LatestCram

LatestCram

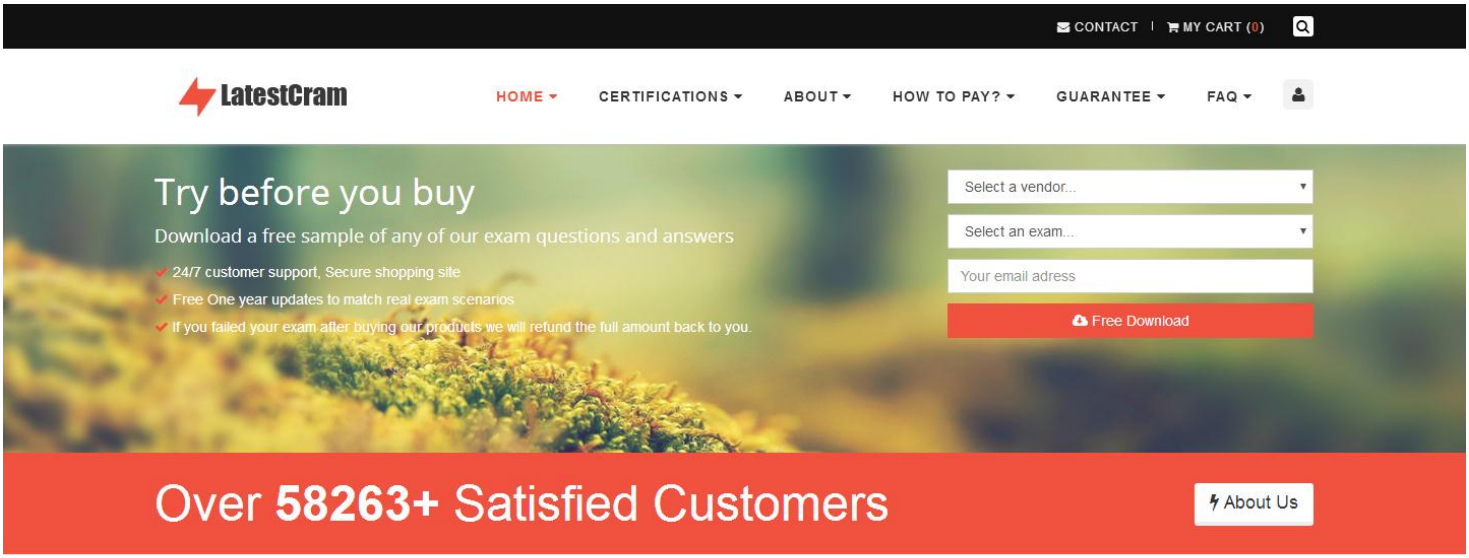

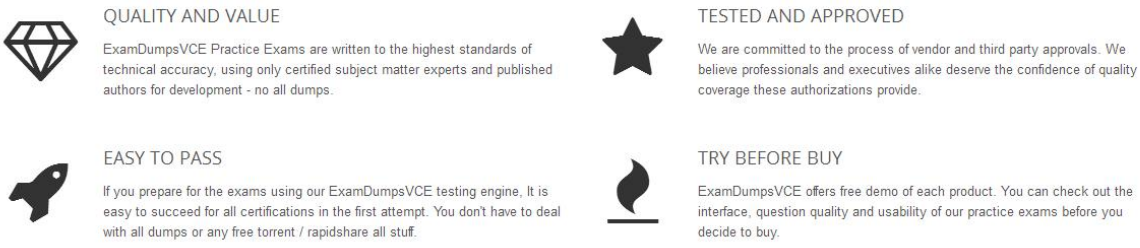

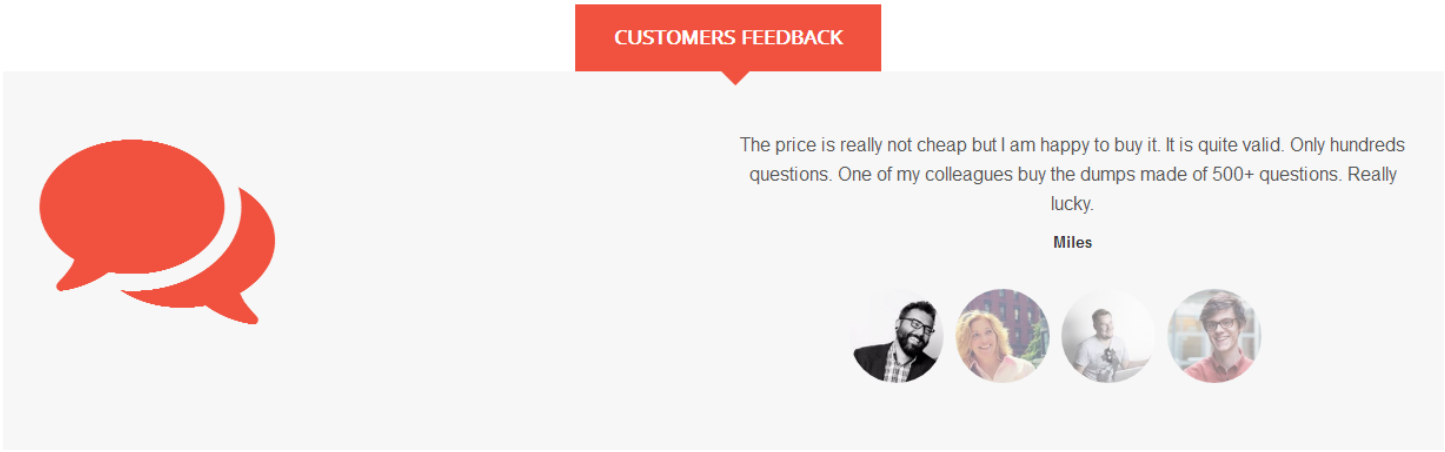

http://www.latestcram.com

Reliable Exam Braindumps & Valid Latest Questions & Right Exam Cram

Exam : **AWS-Solutions-Architect-Associate**

Title : AWS Certified Solutions Architect - Associate (SAA-C03)

Vendor : Amazon

Version : DEMO

NO.1 A company runs an ecommerce platform with a monolithic architecture on Amazon EC2 instances. The platform runs web and API services. The company wants to decouple the architecture and enhance scalability.
The company also wants the ability to track orders and reprocess any failed orders.
Which solution will meet these requirements?

A. Send orders to an Amazon Simple Queue Service (Amazon SQS) queue. Configure AWS Lambda functions to consume the queue and process orders. Implement an SQS dead-letter queue.

B. Send orders to an Amazon Simple Queue Service (Amazon SQS) queue. Configure Amazon Elastic Container Service (Amazon ECS) tasks to consume the queue. Implement SQS visibility timeout.

C. Use Amazon Kinesis Data Streams to queue orders. Use AWS Lambda functions to consume the data stream. Configure Amazon S3 to track and reprocess failed orders.

D. Send orders to an Amazon Simple Queue Service (Amazon SQS) queue. Configure AWS Lambda functions to consume the queue and process orders. Configure the Lambda functions to use SQS long polling.

Answer: A
Explanation:
* To decouple the monolith and enhance scalability, AWS best practice is to introduce an asynchronous message queue, such as Amazon SQS, between the web/API tier and the order-processing logic.
* AWS Lambda functions consuming from the SQS queue provide serverless, auto-scaling processing without managing servers.
* To track and reprocess failed orders, SQS supports dead-letter queues (DLQs). Messages that cannot be processed successfully after a configurable number of attempts are automatically moved to the DLQ, where operations teams or automated processes can inspect and reprocess them.
Why others are not correct:
* B: ECS tasks can consume an SQS queue, but this requires managing container infrastructure and does not inherently provide as simple reprocessing/visibility as combining Lambda with a DLQ. Visibility timeout is not a tracking or archival mechanism.
* C: Kinesis is a streaming service designed for ordered event streams, not primarily for order-queue semantics and DLQs; SQS is simpler and purpose-built for this pattern.
* D: Long polling reduces empty responses and API calls but does nothing for tracking or reprocessing failed messages; without a DLQ, failed orders are harder to manage

NO.2 A company has an application that runs on a single Amazon EC2 instance. The application uses a MySQL database that runs on the same EC2 instance. The company needs a highly available and automatically scalable solution to handle increased traffic.
Which solution will meet these requirements?

A. Deploy the application to EC2 instances that run in an Auto Scaling group behind an Application Load Balancer. Create an Amazon Redshift cluster that has multiple MySQL-compatible nodes.

B. Deploy the application to EC2 instances that are configured as a target group behind an Application Load Balancer. Create an Amazon RDS for MySQL cluster that has multiple instances.

C. Deploy the application to EC2 instances that run in an Auto Scaling group behind an Application Load Balancer. Create an Amazon Aurora Serverless MySQL cluster for the database layer.

D. Deploy the application to EC2 instances that are configured as a target group behind an Application Load Balancer. Create an Amazon ElastiCache (Redis OSS) cluster that uses the MySQL

connector.

Answer: C

Explanation:

Amazon Aurora Serverless is a fully managed, MySQL-compatible database that automatically scales based on demand and provides high availability. Combining this with EC2 Auto Scaling and an Application Load Balancer achieves both application and database high availability and scalability.

Reference Extract:

"Aurora Serverless automatically starts up, shuts down, and scales capacity based on your application's needs, providing a cost-effective, highly available database solution." Source: AWS Certified Solutions Architect - Official Study Guide, Aurora Serverless and Scaling section.

NO.3 A solutions architect is designing a multi-Region disaster recovery (DR) strategy for a company. The company runs an application on Amazon EC2 instances in Auto Scaling groups that are behind an Application Load Balancer (ALB). The company hosts the application in the company's primary and secondary AWS Regions.

The application must respond to DNS queries from the secondary Region if the primary Region fails. Only one Region must serve traffic at a time.

Which solution will meet these requirements?

A. Create an outbound endpoint in Amazon Route 53 Resolver. Create forwarding rules that determine how queries will be forwarded to DNS resolvers on the network. Associate the rules with VPCs in each Region.

B. Create primary and secondary DNS records in Amazon Route 53. Configure health checks and a failover routing policy.

C. Create a traffic policy in Amazon Route 53. Use a geolocation routing policy and a value type of ELB Application Load Balancer.

D. Create an Amazon Route 53 profile. Associate DNS resources to the profile. Associate the profile with VPCs in each Region.

Answer: B

Explanation:

Amazon Route 53 supports failover routing policies, which use health checks to route DNS queries to a secondary Region only if the primary endpoint fails. This design ensures only one Region is active for traffic at any given time. This is the recommended architecture for active-passive, multi-Region DR strategies.

AWS Documentation Extract:

"Failover routing lets you route traffic to a primary resource, such as a web server in one Region, and a secondary resource in another Region. If the primary fails, Route 53 can route traffic to the secondary resource automatically." (Source: Amazon Route 53 documentation, Routing Policy Types)

A, D: These options do not configure DNS failover for external users.

C: Geolocation routing is for regional distribution, not DR failover.

Reference: AWS Certified Solutions Architect - Official Study Guide, Multi-Region DR and Route 53.

NO.4 A company is developing an application using Amazon Aurora MySQL. The team will frequently make schema changes to test new features without affecting production. After testing, changes must be promoted to production with minimal downtime.

Which solution meets these requirements?

A. Create a staging Aurora cluster based on the existing cluster. Test schema changes on the staging cluster.

B. Create a read replica, modify its schema, and then promote it to primary.

C. Create an Aurora MySQL blue/green deployment. Make schema changes in the staging environment and switch traffic after testing.

D. Replicate the Aurora database to DynamoDB, apply schema changes, and switch the application to DynamoDB.

Answer: C

Explanation:

Aurora blue/green deployments are specifically designed for safe schema changes, zero-downtime updates, and production isolation.

The staging (green) environment can receive schema changes without affecting production (blue).

After validation, you perform a fast, minimally disruptive switchover that updates production.

Read replicas (Option B) do not allow schema changes. Creating an independent staging cluster (Option A) does not provide automated, low-downtime cutover. DynamoDB (Option D) is not compatible with MySQL schemas.

NO.5 A company has an application that receives and processes purchase orders. The application supports only XML data. The company needs to configure the application to accept orders in JSON format. The company does not want to modify the application.

A solutions architect is using an Amazon API Gateway HTTP API to create a new purchase order API. The solutions architect needs to modify the application DNS record to point to the new HTTP API.

A. Use an HTTP proxy integration to pass XML requests to the application. For JSON requests, use API Gateway mappings to convert the purchase orders to XML. Use an AWS Lambda function that is integrated with API Gateway to call the application.

B. Use an HTTP proxy integration to pass XML requests to the application. For JSON requests, use an AWS Lambda function that is integrated with API Gateway to convert the purchase orders from JSON to XML and to call the application.

C. Use an HTTP custom integration to pass XML requests to the application. For JSON requests, use API Gateway mappings to convert the purchase orders to XML. Use an AWS Lambda function that is integrated with API Gateway to call the application.

D. Use an HTTP custom integration to pass XML requests to the application. For JSON requests, use an AWS Lambda function that is integrated with API Gateway to convert the purchase orders to JSON and to call the application.

Answer: B

Explanation:

Why Option B is Correct:

HTTP Proxy Integration: Passes XML requests directly to the application, which already supports XML.

JSON Conversion: An AWS Lambda function converts JSON requests to XML and calls the application.

API Gateway: Acts as a front end to handle JSON requests and integrates seamlessly with Lambda for the transformation process.

Why Other Options Are Not Ideal:

Option A: Suggests using API Gateway mappings to convert JSON to XML. API Gateway mapping templates are limited in functionality and are not ideal for complex transformations.

Option C and D: Use HTTP custom integration unnecessarily, which adds complexity without

additional benefits.
AWS References:
Amazon API Gateway Integration:AWS Documentation - API Gateway Integration AWS Lambda:AWS Documentation - Lambda

NO.6 A company is designing an application to connect AWS Lambda functions to an Amazon RDS for MySQL DB instance. The DB instance manages many connections. The company needs to modify the application to improve connectivity and recovery.
Which solution will meet these requirements with the LEAST operational overhead?
A. Use Amazon RDS Proxy for connection pooling. Modify the application to use the RDS Proxy for connections to the DB instance.
B. Create a new RDS instance for connection pooling. Modify the application to use the new RDS instance for connectivity.
C. Create read replicas to distribute the load of the DB instance. Create a Network Load Balancer to distribute the load across the read replicas.
D. Migrate the RDS for MySQL DB instance to Amazon Aurora MySQL to increase DB instance performance.
Answer: A
Explanation:
Amazon RDS Proxy helps manage thousands of concurrent database connections by pooling and reusing them efficiently. It is especially useful for serverless applications like AWS Lambda that can open numerous connections quickly, potentially overwhelming the database. Using RDS Proxy reduces connection management overhead and improves fault tolerance.
Reference: AWS Documentation - Amazon RDS Proxy

NO.7 A company is setting up a development environment on AWS for a team of developers. The team needs to access multiple Amazon S3 buckets to store project data. The team also needs to use Amazon EC2 to run development instances.
The company needs to ensure that the developers have access only to specific Amazon S3 buckets and EC2 instances. Access permissions must be assigned according to each developer's role on the team. The company wants to minimize the use of permanent credentials and to ensure access is securely managed according to the principle of least privilege.
Which solution will meet these requirements?
A. Create IAM roles that have administrative-level permissions for Amazon S3 and Amazon EC2. Require developers to sign in by using Amazon Cognito to access Amazon S3 and Amazon EC2.
B. Create IAM roles that have fine-grained permissions for Amazon S3 and Amazon EC2. Configure AWS IAM Identity Center to manage credentials for the developers.
C. Create IAM users that have programmatic access to Amazon S3 and Amazon EC2. Generate individual access keys for each developer to access Amazon S3 and Amazon EC2.
D. Create a VPC endpoint for Amazon S3. Require developers to access Amazon EC2 instances and Amazon S3 buckets through a bastion host.
Answer: B
Explanation:
The most secure and manageable way to provide developers with temporary, least-privilege access is by using AWS IAM Identity Center (formerly AWS SSO). IAM Identity Center allows assigning IAM

roles with scoped permissions based on the developer's team role. This ensures no permanent credentials are required and minimizes risk.
Option B enables role-based access with centralized identity and access management, making it the most secure and scalable solution for managing developer permissions.

NO.8 A solutions architect needs to secure an Amazon API Gateway REST API. Users need to be able to log in to the API by using common external social identity providers (IdPs). The social IdPs must use standard authentication protocols such as SAML or OpenID Connect (OIDC). The solutions architect needs to protect the API against attempts to exploit application vulnerabilities.
Which combination of steps will meet these security requirements? (Select TWO.)
A. Create an AWS WAF web ACL that is associated with the REST API. Add the appropriate managed rules to the ACL.
B. Subscribe to AWS Shield Advanced. Enable DDoS protection. Associate Shield Advanced with the REST API.
C. Create an Amazon Cognito user pool with a federation to the social IdPs. Integrate the user pool with the REST API.
D. Create an API key in API Gateway. Associate the API key with the REST API.
E. Create an IP address filter in AWS WAF that allows only the social IdPs. Associate the filter with the web ACL and the API.
Answer: A C
Explanation:
Step A:AWS WAF with managed rules protects the API against application-layer attacks, such as SQL injection and cross-site scripting (XSS).
Step C:Amazon Cognito provides secure authentication and supports federation with social IdPs using OIDC or SAML. It integrates seamlessly with API Gateway.
Option B:AWS Shield Advanced provides DDoS protection, which is not explicitly required in this scenario.
Option D:API keys provide identification, not authentication, and are insufficient for this use case.
Option E:IP filters in WAF are overly restrictive for federated authentication scenarios.
AWS Documentation References:
Amazon Cognito Federation
AWS WAF Managed Rules

NO.9 A retail company runs its application on AWS. The application uses Amazon EC2 for web servers, Amazon RDS for database services, and Amazon CloudFront for global content distribution. The company needs a solution to mitigate DDoS attacks.
Which solution will meet this requirement?
A. Implement AWS WAF custom rules to limit the length of query requests. Configure CloudFront to work with AWS WAF.
B. Enable AWS Shield Advanced. Configure CloudFront to work with Shield Advanced.
C. Use Amazon Inspector to scan the EC2 instances. Enable Amazon GuardDuty.
D. Enable Amazon Macie. Configure CloudFront Origin Shield.
Answer: B
Explanation:
AWS Shield Advanced provides advanced DDoS protection for AWS workloads, including EC2,

CloudFront, and RDS. When integrated with CloudFront, Shield Advanced offers comprehensive detection and mitigation against large and sophisticated DDoS attacks, along with 24x7 access to the AWS DDoS Response Team (DRT). AWS WAF provides application-level protection, but for complete DDoS mitigation, Shield Advanced is the recommended solution.
Reference Extract from AWS Documentation / Study Guide:
"AWS Shield Advanced provides expanded DDoS attack protection for applications running on AWS. It offers always-on detection and automatic inline mitigations that minimize application downtime and latency." Source: AWS Certified Solutions Architect - Official Study Guide, Security and DDoS Protection section.

NO.10 A company is migrating a daily Microsoft Windows batch job from the company's on-premises environment to AWS. The current batch job runs for up to 1 hour. The company wants to modernize the batch job process for the cloud environment.
Which solution will meet these requirements with the LEAST operational overhead?
A. Create a fleet of Amazon EC2 instances in an Auto Scaling group to handle the Windows batch job processing.
B. Implement an AWS Lambda function to process the Windows batch job. Use an Amazon EventBridge rule to invoke the Lambda function.
C. Use AWS Fargate to deploy the Windows batch job as a container. Use AWS Batch to manage the batch job processing.
D. Use Amazon Elastic Kubernetes Service (Amazon EKS) on Amazon EC2 instances to orchestrate Windows containers for the batch job processing.
Answer: C
Explanation:
AWS Batch supports Windows-based jobs and automates provisioning and scaling of compute environments.
Paired with AWS Fargate, it removes the need to manage infrastructure. This solution requires the least operational overhead and is cloud-native, providing flexibility and scalability.
Reference: AWS Documentation - AWS Batch with Fargate for Windows Workloads

NO.11 Question:
A machine learning (ML) team is building an application that uses data that is in an Amazon S3 bucket. The ML team needs a storage solution for its model training workflow on AWS. The ML team requires high- performance storage that supports frequent access to training datasets. The storage solution must integrate natively with Amazon S3. Which solution will meet these requirements with the LEAST operational overhead?
Options:
A. Use Amazon Elastic Block Store (Amazon EBS) volumes to provide high-performance storage. Use AWS DataSync to migrate data from the S3 bucket to EBS volumes.
B. Use Amazon EC2 ML instances to provide high-performance storage. Store training data on Amazon EBS volumes. Use the S3 Copy API to copy data from the S3 bucket to EBS volumes.
C. Use Amazon FSx for Lustre to provide high-performance storage. Store training datasets in Amazon S3 Standard storage.
D. Use Amazon EMR to provide high-performance storage. Store training datasets in Amazon S3 Glacier Instant Retrieval storage.

Answer: C
Explanation:
Amazon FSx for Lustre is a high-performance file system optimized for fast processing of workloads such as machine learning, high-performance computing (HPC), and video processing. It integrates natively with Amazon S3, allowing you to:
Access S3 Data: FSx for Lustre can be linked to an S3 bucket, presenting S3 objects as files in the file system.
High Performance: It provides sub-millisecond latencies, high throughput, and millions of IOPS, which are ideal for ML workloads.Amazon Web Services, Inc.
Minimal Operational Overhead: Being a fully managed service, it reduces the complexity of setting up and managing high-performance file systems.
References:
Amazon FSx for Lustre - High-Performance File System Integrated with S3Amazon Web Services, Inc.
What is Amazon FSx for Lustre?

NO.12 A company plans to use an Amazon S3 bucket to archive backup data. Regulations require the company to retain the backup data for 7 years.
During the retention period, the company must prevent users, including administrators, from deleting the data.
The company can delete the data after 7 years.
Which solution will meet these requirements?
A. Create an S3 bucket policy that denies delete operations for 7 years. Create an S3 Lifecycle policy to delete the data after 7 years.
B. Create an S3 Object Lock default retention policy that retains data for 7 years in governance mode .
Create an S3 Lifecycle policy to delete the data after 7 years.
C. Create an S3 Object Lock default retention policy that retains data for 7 years in compliance mode.
Create an S3 Lifecycle policy to delete the data after 7 years.
D. Create an S3 Batch Operations job to set a legal hold on each object for 7 years. Create an S3 Lifecycle policy to delete the data after 7 years.
Answer: C
Explanation:
Comprehensive and Detailed Step-by-Step Explanation:
The requirement is toprevent data deletion by any user, including administrators, for 7 years while allowing automatic deletion afterward.
S3 Object Lock in Compliance Mode (Correct Choice - C)
Compliance mode ensures that even the root user cannot delete or modify the objects during the retention period.
After 7 years, the S3 Lifecycle policy automatically deletes the objects.
This meets bothimmutability and automatic deletionrequirements.
Governance Mode (Option B - Incorrect)
Governance mode prevents deletion,but administrators can override it.
The requirement explicitly states thateven administrators must not be able to delete the data.
S3 Bucket Policy (Option A - Incorrect)
An S3 bucket policy candeny deletes, but policies can be modified at any time by administrators.

It does not enforce strict retention like Object Lock.
S3 Batch Operations Job (Option D - Incorrect)
A legal hold does not have an automatic expiration.
Legal holds must be manually removed, which is not efficient.
Why Option C is Correct:
S3 Object Lock in Compliance Mode prevents deletion by all users, including administrators.
The S3 Lifecycle policy deletes the data automatically after 7 years, reducing operational overhead.
References:
S3 Object Lock Compliance Mode
S3 Lifecycle Policies

NO.13 A solutions architect needs to optimize a large data analytics job that runs on an Amazon EMR cluster. The job takes 13 hours to finish. The cluster has multiple core nodes and worker nodes deployed on large, compute-optimized instances.
After reviewing EMR logs, the solutions architect discovers that several nodes are idle for more than 5 hours while the job is running. The solutions architect needs to optimize cluster performance.
Which solution will meet this requirement MOST cost-effectively?

A. Increase the number of core nodes to ensure there is enough processing power to handle the analytics job without any idle time.

B. Use the EMR managed scaling feature to automatically resize the cluster based on workload.

C. Migrate the analytics job to a set of AWS Lambda functions. Configure reserved concurrency for the functions.

D. Migrate the analytics job core nodes to a memory-optimized instance type to reduce the total job runtime.

Answer: B
Explanation:
EMR managed scaling dynamically resizes the cluster by adding or removing nodes based on the workload.
This feature helps minimize idle time and reduces costs by scaling the cluster to meet processing demands efficiently.
Option A:Increasing the number of core nodes might increase idle time further, as it does not address the root cause of underutilization.
Option C:Migrating the job to Lambda is infeasible for large analytics jobs due to resource and runtime constraints.
Option D:Changing to memory-optimized instances may not necessarily reduce idle time or optimize costs.
AWS Documentation References:
EMR Managed Scaling

NO.14 A company is using AWS DataSync to migrate millions of files from an on-premises system to AWS. The files are 10 KB in size on average.
The company wants to use Amazon S3 for file storage. For the first year after the migration the files will be accessed once or twice and must be immediately available. After 1 year the files must be archived for at least
7 years.
Which solution will meet these requirements MOST cost-effectively?

A. Use an archive tool lo group the files into large objects. Use DataSync to migrate the objects. Store the objects in S3 Glacier Instant Retrieval for the first year. Use a lifecycle configuration to transition the files to S3 Glacier Deep Archive after 1 year with a retention period of 7 years.
B. Use an archive tool to group the files into large objects. Use DataSync to copy the objects to S3 Standard-Infrequent Access (S3 Standard-IA). Use a lifecycle configuration to transition the files to S3 Glacier Instant Retrieval after 1 year with a retention period of 7 years.
C. Configure the destination storage class for the files as S3 Glacier Instant. Retrieval Use a lifecycle policy to transition the files to S3 Glacier Flexible Retrieval after 1 year with a retention period of 7 years.
D. Configure a DataSync task to transfer the files to S3 Standard-Infrequent Access (S3 Standard-IA) Use a lifecycle configuration to transition the files to S3. Deep Archive after 1 year with a retention period of 7 years.
Answer: A

NO.15 A company is using an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. The company must ensure that Kubernetes service accounts in the EKS cluster have secure and granular access to specific AWS resources by using IAM roles for service accounts (IRSA).
Which combination of solutions will meet these requirements? (Select TWO.)
A. Create an IAM policy that defines the required permissions. Attach the policy directly to the IAM role of the EKS nodes.
B. Implement network policies within the EKS cluster to prevent Kubernetes service accounts from accessing specific AWS services.
C. Modify the EKS cluster's IAM role to include permissions for each Kubernetes service account. Ensure a one-to-one mapping between IAM roles and Kubernetes roles.
D. Define an IAM role that includes the necessary permissions. Annotate the Kubernetes service accounts with the Amazon Resource Name (ARN) of the IAM role.
E. Set up a trust relationship between the IAM roles for the service accounts and an OpenID Connect (OIDC) identity provider.
Answer: D E
Explanation:
IAM Roles for Service Accounts (IRSA): IRSA allows you to associate an IAM role with a Kubernetes service account. This enables pods to assume the IAM role and access AWS resources securely.
Annotating Service Accounts: By annotating Kubernetes service accounts with the ARN of the IAM role, you establish the association required for IRSA.
OIDC Identity Provider: EKS clusters use OpenID Connect (OIDC) to authenticate service accounts. Setting up a trust relationship between the IAM role and the OIDC provider allows the Kubernetes service account to assume the IAM role.

NO.16 A company has an application that serves clients that are deployed in more than 20.000 retail storefront locations around the world. The application consists of backend web services that are exposed over HTTPS on port 443 The application is hosted on Amazon EC2 Instances behind an Application Load Balancer (ALB).
The retail locations communicate with the web application over the public internet. The company allows each retail location to register the IP address that the retail location has been allocated by its local ISP.

The company's security team recommends to increase the security of the application endpoint by restricting access to only the IP addresses registered by the retail locations.
What should a solutions architect do to meet these requirements?

A. Associate an AWS WAF web ACL with the ALB Use IP rule sets on the ALB to filter traffic Update the IP addresses in the rule to Include the registered IP addresses

B. Deploy AWS Firewall Manager to manage the ALB. Configure firewall rules to restrict traffic to the ALB Modify the firewall rules to include the registered IP addresses.

C. Store the IP addresses in an Amazon DynamoDB table. Configure an AWS Lambda authorization function on the ALB to validate that incoming requests are from the registered IP addresses.

D. Configure the network ACL on the subnet that contains the public interface of the ALB Update the ingress rules on the network ACL with entries for each of the registered IP addresses.

Answer: A

Explanation:
AWS WAF (Web Application Firewall): AWS WAF allows you to create custom rules to block or allow web requests based on conditions that you specify.
Web ACL (Access Control List):
Create a web ACL and associate it with the ALB.
Use IP rule sets to specify the IP addresses of the retail locations that are allowed to access the application.
Security and Flexibility:
AWS WAF provides a scalable way to manage access control, ensuring that only traffic from registered IP addresses is allowed.
You can dynamically update the IP rule sets to add or remove IP addresses as needed.
Operational Simplicity: Using AWS WAF with a web ACL is straightforward and integrates seamlessly with the ALB, providing an efficient solution for managing access control based on IP addresses.
References:
AWS WAF
How AWS WAF Works

NO.17 A company is designing a secure solution to grant access to its Amazon RDS for PostgreSQL database.
Applications that run on Amazon EC2 instances must be able to securely authenticate to the database without storing long-term credentials.
Which solution will meet these requirements?

A. Enable RDS IAM authentication and configure AWS Secrets Manager to store database credentials. Configure applications to retrieve credentials at runtime.

B. Configure a custom IAM policy for the database that allows access from the EC2 instances' IP addresses. Configure applications to use a static password to authenticate to the database.

C. Set up an IAM user for each application. Store the access key ID and secret access key in the EC2 instances' environment variables. Grant the IAM users permission to the database.

D. Use IAM roles to assign permissions to the EC2 instances. Configure the applications to obtain a token from the RDS database to authenticate by using IAM authentication.

Answer: D

Explanation:
For Amazon RDS for PostgreSQL, AWS provides IAM database authentication. With this feature,

applications do not use stored long-term usernames and passwords. Instead, they use temporary authentication tokens that are generated by AWS and validated by the RDS database.
The AWS best practice pattern is:
* Attach an IAM role to the EC2 instances (instance profile).
* Grant that role the necessary permissions (for example, rds-db:connect) to the specific RDS database user.
* The application running on the EC2 instance uses the role's temporary credentials to call the RDS token-generation API and obtain a short-lived authentication token.
* The application then uses this token as the password when connecting to RDS for PostgreSQL.
This removes the need to store long-term credentials in the application or on the instance and uses IAM roles with temporary credentials, aligning with the security requirement.
Option A still relies on stored credentials (even if in Secrets Manager), which are long-lived and rotated but not token-based per-connection IAM authentication.
Option B uses static passwords and IP-based access, which does not meet the "no long-term credentials" requirement.
Option C stores long-term IAM user keys on the instances, which is explicitly against best practices and does not directly integrate with RDS authentication.

NO.18 The lead member of a DevOps team creates an AWS account. A DevOps engineer shares the account credentials with a solutions architect through a password manager application.
The solutions architect needs to secure the root user for the new account.
Which actions will meet this requirement? (Select TWO.)
A. Update the root user password to a new, strong password.
B. Secure the root user account by using a virtual multi-factor authentication (MFA) device.
C. Create an IAM user for each member of the DevOps team. Assign the AdministratorAccess AWS managed policy to each IAM user.
D. Create root user access keys. Save the keys as a new parameter in AWS Systems Manager Parameter Store.
E. Update the IAM role for the root user to ensure the root user can use only approved services.
Answer: A B
Explanation:
Securing the root user account requires setting a strong password and enabling multi-factor authentication (MFA). AWS recommends never sharing the root user credentials, setting up individual IAM users for everyday operations, and always protecting the root user with MFA for maximum security.
Reference Extract:
"AWS recommends securing the root user with a strong password and enabling multi-factor authentication (MFA). Do not use or share root credentials for everyday tasks." Source: AWS Certified Solutions Architect - Official Study Guide, IAM and Security Best Practices section.

NO.19 A company is developing a serverless web application that gives users the ability to interact with real-time analytics from online games. The data from the games must be streamed in real time. The company needs a durable, low-latency database option for user data. The company does not know how many users will use the application. Any design considerations must provide response times of single-digit milliseconds as the application scales.

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

A. Amazon CloudFront

B. Amazon DynamoDB

C. Amazon Kinesis

D. Amazon RDS

E. AWS Global Accelerator

Answer: B C

Explanation:

Amazon Kinesis allows real-time ingestion of game events at scale, while Amazon DynamoDB provides millisecond-latency access to user data, automatically scaling with demand. This combination ensures real- time processing and fast data retrieval without managing infrastructure.

Reference: AWS Documentation - Real-Time Processing with Kinesis and Low-Latency Databases with DynamoDB

NO.20 A company runs multiple applications in multiple AWS accounts within the same organization in AWS Organizations. A content management system (CMS) runs on Amazon EC2 instances in a VPC. The CMS needs to access shared files from an Amazon Elastic File System (Amazon EFS) file system that is deployed in a separate AWS account. The EFS account is in a separate VPC.

Which solution will meet this requirement?

A. Mount the EFS file system on the EC2 instances by using the EFS Elastic IP address.

B. Enable VPC sharing between the two accounts. Use the EFS mount helper to mount the file system on the EC2 instances. Redeploy the EFS file system in a shared subnet.

C. Configure AWS Systems Manager Run Command to mount the EFS file system on the EC2 instances.

D. Install the amazon-efs-utils package on the EC2 instances. Add the mount target in the efs-config file.Mount the EFS file system by using the EFS access point.

Answer: D

Explanation:

To access an EFS file system across accounts and VPCs, the EFS must be mounted using VPC peering or AWS Transit Gateway, and the EC2 instances must use the amazon-efs-utils package with the correct mount target or access point.

Using an EFS access point simplifies access management, especially across accounts, by providing a POSIX identity and access policy layer.

VPC sharing doesn't support EFS directly unless the subnet and resources are shared properly, which requires redeployment. Therefore, option D is the most complete and correct.

NO.21 A company runs its critical storage application in the AWS Cloud. The application uses Amazon S3 in two AWS Regions. The company wants the application to send remote user data to the nearest S3 bucket with no public network congestion. The company also wants the application to fail over with the least amount of management of Amazon S3.

Which solution will meet these requirements?

A. Implement an active-active design between the two Regions. Configure the application to use the regional S3 endpoints closest to the user.

B. Use an active-passive configuration with S3 Multi-Region Access Points. Create a global endpoint for each of the Regions.

C. Send user data to the regional S3 endpoints closest to the user. Configure an S3 cross-account replication rule to keep the S3 buckets synchronized.

D. Set up Amazon S3 to use Multi-Region Access Points in an active-active configuration with a single global endpoint. Configure S3 Cross-Region Replication.

Answer: D

Explanation:

AWS S3 Multi-Region Access Points enable customers to use a single global endpoint for S3 bucket access across multiple AWS Regions, providing automatic routing to the nearest Region. This reduces public network congestion by directing user data to the closest S3 bucket and supports high availability with active- active configuration.

Cross-Region Replication ensures data is replicated between buckets in different Regions, meeting the failover and resilience requirements with minimal management overhead.

Option D aligns best with AWS's recommended approach to resilient, low-latency, and simplified multi- Region S3 access.

Option A lacks the global endpoint and automatic failover. Option B incorrectly describes Multi-Region Access Points configuration and suggests global endpoints per Region, which is contradictory.

Option C's cross-account replication adds complexity and does not provide a single global endpoint.

References:

AWS Well-Architected Framework - Reliability Pillar (https://d1.awsstatic.com/whitepapers/architecture/AWS_Well-Architected_Framework.pdf)

Amazon S3 Multi-Region Access Points (https://docs.aws.amazon.com/AmazonS3/latest/userguide/MultiRegionAccessPoints.html)

S3 Cross-Region Replication (https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication.html)

NO.22 A company has an online gaming application that has TCP and UDP multiplayer gaming capabilities.

The company uses Amazon Route 53 to point the application traffic to multiple Network Load Balancers (NLBs) in different AWS Regions. The company needs to improve application performance and decrease latency for the online game in preparation for user growth.

Which solution will meet these requirements?

A. Add an Amazon CloudFront distribution in front of the NLBs. Increase the Cache-Control: max-age parameter.

B. Replace the NLBs with Application Load Balancers (ALBs). Configure Route 53 to use latency-based routing.

C. Add AWS Global Accelerator in front of the NLBs. Configure a Global Accelerator endpoint to use the correct listener ports.

D. Add an Amazon API Gateway endpoint behind the NLBs. Enable API caching. Override method caching for the different stages.

Answer: C

Explanation:

AWS Global Accelerator is designed to improve the availability and performance of applications with global users by using the AWS global network. It provides static anycast IP addresses and routes user traffic over the AWS edge network to the optimal AWS Region and endpoint based on health,

geography, and routing policies. Global Accelerator supports both TCP and UDP traffic and can have Network Load Balancers as endpoints.
For latency-sensitive workloads such as multiplayer gaming, Global Accelerator reduces latency and jitter compared to internet-based routing and handles Regional failover quickly.
CloudFront (Option A) is optimized for HTTP/HTTPS content caching and is not appropriate for arbitrary TCP/UDP gaming traffic. Application Load Balancers (Option B) do not support UDP traffic. API Gateway (Option D) is for HTTP APIs and is not suitable for raw TCP/UDP game traffic.

NO.23 A company uses Amazon S3 to host its static website. The company wants to add a contact form to the webpage. The contact form will have dynamic server-side components for users to input their name, email address, phone number, and user message.
The company expects fewer than 100 site visits each month. The contact form must notify the company by email when a customer fills out the form.
Which solution will meet these requirements MOST cost-effectively?
A. Host the dynamic contact form in Amazon Elastic Container Service (Amazon ECS). Set up Amazon Simple Email Service (Amazon SES) to connect to a third-party email provider.
B. Create an Amazon API Gateway endpoint that returns the contact form from an AWS Lambda function.
Configure another Lambda function on the API Gateway to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic.
C. Host the website by using AWS Amplify Hosting for static content and dynamic content. Use server- side scripting to build the contact form. Configure Amazon Simple Queue Service (Amazon SQS) to deliver the message to the company.
D. Migrate the website from Amazon S3 to Amazon EC2 instances that run Windows Server. Use Internet Information Services (IIS) for Windows Server to host the webpage. Use client-side scripting to build the contact form. Integrate the form with Amazon WorkMail.
Answer: B
Explanation:
Using API Gateway and Lambda enables serverless handling of form submissions with minimal cost and infrastructure. When coupled with Amazon SNS, it allows instant email notifications without running servers, making it ideal for low-traffic workloads.
Reference: AWS Documentation - Serverless Contact Form with API Gateway, Lambda, and SNS

NO.24 A company stores data in an on-premises Oracle relational database. The company needs to make the data available in Amazon Aurora PostgreSQL for analysis The company uses an AWS Site-to-Site VPN connection to connect its on-premises network to AWS.
The company must capture the changes that occur to the source database during the migration to Aurora PostgreSQL.
Which solution will meet these requirements?
A. Use the AWS Schema Conversion Tool (AWS SCT) to convert the Oracle schema to Aurora PostgreSQL schema. Use the AWS Database Migration Service (AWS DMS) full-load migration task to migrate the data.
B. Use AWS DataSync to migrate the data to an Amazon S3 bucket. Import the S3 data to Aurora PostgreSQL by using the Aurora PostgreSQL aws_s3 extension.
C. Use the AWS Schema Conversion Tool (AWS SCT) to convert the Oracle schema to Aurora

PostgreSQL schema. Use AWS Database Migration Service (AWS DMS) to migrate the existing data and replicate the ongoing changes.

D. Use an AWS Snowball device to migrate the data to an Amazon S3 bucket. Import the S3 data to Aurora PostgreSQL by using the Aurora PostgreSQL aws_s3 extension.

Answer: C

Explanation:

For the migration of data from an on-premises Oracle database to Amazon Aurora PostgreSQL, this solution effectively handles schema conversion, data migration, and ongoing data replication.

AWS Schema Conversion Tool (SCT): SCT is used to convert the Oracle database schema to a format compatible with Aurora PostgreSQL. This tool automatically converts the database schema and code objects, like stored procedures, to the target database engine.

AWS Database Migration Service (DMS): DMS is employed to perform the data migration. It supports both full-load migrations (for initial data transfer) and continuous replication of ongoing changes (Change Data Capture, or CDC). This ensures that any updates to the Oracle database during the migration are captured and applied to the Aurora PostgreSQL database, minimizing downtime.

Why Not Other Options?:

Option A (SCT + DMS full-load only): This option does not capture ongoing changes, which is crucial for a live database migration to ensure data consistency.

Option B (DataSync + S3): AWS DataSync is more suited for file transfers rather than database migrations, and it doesn't support ongoing change replication.

Option D (Snowball + S3): Snowball is typically used for large-scale data transfers that don't require continuous synchronization, making it less suitable for this scenario where ongoing changes must be captured.

AWS References:

AWS Schema Conversion Tool- Guidance on using SCT for database schema conversions.

AWS Database Migration Service- Detailed documentation on using DMS for data migrations and ongoing replication.

NO.25 A company that uses AWS Organizations runs 150 applications across 30 different AWS accounts. The company used AWS Cost and Usage Report to create a new report in the management account. The report is delivered to an Amazon S3 bucket that is replicated to a bucket in the data collection account.

The company's senior leadership wants to view a custom dashboard that provides NAT gateway costs each day starting at the beginning of the current month.

Which solution will meet these requirements?

A. Share an Amazon QuickSight dashboard that includes the requested table visual. Configure QuickSight to use AWS DataSync to query the new report.

B. Share an Amazon QuickSight dashboard that includes the requested table visual. Configure QuickSight to use Amazon Athena to query the new report.

C. Share an Amazon CloudWatch dashboard that includes the requested table visual. Configure CloudWatch to use AWS DataSync to query the new report.

D. Share an Amazon CloudWatch dashboard that includes the requested table visual. Configure CloudWatch to use Amazon Athena to query the new report.

Answer: B

Explanation:

The AWS Cost and Usage Report (CUR) delivers detailed, line-item billing data to Amazon S3. AWS recommends querying CUR with Amazon Athena by creating external tables over the CUR S3 location (partitioned by time) to produce daily cost aggregations such as NAT Gateway (EC2:NatGateway) usage and cost. Amazon QuickSight natively connects to Athena as a data source to build and share dashboards with visuals (tables, time series) filtered from the start of the current month. DataSync (A, C) is a file transfer service and cannot query data. CloudWatch dashboards (C, D) visualize metrics/logs, not CUR datasets.
Therefore, using Athena to query CUR and QuickSight to present a daily NAT gateway cost dashboard is the most direct and operationally efficient approach.
References: CUR - querying with Amazon Athena; QuickSight - Athena data source; Cost categories an d service/usage type fields for NAT Gateway; AWS Cost Management best practices.

NO.26 A company launches a new web application that uses an Amazon Aurora PostgreSQL database. The company wants to add new features to the application that rely on AI. The company requires vector storage capability to use AI tools.
Which solution will meet this requirement MOST cost-effectively?
A. Use Amazon OpenSearch Service to create an OpenSearch service. Configure the application to write vector embeddings to a vector index.
B. Create an Amazon DocumentDB cluster. Configure the application to write vector embeddings to a vector index.
C. Create an Amazon Neptune ML cluster. Configure the application to write vector embeddings to a vector graph.
D. Install the pgvector extension on the Aurora PostgreSQL database. Configure the application to write vector embeddings to a vector table.
Answer: D
Explanation:
Aurora PostgreSQL supports the pgvector extension, which allows storage and querying of vector embeddings directly inside the database. This eliminates the need for external vector databases and provides cost-effective and performant integration for AI workloads.
Reference: AWS Documentation - Amazon Aurora PostgreSQL and pgvector Support

NO.27 A company uses AWS Lake Formation to govern its S3 data lake. It wants to visualize data in QuickSight by joining S3 data with Aurora MySQL operational data. The marketing team must see only specific columns.
Which solution provides column-level authorization with the least operational overhead?
A. Use EMR to ingest database data into SPICE with only required columns.
B. Use AWS Glue Studio to ingest database data into S3 and use IAM policies for column control.
C. Use AWS Glue Elastic Views to create materialized S3 views with column restrictions.
D. Use a Lake Formation blueprint to ingest database data to S3. Use Lake Formation for column-level access control. Use Athena as the QuickSight data source.
Answer: D
Explanation:
AWS Lake Formation provides fine-grained (column-level) access control for data stored in S3. Using a Lake Formation blueprint ensures database ingestion is automated and governed.
QuickSight can query Athena, and Athena honors Lake Formation permissions, enforcing column-

level controls automatically.
Options A, B, and C rely on manual filtering or IAM policies, which cannot enforce column-level authorization for SQL queries.

NO.28 A company runs an application in a VPC on AWS. The company's on-premises data center has a DNS server.
The data center is connected to AWS through an AWS Direct Connect connection with a private virtual interface (VIF). The on-premises DNS server needs to resolve the DNS name of the application in the VPC.
A. Set up AWS Verified Access endpoints in the VPC. Configure DNS forwarding rules in Verified Access. Configure the on-premises DNS server to forward DNS queries through the Verified Access endpoints.
B. Configure the Direct Connect connection to enable DNS resolution between the on-premises DNS server and the application in the VPC.
C. Create an Amazon Route 53 Resolver outbound endpoint and a Resolver rule in the VPC. Configure the on-premises DNS server to send requests for the application to the outbound endpoint.
D. Create an Amazon Route 53 Resolver inbound endpoint in the VPC. Configure the on-premises DNS server to send requests for the application to the inbound endpoint.
Answer: D
Explanation:
When on-premises DNS servers need to resolve private DNS names in a VPC, the correct pattern is to create a Route 53 Resolver inbound endpoint. The inbound endpoint allows DNS queries to flow from the on-premises environment into the VPC, where Route 53 can resolve VPC-specific names (such as private hosted zones or private resource records). Outbound endpoints (C) are for sending VPC DNS queries to on-premises, not the reverse. Verified Access (A) is unrelated to DNS resolution. Direct Connect (B) provides network connectivity but does not provide DNS forwarding capabilities.
Therefore, option D is the correct design.
References:* Amazon Route 53 Resolver Developer Guide - Inbound and outbound endpoints* AWS Well- Architected Framework - Security Pillar: Hybrid DNS integration

NO.29 A global ecommerce company is designing a three-tier application on AWS. The application includes a web tier that serves static content, an application tier that handles business logic, and a database tier that stores product information and user data. The application interacts with a relational database.
The company needs a highly available application architecture to serve global users with low latency, with the least operational overhead.
Which solution will meet these requirements?
A. Deploy Amazon EC2 instances in an Auto Scaling group for the application tier and web tier in a single AWS Region. Use an Application Load Balancer to distribute web traffic. Use an Amazon RDS database and Multi-AZ deployments for the database tier.
B. Set up an Amazon CloudFront distribution that uses an Amazon S3 bucket as the origin. Use Amazon Elastic Container Service (Amazon ECS) containers on AWS Fargate to deploy the application tier to each AWS Region where the company operates. Use an Amazon Aurora global database for the database tier.
C. Use an Amazon S3 bucket to store the static web content. Use Amazon EC2 Auto Scaling and EC2

Spot Instances for the application tier. Use Amazon RDS for MySQL with read replicas for the database tier. Use AWS Database Migration Service (AWS DMS) to replicate data to secondary AWS Regions.

D. Use an Amazon S3 bucket to store static web content. Use AWS Lambda functions to handle serverless backend logic in the application tier. Use Amazon API Gateway to invoke the Lambda functions for web requests. Use an Amazon DynamoDB database for the database tier. Deploy the DynamoDB database across multiple AWS Regions.

Answer: B

Explanation:

AWS guidance for global, highly available, low-latency applications using a relational database recommends:

* Amazon CloudFront with Amazon S3 for static content to cache data at edge locations globally and minimize latency for static assets.
* A regional application tier deployed close to users using managed container services such as Amazon ECS on AWS Fargate, which removes the need to manage servers and scales automatically, reducing operational overhead.
* A relational database tier using Amazon Aurora global database, which is purpose-built for globally distributed applications. Aurora global database provides a primary cluster in one Region with low-latency read replicas in secondary Regions and fast cross-Region replication, enabling low-latency reads and high availability for global users.

Option A uses only a single Region, which does not meet the "global users with low latency" requirement.

Option C uses RDS and DMS-based replication, which requires more management and does not provide Aurora's integrated global database features.

Option D replaces the relational database with DynamoDB, which violates the requirement that the application interacts with a relational database.

NO.30 A company needs to design a hybrid network architecture The company's workloads are currently stored in the AWS Cloud and in on-premises data centers The workloads require single-digit latencies to communicate The company uses an AWS Transit Gateway transit gateway to connect multiple VPCs Which combination of steps will meet these requirements MOST cost-effectively? (Select TWO.)

A. Establish an AWS Site-to-Site VPN connection to each VPC.

B. Associate an AWS Direct Connect gateway with the transit gateway that is attached to the VPCs.

C. Establish an AWS Site-to-Site VPN connection to an AWS Direct Connect gateway.

D. Establish an AWS Direct Connect connection. Create a transit virtual interface (VIF) to a Direct Connect gateway.

E. Associate AWS Site-to-Site VPN connections with the transit gateway that is attached to the VPCs

Answer: B D

Explanation:

AWS Direct Connect: Provides a dedicated network connection from your on-premises data center to AWS, ensuring low latency and consistent network performance.

Direct Connect Gateway Association:

Direct Connect Gateway: Acts as a global network transit hub to connect VPCs across different AWS regions.

Association with Transit Gateway: Enables communication between on-premises data centers and multiple VPCs connected to the transit gateway.
Transit Virtual Interface (VIF):
Create Transit VIF: To connect Direct Connect with a transit gateway.
Setup Steps:
Establish a Direct Connect connection.
Create a transit VIF to the Direct Connect gateway.
Associate the Direct Connect gateway with the transit gateway attached to the VPCs.
Cost Efficiency: This combination avoids the recurring costs and potential performance variability of VPN connections, providing a robust, low-latency hybrid network solution.
References:
AWS Direct Connect
Transit Gateway and Direct Connect Gateway

NO.31 A company runs HPC workloads requiring high IOPS.
Which combination of steps will meet these requirements? (Select TWO)
A. Use Amazon EFS as a high-performance file system.
B. Use Amazon FSx for Lustre as a high-performance file system.
C. Create an Auto Scaling group of EC2 instances. Use Reserved Instances. Configure a spread placement group. Use AWS Batch for analytics.
D. Use Mountpoint for Amazon S3 as a high-performance file system.
E. Create an Auto Scaling group of EC2 instances. Use mixed instance types and a cluster placement group. Use Amazon EMR for analytics.
Answer: B E
Explanation:
Option B: FSx for Lustre is designed for HPC workloads with high IOPS.
Option E: A cluster placement group ensures low-latency networking for HPC analytics workloads.
Option A: Amazon EFS is not optimized for HPC.
Option D: Mountpoint for S3 does not meet high IOPS needs.

NO.32 A finance company has a web application that generates credit reports for customers. The company hosts the frontend of the web application on a fleet of Amazon EC2 instances that is associated with an Application Load Balancer (ALB). The application generates reports by running queries on an Amazon RDS for SQL Server database.
The company recently discovered that malicious traffic from around the world is abusing the application by submitting unnecessary requests. The malicious traffic is consuming significant compute resources. The company needs to address the malicious traffic.
Which solution will meet this requirement?
A. Use AWS WAF to create a web ACL. Associate the web ACL with the ALB. Update the web ACL to block IP addresses that are associated with malicious traffic.
B. Use AWS WAF to create a web ACL. Associate the web ACL with the ALB. Use the AWS WAF Bot Control managed rule feature.
C. Set up AWS Shield to protect the ALB and the database.
D. Use AWS WAF to create a web ACL. Associate the web ACL with the ALB. Configure the AWS WAF IP reputation rule.

Answer: B
Explanation:
TheAWS WAF Bot Control managed ruleis designed to automatically detect and mitigate bot traffic. This feature is particularly useful for addressing malicious traffic and conserving compute resources by filtering unnecessary requests at the ALB level.
Option A:Blocking IP addresses manually introduces significant operational overhead and is not scalable against dynamic, worldwide malicious traffic.
Option C:AWS Shield provides DDoS protection, but the scenario does not describe a DDoS attack. WAF is better suited for managing application-layer threats like bot traffic.
Option D:The AWS WAF IP reputation rule helps block traffic from known bad IPs but may not address bot traffic effectively.
AWS Documentation References:
AWS WAF Bot Control
AWS WAF Managed Rules

NO.33 A company is migrating some of its applications to AWS. The company wants to migrate and modernize the applications quickly after it finalizes networking and security strategies. The company has set up an AWS Direct Connect connection in a central network account.
The company expects to have hundreds of AWS accounts and VPCs in the near future. The corporate network must be able to access the resources on AWS seamlessly and also must be able to communicate with all the VPCs. The company also wants to route its cloud resources to the internet through its on-premises data center.
Which combination of steps will meet these requirements? (Select THREE.)
A. Create a Direct Connect gateway in the central account. In each of the accounts, create an association proposal by using the Direct Connect gateway and the account ID for every virtual private gateway.
B. Create a Direct Connect gateway and a transit gateway in the central network account. Attach the transit gateway to the Direct Connect gateway by using a transit VIF.
C. Provision an internet gateway. Attach the internet gateway to subnets. Allow internet traffic through the gateway.
D. Share the transit gateway with other accounts. Attach VPCs to the transit gateway.
E. Provision VPC peering as necessary.
F. Provision only private subnets. Open the necessary route on the transit gateway and customer gateway to allow outbound internet traffic from AWS to flow through NAT services that run in the data center.
Answer: B D F
Explanation:
For a large-scale multi-account AWS environment with many VPCs and centralized Direct Connect, AWS recommends using a Transit Gateway (TGW) architecture combined with a Direct Connect gateway (DXGW). This setup allows scalable, centralized connectivity between on-premises and multiple VPCs across accounts.
Step B: Creating a Direct Connect gateway and Transit Gateway in a central network account and connecting them via a transit VIF enables the on-premises network to access all connected VPCs.
Step D: Sharing the transit gateway with other accounts via AWS Resource Access Manager (RAM) allows the central TGW to attach VPCs in multiple accounts, simplifying multi-account connectivity.

Step F: To route cloud resources' internet traffic back through the on-premises data center (for centralized egress), provisioning only private subnets and routing outbound internet traffic through NAT or firewall services in the data center is necessary. This requires configuring transit gateway and customer gateway routes appropriately.
Option A is partially correct in the use of Direct Connect gateway but association proposals are not scalable for hundreds of VPCs and accounts compared to transit gateway. Option C (internet gateway) is irrelevant here as traffic egress is required via on-premises data center, not directly to the internet. Option E (VPC peering) is not scalable for hundreds of VPCs.
References:
AWS Transit Gateway Overview (https://docs.aws.amazon.com/vpc/latest/tgw/what-is-transit-gateway.html) AWS Direct Connect Gateway (https://docs.aws.amazon.com/directconnect/latest/UserGuide/direct-connect- gateways.html) Centralized Egress Architecture with Transit Gateway (https://aws.amazon.com/blogs/networking-and- content-delivery/how-to-set-up-centralized-egress-with-transit-gateway/) AWS Well-Architected Framework - Reliability Pillar (https://d1.awsstatic.com/whitepapers/architecture /AWS_Well-Architected_Framework.pdf)

NO.34 A company uses AWS Organizations to manage multiple AWS accounts. The company needs a secure, event- driven architecture in which specific Amazon SNS topics in Account A can publish messages to specific Amazon SQS queues in Account B.
Which solution meets these requirements while maintaining least privilege?
A. Create a new IAM role in Account A that can publish to any SQS queue. Share the role ARN with Account B.
B. Add SNS topic ARNs to SQS queue policies in Account B. Configure SNS topics to publish to any queue. Encrypt the queue with an AWS KMS key.
C. Modify the SQS queue policies in Account B to allow only specific SNS topic ARNs from Account A to publish messages. Ensure the SNS topics have publish permissions for the specific queue ARN.
D. Create a shared IAM role across both accounts with permission to publish to all SQS queues. Enable cross-account access.
Answer: C
Explanation:
AWS documentation states that the correct and least-privilege method for cross-account SNS-to-SQS integration is:
* Add specific SNS topic ARNs to the SQS queue policy.
* Allow only those topics to publish messages to the queue.
* Ensure SNS has permission to publish to the specific queue ARN.
This ensures strict scoping and adheres to least privilege.
Options A and D grant overly broad permissions. Option B allows publishing to any queue, which violates least privilege.

NO.35 A company runs an enterprise resource planning (ERP) system on Amazon EC2 instances in a single AWS Region. Users connect to the ERP system by using a public API that is hosted on the EC2 instances.
International users report slow API response times from their data centers.
A solutions architect needs to improve API response times for the international users.
Which solution will meet these requirements MOST cost-effectively?

A. Set up an AWS Direct Connect connection that has a public virtual interface (VIF) to connect each user's data center to the EC2 instances. Create a Direct Connect gateway for the ERP system API to route user API requests.

B. Deploy Amazon API Gateway endpoints in multiple Regions. Use Amazon Route 53 latency-based routing to route requests to the nearest endpoint. Configure a VPC peering connection between the Regions to connect to the ERP system.

C. Set up AWS Global Accelerator. Configure listeners for the necessary ports. Configure endpoint groups for the appropriate Regions to distribute traffic. Create an endpoint in each group for the API.

D. Use AWS Site-to-Site VPN to establish dedicated VPN tunnels between multiple Regions and user networks. Route traffic to the API through the VPN connections.

Answer: C

Explanation:

AWS Global Accelerator improves the performance and availability of applications by directing user traffic through the AWS global network of edge locations using anycast IP addresses. It reduces latency and jitter for global users accessing applications in a single Region.

Why this works:

Global Accelerator routes user requests to the nearest AWS edge location using AWS's high-performance backbone network.

It then forwards traffic to the optimal endpoint - in this case, the public API hosted on EC2.

This is much more cost-effective and requires less operational complexity than deploying and maintaining multiple API Gateway endpoints across regions (Option B), or setting up Direct Connect links for every international location (Option A).

Option C requires no application change and is designed specifically for latency improvement and high availability.

References:

AWS Global Accelerator Documentation

Use Cases for Global Accelerator

Performance Improvements for Global Users