources and services using a graphical user interface4. To access all of a company's AWS accounts, a user needs to use AWS Organizations, which is a service that enables users to centrally manage and govern multiple AWS accounts. AWS Organizations allows users to create groups of accounts and apply policies to them5.

Managing access keys for IAM users - AWS Identity and Access Management What Is the AWS Command Line Interface? - AWS Command Line Interface AWS account root user - AWS Identity and Access Management What Is the AWS Management Console? - AWS Management Console
What Is AWS Organizations? - AWS Organizations

Question: 595

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:

On-Demand Instances are the most flexible and cost-effective pricing model for short-term, experimental, or unpredictable workloads on AWS. On-Demand Instances let you pay only for the resources you use, without any long-term commitments or upfront fees. You can easily start and stop instances as needed, and scale up or down depending on your demand.
Savings Plans, Reserved Instances, and Dedicated Hosts are all pricing models that require a commitment for a certain amount of usage or capacity for a one- or three-year term. These pricing models offer lower prices than On-Demand Instances, but they are not suitable for workloads that only run for 3 to 6 months or have variable usage patterns. Savings Plans and Reserved Instances also offer flexibility to change instance types, sizes, or regions within the same family or pool, while Dedicated Hosts are physical servers that can only run specific instance types.

Question: 596

A user wants to allow applications running on an Amazon EC2 instance to make calls to other AWS services. The access granted must be secure. Which AWS service or feature should be used?

A. Security groups
B. AWS Firewall Manager
C. IAM roles
D. IAM user SSH keys

Answer: C

Explanation:
IAM roles are a secure way to grant permissions to applications running on an Amazon EC2 instance to make calls to other AWS services. IAM roles are entities that have specific permissions policies attached to them. You can create an IAM role and associate it with an EC2 instance when you launch it or later. The applications on the instance can then use the temporary credentials provided by the role to access AWS resources that the role allows. This way, you do not have tostore any long-term credentials or access keys on the instance, which reduces the risk of compromise or misuse12. The other options are not correct, because:
Security groups are virtual firewalls that control the inbound and outbound traffic for your EC2 instances. Security groups do not grant permissions to access other AWS services, but rather filter the network traffic based on rules that you define3.
AWS Firewall Manager is a service that helps you centrally configure and manage firewall rules across your accounts and resources. AWS Firewall Manager works with AWS WAF, AWS Shield Advanced, and Amazon VPC security groups. AWS Firewall Manager does not grant permissions to access other AWS services, but rather helps you enforce consistent security policies across your AWS infrastructure4.
IAM user SSH keys are credentials that allow you to connect to your EC2 instance using SSH. SSH keys do not grant permissions to access other AWS services, but rather authenticate your identity when you log in to your instance5.

Using an IAM role to grant permissions to applications running on Amazon EC2 instances - AWS Identity and Access Management IAM roles for Amazon EC2 - Amazon Elastic Compute Cloud Security groups for your VPC - Amazon Virtual Private Cloud What is AWS Firewall Manager? - AWS Firewall Manager Connecting to your Linux instance using SSH - Amazon Elastic Compute Cloud

Question: 597

A company needs to track the activity in its AWS accounts, and needs to know when an API call is made against its AWS resources. Which AWS tool or service can be used to meet these requirements?

A. Amazon CloudWatch
B. Amazon Inspector
C. AWS CloudTrail
D. AWS IAM

Answer: C

Explanation:
AWS CloudTrail is the service that can be used to meet these requirements. AWS CloudTrail is a service that records AWS API calls for your account and delivers log files to you. The recorded information includes the identity of the API caller, the time of the API call, the source IP address of the API caller, therequest parameters, and the response elements returned by the AWS service1. You can use CloudTrail to track the activity in your AWS accounts, such as who made an API call, when it was made, and what resources were affected. You can also use CloudTrail to monitor thecompliance, security, and governance of your AWS environment2. The other services are not designed to track the activity and API calls in your AWS accounts. Amazon CloudWatch is a servicethat monitors and collects metrics, logs, and events from your AWS resources and applications. You can use CloudWatch to set alarms, visualize data, and automate actions based on predefined thresholds or rules3. Amazon Inspector is a service that helps you improve the security and compliance of your applications running on AWS. Inspector automatically assesses applications for exposure, vulnerabilities, and deviations from best practices4. AWS IAM is a service that enables you to manage access to AWS services and resources securely. IAM allows you to create and manage AWS users and groups, and use permissions to allow and deny their access to AWS resources. Reference: AWS CloudTrail, AWS CloudTrail - Capture AWS API Activity, Amazon CloudWatch, Amazon Inspector, [AWS IAM]

Question: 598

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 a set of tools that allow users to connect with AWS and deploy resources programmatically. AWS SDKs provide libraries, code samples, documentation, and other resources to help users write code that interacts with AWS APIs. AWS SDKs support various programming languages, such as Java, Python, Ruby, .NET, Node.js, Go, and more. AWS SDKs make it easier for users to access AWS services, such as Amazon S3, Amazon EC2, Amazon DynamoDB, AWS Lambda, and more, from their applications. AWS SDKs also handle tasks such as authentication, error handling, retries, and data serialization, so users can focus on their application logic .
The other options are not AWS services or tools that give users the ability to connect with AWS and deploy resources programmatically. Amazon QuickSight is a business intelligence service that lets users create and share interactive dashboards and visualizations1. AWS PrivateLink is a service that enables users to securely access services hosted on AWS in a scalable and cost-effective manner2. AWS Direct Connect is a service that establishes a dedicated network connection between a user's premises and AWS3.

Question: 599

A company has deployed applications on Amazon EC2 instances. The company needs to assess application vulnerabilities and must identify infrastructure deployments that do not meet best practices. Which AWS service can the company use to meet these requirements?

A. AWS Trusted Advisor
B. Amazon Inspector
C. AWSConfig
D. Amazon GuardDuty

Answer: B

Explanation:
Amazon Inspector is a service that provides automated security assessment and management for AWS resources, such as Amazon EC2 instances. Amazon Inspector can scan applications for common vulnerabilities, such as SQL injection, cross-site scripting, and remote code execution. Amazon Inspector can also check the configuration of AWS resources against security best practices, such as the CIS Benchmarks and the AWS Security Best Practices. Amazon Inspector can help customers identify and remediate security issues, comply with security standards, and improve the security posture of their AWS environment12. Reference:
Amazon Inspector
Improved, Automated Vulnerability Management for Cloud Workloads with a New Amazon Inspector | AWS News Blog

Question: 600

Which AWS service or tool can be used to set up a firewall to control traffic going into and coming out of an Amazon VPC subnet?

A. Security group
B. AWS WAF
C. AWS Firewall Manager
D. Network ACL

Answer: D

Explanation:

A network ACL (NACL) is an optional layer of security for your VPC that acts as a firewall for controlling traffic in and out of one or more subnets. You can create a network ACL and associate it with a subnet to apply rules that allow or deny traffic to or from the subnet. Network ACLs are stateless, meaning that they evaluate the source and destination IP addresses for both inbound and outbound traffic. You can also use network ACLs to block IP address ranges that are known to be malicious12.

The other options are not AWS services or tools that can be used to set up a firewall to control traffic going into and coming out of an Amazon VPC subnet. Security groups are another layer of security for your VPC that act as a firewall for your EC2 instances. Security groups are stateful, meaning that they automatically allow return traffic for allowed inbound traffic. Security groups can only filter traffic based on protocols, ports, and source or destination IP addresses, not on IP ranges3. AWS WAF is a web application firewall that helps protect your web applications from common web exploits. AWS WAF can filter web requests based on rules that you define, such as IP addresses, HTTP headers, HTTP body, or URI strings. AWS WAF does not apply to non-web traffic or to traffic within a VPC4. AWS Firewall Manager is a service that helps you centrally configure and manage firewall rules across your accounts and resources in AWS Organizations. You can use Firewall Manager to apply AWS WAF rules, AWS Network Firewall policies, and Amazon VPC security groups across your AWS accounts.

AWS Firewall Manager does not provide a firewall service itself, but rather helps you manage other firewall services

Question: 601

Which action is a security best practice for access to sensitive data that is stored in an Amazon S3 bucket?

A. Enable S3 Cross-Region Replication (CRR) on the S3 bucket.
B. Use IAM roles for applications that require access to the S3 bucket.
C. Configure AWS WAF to prevent unauthorized access to the S3 bucket.
D. Configure Amazon GuardDuty to prevent unauthorized access to the S3 bucket.

Answer: B

Explanation:

Understanding IAM Roles: IAM (Identity and Access Management) roles in AWS are designed to delegate access permissions without sharing long-term security credentials. This means applications and services can use temporary security credentials, which enhances security.

Why IAM Roles are Best Practice:

Least Privilege Principle: By using IAM roles, you can ensure that applications only have the minimum permissions they need to function, reducing the risk of unauthorized access.

Temporary Credentials: Roles provide temporary security credentials, which reduce the risk if they are compromised compared to

long-term access keys.

Automated Rotation: Temporary credentials automatically expire and are rotated, which means you don't have to manage the rotation manually.

How to Implement IAM Roles:

Create an IAM Role: In the AWS Management Console, navigate to IAM, and create a new role.

Choose the type of trusted entity (e.g., EC2, Lambda).

Attach Policies: Attach the necessary policies to the role that define the permissions for accessing the S3 bucket.

Assign Role to Service: Attach the IAM role to your EC2 instances, Lambda functions, or other AWS services that need to access the S3 bucket.

Use AWS SDKs: When accessing S3 from your application, use the AWS SDKs to automatically assume the IAM role and obtain temporary credentials.

AWS Identity and Access Management (IAM)

IAM Roles

Question: 602

Which AWS Well-Architected Framework pillar focuses on structured and streamlined allocation of computing resources?

A. Reliability
B. Operational excellence
C. Performance efficiency
D. Sustainability

Answer: C

Explanation:

Understanding Performance Efficiency: This pillar of the AWS Well-Architected Framework focuses on using computing resources efficiently to meet system requirements and maintain that efficiency as demand changes and technologies evolve.

Key Aspects of Performance Efficiency:

Selection: Choose the right resources for the job. This includes using the most appropriate instance types, storage options, and database services.

Review: Regularly review your architecture to take advantage of the latest AWS services and features, and to ensure you're using the best possible resource for your needs.

Monitoring: Continuously monitor your system performance, gather metrics, and use those metrics to make informed decisions about scaling and performance optimization.

Trade-offs: Understand the trade-offs between various performance-related aspects, such as cost, latency, and durability, and make decisions that align with your business goals.

How to Implement Performance Efficiency:

Use Auto Scaling: Implement Auto Scaling to automatically adjust the number of resources based on the demand.

Choose Appropriate Storage Options: Select the right storage solution (e.g., S3, EBS, or EFS) based on performance and access patterns.

Optimize Networking: Utilize Amazon CloudFront, AWS Global Accelerator, and VPC to optimize your network performance.

Regular Review and Testing: Regularly review your architecture, test performance under various loads, and adjust configurations as needed.

AWS Well-Architected Framework

Performance Efficiency Pillar

Question: 603

A company has deployed a web application to Amazon EC2 instances. The EC2 instances have low usage. Which AWS service or feature should lite company use in rightsized the FC? instances?

A. AWS Config
B. AWS Cost Anomaly Detection
C. AWS Budgets
D. AWS Compute Optimizer

Answer: D

Explanation:

Understanding AWS Compute Optimizer: AWS Compute Optimizer is a service that analyzes the configuration and utilization metrics of your AWS resources. It provides recommendations to help you select the optimal configurations for your workloads.

Why AWS Compute Optimizer for Rightsizing:

Resource Recommendations: It provides specific recommendations to rightsize your EC2 instances by suggesting instance types that match your actual usage patterns.

Cost Efficiency: By optimizing instance sizes, you can reduce costs associated with over-provisioned resources.

Performance Improvement: Ensures that you are using instances that provide the required performance without over-allocating resources.

How to Implement AWS Compute Optimizer:

Enable AWS Compute Optimizer: In the AWS Management Console, navigate to AWS Compute Optimizer and enable it for your account.

Review Recommendations: After a period of monitoring, review the recommendations provided for your EC2 instances.

Implement Changes: Follow the suggestions to resize or change instance types based on the recommendations, ensuring you balance cost savings with performance needs.

AWS Compute Optimizer

Question: 604

Which top-level key performance indicator (KPI) is available in AWS rightsizing recommendations of Cost Optimization?

A. Container modernization opportunities
B. Estimated monthly saving
C. Reserved instances savings
D. Compute savings recommendations

Answer: B

Explanation: Understanding Cost Optimization Recommendations: In AWS, cost optimization involves identifying ways to reduce costs

while maintaining or improving performance and capacity.
Top-Level KPI - Estimated Monthly Saving:
Definition: This KPI provides an estimate of how much you can save per month by following the recommended actions.
Importance: It helps you quantify the potential cost savings from rightsizing, purchasing reserved instances, or optimizing resource usage.
Decision-Making: Provides a clear financial benefit to justify changes in your resource configurations. How to Use Estimated Monthly Saving:
Access Recommendations: Navigate to the AWS Cost Management Console to view rightsizing recommendations.
Review Savings Estimates: Look at the estimated monthly savings for each recommendation to understand the potential financial impact.
Implement Recommendations: Prioritize actions based on the savings estimates to maximize cost reduction.
AWS Cost Management
AWS Rightsizing Recommendations

Question: 605

Which fully managed AWS service assists with the creation, testing, and management of custom Amazon EC? images?

A. EC2 Image Builder
B. Amazon Machine Image (AMI)
C. AWS Launch Wizard
D. AWS Elastic Beanstalk

Answer: A

Explanation:
Understanding EC2 Image Builder: EC2 Image Builder is a fully managed service that simplifies the creation, maintenance, validation, and testing of Amazon Machine Images (AMIs).
Why Use EC2 Image Builder:
Automation: Automates the creation and management of AMIs, reducing manual efforts and the risk of errors.
Customization: Allows you to customize the images to include necessary software, configurations, and security settings.
Compliance: Ensures that the images comply with your security and operational standards through continuous monitoring and testing.
How to Implement EC2 Image Builder:
Create a Recipe: Define an image recipe specifying the base image and components to be included. Build Pipeline: Set up an image pipeline that automates the building and testing of the AMI based on a schedule or trigger.
Distribute Images: Use the produced AMIs across multiple AWS regions and accounts as needed.
EC2 Image Builder

Question: 606

A company plans to use an Amazon Snowball Edge device la transfer files to the AWS Cloud. Which activities related in a Snowball device are available to the company at no cost?

A. Use of the Snowball Edge appliance for a 10-day period
B. The transfer of data out of Amazon S3 and 10 the Snowball Edge appliance
C. The transfer of data from the Snowball Edge appliance into Amazon S3
D. Daily use of the Snowball Edge appliance after 10 days

Answer: A

Explanation:
Understanding Amazon Snowball Edge: Amazon Snowball Edge is a data migration and edge computing device that helps in transferring large amounts of data to and from AWS. No-Cost Activities:
Transfer into S3: AWS does not charge for transferring data from the Snowball Edge device into Amazon S3. This is part of the service to facilitate easy data migration into the AWS cloud.
Other Costs: The use of the Snowball Edge appliance for up to 10 days is included in the service fee. However, any usage beyond the 10-day period may incur additional costs.
How to Utilize Snowball Edge:
Order the Device: Use the AWS Management Console to order a Snowball Edge device.
Load Data: Transfer your data onto the Snowball Edge appliance.
Ship Back: Ship the device back to AWS, where the data will be uploaded to Amazon S3 at no additional cost.
Amazon Snowball Edge Pricing
Amazon Snowball Edge Documentation

Question: 607

Which AWS Cloud benefit is shown by an architecture's ability to withstand failures with minimal downtime?

A. Agility
B. Elasticity
C. Scalability
D. High availability

Answer: D

Explanation:
Understanding High Availability: High availability (HA) refers to systems that are durable and likely to operate continuously without failure for a long time. HA ensures that an architecture can withstand failures with minimal downtime.
Importance of High Availability:
Redundancy: Systems are designed with redundancy to prevent single points of failure.
Fault Tolerance: Ensures that failures do not result in significant downtime, maintaining service continuity.
Automated Recovery: Utilizes automated recovery mechanisms to quickly restore services in the event of a failure.
AWS Services for High Availability:
Multi-AZ Deployments: Services like RDS, DynamoDB, and others support Multi-AZ deployments for fault tolerance.
Elastic Load Balancing: Distributes traffic across multiple instances or availability zones to ensure no single point of failure.
Auto Scaling: Automatically adjusts the number of instances based on demand, ensuring availability even during traffic spikes.

AWS Well-Architected Framework: Reliability

Question: 608

How does AWS Cloud computing help businesses reduce costs? (Select TWO.)

A. AWS changes the name prices for servicers in every AWS Region.
B. AWS enables capacity in be adjusted un demand.
C. AWS offers discounts tor Amazon LC2 instances that remain Idle tor more man 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:
Reducing Costs with AWS Cloud:
Capacity Adjustment (B):

Elasticity: AWS allows you to scale your resources up or down based on demand, which means you only pay for what you use. This reduces the cost of over-provisioning resources.

Auto Scaling: Automatically adjusts compute capacity based on usage, ensuring cost efficiency. Eliminating On-Premises Costs (E): No Infrastructure Maintenance: By using AWS, businesses do not need to invest in physical infrastructure or handle maintenance, reducing both capital and operational expenditures. Managed Services: AWS offers managed services that reduce the need for in-house technical staff to manage and maintain infrastructure.

AWS Cloud Economics Center
AWS Benefits

Question: 609

Which AWS resource can help a company reduce Its costs in exchange for a usage commitment when using Amazon EC2 instances?

A. Compute Savings Plans
B. Auto Stalling group
C. On-Demand Instance
D. EC2 instance store

Answer: A

Explanation:

Question: 610

Which option is an environment that consists of one or more data centers?

A. Amazon CloudFront
B. Availability Zone

C. VPC
D. AWS Outposts

Answer: B

Explanation:
Understanding Availability Zones (AZs): An Availability Zone is a distinct location within an AWS region that is engineered to be isolated from failures in other AZs.
Characteristics of Availability Zones:
Data Centers: Each AZ consists of one or more discrete data centers with redundant power, networking, and connectivity.
High Availability: AZs are designed for high availability, providing low-latency network connections to other zones in the same region.
Fault Isolation: They provide fault isolation and are used to deploy applications and services to ensure high availability and reliability.
Use Cases for Availability Zones:
Multi-AZ Deployments: For services like RDS, deploying in multiple AZs ensures fault tolerance. Disaster Recovery: Setting up resources in multiple AZs helps in quick recovery from failures. Load Balancing: Distributing traffic across AZs using Elastic Load Balancing ensures optimal performance and availability.

AWS Global Infrastructure Understanding AWS Regions and Availability Zones

Question: 611

A company wants to set up its workloads to perform their intended functions and recover quickly from failure. Which pillar of the AWS Well-Architected Framework aligns with these goals?

A. Performance efficiency
B. Sustainability
C. Reliability
D. Security

Answer: C

Explanation:
Understanding the Reliability Pillar: The Reliability pillar of the AWS Well-Architected Framework focuses on the ability of a system to recover from infrastructure or service disruptions, dynamically acquire computing resources to meet demand, and mitigate disruptions such as misconfigurations or transient network issues.
Key Concepts of Reliability:
Foundations: Ensure a solid foundation on which to build, including AWS account management, limits, and networking.
Change Management: Manage changes in automation to ensure systems remain reliable during modifications.
Failure Management: Design systems to detect failures and automatically recover from them.
How to Align with Reliability Pillar:
Implement Multi-AZ Deployments: Deploy applications across multiple Availability Zones to ensure fault tolerance.
Use Auto Scaling: Automatically adjust resources to maintain system performance during demand fluctuations.
Monitor and Respond: Implement monitoring and alerting mechanisms using services like
CloudWatch to detect and respond to issues proactively.

AWS Well-Architected Framework: Reliability Pillar

Question: 612

A company Is designing its AWS workloads so that components can be updated regularly and so that changes can be made in small, reversible increments.
Which pillar of the AWS Well-Architected Framework does this design support?

A. Security
B. Performance efficiency
C. Operational excellence
D. Reliability

Answer: C

Explanation:
Understanding Operational Excellence: The Operational Excellence pillar of the AWS Well- Architected Framework focuses on running and monitoring systems to deliver business value and continuously improving supporting processes and procedures.
Key Concepts of Operational Excellence:
Small, Reversible Changes: Making changes in small, incremental steps allows for easier troubleshooting and rollback if issues arise.
Regular Updates: Regularly updating components ensures that systems stay up-to-date with the latest features, security patches, and performance improvements.
Automation: Implementing automation for deployments, updates, and monitoring to reduce human error and increase efficiency.
Continuous Improvement: Encouraging continuous learning and process improvement to enhance operational processes.
Implementing Operational Excellence:
Deployment Automation: Use CI/CD pipelines to automate deployments and ensure that changes can be rolled back if necessary.
Monitoring and Logging: Implement comprehensive monitoring and logging to track system health and performance.
Incident Response: Develop a robust incident response plan to handle issues quickly and efficiently. Documentation and Training: Maintain thorough documentation and provide training to ensure teams can effectively manage and improve operations.

AWS Well-Architected Framework: Operational Excellence Pillar

Question: 613

Which Amazon S3 storage class is MOST cost-effective for unknown access patterns?

A. S3 Standard
B. S3 Standard-Infrequent Access (S3 Standard-IA)
C. S3 One Zone-Infrequent Access (S3 One Zone-IA)
D. S3 Intelligent-Tiering

Answer: D

Explanation:

Understanding S3 Intelligent-Tiering: S3 Intelligent-Tiering is designed to optimize costs by automatically moving data to the most cost-effective access tier based on changing access patterns. It is ideal for data with unknown or unpredictable access patterns.

Why S3 Intelligent-Tiering is Cost-Effective:

Automatic Tiering: Moves data between two access tiers (frequent and infrequent access) based on changing access patterns, optimizing storage costs without performance impact.

No Retrieval Fees: Unlike other storage classes, there are no retrieval fees in Intelligent-Tiering, making it cost-effective for data with unpredictable access patterns.

Monitoring and Automation: Automatically monitors access patterns and transitions data, reducing the need for manual intervention.

When to Use S3 Intelligent-Tiering:

Unpredictable Access Patterns: Ideal for datasets where the access frequency cannot be determined or changes frequently.

Cost Optimization: For organizations looking to minimize storage costs without sacrificing performance or requiring manual intervention to move data between tiers.

Amazon S3 Intelligent-Tiering

Amazon S3 Storage Classes

Question: 614

An ecommerce company has deployed a new web application on Amazon EC2 Instances. The company wants to distribute incoming HTTP traffic evenly across all running instances. Which AWS service or resource will meet this requirement?

A. Amazon EC2 Auto Scaling
B. Application Load Balancer
C. Gateway Load Balancer
D. Network Load Balancer

Answer: B

Explanation:

AnApplication Load Balancer (ALB)is the best choice for distributing incoming HTTP/HTTPS traffic evenly across multiple Amazon EC2 instances. It operates at theapplication layer (Layer 7 of the OSI model) and is specifically designed to handle HTTP and HTTPS traffic, which is ideal for web applications.

Here is why the ALB is the correct choice:

Layer 7 Load Balancing: The ALB works at the application layer and provides advanced routing capabilities based on content. It can inspect the incoming HTTP requests and make decisions on how to route traffic to various backend targets, which include Amazon EC2 instances, containers, or Lambda functions. This is particularly useful for web applications where you need to make routing decisions based on HTTP headers, paths, or query strings.

HTTP and HTTPS Support: The ALB natively supports HTTP and HTTPS protocols, making it the ideal load balancer for web-based applications. It can efficiently manage and route these types of traffic and handle tasks such as SSL/TLS termination.

Health Checks: The ALB can continuously monitor the health of the registered EC2 instances and only route traffic to healthy instances. This ensures high availability and reliability of the web application. Path-based and Host-based Routing: The ALB can route traffic based on the

URL path or host header. This feature allows the same load balancer to serve multiple applications hosted on different domains or subdomains.

Integration with Auto Scaling: The ALB can integrate seamlessly with Amazon EC2 Auto Scaling. As the number of EC2 instances increases or decreases, the ALB automatically includes the new instances in its traffic distribution pool, ensuring even distribution of incoming requests. WebSocket Support: It also supports WebSocket and HTTP/2 protocols, which are essential for modern web applications that require real-time, bidirectional communication.

Why other options are not suitable:

A . Amazon EC2 Auto Scaling: This service is used to automatically scale the number of EC2 instances up or down based on specified conditions. However, it does not provide load balancing capabilities. It works well with load balancers but does not handle the distribution of incoming traffic by itself. C . Gateway Load Balancer: This is designed to distribute traffic to virtual appliances like firewalls, IDS/IPS systems, or deep packet inspection systems. It operates at Layer 3 (Network Layer) and is not ideal for distributing HTTP/HTTPS traffic to EC2 instances.

D . Network Load Balancer: This load balancer operates at Layer 4 (Transport Layer) and is designed to handle millions of requests per second while maintaining ultra-low latencies. It is best suited for TCP, UDP, and TLS traffic but does not provide advanced Layer 7 routing features required for HTTP/HTTPS traffic.

Reference:

AWS Application Load Balancer Documentation Comparison of Elastic Load Balancing Options

Question: 615

Which AWS Support plan provides the full set to AWS Trusted Advisor checks at the LOWEST cost?

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

Answer: B

Explanation:

AWS Trusted Advisor is a tool that helps customers optimize their AWS environment by providing real-time guidance in five key areas: cost optimization, performance, security, fault tolerance, and service limits. However, the availability of the full set of Trusted Advisor checks depends on the AWS Support plan chosen.

AWS Developer Support: This plan provides access to only seven core Trusted Advisor checks. It is designed for developers experimenting or testing in AWS and does not offer the full set of Trusted Advisor checks.

AWS Business Support: This plan is the lowest-cost support plan that provides access to thefull set of AWS Trusted Advisor checks. Business Support is intended for production workloads, providing a broader range of checks, 24x7 access to Cloud Support Engineers, and more extensive support features.

AWS Enterprise On-Ramp Support: This plan offers access to all Trusted Advisor checks as well but is more expensive than the Business Support plan. It is designed for customers running production workloads and needing additional technical support but does not require the full level of engagement that comes with Enterprise Support.

AWS Enterprise Support: This is the most comprehensive and highest-cost support plan, providing access to all Trusted Advisor checks and a full range of AWS Support resources, including a Technical Account Manager (TAM), account management, concierge support, and more.

Conclusion:

The correct answer isB. AWS Business Support, as it provides access to the full set of AWS Trusted Advisor checks at thelowest cost. AWS Developer Support does not offer the complete checks, and both AWS Enterprise On-Ramp and AWS Enterprise Support are higher-cost plans that also provide the full checks.

AWS Cloud Reference:
AWS Support Plans
AWS Trusted Advisor

Question: 616

A company is requesting Payment Card Industry (PCI) reports that validate the operating effectiveness of AWS security controls. How should the company obtain these reports?

A. Contact AWS Support
B. Download reports from AWS Artifact.
C. Download reports from AWS Security Hub.
D. Contact an AWS technical account manager (TAM).

Answer: B

Explanation:
AWS Artifact is a service provided by AWS that offers on-demand access to AWS compliance reports, including the Payment Card Industry (PCI) reports. It is the primary tool for retrieving compliance reports such as Service Organization Control (SOC) reports, ISO certifications, and Payment Card Industry Data Security Standard (PCI DSS) reports.
To obtain these reports:
The company should log into the AWS Management Console and navigate to AWS Artifact.
From there, they can select and download the necessary compliance reports.
Why other options are not suitable:
A . Contact AWS Support: AWS Support is not needed to obtain these reports; they are readily available through AWS Artifact.
C . Download reports from AWS Security Hub: AWS Security Hub is a service that provides a comprehensive view of security alerts and compliance status, but it does not host or provide compliance reports like PCI DSS.
D . Contact an AWS technical account manager (TAM): While a TAM may assist in various AWS- related queries, they are not required to obtain PCI reports. AWS Artifact is designed for this purpose.
Reference:
AWS Artifact Documentation

Question: 617

A company has data lakes designed for high performance computing (HPC) workloads. Which Amazon EC2 instance type should the company use to meet these requirements?

A. General purpose instances
B. Compute optimized instances
C. Memory optimized instances
D. Storage optimized instances

Answer: B

Explanation:

For high performance computing (HPC) workloads, compute resources play a critical role in delivering the necessary processing power and efficiency. HPC workloads are typically computationally intensive, often requiring a large number of CPU cycles to solve complex problems. These workloads benefit most from instances that provide powerful processors and high clock speeds, which is whyCompute optimized instances (Answer B)are the best choice in this scenario. Why Compute Optimized Instances (C Instances)?

Designed for Compute-Intensive Tasks: Compute optimized instances in Amazon EC2, such as the C6i or C5 series, are designed to offer high compute performance, low cost, and consistent CPU power. These instances are ideal for workloads like HPC, which require a high level of processing per second. High Performance CPUs: The compute optimized instance family typically uses the latest-generation processors, such as AWS Graviton2 or Intel Xeon Scalable processors, which provide a higher number of virtual CPUs (vCPUs) and increased clock speeds compared to other instance types. This matches the need for HPC workloads to maximize throughput and minimize compute times.

Use Case Alignment: HPC workloads such as genomic research, computational fluid dynamics (CFD), financial modeling, and 3D rendering require high levels of CPU-bound tasks. Compute optimized instances provide the best CPU-to-memory ratio to handle these efficiently, leading to faster processing times and cost efficiency.

Comparison with Other Instance Types:

A . General Purpose Instances: These are versatile and balanced instances (e.g., T3 or M6i) that are suitable for various workloads but do not provide the specialized compute performance required for HPC. They offer a balanced mix of compute, memory, and networking but are not optimal for HPC workloads where computational power is critical.

C . Memory Optimized Instances: While these instances (e.g., R5, X1) are ideal for memory-intensive workloads such as in-memory databases (e.g., SAP HANA) or real-time data analytics, they do not provide the specialized compute power necessary for HPC tasks that require heavy CPU processing. D . Storage Optimized Instances: These instances (e.g., I3, D3) are designed for workloads that need high disk throughput, like big data or transactional databases. While these are excellent for storageheavy applications, they are not optimized for compute-intensive HPC workloads.

Amazon EC2 Compute Optimized Family (C Instances)

C6i Instances: Based on 3rd Gen Intel Xeon Scalable processors, C6i instances offer up to 15% better price/performance compared to previous generation C5 instances. These are ideal for high compute and HPC workloads.

C5 Instances: These are built for compute-intensive workloads like batch processing, distributed analytics, and high-performance web servers. They offer a high level of sustained CPU performance.

AWS Reference Links:

Amazon EC2 Instance Types

Amazon EC2 Compute Optimized Instances

HPC on AWS

In conclusion,Compute optimized instances (B)are the best choice for HPC workloads due to their high compute performance, optimized CPU architecture, and suitability for computationally intensive tasks.

Question: 618

A company needs to perform an audit of recent AWS account activity. The audit will investigate who initiated an event and what actions were performed.

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

A. AWS Config
B. Amazon Rekognition
C. AWS CloudTrail
D. Amazon Simple Notification Service (Amazon SNS)

Answer: C

Explanation:

AWS CloudTrail is a service that enables governance, compliance, and operational and risk auditing of an AWS account. It logs, continuously monitors, and retains account activity related to actions across an AWS infrastructure.

For auditing purposes:

CloudTrail records AWS API calls made in the account, including details about who made the request, the services used, the actions performed, and the response elements returned by AWS.

This information is critical for understanding user activity, detecting anomalous behavior, and performing security analysis and compliance auditing.

Why other options are not suitable:

A . AWS Config: AWS Config provides a detailed view of the configuration of AWS resources, including how resources are related and their compliance with internal policies, but it does not provide a comprehensive audit trail of user actions.

B . Amazon Rekognition: A service for image and video analysis, not relevant to auditing AWS account activity.

D . Amazon SNS: A notification service for sending alerts and messages, not used for auditing purposes.

Reference:

AWS CloudTrail Documentation

Question: 619

Which of the following can the AWS Pricing Calculator do?

A. Project monthly AWS costs.
B. Calculate historical AWS costs.
C. Provide in-depth information about AWS pricing strategies.
D. Provide users with access to their monthly bills.

Answer: A

Explanation:

TheAWS Pricing Calculatoris a web-based tool that allows customers to estimate their monthly AWS costs by configuring and projecting the costs of different AWS services. The calculator is specifically designed to help customers plan their AWS spending based on their specific architecture and usage patterns. The other options are not correct because:

B . Calculate historical AWS costs: This is incorrect as the AWS Pricing Calculator does not track or calculate historical costs. For historical costs, AWS Cost Explorer is used.

C . Provide in-depth information about AWS pricing strategies: While the AWS Pricing Calculator provides cost estimations, it does not provide detailed insights into AWS pricing strategies.

D . Provide users with access to their monthly bills: This is incorrect; AWS Billing and Cost Management provides access to actual billing information.

AWS Cloud Reference:

AWS Pricing Calculator

Question: 620

A company is using AWS for all its IT Infrastructure. The company's developers are allowed to deploy applications on their own. The developers want to deploy their applications without having to provision the infrastructure themselves.

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

A. AWS Cloud Formation

B. AWS CodeBuild
C. AWS Elastic Beanstalk
D. AWS CodeDeploy

Answer: C

Explanation:

AWS Elastic Beanstalkis a fully managed service designed for developers who want to deploy and manage their applications without having to provision and manage the underlying infrastructure themselves. Developers can simply upload their code, and Elastic Beanstalk automatically handles the deployment, including provisioning the necessary resources (such as EC2 instances, load balancers, and auto-scaling).

A . AWS CloudFormation: Incorrect, as it is an infrastructure-as-code service for defining and provisioning AWS resources but does not directly deploy applications.

B . AWS CodeBuild: Incorrect, as it is a service for building and testing code, not for deploying applications.

D . AWS CodeDeploy: Incorrect, as it is specifically designed for automating software deployments to a variety of compute services, including EC2, but it does not manage the underlying infrastructure. AWS Cloud Reference:

AWS Elastic Beanstalk

Question: 621

Which AWS service or feature requires an Internet service provider (ISP) and a colocation facility to be Implemented?

A. AWS VPN
B. Amazon Conned
C. AWS Direct Connect
D. Internet gateway

Answer: C

Explanation:

AWS Direct Connectis a cloud service solution that makes it easy to establish a dedicated network connection from your premises to AWS. This requires the use of an Internet Service Provider (ISP) and a colocation facility to connect to the Direct Connect location. It provides a private, high-speed, low-latency connection that does not go over the public internet.

A . AWS VPN: Incorrect, as AWS VPN establishes secure connections over the internet and does not necessarily require a colocation facility.

B . Amazon Connect: Incorrect, as it is a cloud-based contact center service.

D . Internet Gateway: Incorrect, as it is a horizontally scaled, redundant, and highly available VPC component that allows communication between instances in your VPC and the internet.

AWS Cloud Reference:

AWS Direct Connect

Question: 622

A company needs a firewall that will control network connections to and from a single Amazon EC2 instance. This firewall will not control network connections to and from other instances that are in the same subnet.
Which AWS service or feature can the company use to meet these requirements?

A. Network ACL
B. AWS WAF
C. Route table
D. Security group

Answer: D

Explanation:
ASecurity Groupacts as a virtual firewall for your Amazon EC2 instances to control inbound and outbound traffic. It provides granular control over network connections to and from a specific EC2 instance or set of instances. Unlike Network ACLs, which operate at the subnet level, Security Groups operate at the instance level, allowing control over network traffic for individual instances.
A . Network ACL: Incorrect, as it controls traffic at the subnet level and not for individual instances.
B . AWS WAF: Incorrect, as AWS WAF is a web application firewall that helps protect web applications from common web exploits but is not designed for controlling instance-level traffic.
C . Route table: Incorrect, as route tables are used for network routing within a VPC and do not act as firewalls.
AWS Cloud Reference:
AWS Security Groups

Question: 623

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:
When a company moves from on-premises IT architecture to the AWS Cloud, it gains several benefits:
A . Reduced or eliminated tasks for hardware troubleshooting, capacity planning, and procurement: In an on-premises environment, companies must maintain their own hardware, which involves procuring servers, configuring them, managing capacity, and troubleshooting hardware failures. Moving to the AWS Cloud eliminates or greatly reduces these tasks since AWS is responsible for the underlying infrastructure.
E . Faster deployment of new features and applications: AWS provides scalable resources and automation tools that allow companies to deploy applications faster. Services like AWS Elastic Beanstalk, AWS CloudFormation, and AWS Lambda help streamline the deployment process, enabling quicker time-to-market for new features and applications.

Why other options are not suitable:

B . Elimination of the need for trained IT staff: While the cloud reduces certain operational burdens, it does not eliminate the need for skilled IT professionals to manage cloud services, ensure proper configurations, handle security, and manage applications.

C . Automatic security configuration of all applications that are migrated to the cloud: AWS provides security tools and services, but security configurations and management still require input from the customer to tailor them to specific application requirements.

D . Elimination of the need for disaster recovery planning: While AWS offers robust disaster recovery options, companies must still plan and implement disaster recovery strategies, such as backups and multi-region deployments.

Reference:

Benefits of Cloud Computing on AWS

Question: 624

Which guidelines are best practices for using AWS Identity and Access Management (1AM)? (Select TWO.)

A. Share access keys.
B. Create individual 1AM users.
C. Use inline policies instead of customer managed policies.
D. Grant maximum privileges to 1AM users.
E. Use groups to assign permissions to 1AM users.

Answer: B,E

Explanation:

Best practices for using AWS Identity and Access Management (IAM) include:

B . Create individual IAM users: Each user should have their own IAM credentials to ensure accountability, control, and traceability. Sharing credentials can lead to security risks and difficulty in auditing.

E . Use groups to assign permissions to IAM users: Assigning permissions through IAM groups simplifies permission management. You can assign the necessary permissions to the group, and then add or remove users from the group as needed, rather than managing permissions for each user individually.

Why other options are not suitable:

A . Share access keys: Sharing access keys is a security risk and violates AWS security best practices. Each user should have their own credentials.

C . Use inline policies instead of customer-managed policies: Customer-managed policies are preferred over inline policies because they offer better control, reusability, and versioning.

D . Grant maximum privileges to IAM users: Granting the least privilege necessary is a best practice to reduce the risk of accidental or malicious actions.

Reference:

AWS IAM Best Practices

Question: 625

Which AWS service supports user sign-up functionality and authentication to mobile and web applications?

A. Amazon Cognito
B. AWS Config

C. Amazon GuardDuty
D. AWS Systems Manager

Answer: A

Explanation:
Amazon Cognitois an AWS service that provides user sign-up, sign-in, and access control to web and mobile applications. It supports authentication for different identity providers, including social identity providers (such as Google, Facebook, and Apple), enterprise identity providers via SAML 2.0, and its own user pools.
Amazon Cognito offers:
User Poolsfor managing user registration, authentication, and account recovery. Federated Identitiesfor managing user access from external identity providers.
Why other options are not suitable:
B . AWS Config: A service for tracking resource configuration changes, not related to user authentication.
C . Amazon GuardDuty: A threat detection service, not related to user sign-up or authentication. D . AWS Systems Manager: A service to manage AWS resources, but it does not provide user authentication.
Reference:
Amazon Cognito Documentation

Question: 626

What is the MOST secure way to store passwords on AWS?

A. Store passwords in an Amazon S3 bucket.
B. Store passwords as AWS CloudFormation parameters
C. Store passwords in AWS Storage Gateway.
D. Store passwords in AWS Secrets Manager.

Answer: D

Explanation:
AWS Secrets Manageris the most secure way to store and manage sensitive information, such as passwords, database credentials, and API keys. Secrets Manager allows you to:
Securely store and rotate secrets.
Automatically manage secret rotation without disrupting applications.
Integrate with AWS services and third-party applications to retrieve secrets securely.
Why other options are not suitable:
A . Store passwords in an Amazon S3 bucket: Although S3 is secure, it is not designed for secret management. It lacks built-in secret rotation and fine-grained access control for sensitive data. B . Store passwords as AWS CloudFormation parameters: CloudFormation is used for managing infrastructure as code, not for securely storing passwords.
C . Store passwords in AWS Storage Gateway: A service for hybrid storage integration, not suitable for storing secrets or passwords.
Reference:
AWS Secrets Manager Documentation

Question: 627

A company plans to migrate its custom marketing application and order-processing application to AWS. The company needs to deploy the applications on different types of instances with various configurations of CPU, memory, storage, and networking capacity. Which AWS service should the company use to meet these requirements?

A. AWS Lambda
B. Amazon Cognito
C. Amazon Athena
D. Amazon EC2

Answer: D

Explanation:

Amazon EC2 (Elastic Compute Cloud)provides scalable computing capacity in the AWS Cloud, allowing customers to run virtual servers (instances) with different configurations of CPU, memory, storage, and networking capacity. This flexibility is ideal for applications that require specific infrastructure configurations, such as custom marketing and order-processing applications. A . AWS Lambda: Incorrect, as it is a serverless compute service that automatically manages the computing resources needed to run code but does not offer the flexibility of choosing different instance types.

B . Amazon Cognito: Incorrect, as it is used for user authentication and authorization, not for deploying applications.

C . Amazon Athena: Incorrect, as it is an interactive query service for analyzing data in Amazon S3 using standard SQL.

AWS Cloud Reference:

Amazon EC2

Question: 628

Which AWS service or feature gives users the ability to connect VPCs and on-premises networks to a central hub?

A. Virtual private gateway
B. AWS Transit Gateway
C. Internet gateway
D. Customer gateway

Answer: B

Explanation:

AWS Transit Gatewayis a network transit hub that customers can use to connect their Amazon VPCs and on-premises networks to a central hub. This service simplifies network management and reduces operational overhead by enabling a single gateway for managing multiple network connections. It facilitates seamless integration and routing between VPCs and on-premises networks. A . Virtual private gateway: Incorrect, as it is used to connect a single VPC to an on-premises network through a VPN connection.

C . Internet gateway: Incorrect, as it provides internet access for instances in a VPC but does not connect multiple networks.

D . Customer gateway: Incorrect, as it represents the on-premises device or software application that connects to AWS, but it does not provide a central hub.

AWS Cloud Reference:
AWS Transit Gateway

Question: 629

A company hosts its website on Amazon EC2 instances. The company needs to ensure that the website reaches a global audience and provides minimum latency to users.
Which AWS service should the company use to meet these requirements?

A. Amazon Route 53
B. Amazon CloudFront
C. Elastic Load Balancing
D. AWS Lambda

Answer: B

Explanation:

Amazon CloudFrontis a content delivery network (CDN) that helps deliver your website content globally with low latency by caching copies of your website content at edge locations around the world. This helps ensure that users receive content from the edge location closest to them, thereby reducing latency and improving user experience.

A . Amazon Route 53: Incorrect, as it is a DNS web service that routes users to the appropriate endpoint, but it does not cache content or reduce latency.

C . Elastic Load Balancing: Incorrect, as it distributes incoming application or network traffic across multiple targets, but does not cache content globally.

D . AWS Lambda: Incorrect, as it is a serverless compute service, not intended for content delivery. AWS Cloud Reference:
Amazon CloudFront

Question: 630

A company wants to build graph queries for real-time fraud pattern detection.
Which AWS service will meet this requirement?

A. Amazon Neptune
B. Amazon DynamoDB
C. Amazon Timestream
D. Amazon Forecast

Answer: A

Explanation:

Amazon Neptuneis a fully managed graph database service optimized for storing and querying highly connected data. It supports popular graph models such as Property Graph and W3C's RDF, making it ideal for building graph queries for real-time fraud pattern detection and other applications that require complex relationships and data traversal.

B . Amazon DynamoDB: Incorrect, as it is a NoSQL database service that is not optimized for graph queries.

C . Amazon Timestream: Incorrect, as it is a time-series database service designed for storing and analyzing time-series data, not graph data.

D . Amazon Forecast: Incorrect, as it is a fully managed service that provides time-series forecasting capabilities, not for graph queries.

AWS Cloud Reference:

Amazon Neptune

Question: 631

A company wants to deploy a web application as a containerized application. The company wants to use a managed service that can automatically create container images from source code and deploy the containerized application.

Which AWS service will meet these requirements?

A. AWS Elastic Beanstalk
B. Amazon Elastic Container Service (Amazon ECS)
C. AWS App Runner
D. Amazon EC2

Answer: C

Explanation:

AWS App Runneris a fully managed service that makes it easy for developers to quickly deploy containerized web applications and APIs at scale. It can automatically build container images from source code or directly from a container registry, and then deploy the application without requiring deep container knowledge or expertise.

AWS App Runner meets the requirements of:

Automatically creating container images from source code.

Managing the deployment of the containerized application with minimal operational overhead. Why other options are not suitable:

A . AWS Elastic Beanstalk: While it simplifies deploying and scaling web applications, it does not automatically create container images from source code.

B . Amazon Elastic Container Service (Amazon ECS): ECS is a container orchestration service that requires more manual setup for creating and managing container images and deployments.

D . Amazon EC2: This service provides virtual servers but does not offer managed container image creation or deployment.

Reference:

AWS App Runner Documentation

Question: 632

A company has moved all its infrastructure to the AWS Cloud. To plan ahead for each quarter, the finance team wants to track the cost and usage data of all resources from previous months. The finance team wants to automatically generate reports that contains the data.

Which AWS service or feature should the finance team use to meet these requirements?

A. Amazon Detective
B. AWS Pricing Calculator
C. AWS Budgets
D. AWS Savings Plans

Answer: C

Explanation:

AWS Budgetsallows users to set custom cost and usage budgets and receive notifications when they exceed their thresholds. It provides detailed reports on cost and usage data for the past and current months, enabling the finance team to track and analyze spending.

AWS Budgets can automatically generate cost and usage reports, which can help the finance team plan ahead for each quarter based on historical data.

Why other options are not suitable:

A . Amazon Detective: A security service for analyzing and investigating AWS account activity for security purposes, not cost tracking.

B . AWS Pricing Calculator: A tool to estimate costs based on expected usage, not for tracking actual past usage.

D . AWS Savings Plans: An offering to save costs on AWS usage; it does not provide cost tracking or reporting features.

Reference:

AWS Budgets Documentation

Question: 633

Which AWS Cloud Adoption Framework (AWS CAF) perspective focuses on real-time insights and answers questions about strategy?

A. Operations
B. People
C. Business
D. Platform

Answer: C

Explanation:

TheBusiness perspectiveof the AWS Cloud Adoption Framework (AWS CAF) focuses on ensuring that cloud investments align with business strategies and that the organization achieves business value from cloud adoption. It provides real-time insights into the organization's strategic goals, financial objectives, and metrics.

This perspective helps answer questions about strategy, ensuring that cloud adoption aligns with the organization's long-term business goals.

Why other options are not suitable:

A . Operations: Focuses on operating cloud environments effectively.

B . People: Focuses on human resources and organizational structure.

D . Platform: Focuses on infrastructure and architecture. Reference:

AWS Cloud Adoption Framework (AWS CAF)

Question: 634

A company wants to migrate a company's on-premises container Infrastructure to the AWS Cloud. The company wants to prevent unplanned administration and operation cost and adapt to a serverless architecture.

Which AWS service will meet these requirements?

A. Amazon Connect
B. AWS Fargate
C. Amazon Lightsail
D. Amazon EC2

Answer: B

Explanation:
AWS Fargateis a serverless compute engine for containers that allows users to run containers without having to manage the underlying infrastructure. Fargate eliminates the need for managing servers and reduces operational overhead, providing a fully managed, serverless approach to containerized applications. It helps avoid unplanned administration and operational costs and is ideal for companies migrating from on-premises container infrastructure.
Why other options are not suitable:
A . Amazon Connect: A service for cloud-based contact centers, not for container management.
C . Amazon Lightsail: Simplifies virtual private server (VPS) management but is not serverless or specialized for containers.
D . Amazon EC2: Provides virtual servers but requires more manual administration and is not serverless.
Reference:
AWS Fargate Documentation

Question: 635

Which AWS service is designed to help users handle large amounts of data in a data warehouse environment?

A. Amazon RDS
B. Amazon DynamoDB
C. Amazon Redshift
D. Amazon Aurora

Answer: C

Explanation:
Amazon Redshiftis a fully managed data warehouse service that enables users to analyze large amounts of data quickly and cost-effectively. It is designed specifically for online analytical processing (OLAP) and is optimized for complex queries against large datasets. Amazon Redshift uses columnar storage, data compression, and massively parallel processing (MPP) to handle petabyte-scale data warehouse environments.
A . Amazon RDS: Incorrect, as it is a managed relational database service for online transaction processing (OLTP) workloads, not specifically designed for data warehousing.
B . Amazon DynamoDB: Incorrect, as it is a NoSQL database service for fast and flexible data storage, not a data warehouse.
D . Amazon Aurora: Incorrect, as it is a MySQL- and PostgreSQL-compatible relational database designed for high performance and availability for OLTP workloads, not data warehousing.
AWS Cloud Reference:
Amazon Redshift

Question: 636

A company wants to use a serverless compute service for an application. Which AWS service will meet this requirement?

A. AWS Lambda
B. AWS Cloud Formation
C. AWS Elastic Beanstalk
D. Elastic Load Balancing

Answer: A

Explanation:

AWS Lambdais a serverless compute service that allows users to run code in response to events without provisioning or managing servers. It automatically scales the application by running code only when needed, and users are charged only for the compute time consumed. This service is ideal for applications that require event-driven compute functions.

B . AWS CloudFormation: Incorrect, as it is an infrastructure-as-code service that helps users automate the deployment of AWS resources, not a serverless compute service.

C . AWS Elastic Beanstalk: Incorrect, as it is a Platform-as-a-Service (PaaS) that still involves managing servers, even though it abstracts much of the infrastructure management.

D . Elastic Load Balancing: Incorrect, as it is a service for distributing incoming application or network traffic across multiple targets, not a serverless compute service.

AWS Cloud Reference:

AWS Lambda

Question: 637

A company wants to migrate critical on-premises production systems to Amazon EC2 instances. The production instances will be used for at least 3 years. The company wants a pricing option that will minimize cost.
Which solution will meet these requirements?

A. On-Demand Instances
B. Reserved Instances
C. Spot Instances
D. AWS Free Tier

Answer: B

Explanation:

Reserved Instances (RIs)offer a significant discount (up to 75%) compared to On-Demand pricing in exchange for committing to use AWS for a period of 1 or 3 years. This pricing model is ideal for applications with predictable usage patterns or long-term commitments, such as critical production systems that will be running for at least 3 years.

A . On-Demand Instances: Incorrect, as they provide flexibility without commitment but are more expensive over the long term.

C . Spot Instances: Incorrect, as they offer the lowest cost for non-critical, interruptible workloads but are not suitable for critical production systems due to their potential for sudden termination.

D . AWS Free Tier: Incorrect, as it provides limited free usage for certain AWS services for a short period (typically one year) and is not intended for long-term production workloads.

AWS Cloud Reference:
Amazon EC2 Reserved Instances

Question: 638

An independent software vendor wants to deliver and share its custom Amazon Machine images (AMIs) to prospective customers. Which AWS service will meet these requirements?

A. AWS Marketplace
B. AWS Data Exchange
C. Amazon EC2
D. AWS Organizations

Answer: A

Explanation:
AWS Marketplaceis an online store where independent software vendors (ISVs) can list and sell their software products, including custom Amazon Machine Images (AMIs). It allows vendors to share, distribute, and monetize their AMIs with a broad audience of AWS customers.
B . AWS Data Exchange: Incorrect, as it is a service for finding, subscribing to, and using third-party data in the cloud, not for delivering custom AMIs.
C . Amazon EC2: Incorrect, as it is the service for running instances, but it does not provide a marketplace for distributing AMIs.
D . AWS Organizations: Incorrect, as it is a service for managing multiple AWS accounts, not for distributing software products like AMIs.
AWS Cloud Reference:
AWS Marketplace

Question: 639

A company uses a third-party identity provider (IdP). The company wants to provide its employees with access to AWS accounts and services without requiring another set of login credentials. Which AWS service will meet this requirement?

A. AWS Directory Service
B. Amazon Cognito
C. AWS IAM Identity Center
D. AWS Resource Access Manager (AWS RAM)

Answer: C

Explanation:
AWS IAM Identity Center (formerly AWS Single Sign-On or AWS SSO)provides a single sign-on experience for accessing AWS accounts and applications by integrating with third-party identity providers (IdPs) like Microsoft Active Directory, Okta, or any SAML 2.0-compliant IdP. This service allows employees to log in once using their existing corporate credentials managed by the third- party IdP and gain access to multiple AWS accounts and services without needing separate AWS credentials.

Why other options are not suitable:

A . AWS Directory Service: Provides a managed Microsoft Active Directory, but does not directly support single sign-on integration with third-party IdPs.

B . Amazon Cognito: Primarily used for managing authentication for web and mobile apps, not for integrating third-party IdPs for AWS management access.

D . AWS Resource Access Manager (AWS RAM): Used for sharing AWS resources across accounts, not for identity and access management.

Reference:

AWS IAM Identity Center Documentation

Question: 640

A company runs an application on AWS that performs batch jobs. The application is fault-tolerant and can handle interruptions. The company wants to optimize the cost to run the application.

Which AWS offering will meet these requirements?

A. Amazon Macie
B. Amazon Neptune
C. Amazon EC2 Spot Instances
D. Amazon EC2 On-Demand Instances

Answer: C

Explanation:

Amazon EC2 Spot Instancesoffer spare AWS compute capacity at a significantly reduced cost compared to On-Demand Instances. Spot Instances are ideal for fault-tolerant and flexibleworkloads that can tolerate interruptions, such as batch jobs, data processing, or large-scale computations. These instances may be interrupted by AWS if there is a demand for capacity, but they provide the best cost optimization for workloads that can handle such interruptions.

Why other options are not suitable:

A . Amazon Macie: A data security and privacy service, not relevant to running batch jobs.

B . Amazon Neptune: A graph database service, not relevant to compute optimization.

D . Amazon EC2 On-Demand Instances: Provide flexible compute capacity but at a higher cost than Spot Instances, which are more suitable for cost optimization.

Reference:

Amazon EC2 Spot Instances Documentation

Question: 641

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:

Users can change their own IAM user password using:

AWS Management Console: The graphical interface allows users to navigate to the "My Security Credentials" page and change their password.

AWS Command Line Interface (CLI): Users with the appropriate permissions can use theaws iam change-passwordcommand to change their password.

Why other options are not suitable:

B . AWS Key Management Service (AWS KMS): Used for creating and managing encryption keys, not for managing user passwords.

D . AWS Resource Access Manager (AWS RAM): Used for resource sharing between accounts, not for password management.

E . AWS Secrets Manager: A service for securely storing and managing secrets, not for changing IAM user passwords.

Reference:

AWS IAM User Guide - Changing Your Own Password

Question: 642

A company is using multiple AWS accounts for different business teams. The finance team wants to receive one bill for all of the company's accounts.

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

A. AWS Organizations
B. AWS Trusted Advisor
C. Cost Explorer
D. AWS Budgets

Answer: A

Explanation:

AWS Organizationsis a service that helps users centrally manage and govern multiple AWS accounts. With AWS Organizations, a company can consolidate billing and receive a single bill for all AWS accounts under an organization, making it easier for the finance team to track costs.

Why other options are not suitable:

B . AWS Trusted Advisor: Provides real-time guidance to help optimize AWS resources, not for consolidated billing.

C . Cost Explorer: A tool for visualizing and managing AWS costs and usage, but it does not consolidate billing.

D . AWS Budgets: Allows setting custom budgets and alerts but does not consolidate billing across accounts.

Reference:

AWS Organizations Documentation

Question: 643

A company has an application that produces unstructured data continuously. The company needs to store the data so that the data is durable and easy to query.

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

A. Amazon RDS
B. Amazon Aurora
C. Amazon QuickSight
D. Amazon DynamoDB

Answer: D

Explanation: Amazon DynamoDBis a fully managed NoSQL database service that provides fast and predictable performance with seamless scalability. It is designed to handle unstructured data, offers high durability, and provides easy querying capabilities using key-value or document data models, making it suitable for applications that continuously produce unstructured data.
Why other options are not suitable:
A . Amazon RDS: A relational database service that is more suited for structured data and SQL queries.
B . Amazon Aurora: A MySQL- and PostgreSQL-compatible relational database, also more suited for structured data.
C . Amazon QuickSight: A business intelligence service for data visualization, not for storing data. Reference:
Amazon DynamoDB Documentation

Question: 644

Which pillar of the AWS Well-Architected Framework focuses on the ability to recover automatically from service Interruptions?

A. Security
B. Performance efficiency
C. Operational excellence
D. Reliability

Answer: D

Explanation:

TheReliabilitypillar of the AWS Well-Architected Framework focuses on the ability of a system to recover from failures and disruptions, dynamically acquire computing resources to meet demand, and mitigate disruptions such as misconfigurations or network issues. This includes the ability to automatically recover from service interruptions and implement mechanisms that anticipate, respond to, and prevent failures.
A . Security: Incorrect, as it focuses on protecting information, systems, and assets while delivering business value.
B . Performance efficiency: Incorrect, as it focuses on using IT and computing resources efficiently. C . Operational excellence: Incorrect, as it focuses on operations in the cloud, including monitoring and managing systems.
AWS Cloud Reference:
AWS Well-Architected Framework - Reliability Pillar

Question: 645

A company's application is running on Amazon EC2 instances. The company is planning a partial migration to a serverless architecture in the next year and wants to pay for resources up front. Which AWS purchasing option will optimize the company's costs?
A. Convertible Reserved Instances

B. Spot Instances
C. EC2 Instance Savings Plans
D. Compute Savings Plan

Answer: D

Explanation:
Compute Savings Plans provide the most flexibility and help to reduce your costs by up to 66%. These plans automatically apply to EC2 instance usage regardless of instance family, size, AZ, Region, OS or tenancy, and also apply to Fargate or Lambda usage. For example, with Compute Savings Plans, you can change from C4 to M5 instances, shift a workload from EU (Ireland) to EU (London), or move a workload from EC2 to Fargate or Lambda at any time and automatically continue to pay the Savings Plans price.
https://aws.amazon.com/savingsplans/compute-pricing/

Question: 646

Which benefit is always free of charge with AWS, regardless of a user's AWS Support plan?

A. AWS Developer Support
B. AWS Developer Forums
C. Programmatic case management
D. AWS technical account manager (TAM)

Answer: B

Explanation:
TheAWS Developer Forumsare always free to use, regardless of the user's AWS Support plan. They provide a platform for users to ask questions, share knowledge, and collaborate with the AWS community.
A . AWS Developer Support: Incorrect, as it is a paid support plan option.
C . Programmatic case management: Incorrect, as this feature is available only with certain AWS Support plans (Business and Enterprise).
D . AWS Technical Account Manager (TAM): Incorrect, as a TAM is provided only under the AWS Enterprise Support plan, which is a paid support option.
AWS Cloud Reference:
AWS Support Plans

Question: 647

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 QuickSightis a scalable, serverless, machine learning-powered business intelligence (BI) service built for the cloud. It allows users to create and publish interactive dashboards with insights powered by machine learning, such as anomaly detection and forecasting. This makes it the best option for creating interactive BI dashboards that require machine learning insights.

A . AWS Glue Studio: Incorrect, as it is an ETL (extract, transform, load) tool for preparing and transforming data, not for creating BI dashboards.

C . Amazon Redshift: Incorrect, as it is a data warehousing service that can be used with BI tools but is not specifically designed for creating and publishing dashboards.

D . Amazon Athena: Incorrect, as it is an interactive query service for analyzing data in Amazon S3 using standard SQL, not for creating BI dashboards.

AWS Cloud Reference:

Amazon QuickSight

Question: 648

A company wants to store its files in the AWS Cloud. Users need to be able to download these files directly using a public URL.
Which AWS service or feature will meet this requirement?

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

Answer: D

Explanation:

Amazon S3 (Simple Storage Service)is a scalable object storage service that allows users to store and retrieve any amount of data at any time from anywhere on the web. S3 provides the ability to make objects (files) publicly accessible via a public URL, which allows users to download files directly. S3 is highly durable, scalable, and secure, making it an ideal service for storing files that need to be accessed publicly.

Why other options are not suitable:

A . Amazon Redshift: A data warehouse service designed for complex queries and analytics, not for storing and publicly sharing files.

B . Amazon Elastic Block Store (Amazon EBS): Provides block storage for use with EC2 instances, but does not support public URLs or direct file access from the web.

C . Amazon Elastic File System (Amazon EFS): A managed file storage service for use with EC2 instances, not designed for public access via URLs.

Reference:

Amazon S3 Documentation

Question: 649

Which AWS service allows for file sharing between multiple Amazon EC2 Instances?

A. AWS Direct Connect
B. AWS Snowball Edge
C. AWS Backup
D. Amazon Elastic File System (Amazon EFS)

Answer: D

Explanation:
Amazon Elastic File System (Amazon EFS)is a scalable, fully managed, shared file storage service that is accessible from multiple Amazon EC2 instances. EFS is designed to provide highly available and durable file storage that can be mounted across multiple instances, making it ideal for shared file access.
Why other options are not suitable:
A . AWS Direct Connect: A network service to establish a dedicated network connection to AWS, not a file-sharing solution.
B . AWS Snowball Edge: A data transfer device for migrating data to and from AWS, not for ongoing file sharing.
C . AWS Backup: A service for centralized backup management, not for sharing files between EC2 instances.
Reference:
Amazon EFS Documentation

Question: 650

A company wants to migrate its on-premises infrastructure to the AWS Cloud. Which advantage of cloud computing will help the company reduce upfront costs?

A. Go global in minutes
B. Increase speed and agility
C. Benefit from massive economies of scale
D. Trade fixed expense for variable expense

Answer: D

Explanation: One of the key advantages of cloud computing is the ability totrade fixed expenses (like data centers and physical servers) for variable expenses. With AWS, companies pay only for the compute power, storage, and other resources they use, rather than investing heavily in on-premises infrastructure upfront. This reduces capital expenditure and helps lower overall costs.
Why other options are not suitable:
A . Go global in minutes: Refers to the ability to quickly deploy applications around the world, but does not directly reduce upfront costs.
B . Increase speed and agility: Refers to the ability to quickly scale and deploy resources, but is not directly related to cost reduction.
C . Benefit from massive economies of scale: Refers to AWS's ability to lower costs by leveraging its scale, but the primary advantage for reducing upfront costs is moving from fixed to variable expenses.
Reference:
AWS Cloud Economics

Question: 651

Which AWS service or tool gives a company the ability to release application changes in an automated way?

A. Amazon AppFlow
B. AWS CodeDeploy
C. AWS PrivateLink
D. Amazon EKS Distro

Answer: B

Explanation:
AWS CodeDeployis a fully managed deployment service that automates software deployments to a variety of compute services such as Amazon EC2, AWS Lambda, and on-premises servers. CodeDeploy makes it easy to release new features, helps avoid downtime during application deployment, and handles the complexity of updating applications in a scalable and reliable manner. Why other options are not suitable:

A . Amazon AppFlow: A service to automate data flows between AWS and SaaS applications, not for application deployment.

C . AWS PrivateLink: A service to securely access AWS services from your virtual private cloud (VPC), not related to application deployments.

D . Amazon EKS Distro: A Kubernetes distribution for running Kubernetes clusters, not specifically designed for automated deployments.

Reference:
AWS CodeDeploy Documentation

Question: 652

A company has multiple SQL-based databases located in a data center. The company needs to migrate all database servers to the AWS Cloud to reduce the cost of operating physical servers.
Which AWS service or resource will meet these requirements with the LEAST operational overhead?

A. Amazon EC2 instances
B. Amazon RDS
C. Amazon DynamoDB
D. OpenSearch

Answer: B

Explanation:
Amazon RDS (Relational Database Service)is a managed database service that simplifies the setup, operation, and scaling of relational databases in the AWS Cloud. It supports several SQL-based databases, such as MySQL, PostgreSQL, Oracle, SQL Server, and MariaDB, while automating common administrative tasks such as backups, patch management, and scaling. This significantly reduces the operational overhead compared to managing databases on Amazon EC2 instances, which would require manual management of the operating system, database software, backups, and maintenance.

A . Amazon EC2 instances: Incorrect, as this would involve more operational overhead, including manual management of the database server.

C . Amazon DynamoDB: Incorrect, as it is a NoSQL database service, not suitable for migrating SQLbased databases.

D . OpenSearch: Incorrect, as it is used for search and analytics workloads, not for hosting SQL-based databases.

AWS Cloud Reference:

Amazon RDS

Question: 653

Which AWS service or feature provides a firewall at the subnet level within a VPC?

A. Security group
B. Network ACL
C. Elastic network interface
D. AWS WAF

Answer: B

Explanation:

ANetwork ACL (Access Control List)is a stateless firewall that controls inbound and outbound traffic at the subnet level within a VPC. It provides an additional layer of security to the VPC by allowing or denying traffic to and from a subnet based on defined rules.

A . Security group: Incorrect, as security groups act as a firewall at the instance level, not the subnet level.

C . Elastic network interface: Incorrect, as it is a virtual network interface that you can attach to an instance, not a firewall feature.

D . AWS WAF: Incorrect, as it is a web application firewall that protects web applications from common exploits, not for subnet-level protection.

AWS Cloud Reference: Network ACLs

Question: 654

A company wants to add a conversational chatbot to its website. Which AWS service can the company use to meet this requirement?

A. Amazon Textract
B. Amazon Lex
C. AWS Glue
D. Amazon Rekognition

Answer: B

Explanation:

Amazon Lexis a service for building conversational interfaces into any application using voice and text. It provides the tools necessary to create, test, and publish intelligent chatbots that can be integrated with web applications to provide a conversational experience.

A . Amazon Textract: Incorrect, as it is a service for extracting text and data from documents, not for creating chatbots.

C . AWS Glue: Incorrect, as it is a data integration service used for preparing and transforming data. D . Amazon Rekognition: Incorrect, as it is an image and video analysis service.

AWS Cloud Reference:

Amazon Lex

Question: 655

A company has batch workloads that need to run for short periods of time on Amazon EC2. The workloads can handle interruptions and

can start again from where they ended.

What is the MOST cost-effective EC2 instance purchasing option to meet these requirements?

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

Answer: B

Explanation:

Spot Instancesare the most cost-effective EC2 instance purchasing option for batch workloads that can handle interruptions and can be restarted from where they left off. Spot Instances offer significant cost savings (up to 90%) over On-Demand Instances by using spare EC2 capacity. They are ideal for workloads that are fault-tolerant and flexible in terms of timing.

A . Reserved Instances: Incorrect, as they require a long-term commitment and do not offer the flexibility needed for short-term, interruptible workloads.

C . Dedicated Instances: Incorrect, as they are designed to run on hardware dedicated to a single customer, which is more expensive.

D . On-Demand Instances: Incorrect, as they are more expensive than Spot Instances and are used for applications that require immediate or predictable capacity.

AWS Cloud Reference:

Amazon EC2 Spot Instances

Question: 656

An AWS user wants to proactively detect when an instance or account might be compromised or if there are threats from attacks. Which AWS service should the user choose?

A. Amazon GuardDuty
B. AWS WAF
C. AWS Shield
D. Amazon Inspector

Answer: A

Explanation:

Amazon GuardDutyis a threat detection service that continuously monitors AWS accounts and workloads for malicious activity and unauthorized behavior to help protect your AWS resources. It uses machine learning, anomaly detection, and integrated threat intelligence to detect when an instance or account might be compromised or if there are threats from attacks.

B . AWS WAF: Incorrect, as it is a web application firewall that protects against common web exploits but does not provide comprehensive threat detection.

C . AWS Shield: Incorrect, as it provides protection against DDoS attacks but does not detect compromises within AWS accounts.

D . Amazon Inspector: Incorrect, as it is a service that helps improve the security and compliance of applications deployed on AWS by assessing for vulnerabilities, not for threat detection.

AWS Cloud Reference:

Amazon GuardDuty

Question: 657

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

A. Patch the Amazon DynamoDB operating system.
B. Secure Amazon CloudFront edge locations by allowing physical access according to the principle of least privilege.
C. Protect the hardware that runs AWS services.
D. Use AWS Identity and Access Management (1AM) according to the principle of least privilege.

Answer: D

Explanation:
According to theAWS Shared Responsibility Model, AWS is responsible for the security "of" the cloud (such as protecting the infrastructure, including hardware, software, networking, and facilities that run AWS Cloud services). In contrast, customers are responsible for security "in" the cloud. This includes configuring and using AWS services securely.
D . Use AWS Identity and Access Management (IAM) according to the principle of least privilegeis a customer's responsibility. Customers must manage their credentials, control access to resources, and ensure that IAM policies follow the principle of least privilege, which means granting only the permissions necessary to perform a task.
Why other options are not suitable:
A . Patch the Amazon DynamoDB operating system: AWS is responsible for managing and patching the infrastructure for managed services like DynamoDB.
B . Secure Amazon CloudFront edge locations by allowing physical access according to the principle of least privilege: AWS handles physical security of its edge locations.
C . Protect the hardware that runs AWS services: AWS is responsible for protecting the physical hardware that runs AWS services.
Reference:
AWS Shared Responsibility Model

Question: 658

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:
AVPC (Virtual Private Cloud)in AWS is a logically isolated network that you define in the AWS Cloud. Within a VPC, you can create subnets, route tables, network gateways, and more.
D . Internet Gateway: An internet gateway is a component that allows communication between resources in a VPC and the internet.
E . Subnet: A subnet is a range of IP addresses in your VPC. Subnets can be public or private and are essential for organizing

resources within a VPC.

Why other options are not suitable:

A . Amazon API Gateway: Used for creating and managing APIs, not a direct component of a VPC.

B . Amazon S3 buckets and objects: Amazon S3 is a storage service; its resources are globally accessible and not confined to a VPC.

C . AWS Storage Gateway: A hybrid cloud storage service, not a core component of a VPC. Reference:

Amazon VPC Documentation

Question: 659

A company wants to build, tram, and deploy machine learning (ML) models. Which AWS service can the company use to meet this requirement?

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

Answer: D

Explanation:

Amazon SageMakeris a fully managed service that provides every developer and data scientist with the ability to build, train, and deploy machine learning (ML) models quickly. It offers integrated Jupyter notebooks for data exploration, pre-built algorithms, and automatic model tuning to optimize hyperparameters.

Why other options are not suitable:

A . Amazon Personalize: A managed machine learning service that provides personalized recommendations, but does not offer the full suite of ML tools for building, training, and deploying models.

B . Amazon Comprehend: A natural language processing (NLP) service, not a general-purpose ML model building and deployment service.

C . Amazon Forecast: A managed service specifically for creating accurate forecasts using machine learning, not for general-purpose ML model development.

Reference:

Amazon SageMaker Documentation

Question: 660

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

A. Patch a guest operating system that is deployed on an Amazon EC2 instance.
B. Control physical access to an AWS data center
C. Control access to AWS underlying hardware.
D. Patch a host operating system that is deployed on Amazon S3.

Answer: A

Explanation:
According to theAWS Shared Responsibility Model, the customer is responsible for managing the guest operating system (including patches and updates) and any applications that they run on their Amazon EC2 instances.
Why other options are not suitable:

B . Control physical access to an AWS data center: AWS is responsible for the physical security of its data centers.
C . Control access to AWS underlying hardware: AWS manages the hardware infrastructure.
D . Patch a host operating system that is deployed on Amazon S3: Amazon S3 is a managed storage service, and AWS manages the underlying infrastructure, including the operating system. Reference:
AWS Shared Responsibility Model

Question: 661

A company wants to design its cloud architecture so that it can support development innovations, and continuously improve processes and procedures.
This is an example of which pillar of the AWS Well-Architected Framework?

A. Security
B. Performance efficiency
C. Operational excellence
D. Reliability

Answer: D

Explanation:
"The Reliability pillar encompasses the ability of a workload to perform its intended function correctly and consistently when it's expected to. This includes the ability to operate and test the workload through its total lifecycle."
https://docs.aws.amazon.com/wellarchitected/latest/framework/reliability.html

Question: 662

A company runs its production workload in the AWS Cloud. The company needs to choose one of the AWS Support Plans.
Which of the AWS Support Plans will meet these requirements at the LOWEST cost?

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

Answer: A

Explanation:
TheAWS Developer Supportplan is the lowest-cost AWS Support plan that provides basic guidance and technical support for running

production workloads. It offers 24/7 access to AWS customer service, documentation, whitepapers, and access to a limited set of Trusted Advisor checks. This plan is ideal for non-critical workloads or early development phases but provides lower levels of support compared to the Business, Enterprise On-Ramp, and Enterprise plans.

B . Enterprise On-Ramp: Incorrect, as it is a higher-cost support plan designed for production workloads needing more guidance and technical support.

C . Enterprise: Incorrect, as it is the most expensive support plan, providing a Technical Account Manager and other premium support features.

D . Business: Incorrect, as it provides comprehensive support for production workloads but is more expensive than the Developer plan.

AWS Cloud Reference:

AWS Support Plans

Question: 663

A company is planning to move data backups to the AWS Cloud. The company needs to replace on premises storage with storage that is cloud-based but locally cached.

Which AWS service meets these requirements?

A. AWS Storage Gateway
B. AWS Snowcone
C. AWS Backup
D. Amazon Elastic File System (Amazon EFS)

Answer: A

Explanation:

AWS Storage Gatewayis a hybrid cloud storage service that provides on-premises access to virtually unlimited cloud storage. TheFile Gatewayconfiguration allows for on-premises applications to store data in Amazon S3 using NFS and SMB protocols, while keeping frequently accessed data cached locally. This meets the company's requirement for cloud-based storage that is locally cached. B . AWS Snowcone: Incorrect, as it is a portable edge computing and storage device, not a storage service that provides local caching for cloud storage.

C . AWS Backup: Incorrect, as it is a centralized backup service to automate data protection across AWS services but does not provide local caching.

D . Amazon Elastic File System (Amazon EFS): Incorrect, as it provides scalable file storage for use with Amazon EC2 but does not offer local caching for on-premises storage needs.

AWS Cloud Reference:

AWS Storage Gateway

Question: 664

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:

AnIAM useris created when the company needs to provide unique credentials (username and password) to individuals who need access to the AWS Management Console or programmatic access (using access keys) to AWS services.

B . When the company creates AWS access credentials for individuals: Correct, as an IAM user is created to provide credentials for specific individuals.

D . When the company needs to add users to IAM groups: Correct, as IAM users can be added to groups to apply permissions and policies at a group level.

A . When an application that runs on Amazon EC2 instances requires access to other AWS services: Incorrect, as an IAM role is more appropriate for applications running on EC2 to assume temporary credentials.

C . When the company creates an application that runs on a mobile phone that makes requests to AWS: Incorrect, as using Cognito or a role with temporary credentials is more suitable.

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: Incorrect, as this use case typically involves IAM roles combined with AWS Single Sign-On (SSO).

AWS Cloud Reference:

IAM Users and Groups

IAM Roles

Question: 665

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

A. Secure the virilization layer.

B. Encrypt data and maintain data integrity.

C. Patch the Amazon RDS operating system.

D. Maintain identity and access management controls.

E. Secure Availability Zones.

Answer: B,D

Explanation:

Under theAWS Shared Responsibility Model, AWS manages securityofthe cloud (such as physical infrastructure and virtualization), while customers are responsible for securityinthe cloud. This means that customers are responsible for:

B . Encrypt data and maintain data integrity: Correct, as customers are responsible for securing their data, including encryption and maintaining its integrity.

D . Maintain identity and access management controls: Correct, as customers are responsible for managing access to their AWS resources, including creating and managing IAM users, roles, and permissions.

A . Secure the virtualization layer: Incorrect, as AWS is responsible for securing the underlying virtualization layer.

C . Patch the Amazon RDS operating system: Incorrect, as AWS handles patching and maintenance of the managed service's underlying infrastructure.

E . Secure Availability Zones: Incorrect, as AWS is responsible for securing the physical infrastructure, including Availability Zones.

AWS Cloud Reference:
AWS Shared Responsibility Model

Question: 666

A company needs stateless network filtering for its VPC.
Which AWS service, tool, or feature will meet this requirement?

A. AWS PrivateLink
B. Security group
C. Network access control list (ACL)
D. AWS WAF

Answer: C

Explanation:
A Network Access Control List (ACL) is a stateless network filtering mechanism provided by AWS for controlling traffic in and out of subnets within a VPC. Unlike security groups, which are stateful, network ACLs are stateless. This means that they do not automatically allow responses to inbound traffic unlessexplicitly specified. Network ACLs allow you to set rules for both inbound and outbound traffic, making them suitable for stateless filtering. Security groups, on the other hand, are stateful, while AWS WAF is primarily for web application-level security. AWS PrivateLink is used for privately connecting VPCs to AWS services without using an internet gateway. Therefore, for stateless network filtering, Network ACL is the correct choice.

Question: 667

Which AWS compute service gives users the ability to securely and reliably run containers at scale?

A. Amazon Elastic Container Service (Amazon ECS)
B. Amazon Aurora
C. Amazon Athena
D. Amazon Polly

Answer: A

Explanation:
Amazon Elastic Container Service (Amazon ECS) is a highly scalable, fast container management service that allows you to run, stop, and manage Docker containers on a cluster of Amazon EC2 instances or a serverless environment powered by AWS Fargate. It provides secure, reliable, and scalable operations for containers, allowing users to deploy containers in a highly secure and reliable manner. Amazon Aurora is a relational database service, Amazon Athena is an interactive query service, and Amazon Polly is a text-to-speech service. Therefore, the correct answer is Amazon ECS. Reference:
AWS Documentation on Amazon ECS

Question: 668

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

A. Users can implement all AWS services in seconds.

B. AWS assumes all responsibility for the security of infrastructure and applications.
C. Users experience increased speed and agility.
D. Users benefit from massive economies of scale.
E. Users can move hardware from their data center to the AWS Cloud.

Answer: C,D

Explanation:

Moving to the AWS Cloud offers several advantages, including increased speed and agility, which allows users to experiment, innovate, and iterate faster by using the global AWS infrastructure. Additionally, AWS offers massive economies of scale due to its large customer base, leading to lower pay-as-you-go prices. AWS does not assume all responsibilities for the security of infrastructure and applications; it follows a shared responsibility model. Also, not all AWS services can be implemented in seconds, and physical hardware cannot be moved from a user's data center to the AWS Cloud. Reference:
AWS Cloud Benefits

Question: 669

A company needs to provision uninterruptible Amazon EC2 instances, when needed, and pay for compute capacity by the second. Which EC2 instance purchasing option will meet these requirements?

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

Answer: C

Explanation:
On-Demand Instances are ideal for workloads that require flexibility, offering compute capacity with no upfront cost, allowing users to pay by the second for the instances they use. This makes it suitable for applications with short-term, spiky, or unpredictable workloads that cannot be interrupted. Reserved Instances require long-term commitments, Spot Instances can be interrupted, and Dedicated Instances are for specific hardware requirements, making On-Demand the correct choice for uninterruptible and flexible instance usage.
Reference:
AWS Documentation on EC2 On-Demand Instances

Question: 670

An administrator observed that multiple AWS resources were deleted yesterday.
Which AWS service will help identify the cause and determine which user deleted the resources?

A. AWS CtoudTrail
B. Amazon Inspector
C. Amazon GuardDuty
D. AWS Trusted Advisor

Answer: A

Explanation:
AWS CloudTrailis a service that enables governance, compliance, and operational and risk auditing of your AWS account. With CloudTrail, you can log, continuously monitor, and retain account activity related to actions across your AWS infrastructure. CloudTrail logs provide a history of AWS API calls for your account, including those made by the AWS Management Console, AWS SDKs, command-line tools, and other AWS services. In this case, AWS CloudTrail will help the administrator identify which user deleted the resources by reviewing the event history that records details such as which user performed the action, the time of the action, and which resources were affected.
B . Amazon Inspector: Incorrect, as it is a security assessment service that helps identify vulnerabilities and deviations from best practices, not for tracking user activity.
C . Amazon GuardDuty: Incorrect, as it is a threat detection service that monitors malicious activity and unauthorized behavior, not specifically for tracking changes made by users.
D . AWS Trusted Advisor: Incorrect, as it provides best practices and guidance for cost optimization, security, fault tolerance, and performance, not for logging user actions.
AWS Cloud Reference:
AWS CloudTrail

Question: 671

A company plans to launch an ecommerce website that contains many images for a product catalog. The company wants to keep the cost of running the website within a specific budget.
Which AWS service or tool should the company use to monitor the ongoing costs of the website?

A. AWS Cost Explorer
B. AWS SDKs
C. EC2 Image Builder
D. AWS CloudFormation

Answer: A

Explanation:
AWS Cost Exploreris a tool that enables you to view and analyze your costs and usage. It provides an easy-to-use interface to visualize, understand, and manage AWS costs and usage over time. Cost Explorer allows the company to monitor ongoing costs, set budget alerts, and analyze cost drivers, which is ideal for keeping the cost of running an ecommerce website within a specific budget. B . AWS SDKs: Incorrect, as they are software development kits that allow integration with AWS services but do not provide cost monitoring or management capabilities.
C . EC2 Image Builder: Incorrect, as it is a service for automating the creation of virtual machine images, not for monitoring costs.
D . AWS CloudFormation: Incorrect, as it is used for provisioning and managing infrastructure as code, not for cost monitoring.
AWS Cloud Reference:
AWS Cost Explorer

Question: 672

Which of the following is a customer responsibility according to the AWS shared responsibility model?

A. Apply security patches for Amazon S3 infrastructure devices.

B. Provide physical security for AWS datacenters.
C. Install operating system updates on Lambda@Edge.
D. Implement multi-factor authentication (MFA) for 1AM user accounts.

Answer: D

Explanation:
According to theAWS Shared Responsibility Model, customers are responsible for managing their data, classifying their assets, and using AWS's identity and access management tools to apply the appropriate permissions. Implementingmulti-factor authentication (MFA) for IAM user accountsis a customer responsibility to enhance the security of their accounts and protect against unauthorized access.
A . Apply security patches for Amazon S3 infrastructure devices: Incorrect, as AWS is responsible for managing and patching the underlying infrastructure, including storage devices.
B . Provide physical security for AWS data centers: Incorrect, as AWS is responsible for the physical security of its data centers.
C . Install operating system updates on Lambda@Edge: Incorrect, as AWS manages the underlying infrastructure for Lambda@Edge, including operating system updates.
AWS Cloud Reference:
AWS Shared Responsibility Model

Question: 673

Which AWS service can be used at no additional cost?

A. Amazon SageMaker
B. AWS Config
C. AWS Organizations
D. Amazon CloudWatch

Answer: C

Explanation:
AWS Organizationsis a no-cost service that helps you centrally manage and govern your environment as you grow and scale your AWS resources. With AWS Organizations, you can create new AWS accounts, invite existing accounts to join your organization, and apply policies to groups of accounts for governance. While some features within AWS services incur additional costs, using AWS Organizations itself does not add any direct costs.
A . Amazon SageMaker: Incorrect, as it is a fully managed service for building, training, and deploying machine learning models, and it is not free.
B . AWS Config: Incorrect, as while it offers some free-tier usage, it generally incurs charges for recording and evaluating configuration changes.
D . Amazon CloudWatch: Incorrect, as certain CloudWatch metrics, custom metrics, and alarms have associated costs beyond the free tier.
AWS Cloud Reference:
AWS Organizations

Question: 674

A company migrated its systems to the AWS Cloud. The systems are rightsized, and a security review did not reveal any issues. The company must ensure that additional developments, integrations, changes, and system usage growth do not jeopardize this optimized AWS infrastructure.

Which AWS service should the company use to report ongoing optimization and security?

A. AWS Trusted Advisor
B. AWS Health Dashboard
C. Amazon Connect
D. AWS Systems Manager

Answer: A

Explanation:

AWS Trusted Advisoris a service that provides real-time guidance to help you provision your resources following AWS best practices. It offers checks in five categories: cost optimization, performance, security, fault tolerance, and service limits. By continuously monitoring the AWS environment, AWS Trusted Advisor helps ensure that additional developments, integrations, changes, and growth do not jeopardize the optimized infrastructure by providing ongoing optimization and security recommendations.

B . AWS Health Dashboard: Incorrect, as it provides personalized view of service health and status updates but does not offer continuous optimization and security advice.

C . Amazon Connect: Incorrect, as it is a contact center service, not related to infrastructure optimization or security.

D . AWS Systems Manager: Incorrect, as it provides operational insights and management for AWS resources but is not specifically focused on ongoing optimization and security checks.

AWS Cloud Reference:

AWS Trusted Advisor

Question: 675

Which AWS service can migrate Amazon EC2 instances from one AWS Region to another?

A. AWS Application Migration Service
B. AWS Database Migration Service (AWS DMS)
C. AWS DataSync
D. AWS Migration Hub

Answer: A

Explanation:

AWS Application Migration Service (AWS MGN) is the primary service for migrating Amazon EC2 instances and on-premises servers to AWS, and it supports migration from one AWS Region to another. It automates the process of replicating and converting the source servers (including their applications, configurations, and data) to run natively on AWS. AWS Database Migration Service (DMS) is for migrating databases, AWS DataSync is for data transfer, and AWS Migration Hub provides a single location to track migration tasks across multiple AWS services but does not perform the migration itself.

Question: 676

Which of the following is an advantage that the AWS Cloud provides to users?

A. Users eliminate the need to guess about infrastructure capacity requirements.
B. Users decrease their variable costs by maintaining sole ownership of IT hardware.
C. Users maintain control of underlying IT infrastructure hardware.
D. Users maintain control of operating systems for managed services.

Answer: A

Explanation:
One of the primary advantages of the AWS Cloud is the ability to provision resources on demand. Users no longer need to over-provision or guess the capacity needed for infrastructure, which helps optimizecosts and improve agility. AWS handles infrastructure scaling dynamically based on the actual usage. Options B, C, and D are incorrect as users do not maintain sole ownership of IT hardware or control of underlying infrastructure, and while they do maintain some control over the operating system in some cases, it is not an advantage specific to the cloud.

Question: 677

A company wants a cost-effective option when running its applications in an Amazon EC2 instance for short time periods. The applications can be interrupted.
Which EC2 instance type will meet these requirements?

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

Answer: A

Explanation:
Spot Instances provide a cost-effective option for running Amazon EC2 instances for workloads that can tolerate interruptions. They allow you to use unused EC2 capacity in the AWS Cloud at a discounted price, making them suitable for applications that are flexible in terms of start and stop times. On-Demand Instances are more expensive and do not provide discounts for short-term, interruptible workloads, while Reserved and Dedicated Instances are for long-term or specific hardware needs.

Question: 678

A company runs an uninterruptible Amazon EC2 workload on AWS 24 hours a day. 7 days a week. The company will require the same instance family and instance type to run the workload for the next 12 months.
Which combination of purchasing options should the company choose to MOST optimize costs? (Select TWO.)

A. Standard Reserved Instance
B. Convertible Reserved Instance
C. Compute Savings Plan

D. Spot Instance
E. All Upfront payment

Answer: A,E

Explanation:

For workloads running 24/7 for a year, Standard Reserved Instances provide a significant discount compared to On-Demand pricing. Choosing the "All Upfront" payment option maximizes the cost savings as AWS offers the highest discount for upfront payments. Convertible Reserved Instances provide flexibility to change the instance type but usually at a slightly higher cost than Standard Reserved Instances. Compute Savings Plans offer cost savings, but in this scenario, the best optimization would be a combination of Standard Reserved Instances with All Upfront payment. Spot Instances are not suitable due to their interruptible nature.

Reference:
AWS EC2 Reserved Instances
AWS Savings Plans

Question: 679

What is the recommended use case for Amazon EC2 On-Demand Instances?

A. A steady-state workload that requires a particular EC2 instance configuration for a long period of time
B. A workload that can be interrupted for a project that requires the lowest possible cost
C. An unpredictable workload that does not require a long-term commitment
D. A workload that is expected to run for longer than 1 year

Answer: C

Explanation:

Amazon EC2 On-Demand Instances are ideal for unpredictable workloads that do not require a longterm commitment. This option allows users to pay for compute capacity by the hour or second, with no long-term commitments. It is suitable for short-term, spiky, or unpredictable workloads that cannot be interrupted and require flexibility. Option A is better suited for Reserved Instances, Option B is ideal for Spot Instances, and Option D is more suited for Reserved Instances due to cost optimization for long-term use.
"There is no long-term commitment required when you purchase On-Demand Instances. You pay only for the seconds that your On-Demand Instances are in the running state, with a 60-second minimum. The price per second for a running On-Demand Instance is fixed, and is listed on the Amazon EC2 Pricing, On-Demand Pricing page. We recommend that you use On-Demand Instances for applications with short-term, irregular workloads that cannot be interrupted."https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-on-demand- instances.html

Question: 680

Which advantage of cloud computing allows users to scale resources up and down based on the amount of load that an application supports?

A. Go global in minutes
B. Stop guessing capacity
C. Benefit from massive economies of scale
D. Trade fixed expense for variable expense

Answer: B

Explanation:
The ability to scale resources up and down based on application load is tied to the advantage of "Stop guessing capacity." This feature allows users to scale resources automatically to match the demand, thereby avoiding over-provisioning or under-provisioning of resources. "Go global in minutes," "Benefit from massive economies of scale," and "Trade fixed expense for variable expense" are other advantages of cloud computing, but they do not specifically address scaling resources up and down based on load.

Question: 681

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 (DMS) and AWS Schema Conversion Tool are the primary services for migrating a commercial relational database to an Amazon-managed open-source database. AWS DMS helps migrate the data, while the AWS Schema Conversion Tool converts the database schema from the source database format to the target format, including SQL code. AWS SDKs are for software development, AWS Systems Manager is for operational management, and Amazon EMR is for big data processing, which are not relevant to this use case.

Question: 682

A company has a workload that will run continuously for 1 year. The workload cannot tolerate service interruptions.
Which Amazon EC2 purchasing option will be MOST cost-effective?

A. All Upfront Reserved Instances
B. Partial Upfront Reserved Instances
C. Dedicated Instances
D. On-Demand Instances

Answer: A

Explanation:
All Upfront Reserved Instances offer the most cost-effective solution for a workload that will run continuously for one year without interruption. By paying upfront, the user receives the maximum discount over the On-Demand pricing model. Partial Upfront Reserved Instances and Dedicated Instances are more expensive than All Upfront Reserved Instances. On-Demand Instances are not cost-effective for continuous long-term workloads due to their higher hourly rates.

Question: 683

A company wants to discover, prepare, move, and integrate data from multiple sources for data analytics and machine learning. Which AWS serverless data integration service should the company use to meet these requirements?

A. AWS Glue
B. AWS Data Exchange
C. Amazon Athena
D. Amazon EMR

Answer: A

Explanation:
AWS Glue is a serverless data integration service designed to discover, prepare, move, and integrate data from multiple sources for data analytics and machine learning purposes. It provides a managed ETL (Extract, Transform, Load) service that is ideal for preparing and transforming data for analytics. AWS Data Exchange is used for finding and subscribing to third-party data, Amazon Athena is for querying data stored in Amazon S3 using SQL, and Amazon EMR is for big data processing using Apache Hadoop and Spark, but AWS Glue is specifically designed for data integration and preparation tasks.

Question: 684

Which AWS service or tool provides a visualization of historical AWS spending patterns and projections of future AWS costs?

A. AWS Cos! and Usage Report
B. AWS Budgets
C. Cost Explorer
D. Amazon CloudWatch

Answer: C

Explanation:
AWS Cost Explorer provides a visualization of historical AWS spending patterns and allows users to project future costs based on past usage. It offers advanced filtering and grouping features, enabling users to analyze costs and usage at a granular level. The AWS Cost and Usage Report provides detailed AWS cost and usage data but does not offer visualization or future cost projections. AWS Budgets is used for setting custom cost and usage budgets and receiving alerts. Amazon CloudWatch is for monitoring AWS resources and applications, not for cost management.

Question: 685

Which AWS service integrates with other AWS services to provide the ability to encrypt data at rest?

A. AWS Key Management Service (AWS KMS)
B. AWS Certificate Manager (ACM)
C. AWS Identity and Access Management (1AM)
D. AWS Security Hub

Answer: A

Explanation:
AWS Key Management Service (AWS KMS) is designed to integrate with various AWS services to encrypt data at rest. It provides a secure and highly available service to create, control, and manage encryption keys used to encrypt your data. AWS Certificate Manager (ACM) is for managing SSL/TLS certificates, AWS Identity and Access Management (IAM) is for managing user access and permissions, and AWS Security Hub is for security monitoring and compliance, but none of these services provide data encryption at rest like AWS KMS.

Question: 686

A company has an AWS Business Support plan. The company needs to gain access to the AWS DDoS Response Team (DRT) to help mitigate DDoS events.
Which AWS service or resource must the company use to meet these requirements?

A. AWS Shield Standard
B. AWS Enterprise Support
C. AWS WAF
D. AWS Shield Advanced

Answer: D

AWS Shield Advanced provides enhanced protection against DDoS attacks and includes access to the AWS DDoS Response Team (DRT) to help mitigate complex DDoS events. AWS Shield Standard offers basic DDoS protection, which is included with AWS services, but does not provide access to the DRT. AWS WAF is a web application firewall, and AWS Enterprise Support is a premium support plan but does not specifically provide DDoS mitigation services or access to the DRT.

Question: 687

To assist companies with Payment Card Industry Data Security Standard (PCI DSS) compliance in the cloud. AWS provides:

A. physical inspections of data centers by appointment.
B. required PCI compliance certifications for any application running on AWS.
C. an AWS Attestation of Compliance (AOC) report for specific AWS services.
D. professional PCI compliance services.

Answer: C

Explanation:
AWS provides an Attestation of Compliance (AOC) report for specific AWS services to assist companies in achieving Payment Card Industry Data Security Standard (PCI DSS) compliance in the cloud. This report demonstrates that AWS services meet the necessary PCI DSS requirements. AWS does not offer physical inspections of data centers by appointment, nor does it provide certifications for any application running on AWS. Additionally, AWS does not provide professional PCI compliance services; companies must manage their PCI compliance in their environment.

Question: 688

Which programming languages does AWS Cloud Development Kit (AWS CDK) currently support? (Select TWO.)

A. Python
B. Swift
C. TypeScript
D. Ruby
E. PHP

Answer: A,C

Explanation:
TheAWS Cloud Development Kit (AWS CDK)currently supports multiple programming languages, includingPythonandTypeScript. These languages allow developers to define cloud infrastructure using familiar programming constructs. Python and TypeScript are among the first languages supported by AWS CDK, which also supports Java, C#, and JavaScript. This enables developers to use their existing programming skills and tools to define cloud infrastructure in code.
B . Swift: Incorrect, as Swift is not currently supported by AWS CDK.
D . Ruby: Incorrect, as Ruby is not currently supported by AWS CDK.
E . PHP: Incorrect, as PHP is not currently supported by AWS CDK.

AWS Cloud Reference:

AWS Cloud Development Kit (AWS CDK)

Question: 689

Which AWS Cloud design principle is a company using when the company implements AWS CloudTrail?

A. Activate traceability.
B. Use serverless compute architectures.
C. Perform operations as code.
D. Go global in minutes.

Answer: A

Explanation:
By implementingAWS CloudTrail, a company is adhering to the AWS Cloud design principle ofactivating traceability. AWS CloudTrail provides detailed logs of all API calls made in an AWS account, which helps monitor, troubleshoot, and detect unusual activity, thereby improving security and compliance. This supports the principle of "activating traceability" by enabling continuous monitoring and auditing of all actions and changes within the AWS environment.
B . Use serverless compute architectures: Incorrect, as this principle encourages the use of managed services that handle infrastructure, such as AWS Lambda, and is not directly related to CloudTrail.
C . Perform operations as code: Incorrect, as this principle emphasizes the use of code and automation for infrastructure management.
D . Go global in minutes: Incorrect, as this principle relates to the global deployment of applications and services.
AWS Cloud Reference:
AWS Well-Architected Framework
AWS CloudTrail

Question: 690

A company stores data in an Amazon S3 bucket. Which task is the responsibility of AWS?

A. Configure an S3 Lifecycle policy.
B. Activate S3 Versioning.
C. Configure S3 bucket policies.
D. Protect the infrastructure that supports S3 storage.

Answer: D

Explanation:
According to theAWS Shared Responsibility Model, AWS is responsible for protecting the infrastructure that runs all of the services offered in the AWS Cloud, including Amazon S3. This infrastructure includes hardware, software, networking, and facilities that run AWS services.
A . Configure an S3 Lifecycle policy: Incorrect, as configuring S3 Lifecycle policies to manage object lifecycle (e.g., transitioning objects to different storage classes or deleting them after a certain period) is the customer's responsibility.
B . Activate S3 Versioning: Incorrect, as enabling S3 Versioning is a customer responsibility for managing data protection.
C . Configure S3 bucket policies: Incorrect, as setting and managing S3 bucket policies to control access is the customer's responsibility.
AWS Cloud Reference:
AWS Shared Responsibility Model Amazon S3

Question: 691

A company is building a new application on AWS. The company needs the application to remain available if an individual application component fails.
Which design principle should the company use to meet this requirement?

A. Disposable resources

B. Automation
C. Rightsizing
D. Loose coupling

Answer: D

Explanation:
Loose coupling is a design principle that involves reducing dependencies between application components so that they can operate independently. This approach ensures that the failure of one component does not affect the availability of the others, thereby improving the application's fault tolerance and resilience. Disposable resources, automation, and rightsizing are valuable principles in cloud architecture, but they do not directly address the requirement of remaining available despite the failure of an individual component like loose coupling does.

Reference:
AWS Well-Architected Framework - Design Principles

Question: 692

Under the AWS shared responsibility model, which of the following is a responsibility of the customer?

A. Shred disk drives before they leave a data center.
B. Prevent customers from gathering packets or collecting traffic at the hypervisor level.
C. Patch the guest operating system with the latest security patches.
D. Maintain security systems that provide physical monitoring of data centers.

Answer: C

Explanation:
Under the AWS shared responsibility model, AWS is responsible for the security "of" the cloud, which includes the physical infrastructure, networking, and hypervisor layer. The customer, however, is responsible for security "in" the cloud, which includes managing the security of their data, patching and maintaining their guest operating system and applications, and managing identity and access. The responsibilities of shredding disk drives, preventing packet capture at the hypervisor level, and physical monitoring are handled by AWS as part of its responsibility for security "of" the cloud.

Question: 693

A company is planning to migrate applications to the AWS Cloud. During a system audit, the company finds that its content management system (CMS) application is incompatible with cloud environments.
Which migration strategies will help the company to migrate the CMS application with the LEAST effort? (Select TWO.)

A. Retire
B. Rehost
C. Repurchase
D. Replatform
E. Refactor

Answer: B,C

Explanation:
Rehosting ("lift and shift") is a migration strategy where applications are moved to the cloud with minimal changes, making it the least effort-intensive method for applications incompatible with the cloud. Repurchasing involves moving to a different product, often a SaaS solution, which can also minimize migration effort by avoiding the need for application-level changes. Retiring, replatforming, and refactoring require significant effort either in terms of analyzing and shutting down the application, making changes to the underlying platform, or redesigning the application architecture, respectively.

Reference:
AWS Migration Strategies

Question: 694

A company wants to run its application on Amazon EC2 instances. The company needs to keep the application on-premises to meet a compliance requirement. Which AWS offering will meet these requirements?

A. Dedicated Instances
B. Amazon CloudFront
C. AWS Fargate
D. AWS Outposts

Answer: D

Explanation:
AWS Outposts is an AWS offering that brings AWS infrastructure and services to a customer's onpremises location. It allows companies to run AWS services locally while meeting any regulatory or compliance requirements to keep data or applications on-premises. Dedicated Instances are EC2 instances that run on hardware dedicated to a single customer but are still within AWS data centers. Amazon CloudFront is a CDN service, and AWS Fargate is a serverless compute engine for containers, neither of which meets the requirement for running an application on-premises.

Reference:
AWS Outposts

Question: 695

Which AWS Support plan provides the full set of AWS Trusted Advisor checks at the LOWEST cost?

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

Answer: B

Explanation:

The AWS Support plan that provides the full set of AWS Trusted Advisor checks at the lowest cost is theAWS Business Supportplan. The AWS Business Support plan includes access to the complete set of Trusted Advisor checks, which cover areas such as cost optimization, security, performance, fault tolerance, and service limits. This plan is specifically designed to support production workloads and includes 24/7 access to cloud support engineers, response times for impaired systems, and other enhanced technical support features.

AWS Developer Support, while more affordable, only provides limited Trusted Advisor checks, specifically around Service Limits and basic Security checks. Full access to all Trusted Advisor checks is only available with Business Support and higher-tier plans, such as Enterprise On-Ramp and Enterprise Support

Question: 696

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 Pile System (Amazon EFS)

Answer: D

Explanation:

Amazon Elastic File System (Amazon EFS) is a fully managed, scalable, and serverless NFS (Network File System) file system specifically designed for use with AWS services and on-premises resources. It enables companies to create and configure file systems that can be accessed from multiple Amazon EC2instances simultaneously, making it ideal for use cases that require shared file storage for AWS compute services.

Why Amazon EFS Fits the Requirements:

Managed Service: Amazon EFS is a fully managed file storage service that simplifies the process of setting up and managing NFS file systems.

Scalability and Elasticity: EFS automatically scales to accommodate the storage needs of applications, without the need to provision or manage storage capacity.

NFS Compatibility: Amazon EFS natively supports the NFSv4 protocol, making it compatible with a wide range of applications and workloads that require NFS access.

Integration with AWS Compute Services: EFS integrates seamlessly with Amazon EC2 and other AWS services, providing a shared file storage solution across multiple instances and services within the AWS cloud environment.

Why Other Options Do Not Fit:

A . Amazon Elastic Block Store (Amazon EBS): While EBS provides block-level storage that can be attached to individual EC2 instances, it is not a file system, nor does it provide managed NFS file storage capabilities. EBS is designed for single-instance access rather than shared file access across multiple instances.

B . AWS Storage Gateway Tape Gateway: Tape Gateway is designed for archival purposes and allows companies to store virtual tape backups in Amazon S3 or Glacier. It does not support NFS file storage and is not intended for regular compute access.

C . Amazon S3 Glacier Flexible Retrieval: Amazon S3 Glacier is optimized for data archiving and longterm storage of infrequently accessed data, but it does not provide NFS file system capabilities, nor is it suitable for high-performance access needs associated with compute

services.
For more details, you can refer to theAWS Cloud Practitioner Essentialscontent, specifically the modules onStorage Serviceswhere Amazon EFS is covered as the managed NFS file storage solution offered by AWS.

Question: 697

A company wants to connect its supported AWS services and VPCs. The company does not wa Which AWS service will meet these requirements?

A. Amazon Inspector
B. AWS PrivateLink
C. Amazon Connect
D. AWS Internet Gateway

Answer: B

Explanation:
AWS PrivateLink provides private connectivity between VPCs, AWS services, and on-premises applications without exposing traffic to the public internet. It enables private access to AWS services by creating private endpoints within a VPC, which enhances security and simplifies network architecture by keeping all communication on the AWS network. This service is ideal for scenarios where a company wants to connect AWS services and VPCs privately.

Question: 698

A company wants to implement detailed tracking of its cloud costs by department and project. Which AWS feature or service should the company use?

A. Consolidated billing
B. Cost allocation tags
C. AWS Marketplace
D. AWS Budgets

Answer: B

Explanation:
Cost allocation tags are an AWS feature that allows for detailed tracking of cloud costs by tagging resources. These tags can be customized to represent departments, projects, or cost centers, enabling organizations to allocate and monitor expenses accurately. By using cost allocation tags, companies can create reports that break down costs by tag, providing granular insights into their cloud spending. This feature is critical for achieving detailed cost management and is a core part of AWS's billing and cost management offerings.

Question: 699

A company wants to manage its cloud resources by using infrastructure as code (IaC) template Which AWS service should the company use to meet these requirements?

A. AWS Artifact
B. AWS Resource Explorer
C. AWS License Manager
D. AWS Service Catalog

Answer: D

Explanation:
AWS Service Catalog enables companies to create, manage, and distribute catalogs of approved infrastructure as code (IaC) templates and software configurations. This service supports infrastructure automation by allowing users to deploy resources using predefined templates, which align with the company's policies and best practices. It helps in managing cloud resources at scale by utilizing IaC principles, enabling consistent deployments, and enforcing governance across AWS resources. AWS Service Catalog is an excellent choice for organizations implementing IaC.

Question: 700

Which AWS Support plans provide access to an AWS technical account manager (TAM)? (Select)

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

Answer: D,E

Explanation:
AWS Technical Account Managers (TAMs) are available to customers who subscribe to either theAWS Enterprise On-RamporAWS Enterprise Supportplans. TAMs provide proactive guidance and ongoing support, assisting with architecture reviews, operational performance improvements, and business reviews. TheBasicandDeveloperSupport plans do not include access to TAMs, and whileBusiness Supportoffers enhanced technical support, it does not include TAMs either.

Question: 701

A company needs to convert video files and audio files to a format that will play on smartphones. Which AWS service will meet this requirement?

A. Amazon Comprehend
B. Amazon Rekognition
C. Amazon Elastic Transcoder
D. Amazon Polly

Answer: C

Explanation:

Amazon Elastic Transcoder is a media transcoding service that enables companies to convert video and audio files into formats optimized for playback on various devices, including smartphones. It automates the transcoding process and supports a wide array of video and audio formats, making it ideal for converting files into mobile-friendly formats. Services like Amazon Comprehend, Rekognition, and Polly do not perform media transcoding functions.

Topic 4, Exam Pool D

Question: 702

Which AWS Cloud Adoption Framework (AWS CAF) perspective focuses on organizing an inventory of data products in a data catalog?

A. Operations
B. Governance
C. Business
D. Platform

Answer: B

Explanation:

In the AWS Cloud Adoption Framework (AWS CAF), theGovernanceperspective focuses on aligning IT strategy and goals with business processes, which includes managing data assets, setting up an inventory of data products, and ensuring that data cataloging and metadata management are in place. This perspective is crucial for organizing data products and services, which aligns with building a comprehensive data catalog. Other perspectives like Operations, Business, and Platform do not specifically address the management of a data catalog.

Question: 703

Which AWS service can a company use to manage encryption keys in the cloud?

A. AWS License Manager
B. AWS Certificate Manager (ACM)
C. AWS CloudHSM
D. AWS Directory Service

Answer: C

Explanation:

AWS CloudHSM provides hardware-based key management to manage and protect encryption keys in the AWS Cloud. It allows customers to generate and use their own encryption keys while complying withrigorous security requirements. WhileAWS Certificate Manager (ACM)manages SSL/TLS certificates, it does not handle encryption keys independently, andAWS License ManagerandAWS Directory

Serviceare not designed for managing encryption keys. AWS KMS is also relevant for key management but wasn't listed as an option in this question.

Question: 704

Which AWS service supports the deployment and management of applications in the AWS Cloud?

A. Amazon CodeGuru
B. AWS Fargate
C. AWS CodeCommit
D. AWS Elastic Beanstalk

Answer: D

Explanation:
AWS Elastic Beanstalk is a managed service that facilitates the deployment and management of applications in the AWS Cloud. It supports multiple programming languages and frameworks, allowing users to deploy web applications without managing the underlying infrastructure. Elastic Beanstalk automatically handles deployment, capacity provisioning, load balancing, and autoscaling. The other services listed, like CodeGuru, Fargate, and CodeCommit, do not provide full application deployment and management capabilities.

Question: 705

A company wants to enhance security by launching a third-party ISP intrusion detection system from its AWS account.
Which AWS service or resource should the company use to meet this requirement?

A. AWS Security Hub
B. AWS Marketplace
C. AWS Quick Starts
D. AWS Security Center

Answer: B

Explanation:
AWS Marketplace is a digital catalog that offers third-party software, including security solutions like intrusion detection systems, which can be deployed directly within an AWS environment. By using AWS Marketplace, companies can find, purchase, and deploy ISP intrusion detection systems or other security tools that meet their specific requirements. Other services, such as AWS Security Hub and Quick Starts, do not directly provide third-party application deployment.

Question: 706

A company migrated to the AWS Cloud. Now the company pays for services on an as-needed basis. Which advantage of cloud

computing is the company benefiting from?

A. Stop spending money running and maintaining data centers
B. Increase speed and agility
C. Go global in minutes
D. Trade fixed expense for variable expense

Answer: D

Explanation:
One of the main advantages of cloud computing is the ability to trade fixed capital expenditures for variable operating expenses. By moving to the AWS Cloud, the company pays only for the resources it uses on a per-usage basis, reducing the need for large upfront investments in data center infrastructure. This benefit allows for more flexible and cost-effective spending based on actual usage. Other options, such as increased speed and agility or global reach, are also cloud benefits but do not specifically address the shift from fixed to variable expenses.

Question: 707

A company is building a web application using AWS.
Which AWS service will help prevent network layer DDoS attacks against the web application?

A. AWS WAF
B. AWS Firewall Manager
C. Amazon GuardDuty
D. AWS Shield

Answer: D

Explanation:
AWS Shield is a managed Distributed Denial of Service (DDoS) protection service that safeguards AWS applications against network and transport layer attacks. AWS Shield Standard is automatically included with all AWS services and offers protection against common DDoS attacks. For more advanced protection, AWS Shield Advanced provides additional DDoS mitigation measures. Although AWS WAF, Firewall Manager, and GuardDuty contribute to security, they do not specialize in DDoS mitigation at the network layer like AWS Shield does.

Question: 708

A company is moving some of its on-premises IT services to the AWS Cloud. The finance department wants to see the entire bill so it can forecast spending limits.
Which AWS service can the company use to set spending limits and receive notifications if those limits are exceeded?

A. AWS Cost and Usage Reports
B. AWS Budgets
C. AWS Organizations consolidated billing

D. Cost Explorer

Answer: B

Explanation:

AWS Budgets allows organizations to set custom cost and usage budgets and receive notifications when they exceed predefined thresholds. This feature helps the finance department monitor spending and manage budget forecasts effectively. AWS Cost and Usage Reports and Cost Explorer provide detailed billing data but do not include budget notifications. Consolidated billing in AWS Organizations is useful for aggregating billing across accounts but does not provide budget alerts.

Question: 709

Which AWS service gives users on-demand, sell-service access to AWS compliance control reports?

A. AWS Config
B. Amazon GuardDuty
C. AWS Trusted Advisor
D. AWS Artifact

Answer: D

Explanation:

AWS Artifact provides on-demand access to AWS compliance reports and security documentation. It is a self-service portal where customers can download documents like SOC reports, ISO certifications, and other compliance-related materials necessary for meeting regulatory requirements. AWS Config and Trusted Advisor offer security assessments and compliance monitoring, but they do not provide direct access to compliance reports.

Question: 710

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 1AM user and provide AWS Management Console access only.
B. Create an 1AM user and provide programmatic access only.
C. Create an 1AM role and provide AWS Management Console access only.
D. Create an 1AM policy with administrator access and attach it to the 1AM user.
E. Create an 1AM policy with Amazon RDS access and attach it to the 1AM user.

Answer: B,E

Explanation:

To follow the principle of least privilege, the company should create an IAM user with only programmatic access since the access is limited to AWS CLI and SDKs, not the Management Console. Additionally, a custom IAM policy granting specific Amazon RDS permissions should be

created and attached to this user to restrict access solely to necessary actions. Providing programmatic access only ensures adherence to security best practices by limiting access to the required interfaces.

Question: 711

Which AWS service can create a private network connection from on premises to the AWS Cloud?

A. AWS Config
B. Virtual Private Cloud (Amazon VPC)
C. AWS Direct Connect
D. Amazon Route 53

Answer: C

Explanation:
AWS Direct Connect provides a dedicated private network connection from on-premises data centers directly to the AWS Cloud, bypassing the public internet. This setup is ideal for reducing network costs, increasing bandwidth throughput, and providing a more consistent network experience compared to standard internet connections. Other services, such as Amazon VPC, relate to networking but do not establish a private network connection from on-premises to AWS.

Question: 712

A company is running a reporting web server application on Amazon EC2 instances. The application runs once every week and once again at the end of the month. The EC2 instances can be shut down when they are not in use.
What is the MOST cost-effective billing model for this use case?

A. Standard Reserved Instances
B. Convertible Reserved Instances
C. On-Demand Capacity Reservations
D. On-Demand Instances

Answer: D

Explanation:
For a reporting application that runs only periodically, On-Demand Instances are the most costeffective choice because they allow the company to pay only for the compute capacity used, without long-term commitments. Reserved Instances are less flexible due to the need for upfront payment or long-term contracts, which would not be cost-effective given the application's intermittent usage. On-Demand Capacity Reservations would also be more costly, as they hold capacity regardless of usage.

Question: 713

A company is moving Us development and test environments to AWS to increase agility and reduce cost. Because these are not production workloads and the servers are not fully utilized, occasional unavailability is acceptable.
What is the MOST cost-effective Amazon EC2 pricing model that will meet these requirements?

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

Answer: C

Explanation:
Spot Instances offer the lowest cost for Amazon EC2 and are suitable for non-production workloads like development and testing where occasional unavailability is acceptable. Spot Instances take advantage of unused EC2 capacity at a reduced cost, making them ideal for environments that can tolerate interruptions. Reserved or On-Demand Instances would be more expensive for this scenario, and Dedicated Hosts are not cost-effective for non-production environments.

Question: 714

A company wants to use an AWS networking solution that can act as a centralized gateway between multiple VPCs and on-premises networks. Which AWS service or feature will meet this requirement?

A. Gateway VPC endpoint
B. AWS Direct Connect
C. AWS Transit Gateway
D. AWS PrivateLink

Answer: C

Explanation:
AWS Transit Gateway serves as a central hub that enables connectivity between multiple VPCs and on-premises networks. It simplifies network architecture and management by acting as a centralized gateway for traffic flowing between all connected networks. Other options, such as Gateway VPC Endpoints and AWS PrivateLink, do not provide the centralized, scalable connectivity that Transit Gateway offers across multiple VPCs and on-premises environments.

Question: 715

Which AWS service is a fully managed NoSQL database service?

A. Amazon RDS
B. Amazon Redshift
C. Amazon DynamoDB
D. Amazon Aurora

Answer: C

Explanation:

Amazon DynamoDB is a fully managed NoSQL database service that provides fast and predictable performance with seamless scalability. It supports both document and key-value data models and is designed to handle large amounts of data across multiple servers. Other options, like Amazon RDS and Aurora, are managed relational database services, and Amazon Redshift is a data warehousing service.

Question: 716

Which AWS service or tool provides users with a graphical interface that they can use to manage AWS services?

A. AWS Copilot
B. AWS CLI
C. AWS Management Console
D. AWS software development kits (SDKs)

Answer: C

Explanation:
The AWS Management Console provides a web-based graphical user interface (GUI) that allows users to manage AWS services. It is user-friendly and accessible, enabling users to control and configure resources without needing to interact with AWS through code or command-line interfaces. AWS CLI and SDKs are command-line and programming tools, respectively, and do not offer a graphical interface.

Question: 717

Which AWS service or feature allows users to securely store encrypted credentials and retrieve these credentials when required?

A. AWS Encryption SDK
B. AWS Security Hub
C. AWS Secrets Manager
D. AWS Artifact

Answer: C

Explanation:
AWS Secrets Manager is a service designed to securely store and manage access to sensitive information such as API keys, passwords, and database credentials. It allows automatic rotation, secure storage, and fine-grained access control for these credentials, ensuring they are securely managed. The other options do not provide encrypted storage for credentials specifically.

Question: 718

A company wants to migrate its on-premises SQL Server database to the AWS Cloud. The company wants AWS to handle the day-to-day administration of the database. Which AWS service will meet the company's requirements?

A. Amazon EC2 foe Microsoft SQL Server
B. Amazon DynamoDB

C. Amazon RDS
D. Amazon Aurora

Answer: C

Explanation:
Amazon RDS (Relational Database Service) supports SQL Server and other relational databases while handling administrative tasks such as backups, patching, monitoring, and scaling. This allows the company to migrate its SQL Server database without managing the infrastructure manually. Amazon EC2for SQL Server would require more administrative overhead, as the company would be responsible for the maintenance. Amazon DynamoDB and Aurora do not support SQL Server directly.

Question: 719

A company runs a legacy workload in an on-premises data center. The company wants to migrate the workload to AWS. The company does not want to make any changes to the workload.
Which migration strategy should the company use?

A. Repurchase
B. Replatform
C. Rehost
D. Refactor

Answer: C

Explanation:
TheRehostmigration strategy, often referred to as "lift-and-shift," involves moving applications to the cloud with minimal or no modifications. This approach is suitable when a company wants to migrate legacy workloads to AWS without altering them. Other strategies, such asRepurchase,Replatform, andRefactor, involve making changes to the application or adopting different services, which is not aligned with the requirement to avoid modifications.

Question: 720

A company wants to transport 100 TB of data from its data center to AWS without using internet. Which AWS service will meet this requirement?

A. AWS Snowcone
B. AWS Snowball Edge
C. AWS Data Exchange
D. AWS DataSync

Answer: B

Explanation:
AWS Snowball Edge is designed for large-scale data transfer tasks, allowing companies to transport significant amounts of data (up to petabytes) from their data centers to AWS without relying on the internet. With the ability to transfer 100 TB of data securely and efficiently, Snowball Edge is the ideal choice for this use case. AWS Snowcone is for smaller transfers, and AWS Data Exchange and DataSync are not designed for physical data transport.

Question: 721

A company wants to rightsize its Amazon EC2 instances.
Which configuration change will meet this requirement with the LEAST operational overhead?

A. Add EC2 instances in another Availability Zone.
B. Change the size and type of the EC2 instances based on utilization.
C. Convert the payment method from On-Demand to Savings Plans.
D. Reprovision the EC2 instances with a larger instance type.

Answer: B

Explanation:
Rightsizing involves adjusting instance types based on actual utilization to optimize costs and performance. Changing the size and type of EC2 instances according to utilization data achieves this with minimal operational overhead. Adding instances or changing payment methods does not directly address instance resizing, and reprovisioning with larger instances may not be optimal based on actual usage.

Question: 722

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:
Elastic Load Balancing (ELB) distributes incoming traffic across multiple targets, such as EC2 instances, containers, and IP addresses, ensuring better application availability and fault tolerance. While ELB scales to meet traffic demands, it does not inherently scale compute resources themselves, span multiple regions, or balance traffic between internet gateways.

Question: 723

Which action should a company take to improve security in its AWS account?

A. Require multi-factor authentication (MFA) for privileged users.
B. Remove the root user account.
C. Create an access key for the AWS account root user.
D. Create an access key for each privileged user.

Answer: A

Explanation:

Enforcing multi-factor authentication (MFA) for privileged users enhances account security by requiring a second form of authentication. It reduces the risk of unauthorized access, even if credentials are compromised. Removing the root account is not possible, and creating access keys for the root account or privileged users can increase security risks rather than reduce them.

Question: 724

A developer needs to use a standardized template to create copies of a company's AWS architecture for development test, and production environments. Which AWS service should the developer use to meet this requirement?

A. AWS Cloud Map
B. AWS Cloud Formation
C. Amazon CloudFront
D. AWS CloudTrail

Answer: B

Explanation:

AWS CloudFormation allows developers to create, manage, and deploy AWS resources using templates, ensuring a standardized and repeatable infrastructure setup. It's ideal for setting up development, test, and production environments consistently. AWS Cloud Map, CloudFront, and CloudTrail do not provide templated infrastructure management capabilities.

Question: 725

A company wants to test a new application.
Which AWS principle will help the company test the application?

A. Make long-term commitments in exchange for a cost discount.
B. Scale up and down when needed without any long-term commitments.
C. Have total control over the application infrastructure.
D. Manage all of the maintenance tasks associated with the cloud.

Answer: B

Explanation:

AWS provides the flexibility to scale resources up or down on demand, enabling companies to test new applications without making long-term commitments. This principle allows for cost efficiency and agility during testing and development. The other options do not specifically highlight the benefits of flexible, on-demand resource scaling.

Question: 726

Which AWS tool or feature acts as a VPC firewall at the subnet level?

A. Security group
B. Network ACL
C. Traffic Mirroring
D. Internet gateway

Answer: B

Explanation:
Network ACLs (NACLs) are subnet-level firewalls in AWS, controlling inbound and outbound traffic for VPC subnets. They provide an additional layer of security by allowing or denying traffic based on IP protocol, source and destination IP, and port. Security groups operate at the instance level, while Traffic Mirroring and Internet Gateways do not function as firewalls at the subnet level.

Question: 727

Which AWS service uses edge locations to cache content?

A. Amazon Kinesis
B. Amazon Simple Queue Service (Amazon SQS)
C. Amazon CloudFront
D. Amazon Route 53

Answer: C

Explanation:
Amazon CloudFront is a content delivery network (CDN) that uses edge locations to cache content closer to users, reducing latency and improving performance. It supports the delivery of web content, such as videos and images, by caching copies at edge locations around the world. Amazon Kinesis, SQS, and Route 53 do not utilize edge locations for content caching.

Question: 728

Where can users find examples of AWS Cloud solution designs?

A. AWS Marketplace
B. AWS Service Catalog
C. AWS Architecture Center
D. AWS Trusted Advisor

Answer: C

Explanation:

The AWS Architecture Center provides examples, best practices, and architectural blueprints for various AWS Cloud solutions. It includes whitepapers, reference architectures, and detailed diagrams to help users design secure, scalable, and reliable applications on AWS. Other services like AWS Marketplace, Service Catalog, and Trusted Advisor do not offer solution design examples.

Question: 729

What is the total volume of data that can be stored in Amazon S3?

A. 10 PB
B. 50 PB
C. 100 PB
D. Virtually unlimited

Answer: D

Explanation:

Amazon S3 offers virtually unlimited storage, allowing customers to store and retrieve any amount of data. There are no practical limits to the total volume of data that can be stored in S3, making it suitable for applications that require vast amounts of storage. The options of 10 PB, 50 PB, and 100 PB are incorrect as they do not reflect the actual scale of S3.

Question: 730

A company has an on-premises application. The application has processing times of less than 5 minutes and is invoked only a few times each day. The company wants to move the application to the AWS Cloud.

Which AWS service will support this application MOST cost-effectively?

A. Amazon Elastic Container Service (Amazon ECS)
B. AWS Lambda
C. Amazon Elastic Kubernetes Service (Amazon EKS)
D. Amazon EC2

Answer: B

Explanation:

AWS Lambda is a serverless compute service that is ideal for applications with short-running processes and low-frequency invocation, as it only charges for the compute time consumed. Since the application is invoked only a few times a day and runs for less than 5 minutes, Lambda is the most cost-effective option. ECS, EKS, and EC2 are better suited for more persistent workloads and would be more costly in this use case.

Question: 731

Which AWS service can generate information that can be used by external auditors?

A. Amazon Cognito
B. Amazon FSx
C. AWS Config
D. Amazon Inspector

Answer: C

Explanation:
AWS Config enables users to assess, audit, and evaluate the configurations of AWS resources. It provides information that can be used by external auditors to ensure compliance with various regulatory requirements by tracking changes and maintaining configuration history. Amazon Cognito, FSx, and Inspector do not provide detailed configuration tracking for audit purposes.

Question: 732

A company is learning about its responsibilities that are related to the management of Amazon EC2 instances.
Which tasks for EC2 instances are the company's responsibility, according to the AWS shared responsibility model? (Select TWO.)

A. Install and patch the machine hypervisor.
B. Patch the guest operating system.
C. Encrypt data at rest on associated storage.
D. Install the physical hardware and cabling.
E. Provide physical security for the EC2 instances.

Answer: B,C

Explanation:
Under the AWS Shared Responsibility Model, AWS manages the securityofthe cloud, while customers manage securityinthe cloud. For EC2 instances, it is the customer's responsibility to manage the guestoperating system, including patching and encrypting data stored on attached storage volumes. AWS is responsible for the underlying infrastructure, including physical security and hypervisor maintenance.

Question: 733

A company wants to automatically patch its Windows instances that are deployed on Amazon EC2. Which AWS service will meet these requirements?

A. AWS Systems Manager
B. AWS Organizations
C. AWS Control Tower
D. Elastic Load Balancing (ELB)

Answer: A

Explanation:
AWS Systems Manager offers the capability to automate patching for Amazon EC2 instances, including Windows instances, through Patch Manager. It can automate patching tasks, schedule updates, and apply patches based on specified baselines. AWS Organizations and Control Tower focus on account management and governance, while Elastic Load Balancing (ELB) is unrelated to patch management.

Question: 734

Which design principles are included in the reliability pillar of the AWS Well-Architected Framework? (Select TWO.)

A. Automatically recover from failure.
B. Grant everyone access to increase AWS service quotas.
C. Stop guessing capacity.
D. Design applications to run in a single Availability Zone.
E. Plan to increase AWS service quotas first in a secondary AWS Region.

Answer: A,C

Explanation:
The reliability pillar of the AWS Well-Architected Framework emphasizes building systems that can recover from failures and dynamically adjust to meet demand. This includes designing to automatically recover from failures and implementing mechanisms to manage capacity demands, avoiding manual guesswork. The other options do not align with core reliability principles, as operating in a single Availability Zone or preemptively adjusting quotas in a secondary region does not inherently improve reliability.

Question: 735

A company needs to organize its resources and track AWS costs on a detailed level. The company needs to categorize costs by business department, environment, and application. Which solution will meet these requirements'?

A. Access the AWS Cost Management console to organize resources set an AWS budget, and receive notifications of unintentional usage.
B. Use tags to organize the resources. Activate cost allocation tags to track AWS costs on a detailed level.
C. Create Amazon CloudWatch dashboards to visually organize and track costs individually.
D. Access the AWS Billing and Cost Management dashboard to organize and track resource consumption on a detailed level.

Answer: B

Explanation:
AWS recommends using tags to organize resources effectively and to activate cost allocation tags to enable detailed tracking of costs by categories such as department, environment, and application. This approach provides the granularity needed for tracking and managing AWS costs accurately. AWS Cost Management and Billing tools provide insights but do not inherently categorize costs without tags.

Question: 736

Which AWS service helps users plan and track their server and application inventory migration data to AWS?

A. Amazon CloudWatch
B. AWS DataSync
C. AWS Migration Hub
D. AWS Application Migration Service

Answer: C

Explanation:
AWS Migration Hub assists users in planning and tracking the progress of their server and application migrations. It centralizes migration tracking across various AWS services, providing visibility into application inventory and migration status. While AWS Application Migration Service also assists with migrations, Migration Hub is specifically designed for tracking migration data comprehensively.

Question: 737

Which AWS services are serverless? (Select TWO.)

A. AWS Fargate
B. Amazon Managed Streaming for Apache Kafka
C. Amazon EMR
D. Amazon S3
E. Amazon EC2

Answer: A,D

Explanation:
AWS Fargate and Amazon S3 are both serverless services. Fargate allows users to run containers without managing the underlying infrastructure, while S3 provides object storage without the need for provisioning or managing servers. Amazon EC2, Amazon Managed Streaming for Apache Kafka, and Amazon EMR involve server management to varying degrees and are not serverless by nature.

Question: 738

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:
The AWS Cloud is cost-effective for dynamic workloads because of its elasticity, allowing resources to scale up or down based on demand, and its pay-as-you-go pricing, which enables companies to pay only for what they use. Reliability, security, and high availability are also benefits of AWS, but they do not specifically relate to cost-effectiveness in the context of dynamic workloads.

Question: 739

A company plans to migrate to the AWS Cloud. The company wants to gather information about its on-premises data center.
Which AWS service should the company use to meet these requirements?

A. AWS Application Discovery Service
B. AWS DataSync
C. AWS Storage Gateway
D. AWS Database Migration Service (AWS DMS)

Answer: A

Explanation:
AWS Application Discovery Service helps companies gather information about on-premises data centers to support migration planning. It collects data on system configuration, resource utilization, and network dependencies, providing essential insights for migration. AWS DataSync and Storage Gateway are used for data transfer, while AWS DMS is for database migration specifically.

Question: 740

Which AWS service provides on-premises applications with low-latency access to data that is stored in the AWS Cloud?

A. Amazon CloudFront
B. AWS Storage Gateway
C. AWS Backup
D. AWS DataSync

Answer: B

Explanation:
AWS Storage Gateway provides on-premises applications with low-latency access to data stored in AWS by caching frequently accessed data locally. It seamlessly integrates on-premises environments with cloud storage, enabling hybrid storage solutions. AWS DataSync is for data transfer, CloudFront is a content delivery network, and AWS Backup is for backup management, not low-latency access.

Question: 741

A company wants to securely access an Amazon S3 bucket from an Amazon EC2 instance without accessing the internet.
What should the company use to accomplish this goal?

A. VPN connection
B. Internet gateway
C. VPC endpoint
D. NAT gateway

Answer: C

Explanation:
A VPC endpoint enables private connections between an Amazon VPC and AWS services, like Amazon S3, without requiring internet access. This allows secure access to S3 from an EC2 instance within the same VPC, reducing latency and improving security. VPN connections and NAT gateways do not eliminate internet traffic, and an internet gateway would expose the VPC to the public internet.

Question: 742

Which design principle aligns with performance efficiency pillar of the AWS Well-Architected Framework?

A. Using serverless architectures
B. Scaling horizontally
C. Measuring the cost of workloads
D. Using managed services

Answer: B

Explanation:
Scaling horizontally, or adding more instances of resources rather than increasing the size of a single instance, is a core principle of the Performance Efficiency pillar of the AWS Well-Architected Framework. It enables applications to handle increasing loads by distributing traffic across multiple resources. Other options, such as serverless architectures, managed services, and cost measurement, align with other pillars of the framework.

Question: 743

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 a machine learning service that can analyze and recognize objects, people, text, scenes, and activities in images and videos. It is specifically designed for organizing, categorizing, and searching large volumes of images based on various attributes. Amazon

Transcribe is for speech-to-text, Amazon Aurora is a relational database service, and Amazon QuickSight is a business intelligence tool for data visualization.

Question: 744

A company needs to manage multiple logins across AWS accounts within the same organization in AWS Organizations.
Which AWS service should the company use to meet this requirement?

A. Amazon VPC
B. Amazon GuardDuty
C. Amazon Cognito
D. AWS IAM Identity Center

Answer: D

Explanation:
AWS IAM Identity Center (formerly AWS Single Sign-On) provides centralized management of user access to multiple AWS accounts within an organization in AWS Organizations. It allows users to log in once and gain access to all assigned accounts without managing separate logins for each account. Amazon Cognito is generally used for application-level user management and authentication, not for managing access across AWS accounts.

Question: 745

A company plans to perform a one-time migration of a large dataset with millions of files from its onpremises data center to the AWS Cloud.
Which AWS service should the company use for the migration?

A. AWS Database Migration Service (AWS DMS)
B. AWS DataSync
C. AWS Migration Hub
D. AWS Application Migration Service

Answer: B

Explanation:
AWS DataSync is designed for large-scale data transfers, especially involving large datasets with millions of files from on-premises to AWS. It provides fast and efficient transfer capabilities, and supports a one-time migration. AWS DMS is specific to databases, while Migration Hub is for tracking migrations, and Application Migration Service is for continuous replication rather than one-time file migrations.

Question: 746

A company is planning to migrate a monolithic application to AWS. The company wants to modernize the application by splitting it into microservices. The company will deploy the microservices on AWS.

Which migration strategy should the company use?

A. Rehost
B. Repurchase
C. Replatform
D. Refactor

Answer: D

Explanation:
Refactoring involves re-architecting and modifying an application to take advantage of cloud-native features. In this case, the company wants to modernize a monolithic application by breaking it into microservices. This process aligns with the Refactor strategy, which is aimed at modernizing and rearchitecting applications. Rehost, Repurchase, and Replatform do not involve the level of rearchitecting needed to move from a monolithic to a microservices architecture.

Question: 747

A company wants its Amazon EC2 instances to be in different locations but share the same geographic are
a. The company also wants to use multiple power grids and independent networking connectivity for the EC2 instances.
Which solution meets these requirements?

A. Use EC2 instances in multiple edge locations in the same AWS Region.
B. Use EC2 instances in multiple Availability Zones in the same AWS Region.
C. Use EC2 instances in multiple Amazon Connect locations in the same AWS Region
D. Use EC2 instances in multiple AWS Artifact locations in the same AWS Region.

Answer: B

Explanation:
Using multiple Availability Zones within the same AWS Region meets the requirements for having instances in different locations with independent power and networking. Availability Zones are distinct physical locations within a region, each with separate power sources and networking. Edge locations are used for content delivery, and Amazon Connect and AWS Artifact locations are not relevant to EC2 deployment and infrastructure.

Question: 748

A company wants to centrally manage Its employee's access to multiple AWS accounts. Which AWS service or feature should the company use to meet this requirement?

A. AWS Identity and Access Management Access Analyzer
B. AWS Secrets Manager
C. AWS IAM Identity Center
D. AWS Security Token Service (AWS STS)

Answer: C

Explanation:
AWS IAM Identity Center provides centralized access management across multiple AWS accounts, enabling organizations to manage employee access efficiently. It is specifically designed for this purpose within AWS Organizations. AWS STS provides temporary credentials but does not manage multiple account access centrally, while IAM Access Analyzer and Secrets Manager serve different purposes related to access and secret management.

Question: 749

A company has multiple AWS accounts. The company needs to receive a consolidated bill from AWS and must centrally manage security and compliance. Which AWS service or feature should the company use to meet these requirements?

A. AWS Cost and Usage Report
B. AWS Organizations
C. AWS Config
D. AWS Security Hub

Answer: B

Explanation:
AWS Organizations enables consolidated billing across multiple AWS accounts and allows for centralized management of security and compliance policies. It provides account grouping and centralized payment management, making it the optimal choice for a company requiring consolidated billing and centralized governance. AWS Cost and Usage Report only provides billing information, and AWS Config and Security Hub offer monitoring and security insights but do not handle billing consolidation.

Question: 750

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 a scalable Domain Name System (DNS) web service that routes end-user requests to infrastructure running in AWS. It is specifically designed to host domain names and manage DNS records, making it the ideal service for hosting the domain name of a public website. AWS Lambda, CloudFront, and Direct Connect serve different functions and are not DNS hosting services.

Question: 751

A company needs to store infrequently used data for data archives and long-term backups. Which AWS service or storage class will meet these requirements MOST cost-effectively?

A. Amazon FSx for Lustre
B. Amazon Elastic Block Store (Amazon EBS)
C. Amazon Elastic File System (Amazon EFS)
D. Amazon S3 Glacier Flexible Retrieval

Answer: D

Explanation:
Amazon S3 Glacier Flexible Retrieval is optimized for storing infrequently accessed data at a very low cost, making it suitable for data archiving and long-term backups. It provides flexible retrieval options for archived data. Other services like FSx for Lustre, EBS, and EFS are more suited for frequently accessed data and higher-performance use cases, not archival storage.

Question: 752

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 provides checks and recommendations across various categories, including cost optimization, security, and performance. Access to all Trusted Advisor checks is available with AWS Business Support, ensuring the company can follow AWS best practices comprehensively. The Developer Support plan offers limited Trusted Advisor checks, and AWS Health Dashboard provides insights into service health but not specific best practice recommendations.

Question: 753

Which AWS services can host PostgreSQL databases? (Select TWO.)

A. Amazon S3
B. Amazon Aurora
C. Amazon EC2
D. Amazon OpenSearch Service

E. Amazon Elastic File System (Amazon EFS)

Answer: B,C

Explanation:
Amazon Aurora (with PostgreSQL compatibility) and Amazon EC2 can both host PostgreSQL databases. Aurora is a managed database service that provides a fully compatible PostgreSQL environment with automated management. EC2 can also host PostgreSQL, but it requires more manual setup and maintenance. Amazon S3 and EFS do not host databases, and Amazon OpenSearch Service is intended for search and analytics, not relational databases.

Question: 754

A company's gaming application has been gaining popularity. There has been high demand for the gaming application in countries where the company does not currently deploy the application. Which advantage of the AWS Cloud can help the company to deploy the application to more countries around the world?

A. Increase speed and agility
B. Go global in minutes
C. Trade fixed expense for variable expense
D. Benefit from massive economies of scale

Answer: B

Explanation:
The AWS Cloud's global infrastructure enables rapid deployment across multiple geographic regions, allowing companies to extend applications to new markets quickly. This capability to "go global in minutes" is especially valuable for businesses looking to reach customers in new countries without building data centers. Other options, like speed and agility, are advantages but do not specifically address global deployment.

Question: 755

A company has a MariaDB database on premises. The company wants to move the data to the AWS Cloud. Which AWS service will host this database with the LEAST amount of operational overhead?

A. Amazon RDS
B. Amazon Neptune
C. Amazon S3
D. Amazon DynamoDB

Answer: A

Explanation:
Amazon RDS provides a managed database service that supports MariaDB with minimal operational overhead. It handles routine database

tasks such as backups, patching, and scaling, reducing the burden on users. Amazon Neptune and DynamoDB are not compatible with MariaDB, and Amazon S3 is not a relational database service.

Question: 756

A company is building AWS architecture to deliver real-time data feeds from an on-premises data center into an application that runs on AWS. The company needs a consistent network connection with minimal latency.
What should the company use to connect the application and the data center to meet these requirements?

A. AWS Direct Connect
B. Public internet
C. AWS VPN
D. Amazon Connect

Answer: A

Explanation:
AWS Direct Connect provides a dedicated, low-latency, and consistent network connection between on-premises data centers and AWS. This private connection minimizes latency and improves network stability compared to a public internet connection or VPN. AWS Direct Connect is ideal for applications requiring real-time data transfer with minimal latency.

Question: 757

A cloud engineer wants to store data in Amazon S3. The engineer will access some of the data yearly and some of the data dally.
Which S3 storage class will meet these requirements MOST cost-effectively?

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

Answer: D

Explanation:
S3 Intelligent-Tiering is cost-effective for data with unpredictable access patterns. It automatically moves data between two access tiers (frequent and infrequent) based on access patterns, which is suitable for data accessed daily or infrequently. S3 Glacier Deep Archive is for archival data accessed infrequently, and S3 One Zone-IA is for infrequent access data stored in a single availability zone, which may not be ideal for data accessed daily.

Question: 758

A company wants to run a graph query that provides credit card users' names, addresses, and transactions. The company wants the graph to show if the names, addresses, and transactions indicates possible fraud.

Which AWS database service will meet these requirements?

A. Amazon DocumenlDB (with MongoDB compatibility)
B. Amazon Timestream
C. Amazon DynamoDB
D. Amazon Neptune

Answer: D

Explanation:
Amazon Neptune is a managed graph database service that is well-suited for complex queries involving relationships, such as detecting patterns indicative of fraud in credit card transactions. It supports graph query languages like Gremlin and SPARQL, making it ideal for use cases requiring intricate relationship analysis. DocumentDB, Timestream, and DynamoDB are not designed for graph-based queries.

Question: 759

A company needs to block SOL injection attacks.
Which AWS service or feature provides this functionality?

A. AWS WAF
B. Network ACLs
C. Security groups
D. AWS Trusted Advisor

Answer: A

Explanation:
AWS WAF is a web application firewall that helps protect applications from SQL injection attacks and other common web exploits. By defining rules to block, allow, or monitor specific types of requests, AWS WAF provides an effective defense against SQL injection. Network ACLs and Security Groups provide network-level security but do not inspect web traffic for specific attack patterns like SQL injection.

Question: 760

A user wants to invoke an AWS Lambda function when an Amazon EC2 instance enters the "stopping" state.
Which AWS service is appropriate for this use case?

A. Amazon EventBridge
B. AWS Config
C. Amazon Simple Notification Service (Amazon SNS)
D. AWS CloudFormation

Answer: A

Explanation:

Amazon EventBridge (formerly CloudWatch Events) allows users to respond to changes in the state of AWS resources. It can be configured to invoke an AWS Lambda function when an EC2 instance enters the "stopping" state, providing a serverless way to automate responses to changes in EC2 instance states. AWS Config, SNS, and CloudFormation do not provide direct triggering for specific instance state changes.

Question: 761

Which AWS service provides machine learning capability to detect and analyze content in images and videos?

A. Amazon Connect
B. Amazon Lightsail
C. Amazon Personalize
D. Amazon Rekognition

Answer: D

Explanation:
Amazon Rekognition provides machine learning capabilities to analyze images and videos, enabling the detection of objects, people, text, and scenes. It is designed specifically for image and video analysis, making it suitable for various use cases like facial recognition and content moderation. Other services like Amazon Connect, Lightsail, and Personalize do not offer image or video analysis capabilities.

Question: 762

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 enables visualization and analysis of AWS costs and usage over specific time periods. It provides detailed reports, trends, and forecasts for cost management. Consolidated billing and AWS Organizations are useful for billing management across accounts, while AWS Budgets helps set spending thresholds but does not provide visualization tools for cost and usage.

Question: 763

A company purchased Amazon EC2 Standard Reserved Instances (RIs) for a workload in the AWS Cloud. The company needs to move part of the workload to an instance family that does not match the instance family of these Standard RIs.
How can the company take advantage of the Standard RIs that it no longer needs?

A. Contact the AWS Support team, and ask the team to sell the Standard RIs.
B. Sell the Standard RIs on the Amazon EC2 Reserved Instance Marketplace.
C. Sell the Standard RIs as a third-party seller on the AWS Marketplace.
D. Convert the Standard RIs to Savings Plans.

Answer: B

Explanation:
The Amazon EC2 Reserved Instance Marketplace allows customers to sell unused Standard Reserved Instances to other AWS users. This enables companies to recoup some of the costs if they no longer need certain RIs. Standard RIs cannot be converted to Savings Plans, and AWS Support does not facilitate the resale of RIs directly.

Question: 764

Which AWS service provides storage-optimized and compute-optimized device configurations?

A. AWS Snowcone
B. AWS Storage Gateway
C. AWS Snowball Edge
D. AWS DataSync

Answer: C

Explanation:
AWS Snowball Edge offers configurations that are either storage-optimized or compute-optimized, providing flexibility for different data migration and processing needs. It supports local processing and data transfer to AWS, catering to scenarios that require heavy storage or compute resources. AWS Snowcone and DataSync do not offer these optimizations, and Storage Gateway focuses on hybrid cloud storage.

Question: 765

Which AWS service or resource can a company use to deploy AWS WAF rules?

A. Amazon EC2
B. Application Load Balancer
C. AWS Trusted Advisor
D. Network Load Balancer

Answer: B

Explanation:
Application Load Balancer (ALB) integrates with AWS WAF to deploy and manage WAF rules for incoming traffic. ALB can route HTTP and HTTPS traffic and apply WAF rules to protect applications from common web exploits. Network Load Balancer does not support AWS WAF,

and Trusted Advisor does not deploy WAF rules.

Question: 766

Which AWS service gives users the ability to deploy highly repeatable infrastructure configurations?

A. AWS CloudFormation
B. AWS CodeDeploy
C. AWS CodeBuild
D. AWS Systems Manager

Answer: A

Explanation:

AWS CloudFormation allows users to define and deploy infrastructure as code, creating highly repeatable and consistent configurations across environments. It uses templates to automate the provisioning and management of resources. CodeDeploy focuses on application deployment, and Systems Manager offers operational management, but neither provides templated infrastructure deployment at the same level as CloudFormation.

Question: 767

A company will run a predictable compute workload on Amazon EC2 Instances for the next 3 years. The workload is critical for the company. The company wants to optimize costs to run the workload. Which solution will meet these requirements?

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

Answer: C

Explanation:

Savings Plans offer cost savings for predictable, steady workloads over a one or three-year term, with flexibility in instance family and size, making them suitable for critical workloads on EC2 instances that will run for a long term like three years. Spot Instances are more cost-effective but not suitable forcritical, predictable workloads due to potential interruptions. Dedicated Hosts and On-Demand Instances would be more costly.

Question: 768

For which use case are Amazon EC2 On-Demand Instances MOST cost-effective?

A. Compute-intensive video transcoding that can be restarted it necessary
B. An instance in continual use for 1 month to conduct quality assurance tests

C. An instance that runs a web server that will run for 1 year
D. An instance that runs a database that will run for 3 years

Answer: B

Explanation:
On-Demand Instances are most cost-effective for short-term, steady, and unpredictable workloads. Using them for a one-month testing period allows flexibility without a long-term commitment. For long-term workloads (like a year or more), Reserved Instances or Savings Plans would be more costeffective. Spot Instances are better for interruptible, flexible workloads.

Question: 769

A company wants to log in securely to Linux Amazon EC2 instances. How can the company accomplish this goal?

A. Use SSH keys.
B. Use a VPN.
C. Use end-to-end encryption.
D. Use Amazon Route 53.

Answer: A

Explanation:
SSH keys provide secure login for Linux-based Amazon EC2 instances by establishing a secure connection over SSH (Secure Shell), protecting login credentials from interception. VPNs and encryption enhance security in other contexts, but SSH keys are the standard approach for accessing Linux EC2 instances. Amazon Route 53 is unrelated to EC2 instance access.

Question: 770

A company wants to use AWS. The company has stringent requirements about low-latency access to on-premises systems and data residency.
Which AWS service should the company use to design a solution that meets these requirements?

A. AWS Wavelength
B. AWS Transit Gateway
C. AWS Ground Station
D. AWS Outposts

Answer: D

Explanation:
AWS Outposts extend AWS infrastructure and services to on-premises locations, providing low- latency access to AWS resources and ensuring data residency. This service is suitable for hybrid environments that require the same AWS services and infrastructure to be available locally. Wavelength, Transit Gateway, and Ground Station do not specifically address low-latency access to on-premises resources or data residency.

Question: 771

A company has a client that uses an Amazon RDS database. The client requests Information about operating system-level upgrades on the AWS resources that host the RDS database. The company employs a third-party provider to monitor the RDS database.
Who is responsible for upgrading the operating systems for Amazon RDS under the AWS shared responsibility model?

A. The client
B. The company
C. AWS
D. The third-party provider

Answer: C

Explanation:
Under the AWS shared responsibility model, AWS is responsible for managing the underlying infrastructure, including operating system-level updates on managed services like Amazon RDS. Customers are responsible for managing the database instance and configurations, but AWS handles OS updates for the infrastructure supporting RDS.

Question: 772

Which design principle is related to the reliability pillar according to the AWS Well-Architected Framework?

A. Test recovery procedures
B. Experiment more often
C. Go global in minutes
D. Analyze and attribute to expenditure

Answer: A

Explanation:
The reliability pillar of the AWS Well-Architected Framework includes the principle of testing recovery procedures to ensure systems can effectively recover from failures. Regular testing of recovery processes helps verify that systems are resilient and can handle potential disruptions. The other options align with different pillars like cost optimization and operational excellence.

Question: 773

A company is connecting multiple VPCs and on-premises networks. The company needs to use an AWS service as a cloud router to simplify peering relationships.
Which AWS service can the company use to meet this requirement?

A. AWS Direct Connect
B. AWS Transit Gateway
C. Amazon Connect
D. Amazon Route 53

Answer: B

Explanation:

AWS Transit Gateway acts as a cloud router for connecting multiple VPCs and on-premises networks, simplifying network management by creating a hub-and-spoke model for routing traffic. Direct Connect provides a private connection to AWS but does not function as a central router. Amazon

Connect is unrelated, and Route 53 is for DNS services, not VPC connectivity.

Question: 774

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:

AWS Enterprise Support provides strategic planning assistance, infrastructure event management, and real-time support, which are necessary for business-critical applications. Trusted Advisor and APN do not offer direct strategic support, and while AWS Professional Services can assist with complex solutions, Enterprise Support specifically includes ongoing operational support and event management.

Question: 775

A company wants to secure its consumer web application by using SSL/TLS to encrypt traffic. Which AWS service can the company use to meet this goal?

A. AWS WAF
B. AWS Shield
C. Amazon VPC
D. AWS Certificate Manager (ACM)

Answer: D

Explanation:

AWS Certificate Manager (ACM) Overview:
ACM simplifies the process of provisioning, managing, and deploying SSL/TLS certificates.

These certificates are used to secure HTTPS connections, ensuring encrypted communication between clients and servers.

How ACM Meets the Goal:

Provides free public SSL/TLS certificates.

Automates certificate renewals and deployment with services like Elastic Load Balancing and Amazon CloudFront.

Reduces the operational overhead of managing SSL/TLS certificates.

Why Other Options Are Incorrect:

A . AWS WAF:Focuses on application-layer protection (e.g., SQL injection, cross-site scripting) but does not handle SSL/TLS encryption.

B . AWS Shield:Provides protection against DDoS attacks, not encryption.

C . Amazon VPC:Used for networking and security at the infrastructure level but does not manage SSL/TLS certificates.

Reference:

AWS Certificate Manager (ACM) Documentation

Securing Web Applications with ACM

Question: 776

Which AWS service or tool can a company use to set up consolidated billing?

A. AWS Billing and Cost Management console
B. AWS Organizations
C. AWS Cost and Usage Report
D. AWS Systems Manager

Answer: B

Explanation:

AWS Organizations Overview:

AWS Organizations allows a company to manage multiple AWS accounts under a single organization. Through the feature of consolidated billing, a company can receive a single bill for all linked accounts.

Key Features of Consolidated Billing in AWS Organizations:

Aggregates usage across accounts to take advantage of volume pricing discounts.

Provides detailed cost reports for individual linked accounts.

Simplifies payment processing by centralizing billing.

Why Other Options Are Incorrect:

A . AWS Billing and Cost Management console:Used for managing budgets and payments but does not set up consolidated billing.

C . AWS Cost and Usage Report:Provides detailed cost and usage reports but does not set up consolidated billing.

D . AWS Systems Manager:Focuses on operational management, not billing.

Reference:

AWS Organizations Documentation

Question: 777

A company needs to invoke an AWS Step Functions workflow each time an Amazon EC2 instance state changes to RUNNING.

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

A. Amazon SageMaker
B. Amazon Connect
C. Amazon EventBridge
D. AWS Fargate

Answer: C

Explanation:

Amazon EventBridge Overview:

EventBridge is a serverless event bus that enables applications to react to changes in AWS resources. It supports routing events such as EC2 state changes to various targets, including AWS Step Functions.

How EventBridge Meets the Requirement:

EventBridge can capture the EC2 instance state change event and trigger the execution of a Step Functions workflow.

The integration is seamless and supports workflows triggered by multiple event sources.

Why Other Options Are Incorrect:

A . Amazon SageMaker:Used for building, training, and deploying machine learning models; not related to event triggers.

B . Amazon Connect:A cloud-based contact center service; unrelated to event triggers.

D . AWS Fargate:A compute engine for containers; does not manage events or invoke workflows. Reference:

Amazon EventBridge Documentation

Question: 778

A company is preparing for an audit and wants documentation that AWS complies with the Payment Card Industry Data Security Standard (PCI DSS).

Where can the company find this documentation?

A. AWS Artifact
B. AWS Organizations
C. AWS Trusted Advisor
D. AWS Support Center

Answer: A

Explanation:

AWS Artifact Overview:

AWS Artifact is a service that provides on-demand access to AWS compliance documentation and agreements.

It includes reports such as PCI DSS, SOC, and ISO certifications.

How AWS Artifact Meets the Requirement:

The PCI DSS compliance documentation can be downloaded directly from the Artifact console. Artifact helps customers demonstrate compliance to auditors by providing official certifications and attestations.

Why Other Options Are Incorrect:

B . AWS Organizations:Used for managing accounts, not compliance reports.

C . AWS Trusted Advisor:Offers best practice checks but does not provide compliance documentation. D . AWS Support Center:Provides customer support and ticket management but not compliance documentation.

Reference:

AWS Artifact Documentation

Question: 779

A company wants to develop applications that run on AWS. The company's developers need a set of libraries and development tools that are available in multiple programming languages.

Which AWS solution provides these libraries and tools?

A. AWS CodePipeline
B. AWS SDKs
C. Amazon CloudWatch
D. AWS CodeDeploy

Answer: B

Explanation:

AWS SDKs Overview:

AWS Software Development Kits (SDKs) provide libraries and tools for developers to interact with AWS services programmatically.

SDKs are available for multiple programming languages, including Python, Java, JavaScript, and .NET.

How AWS SDKs Meet the Requirement:

Enable developers to integrate AWS services into their applications easily.

Include API clients, authentication helpers, and other utilities specific to AWS services.

Why Other Options Are Incorrect:

A . AWS CodePipeline:Automates CI/CD pipelines but does not provide libraries for development.

C . Amazon CloudWatch:Focuses on monitoring and logging; not a development toolset.

D . AWS CodeDeploy:Automates application deployment but does not include development libraries. Reference:

AWS SDKs Documentation

Question: 780

Which option is a shared responsibility between AWS and its customers under the AWS shared responsibility model?

A. Configuration of Amazon EC2 instance operating systems
B. Application file system server-side encryption
C. Patch management
D. Security of the physical infrastructure

Answer: C

Explanation:

AWS Shared Responsibility Model Overview:

AWS manages securityofthe cloud, including physical infrastructure and foundational services. Customers are responsible for securityinthe cloud, which includes operating system configuration, data encryption, and application patch management.

Why Patch Management Is Shared:

AWS is responsible for patching the underlying infrastructure.

Customers are responsible for patching the operating system and applications they install on their resources (e.g., EC2 instances).

Why Other Options Are Incorrect:

A . Configuration of Amazon EC2 instance operating systems:Fully the customer's responsibility. B . Application file system server-side encryption:Customers configure and manage encryption. D . Security of the physical infrastructure:Fully AWS's responsibility.

Reference:

AWS Shared Responsibility Model

Question: 781

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 Overview:

Amazon Athena is a serverless query service that allows users to analyze data in S3 using standard SQL.

It is commonly used to query AWS Cost and Usage Reports stored in S3.

How It Works for Cost Reports:

Cost and Usage Reports are delivered in a structured format to an S3 bucket.

Athena can query these reports without requiring additional ETL processes.

Why Other Options Are Incorrect:

A . Amazon OpenSearch Service:Designed for search and log analytics, not querying structured data.

C . Amazon Aurora:A relational database service, unsuitable for querying S3 data directly.

D . AWS Glue:Used for ETL tasks but not directly for querying.

Reference:

Analyzing AWS Cost and Usage Reports with Amazon Athena

Question: 782

Which task can an IAM user perform without AWS account root user credentials?

A. Change to a different AWS Support plan.
B. Close an AWS account.
C. View the AWS Billing console.
D. Activate access to the AWS Billing console.

Answer: C

Explanation:

IAM Users and Billing Permissions:

IAM users with appropriate permissions can access and view the AWS Billing console.

The root user is required only for certain account-level tasks, such as changing support plans or closing accounts.

Why Other Options Are Incorrect:

A . Change to a different AWS Support plan:Requires root user credentials.

B . Close an AWS account:Requires root user credentials.

D . Activate access to the AWS Billing console:Requires root user credentials. Reference:

Controlling Access to the Billing Console

Question: 783

A company notices suspicious network activity against an application that is running on a fleet of Amazon EC2 instances. The suspicious activity is coming from a single IP address.

Which AWS service should the company use to block access from this IP address?

A. AWS Shield
B. AWS Config
C. Amazon GuardDuty
D. AWS WAF

Answer: D

Explanation:

AWS WAF Overview:

AWS Web Application Firewall (WAF) allows users to create rules to block or allow traffic based on IP addresses, request patterns, and other conditions.

It is ideal for blocking traffic from a specific IP address.

Why AWS WAF Meets the Requirement:

The company can create a WAF rule to block traffic from the malicious IP address.

WAF integrates with services like Amazon CloudFront, Application Load Balancer, and API Gateway. Why Other Options Are Incorrect:

A . AWS Shield:Protects against DDoS attacks but does not allow custom IP blocking.

B . AWS Config:Monitors resource configurations but does not block IPs.

C . Amazon GuardDuty:Detects threats but does not block traffic directly. Reference:

AWS WAF Documentation

Question: 784

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:

Question: 785

A company's IT administrator needs to configure the AWS CLI for programmatic access to AWS services for the company's employees.
Which combination of credential components must the IT administrator use to meet this requirement? (Select TWO.)

A. A public key
B. A secret access key
C. An IAM role
D. An access key ID
E. A private key

Answer: B,D

Explanation:

Question: 786

Which AWS service or resource can provide discounts on some AWS service costs in exchange for a spending commitment?

A. Amazon Detective
B. AWS Pricing
C. Savings Plans
D. Basic Support

Answer: D

Explanation:

Question: 787

A company is running a key-value NoSQL workload on Amazon EC2 instances. The company needs the workload to have scalability, failover protection, and backup capabilities.
What is the MOST operationally efficient way to meet these requirements?

A. Add additional EC2 instances to the database cluster.

B. Run an identical copy of the database in a second Availability Zone.
C. Migrate the database to Amazon DynamoDB.
D. Migrate the database to a relational database.

Answer: C

Explanation:

Question: 788

Which AWS service gives companies the ability to create graph applications that can analyze billions of relationships between data points in milliseconds?

A. Amazon Redshift
B. Amazon Neptune
C. Amazon DocumentDB (with MongoDB compatibility)
D. Amazon ElastiCache

Answer: B

Explanation:

Question: 789

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

A. Configure identity and access management for applications.
B. Manage encryption options for data that is stored on AWS.
C. Configure security groups for Amazon EC2 instances.
D. Maintain the physical hardware of the infrastructure.

Answer: B

Explanation:

Question: 790

A company needs to reserve a certain amount of Amazon EC2 compute resources in a specific Availability Zone within an AWS Region. Which purchasing option should the company use to meet this requirement?

A. EC2 Instance Savings Plans
B. Compute Savings Plans
C. Regional Reserved Instances
D. Zonal Reserved Instances

Answer: D

Explanation:

Question: 791

A user wants to review all Amazon S3 buckets with ACLs and S3 bucket policies in the S3 console. Which AWS service or resource will meet this requirement?

A. S3 Multi-Region Access Points
B. S3 Storage Lens
C. AWS IAM Identity Center
D. Access Analyzer for S3

Answer: D

Explanation:

Question: 792

A company's workload can recover with minimal downtime when failures occur. Which AWS Cloud benefit does this scenario represent?

A. Agility
B. Elasticity
C. Resiliency
D. Scalability

Answer: C

Explanation:

Question: 793

Which of the following are features of network ACLs as they are used in the AWS Cloud? (Select TWO.)

A. They are stateless.
B. They are stateful.
C. They evaluate all rules before allowing traffic.
D. They process rules in order, starting with the lowest numbered rule, when deciding whether to allow traffic.
E. They operate at the instance level.

Answer: B,D

Explanation:

Question: 794

Which options are benefits of using third-party software from AWS Marketplace? (Select TWO.)

A. The software's data encryption is managed by a third-party vendor.
B. The software has been evaluated by vendors to ensure that it will run on AWS.
C. Users do not need to upgrade to newer software versions.
D. Users do not need to conduct security testing on the software.
E. Users can launch preconfigured software in only a few steps.

Answer: B,E

Explanation:

Question: 795

A company needs to create a portfolio that provides central management of approved IT services. Which AWS service offers this functionality?

A. AWS Service Catalog
B. AWS Control Tower
C. AWS Cloud Map
D. AWS Clean Rooms

Answer: A

Explanation:

Question: 796

Which of the following AWS services are serverless? (Select TWO.)

A. AWS Outposts
B. Amazon EC2
C. Amazon Elastic Kubernetes Service (Amazon EKS)
D. AWS Fargate
E. AWS Lambda

Answer: D,E

Explanation:

Question: 797

Which AWS service allows users to model and provision AWS resources using common programming languages?

A. AWS CloudFormation
B. AWS CodePipeline
C. AWS Cloud Development Kit (AWS CDK)
D. AWS Systems Manager

Answer: C

Explanation:

Question: 798

A company needs a hybrid cloud storage service to connect its on-premises environment to scalable AWS Cloud storage. Which AWS service will meet these requirements?

A. Amazon S3
B. Amazon FSx
C. AWS Storage Gateway
D. AWS Fargate

Answer: C

Explanation:

Question: 799

Which AWS service or feature should a company use between two microservices to ensure that messages are sent and received in exact order?

A. Amazon Simple Email Service (Amazon SES)
B. Amazon Simple Notification Service (Amazon SNS)
C. Amazon S3 Event Notifications
D. Amazon Simple Queue Service (Amazon SQS) FIFO queues

Answer: D

Explanation:

Question: 800

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

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

Answer: A

Explanation:

Question: 801

Which AWS service or tool inspects a user's AWS environment and makes recommendations for cost savings and system performance improvements?

A. Cost Explorer
B. AWS Trusted Advisor
C. Amazon Inspector
D. AWS Budgets

Answer: B

Explanation:

Question: 802

Which AWS service or feature can a company use to create a private, secured, and scalable network environment in the AWS Cloud?

A. Amazon Elastic Container Service (Amazon ECS)
B. Amazon S3
C. Amazon VPC
D. Route tables

Answer: C

Explanation:

Question: 803

A company wants to avoid unnecessary charges and run workloads at the lowest price point. Which pillar of the AWS Well-Architected Framework includes these goals?

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

Answer: D

Explanation:

Question: 804

A company uploads audio and video files to a centralized Amazon S3 bucket from different geographic locations. Which AWS solution will optimize transfer speeds for these files?

A. AWS Global Accelerator
B. S3 Transfer Acceleration
C. AWS Direct Connect
D. Amazon CloudFront

Answer: B

Explanation:

Question: 805

A company wants to control the protection of its AWS resources. The company wants to block SQL injection attacks and cross-site scripting.
Which AWS service or feature meets these requirements?

A. Amazon GuardDuty
B. AWSWAF
C. Security groups
D. AWS Shield

Answer: B

Explanation:

Question: 806

A company is building a business intelligence solution that uses Amazon Redshift. The company wants to use an AWS service to create interactive dashboards and not pay any upfront costs for it. Which service should the company use?

A. Amazon CloudWatch
B. AWS Health Dashboard
C. AWS Service Catalog
D. Amazon QuickSight

Answer: D

Explanation:

Question: 807

Which AWS service or feature can a company use to determine which business unit is using specific

AWS resources?

A. Cost allocation tags
B. Key pairs
C. Amazon Inspector
D. AWS Trusted Advisor

Answer: A

Explanation:

Question: 808

A company is considering a move to the AWS Cloud. The company wants to be able to scale its compute resources as needed to accommodate changing loads.
Which benefit of the AWS Cloud does this scenario describe?

A. Global deployment in minutes
B. Cost savings
C. Agility
D. Elasticity

Answer: D

Explanation:

Question: 809

A company wants to securely rehost databases to AWS with minimal downtime. Which AWS service will meet these requirements?

A. AWS Database Migration Service (AWS DMS)
B. AWS Snow Family
C. AWSDataSync
D. AWS Mainframe Modernization

Answer: A

Explanation:

Question: 810

A company needs to establish a dedicated network connection from on premises to AWS. The connection must provide consistent, low-latency network performance. Which AWS service should the company use to meet this requirement?

A. AWS Direct Connect
B. AWS Site-to-Site VPN
C. AWS Directory Service
D. AWS Transit Gateway

Answer: A

Explanation:

Question: 811

A company is creating a web application that requires a relational database to store customer dat a. Which AWS service should the company use to host the database?

A. Amazon Aurora
B. Amazon DynamoDB
C. Amazon ElastiCache
D. Amazon Redshift

Answer: A

Explanation:

Question: 812

Which AWS service can a company use to build conversational chatbots for customer service?

A. Amazon Lex
B. AWS Amplify
C. Amazon Comprehend
D. Amazon Polly

Answer: A

Explanation:

Question: 813

A company needs to request temporary, limited-privilege credentials for IAM users and for the federated users that the company authenticates.
Which AWS service will provide these credentials?

A. Amazon GuardDuty
B. AWS Key Management Service (AWS KMS)
C. AWS Security Token Service (AWS STS)
D. AWS Identity and Access Management Access Analyzer

Answer: C

Explanation:

Question: 814

A company is running Amazon EC2 instances in a private subnet in a VPC.

Which AWS service or feature can provide the EC2 instances with network connections to the internet?

A. Gateway endpoint
B. NAT gateway
C. Network Load Balancer
D. Amazon Route 53

Answer: B

Explanation:

Question: 815

Treating infrastructure as code in the AWS Cloud allows users to:

A. automate migration of on-premises hardware to AWS data centers.
B. let a third party automate an audit of the AWS infrastructure.
C. turn over application code to AWS so it can run on the AWS infrastructure.
D. automate the infrastructure provisioning process.

Answer: D

Explanation:

Question: 816

A cloud practitioner wants a repeatable way to deploy identical AWS resources by using infrastructure templates. Which AWS service will meet these requirements?

A. AWS CloudFormation
B. AWS Directory Service
C. Amazon Lightsail
D. AWS CodeDeploy

Answer: A

Explanation:

Question: 817

A company needs to mount a file share across multiple Amazon EC2 instances as a mapped drive by using the SMB protocol. Which AWS service will meet these requirements?

A. Amazon FSx for Windows File Server
B. Amazon Elastic File System (Amazon EFS)
C. Amazon S3
D. AWS DataSync

Answer: A

Explanation:

Question: 818

A company uses AWS and has a VPC that includes two public subnets. The company needs to allow and deny specific inbound and outbound traffic for each public subnet.
Which AWS service or tool can the company use to meet this requirement?

A. Network ACL
B. AWSWAF
C. VPC route table entry
D. Security group

Answer: A

Explanation:

Question: 819

A company wants to automatically run operating system command scripts on Amazon EC2 instances. Which AWS service will meet these requirements in the MOST operationally efficient way?

A. AWS Organizations
B. AWS Control Tower
C. AWS Lambda
D. AWS Systems Manager

Answer: D

Explanation:

Question: 820

A company wants to know more about the benefits offered by cloud computing. The company wants to understand the operational advantage of agility.
How does AWS provide agility for users?

A. The ability the ensure high availability by deploying workloads to multiple regions.
B. A pay-as-you-go model for many services and resources
C. The ability to transfer infrastructure management to the AWS Cloud
D. The ability to provision and deprovision resources quickly with minimal effort

Answer: D

Explanation:

Question: 821

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:

Question: 822

A company is undergoing a security audit. The audit includes security validation and compliance validation of the AWS infrastructure and services that the company uses. The auditor needs to locate compliance-related information and must download AWS security and compliance documents. These documents include the System and Organization Control (SOC) reports.

A. AWS Abuse team
B. AWS Artifact
C. AWS Support
D. AWS Config

Answer: B

Explanation:

Question: 823

A company needs to automatically protect its Amazon EC2 instances from distributed denial of service (DDoS) attacks.

A. Network access control list (ACL)
B. AWS Shield
C. Security group

D. Amazon GuardDuty

Answer: B

Explanation:

Question: 824

Which perspective of the AWS Cloud Adoption Framework (AWS CAF) encourages the development of well-architected, cloud-focused applications?

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

Answer: B

Explanation:

Question: 825

A company wants to securely log in to Linux Amazon EC2 instances.

A. Use end-to-end encryption.
B. Use multi-factor authentication (MFA).
C. Use AWS Systems Manager Session Manager.
D. Use AWS Systems Manager State Manager.

Answer: C

Explanation:

Question: 826

A company is designing AWS architecture that will add compute resources when the company needs them. The architecture also includes a disaster recovery plan with automatic failover.

A. Reliability
B. Operational excellence
C. Security
D. Performance efficiency

Answer: A

Explanation:

Question: 827

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.

A. AWS Lambda
B. Amazon EventBridge
C. Elastic Load Balancing (ELB)
D. Amazon EC2 Auto Scaling

Answer: D

Explanation:

Question: 828

A company runs Amazon EC2 instances in a research lab. The instances run for 3 hours each week and cannot be interrupted.

A. Compute Savings Plan
B. On-Demand Instances
C. Convertible Reserved Instances
D. Spot Instances

Answer: B

Explanation:

Question: 829

A company wants to use a template to reliably provision, manage, and update its infrastructure in the AWS Cloud.

A. AWS Lambda
B. AWS CloudFormation
C. AWS Fargate
D. AWS CodeDeploy

Answer: B

Explanation:

Question: 830

A company wants phone, email, and chat access to AWS support engineers. The company also wants the ability to get support for

business-critical systems within 15 minutes.

Which AWS Support plan meets these requirements?

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

Answer: C

Explanation:

Question: 831

A company wants to push VPC Flow Logs to an Amazon S3 bucket. Which action is the company's responsibility?

A. Managing the infrastructure that runs the S3 bucket
B. Managing the data in transit
C. Managing the encryption options on the S3 bucket
D. Managing the operating system updates on the S3 bucket

Answer: C

Explanation:

Question: 832

Which AWS service helps assess the security and compliance of applications that are deployed on Amazon EC2 instances?

A. AWS Security Hub
B. Amazon Inspector
C. Amazon GuardDuty
D. AWS Shield

Answer: B

Explanation:

Question: 833

A company is planning to set up a new application in the AWS Cloud. The company needs a complete estimate of the AWS expenses that the application is likely to incur.

A. AWS Trusted Advisor
B. AWS Cost Explorer
C. AWS Price List API

D. AWS Pricing Calculator

Answer: D

Explanation:

Question: 834

Which of the following is an architectural design principle of the AWS Well-Architected Framework?

A. Loosely couple components
B. Build monolithic systems
C. Scale vertically, not horizontally
D. Use third-party software

Answer: A

Explanation:

Question: 835

A company needs a file-sharing service that supports SMB protocol.

A. Amazon Aurora
B. AWS Config
C. AWS DataSync
D. Amazon FSx for Windows File Server

Answer: D

Explanation:

Question: 836

An Amazon EC2 instance previously used for development is inaccessible and no longer appears in the AWS Management Console. Which AWS service should be used to determine what action made this EC2 instance inaccessible?

A. Amazon CloudWatch Logs
B. AWS Security Hub
C. Amazon Inspector
D. AWS CloudTrail

Answer: D

Explanation:

Question: 837

Which task can a user complete by using AWS Identity and Access Management (IAM)?

A. Validate JSON syntax from an application configuration file.
B. Analyze logs from an Amazon API Gateway call.
C. Filter traffic to or from an Amazon EC2 instance.
D. Grant permissions to applications that run on Amazon EC2 instances.

Answer: D

Explanation:

Question: 838

Which AWS service or tool can be used to capture information about inbound and outbound traffic in an Amazon VPC?

A. Amazon Inspector
B. VPC endpoint services
C. VPC Flow Logs
D. NAT gateway

Answer: C

Explanation:

Question: 839

An online gaming company hosts its gaming application on AWS. The company needs to decrease latency by sending internet traffic to the application through AWS infrastructure.

A. Amazon Route 53
B. AWS VPN
C. AWS Direct Connect
D. AWS Global Accelerator

Answer: D

Explanation:

Question: 840

A company is developing a web application that uses containers. The company needs a repository to store and control access to container images.

A. Amazon Elastic Container Registry (Amazon ECR)
B. Amazon Elastic Kubernetes Service (Amazon EKS)
C. Amazon Elastic Container Service (Amazon ECS)
D. Amazon EC2

Answer: A

Explanation:

Question: 841

A company is moving its on-premises IT services to the AWS Cloud. The company wants to set spending limits and to receive notifications if the limits are exceeded.
Which AWS service or resource will meet these requirements?

A. AWS Budgets
B. AWS Cost and Usage Reports
C. AWS Cost Explorer
D. AWS Organizations consolidated billing

Answer: A

Explanation:

Question: 842

A company needs a database that can quickly read and write large amounts of unstructured data. Which AWS service meets these requirements with the MOST scalability?

A. Amazon RDS
B. Amazon Elastic Block Store (Amazon EBS)
C. Amazon DynamoDB
D. Amazon Lightsail

Answer: C

Explanation:

Question: 843

A company plans to host its data warehouse application on AWS. The company has a machine learning (ML) model and wants to use that model within its data warehouse for data forecasting.

A. Amazon DynamoDB
B. Amazon Redshift ML

C. Amazon Aurora ML

D. Amazon MemoryDB

Answer: B

Explanation:

Question: 844

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:

Question: 845

A company wants to build, train, and deploy machine learning (ML) models. Which AWS service will meet these requirements?

A. Amazon Athena

B. Amazon Comprehend

C. Amazon Polly

D. Amazon SageMaker AI

Answer: D

Explanation:

Question: 846

Which AWS service or feature gives users the ability to run containers and Kubernetes applications without the need to manage scaling?

A. Amazon Elastic Container Registry (Amazon ECR)

B. Elastic network interface

C. AWS Fargate

D. Amazon EC2

Answer: C

Explanation:

Question: 847

A company wants to monitor and block malicious HTTP and HTTPS requests that its Amazon CloudFront distributions receive.
Which AWS service should the company use to meet these requirements?

A. Amazon GuardDuty
B. Amazon Inspector
C. AWS WAF
D. Amazon Detective

Answer: C

Explanation:

Question: 848

Which cloud concept is demonstrated by using AWS Cost Explorer?

A. Rightsizing
B. Reliability
C. Resilience
D. Modernization

Answer: A

Explanation:

Question: 849

Which AWS service provides a cloud-based customer contact center?

A. AWS CloudFormation
B. Amazon CloudWatch
C. Amazon Connect
D. AWS Direct Connect

Answer: C

Explanation:

Question: 850

Which AWS service can a company use to send SMS messages and email messages from applications?

A. AWS Direct Connect
B. Amazon Simple Email Service (Amazon SES)
C. Amazon Simple Notification Service (Amazon SNS)
D. Amazon Simple Queue Service (Amazon SQS)

Answer: C