Exam : **CLF-C02**

Title : AWS Certified Cloud Practitioner

Vendor : Amazon

Version : DEMO

NO.1 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.

NO.2 A company wants to define a central data protection policy that works across AWS services for compute, storage, and database resources.
Which AWS service will meet this requirement?
A. AWS Batch
B. AWS Elastic Disaster Recovery
C. AWS Backup
D. Amazon FSx
Answer: C
Explanation:
The AWS service that will meet this requirement is C. AWS Backup.
AWS Backup is a service that allows you to define a central data protection policy that works across AWS services for compute, storage, and database resources. You can use AWS Backup to create backup plans that specify the frequency, retention, and lifecycle of your backups, and apply them to your AWS resources using tags or resource IDs.AWS Backup supports various AWS services, such as Amazon EC2,Amazon EBS, Amazon RDS, Amazon DynamoDB, Amazon EFS, Amazon FSx, and AWS Storage Gateway12.
AWS Batch is a service that allows you to run batch computing workloads on AWS.AWS Batch does not provide a central data protection policy, but rather enables you to optimize the allocation and utilization of your compute resources3.
AWS Elastic Disaster Recovery is a service that allows you to prepare for and recover from disasters using AWS.AWS Elastic Disaster Recovery does not provide a central data protection policy, but rather helps you minimize downtime and data loss by replicating your applications and data to AWS4.
Amazon FSx is a service that provides fully managed file storage for Windows and Linux applications. Amazon FSx does not provide a central data protection policy, but rather offers features such as encryption, snapshots, backups, and replication to protect your file systems5.
References:

1:AWS Backup - Centralized backup across AWS services3:AWS Batch - Run Batch Computing Jobs o n 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

NO.3 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.

NO.4 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

NO.5 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].

NO.6 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.

NO.7 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 . References: 7: VPC peering - Amazon Virtual Private Cloud, 8: Connect VPCs using VPC peering - Amazon Virtual Private Cloud, 9: Amazon VPC-to- Amazon VPC connectivit y options - Amazon Virtual Private Cloud, : [AWS Transit Gateway - Amazon Web Services], : [Connec t VPCs using AWS Transit Gateway - Amazon Virtual Private Cloud], : [AWS Transit Gateway: Simplify Your Network Architecture]

NO.8 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.
References:
AWS Budgets Documentation

NO.9 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 References:
Amazon CloudFront

NO.10 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 .

NO.11 A company wants to set up a high-speed connection between its data center and its applications that run on AWS. The company must not transfer data over the internet.
Which action should the company take to meet these requirements?
A. Transfer data to AWS by using AWS Snowball.
B. Transfer data to AWS by using AWS Storage Gateway.
C. Set up a VPN connection between the data center and an AWS Region.
D. Set up an AWS Direct Connect connection between the company network and AWS.
Answer: D
Explanation:
AWS Direct Connect is a cloud service solution that makes it easy to establish a dedicated network connection from a customer's premises to AWS. AWS Direct Connect does not involve the public internet, and therefore can reduce network costs, increase bandwidth throughput, and provide a more consistent network experience than internet-based connections. AWS Snowball is a petabyte-scale data transport service that uses secure devices to transfer large amounts of data into and out of the AWS Cloud. AWS Storage Gateway is a hybrid cloud storage service that gives customers on-premises access to virtually unlimited cloud storage. A VPN connection enables customers to establish a secure and private connection between their network and AWS.

NO.12 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. References: Logging AWS Audit Manager API calls with CloudTrail, Logging AWS Account Management API calls using AWS CloudTrail, Review API calls in your AWS account using CloudTrail, Monitor the usage of AWS API calls using Amazon CloudWatch, Which service enables customers to audit API calls in their AWS ...

NO.13 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.

NO.14 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 offline 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.

NO.15 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.

NO.16 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 cost-effective option, but it has no redundancy or fault tolerance.A single-AZ deployment model may not be suitable for a company that wants to deploy an application globally, as it may face challenges such as network latency, regional outages, or regulatory compliance34.

References:

1:AWS Cloud Computing - W3Schools2:Understand the Different Cloud Computing Deployment Models Unit - Trailhead3:Regions and Availability Zones - Amazon Elastic Compute Cloud4:AWS Reference Architecture Diagrams

NO.17 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

NO.18 A company's IT team is managing MySQL database server clusters. The IT team has to patch the database and take backup snapshots of the data in the clusters. The company wants to move this workload to AWS so that these tasks will be completed automatically.
What should the company do to meet these requirements?

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

B. Use Amazon RDS with a MySQL database.

C. Use an AWS CloudFormation template to deploy MySQL database servers on Amazon EC2 instances.

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

Answer: B

Explanation:
Amazon RDS is a service that makes it easy to set up, operate, and scale a relational database in the cloud.
Amazon RDS supports MySQL as one of the database engines. By using Amazon RDS with a MySQL database, the company can offload the tasks of patching the database and taking backup snapshots to AWS.
Amazon RDS automatically patches the database software and operating system of the database instances.
Amazon RDS also automatically backs up the database and retains the backups for a user-defined retention period. The company can also restore the database to any point in time within the retention period. Deploying MySQL database server clusters 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

NO.19 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

NO.20 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.

NO.21 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.

NO.22 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

NO.23 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?

NO.24 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.

NO.25 A company operates a petabyte-scale data warehouse to analyze its data. 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 petabyte-scale 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]

NO.26 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

NO.27 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

NO.28 A company wants to migrate its applications from its on-premises data center to a VPC in the AWS Cloud.
These applications will need to access on-premises resources. Which actions will meet these requirements?
(Select TWO.)
A. Use AWS Service Catalog to identify a list of on-premises resources that can be migrated.
B. Create a VPN connection between an on-premises device and a virtual private gateway in the VPC.
C. Use an Amazon CloudFront distribution and configure it to accelerate content delivery close to the on- premises resources
D. Set up an AWS Direct Connect connection between the on-premises data center and AWS.
E. Use Amazon CloudFront to restrict access to static web content provided through the on-premises web servers.
Answer: B D

NO.29 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.

NO.30 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.

NO.31 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.

NO.32 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.

NO.33 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.

NO.34 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.

NO.35 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.
References:
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