

ActualtestsIT

ActualtestsIT

HOME

ALL VENDORS

★ GUARANTEE

? FAQ

TESTIMONIALS

CART (0)



Try **PDF Demo** before you buy

Online Test Engine: Online Tool, Convenient, easy to study. Instant Online Access. Supports All Web Browsers.

PDF format: Easy to read and print learning materials, our products are available in PDF file format.

Desktop Test Engine: Installable Software Application. Simulates Real Exam Environment. Practice Offline Anytime.

Select a vendor...

Select an test...

Your email address

Free Download Demo

Instant Download

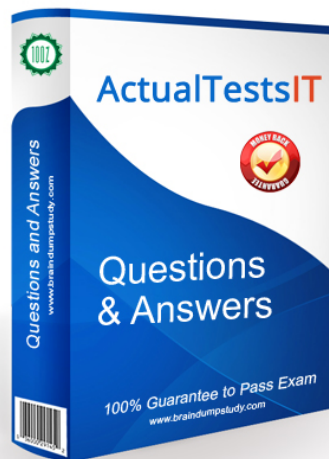


After Payment, our system will send you the products you purchase in mailbox in a minute after payment. If not received within 2 hours, please contact us.

365 Days Free Updates



Free update is available within 365 days after your purchase. After 365 days, you will get 50% discounts for updating.



Money Back Guarantee

Full refund if you fail the corresponding exam in 60 days after purchasing. And Free get any another product.



Security & Privacy

We respect customer privacy. We use McAfee's security service to provide you with utmost security for your personal information & peace of mind.

<http://www.actualtestsit.com/>

Topping Certification Exam Prep, Test Dumps Materials - ActualTestsIT

Exam : **SAP-C01**

Title : **AWS Certified Solutions
Architect - Professional**

Vendor : **Amazon**

Version : **DEMO**

NO.1 A company is running a web application on Amazon EC2 instances in a production AWS account.

The company requires all logs generated from the web application to be copied to a central AWS account (or analysis and archiving). The company's AWS accounts are currently managed independently. Logging agents are configured on the EC2 instances to upload the log files to an Amazon S3 bucket in the central AWS account.

A solutions architect needs to provide access for a solution that will allow the production account to store log files in the central account. The central account also needs to have read access to the log files.

What should the solutions architect do to meet these requirements?

- A.** Create a cross-account role in the central account. Assume the role from the production account when the logs are being copied.
- B.** Create a cross-account role in the production account. Assume the role from the production account when the logs are being copied.
- C.** Create a policy on the S3 bucket with access from only the CIDR range of the EC2 instances in the production account. Use the production account ID as the principal.
- D.** Create a policy on the S3 bucket with the production account ID as the principal. Allow S3 access from a delegated user.

Answer: D

NO.2 A finance company hosts a data lake in Amazon S3. The company receives financial data records over SFTP each night from several third parties. The company runs its own SFTP server on an Amazon EC2 instance in a public subnet of a VPC. After the files are uploaded, they are moved to the data lake by a cron job that runs on the same instance. The SFTP server is reachable on DNS `sftp.examWe.com` through the use of Amazon Route 53.

What should a solutions architect do to improve the reliability and scalability of the SFTP solution?

- A.** Migrate the SFTP server to AWS Transfer for SFTP. Update the DNS record `sftp.example.com` in Route 53 to point to the server endpoint hostname.
- B.** Move the EC2 instance into an Auto Scaling group. Place the EC2 instance behind an Application Load Balancer (ALB). Update the DNS record `sftp.example.com` in Route 53 to point to the ALB.
- C.** Migrate the SFTP server to a file gateway in AWS Storage Gateway. Update the DNS record `sftp.example.com` in Route 53 to point to the file gateway endpoint.
- D.** Place the EC2 instance behind a Network Load Balancer (NLB). Update the DNS record `sftp.example.com` in Route 53 to point to the NLB.

Answer: A

NO.3 A company is running an application distributed over several Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer. The security team requires that all application access attempts be made available for analysis. Information about the client IP address, connection type, and user agent must be included.

Which solution will meet these requirements?

- A.** Enable access logs for the Application Load Balancer, and publish the logs to an Amazon S3 bucket. Have the security team use Amazon Athena to query and analyze the logs.
- B.** Enable Traffic Mirroring and specify all EC2 instance network interfaces as the source. Send all

traffic information through Amazon Kinesis Data Firehose to an Amazon Elastic search Service (Amazon ES) cluster that the security team uses for analysis.

C. Enable VPC Flow Logs for all EC2 instance network interfaces Publish VPC Flow Logs to an Amazon S3 bucket Have the security team use Amazon Athena to query and analyze the logs

D. Enable EC2 detailed monitoring, and include network logs Send all logs through Amazon Kinesis Data Firehose to an Amazon Elasticsearch Service (Amazon ES) cluster that the security team uses for analysis.

Answer: A

Explanation:

<https://docs.aws.amazon.com/elasticloadbalancing/latest/application/load-balancer-access-logs.html>

NO.4 A scientific organization requires the processing of text and picture data stored in an Amazon S3 bucket. The data is gathered from numerous radar stations during a mission's live, time-critical phase. The data is uploaded by the radar stations to the source S3 bucket. The data is preceded with the identification number of the radar station.

In a second account, the business built a destination S3 bucket. To satisfy a compliance target, data must be transferred from the source S3 bucket to the destination S3 bucket. Replication is accomplished by using an S3 replication rule that covers all items in the source S3 bucket.

A single radar station has been recognized as having the most precise data

a. At this radar station, data replication must be completed within 30 minutes of the radar station uploading the items to the source S3 bucket.

What actions should a solutions architect take to ensure that these criteria are met?

A. In the second account, create another S3 bucket to receive data from the radar station with the most accurate data Set up a new replication rule for this new S3 bucket to separate the replication from the other radar stations Monitor the maximum replication time to the destination. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger an alert when the time exceeds the desired threshold

B. Set up an AWS DataSync agent to replicate the prefixed data from the source S3 bucket to the destination S3 bucket. Select to use at available bandwidth on the task, and monitor the task to ensure that it is in the TRANSFERRING status. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger an alert if this status changes.

C. Create a new S3 replication rule on the source S3 bucket that filters for the keys that use the prefix of the radar station with the most accurate data Enable S3 Replication Time Control (S3 RTC) Monitor the maximum replication time to the destination Create an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger an alert when the time exceeds the desired threshold

D. Enable Amazon S3 Transfer Acceleration on the source S3 bucket, and configure the radar station with the most accurate data to use the new endpoint Monitor the S3 destination bucket's TotalRequestLatency metric Create an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger an alert if this status changes

Answer: C

Explanation:

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication-time-control.html>

NO.5 A company uses AWS Cloud Formation to deploy applications within multiple VPCs that are all

attached to a transit gateway. Each VPC that sends traffic to the public internet must send the traffic through a shared services VPC. Each subnet within a VPC uses the default VPC route table, and the traffic is routed to the transit gateway. The transit gateway uses its default route table for any VPC attachment.

A security audit reveals that an Amazon EC2 instance that is deployed within a VPC can communicate with an EC2 instance that is deployed in any of the company's other VPCs. A solutions architect needs to limit the traffic between the VPCs. Each VPC must be able to communicate only with a predefined, limited set of authorized VPCs.

What should the solutions architect do to meet these requirements?

- A.** Create a dedicated transit gateway route table for each VPC attachment. Route traffic only to the authorized VPCs.
- B.** Update the network ACL of each subnet within a VPC to allow outbound traffic only to the authorized VPCs. Remove all deny rules except the default deny rule.
- C.** Update all the security groups that are used within a VPC to deny outbound traffic to security groups that are used within the unauthorized VPCs
- D.** Update the main route table of each VPC to route traffic only to the authorized VPCs through the transit gateway.

Answer: B

NO.6 An enterprise company wants to allow its developers to purchase third-party software through AWS Marketplace. The company uses an AWS Organizations account structure with full features enabled, and has a shared services account in each organizational unit (OU) that will be used by procurement managers. The procurement team's policy indicates that developers should be able to obtain third-party software from an approved list only and use Private Marketplace in AWS Marketplace to achieve this requirement. The procurement team wants administration of Private Marketplace to be restricted to a role named procurement-manager-role, which could be assumed by procurement managers. Other IAM users, groups, roles, and account administrators in the company should be denied Private Marketplace administrative access.

What is the MOST efficient way to design an architecture to meet these requirements?

- A.** Create an IAM role named procurement-manager-role in all AWS accounts in the organization. Add the AdministratorAccess managed policy to the role. Define a permissions boundary with the AWSPublicMarketplaceAdminFullAccess managed policy and attach it to all the developer roles.
- B.** Create an IAM role named procurement-manager-role in the AWS accounts that will be used by developers. Add the AWSPublicMarketplaceAdminFullAccess managed policy to the role. Create ...Organizations to deny permissions to administer Private Marketplace to everyone except the role named procurement-manager-role. Apply the SCP to all the shared services accounts in the.....
- C.** Create an IAM role named procurement-manager-role in all AWS accounts in the organization. Add the PowerUserAccess managed policy to the role. Apply an inline policy to all IAM users and roles in every AWS account to deny permissions on the AWSPublicMarketplaceAdminFullAccess managed policy.
- D.** Create an IAM role named procurement-manager-role in all the shared services accounts in the organization. Add the AWSPublicMarketplaceAdminFullAccess managed policy to the role. Create an organization root-level SCP to deny permissions to administer Private Marketplace to everyone except the role named procurement-manager-role. Create another organization root-level SCP to

deny permissions to create an IAM role named procurement-manager-role to everyone in the organization.

Answer: D

NO.7 A company's AWS architecture currently uses access keys and secret access keys stored on each instance to access AWS services Database credentials are hard-coded on each instance SSH keys for command-line remote access are stored in a secured Amazon S3 bucket The company has asked its solutions architect to improve the security posture of the architecture without adding operational complexity.

Which combination of steps should the solutions architect take to accomplish this? (Select THREE.)

- A.** Use AWS Systems Manager Parameter Store to store database credentials
- B.** Use AWS KMS to store database credentials
- C.** Use AWS Systems Manager Session Manager for remote access
- D.** Use a secure fleet of Amazon EC2 bastion hosts for remote access
- E.** Use AWS Secrets Manager to store access keys and secret access keys
- F.** Use Amazon EC2 instance profiles with an IAM role

Answer: A,C,F

NO.8 A company is running a containerized application in the AWS Cloud. The application is running by using Amazon Elastic Container Service (Amazon ECS) on a set Amazon EC2 instances. The EC2 instances run in an Auto Scaling group.

The company uses Amazon Elastic Container Registry (Amazon ECR) to store its container images When a new image version is uploaded, the new image version receives a unique tag

The company needs a solution that inspects new image versions for common vulnerabilities and exposures The solution must automatically delete new image tags that have Critical or High severity findings The solution also must notify the development team when such a deletion occurs

Which solution meets these requirements?

- A.** Configure periodic image scan on the repository Configure scan results to be added to an Amazon Simple Queue Service (Amazon SQS) queue Invoke an AWS Step Functions state machine when a new message is added to the SQS queue Use the Step Functions state machine to delete the image tag for images that have Critical or High severity findings. Notify the development team by using Amazon Simple Email Service (Amazon SES).
- B.** Schedule an AWS Lambda function to start a manual image scan every hour Configure Amazon EventBridge (Amazon CloudWatch Events) to invoke another Lambda function when a scan is complete. Use the second Lambda function to delete the image tag for images that have Critical or High severity findings. Notify the development team by using Amazon Simple Notification Service (Amazon SNS)
- C.** Configure scan on push on the repository Configure scan results to be pushed to an Amazon Simple Queue Service (Amazon SQS) queue Invoke an AWS Lambda function when a new message is added to the SQS queue Use the Lambda function to delete the image tag for images that have Critical or High severity findings. Notify the development team by using Amazon Simple Email Service (Amazon SES).
- D.** Configure scan on push on the repository. Use Amazon EventBridge (Amazon CloudWatch Events) to invoke an AWS Step Functions state machine when a scan is complete for images that have Critical

or High severity findings Use the Step Functions state machine to delete the image tag for those images and to notify the development team through Amazon Simple Notification Service (Amazon SNS)

Answer: B

NO.9 A company wants to send data from its on-premises systems to Amazon S3 buckets. The company created the S3 buckets in three different accounts. The company must send the data privately without the data traveling across the internet. The company has no existing dedicated connectivity to AWS

Which combination of steps should a solutions architect take to meet these requirements? (Select TWO.)

- A.** Create an Amazon S3 gateway endpoint in the networking account
- B.** Establish a networking account in the AWS Cloud Create a private VPC in the networking account Set up an AWS Direct Connect connection with a public VIF between the on-premises environment and the private VPC
- C.** Create an Amazon S3 interface endpoint in the networking account
- D.** Establish a networking account in the AWS Cloud Create a private VPC in the networking account Set up an AWS Direct Connect connection with a private VIF between the on-premises environment and the private VPC

Answer: A,D

E, Establish a networking account in the AWS Cloud. Create a private VPC in the networking account Peer VPCs from the accounts that host the S3 buckets with the VPC in the network account

NO.10 A company has an organization in AWS Organizations that has a large number of AWS accounts. One of the AWS accounts is designated as a transit account and has a transit gateway that is shared with all of the other AWS accounts AWS Site-to-Site VPN connections are configured between all of the company's global offices and the transit account The company has AWS Config enabled on all of its accounts.

The company's networking team needs to centrally manage a list of internal IP address ranges that belong to the global offices Developers Will reference this list to gain access to applications securely. Which solution meets these requirements with the LEAST amount of operational overhead?

- A.** In the transit account, create a VPC prefix list with all of the internal IP address ranges. Use AWS Resource Access Manager to share the prefix list with all of the other accounts. Use the shared prefix list to configure security group rules in the other accounts.
- B.** In the transit account create a security group with all of the internal IP address ranges. Configure the security groups in the other accounts to reference the transit account's security
- C.** Create a new AWS Config managed rule that contains all of the internal IP address ranges Use the rule to check the security groups in each of the accounts to ensure compliance with the list of IP address ranges. Configure the rule to automatically remediate any noncompliant security group that is detected.
- D.** Create a JSON file that is hosted in Amazon S3 and that lists all of the internal IP address ranges Configure an Amazon Simple Notification Service (Amazon SNS) topic in each of the accounts that can be involved when the JSON file is updated. Subscribe an AWS Lambda function to the SNS topic to update all relevant security group rules with the updated IP address ranges.

Answer: A

group by using a nested security group reference of *<transit-account-id>./sg-1a2b3c4d".

NO.11 A company runs a software-as-a-service (SaaS) application on AWS. The application consists of an AWS Lambda function and an Amazon RDS for MySQL Multi-AZ database. During market events, the application has a much higher workload than normal. Users notice slow response times during the peak periods because of many database connections. The company needs to improve the scalable performance and availability of the database.

Which solution meets these requirements?

- A.** Create an Amazon CloudWatch alarm action that triggers a Lambda function to add an Amazon RDS for MySQL read replica when resource utilization hits a threshold.
- B.** Migrate the database to Amazon Aurora and add a read replica. Add a database connection pool outside of the Lambda handler function.
- C.** Migrate the database to Amazon Aurora and add a read replica. Use Amazon Route 53 weighted records.
- D.** Migrate the database to Amazon Aurora and add an Aurora Replica. Configure Amazon RDS Proxy to manage database connection pools.

Answer: D

NO.12 A gaming company created a game leaderboard by using a Multi-AZ deployment of an Amazon RDS database. The number of users is growing, and the queries to get individual player rankings are getting slower over time. The company expects a surge in users for an upcoming version and wants to optimize the design for scalability and performance.

Which solution will meet these requirements?

- A.** Add a read-only replica to the RDS DB instance. Add an RDS Proxy database proxy.
- B.** Keep the leaderboard data in the RDS DB instance. Provision a Multi-AZ deployment of an Amazon ElastiCache for Redis cluster.
- C.** Migrate the database to Amazon DynamoDB. Store the leaderboard data in different tables. Use Apache HiveQL JOIN statements to build the leaderboard.
- D.** Stream the leaderboard data by using Amazon Kinesis Data Firehose with an Amazon S3 bucket as the destination. Query the S3 bucket by using Amazon Athena for the leaderboard.

Answer: C

NO.13 A company processes environmental data.

a. The company has set up sensors to provide a continuous stream of data from different areas in a city. The data is available in JSON format.

The company wants to use an AWS solution to send the data to a database that does not require fixed schemas for storage. The data must be sent in real time.

Which solution will meet these requirements?

- A.** Use Amazon Kinesis Data Streams to send the data to Amazon DynamoDB.
- B.** Use Amazon Managed Streaming for Apache Kafka (Amazon MSK) to send the data to Amazon Aurora.
- C.** Use Amazon Kinesis Data Firehose to send the data to Amazon Redshift.
- D.** Use Amazon Kinesis Data Firehose to send the data to Amazon Keyspaces (for Apache Cassandra).

Answer: A

NO.14 A development team has created a new flight tracker application that provides near-real-time data to users. The application has a front end that consists of an Application Load Balancer (ALB) in front of two large Amazon EC2 instances in a single Availability Zone. Data is stored in a single Amazon RDS MySQL DB instance. An Amazon Route 53 DNS record points to the ALB.

Management wants the development team to improve the solution to achieve maximum reliability with the least amount of operational overhead.

Which set of actions should the team take?

- A.** Replace the DB instance with Amazon Aurora with Aurora Replicas. Deploy the application to multiple smaller EC2 instances across multiple Availability Zones in an Auto Scaling group behind an ALB.
- B.** Create RDS MySQL read replicas. Deploy the application to multiple AWS Regions. Use a Route 53 latency-based routing policy to route to the application.
- C.** Configure the DB instance as Multi-AZ. Deploy the application to two additional EC2 instances in different Availability Zones behind an ALB.
- D.** Replace the DB instance with Amazon DynamoDB global tables. Deploy the application in multiple AWS Regions. Use a Route 53 latency-based routing policy to route to the application.

Answer: A

Explanation:

Multi AZ ASG + ALB + Aurora = Less over head and automatic scaling

NO.15 A company is planning to migrate an Amazon RDS for Oracle database to an RDS for PostgreSQL DB instance in another AWS account. A solutions architect needs to design a migration strategy that will require no downtime and that will minimize the amount of time necessary to complete the migration. The migration strategy must replicate all existing data and any new data that is created during the migration. The target database must be identical to the source database at completion of the migration process.

All applications currently use an Amazon Route 53 CNAME record as their endpoint for communication with the RDS for Oracle DB instance. The RDS for Oracle DB instance is in a private subnet.

Which combination of steps should the solutions architect take to meet these requirements? (Select THREE)

- A.** Create a new RDS for PostgreSQL DB instance in the target account. Use the AWS Schema Conversion Tool (AWS SCT) to migrate the database schema from the source database to the target database.
- B.** Configure VPC peering between the VPCs in the two AWS accounts to provide connectivity to both DB instances from the target account. Configure the security groups that are attached to each DB instance to allow traffic on the database port from the VPC in the target account.
- C.** Use AWS Database Migration Service (AWS DMS) in the target account to perform a change data capture (CDC) migration from the source database to the target database. When the migration is complete, change the CNAME record to point to the target DB instance endpoint.
- D.** Temporarily allow the source DB instance to be publicly accessible to provide connectivity from the VPC in the target account. Configure the security groups that are attached to each DB instance to allow traffic on the database port from the VPC in the target account.

- E.** Use AWS Database Migration Service (AWS DMS) in the target account to perform a full load plus change data capture (CDC) migration from the source database to the target database. When the migration is complete, change the CNAME record to point to the target DB instance endpoint.
- F.** Use the AWS Schema Conversion Tool (AWS SCT) to create a new RDS for PostgreSQL DB instance in the target account with the schema and initial data from the source database.

Answer: B,E,F

NO.16 A solutions architect uses AWS Organizations to manage several AWS accounts for a company. The full Organizations feature set is activated for the organization. All production AWS accounts exist under an OU that is named "production". Systems operators have full administrative privileges within these accounts by using IAM roles.

The company wants to ensure that security groups in all production accounts do not allow inbound traffic for TCP port 22. All noncompliant security groups must be remediated immediately, and no new rules that allow port 22 can be created.

Which solution will meet these requirements?

- A.** Write an SCP that denies the `CreateSecurityGroup` action with a condition of `ec2:ingress rule with value 22`. Apply the SCP to the 'production' OU.
- B.** Create an Amazon EventBridge (Amazon CloudWatch Events) event bus in the Organizations management account. Create an AWS CloudFormation template to deploy configurations that send `CreateSecurityGroup` events to the event bus from all production accounts. Configure an AWS Lambda function in the management account with permissions to assume a role in all production accounts to describe and modify security groups. Configure the event bus to invoke the Lambda function. Configure the Lambda function to analyze each event for noncompliant security group actions and to automatically remediate any issues.
- C.** Create an AWS CloudFormation template to turn on AWS Config. Activate the `INCOMING_SSH_DISABLED` AWS Config managed rule. Deploy an AWS Lambda function that will run based on AWS Config findings and will remediate noncompliant resources. Deploy the CloudFormation template by using a StackSet that is assigned to the "production" OU. Apply an SCP to the OU to deny modification of the resources that the CloudFormation template provisions.
- D.** Configure an AWS CloudTrail trail for all accounts. Send CloudTrail logs to an Amazon S3 bucket in the Organizations management account. Configure an AWS Lambda function on the management account with permissions to assume a role in all production accounts to describe and modify security groups. Configure Amazon S3 to invoke the Lambda function on every `PutObject` event on the S3 bucket. Configure the Lambda function to analyze each CloudTrail event for noncompliant security group actions and to automatically remediate any issues.

Answer: C

NO.17 A company is running multiple workloads in the AWS Cloud. The company has separate units for software development. The company uses AWS Organizations and federation with SAML to give permissions to developers to manage resources in their AWS accounts. The development units each deploy their production workloads into a common production account. Recently, an incident occurred in the production account in which members of a development unit terminated an EC2 instance that belonged to a different development unit. A solutions architect must create a solution that prevents a similar incident from happening in the future. The solution also must allow developers the possibility to manage the instances used for their workloads.

Which strategy will meet these requirements?

- A.** Create separate OUs in AWS Organizations for each development unit Assign the created OUs to the company AWS accounts Create separate SCPs with a deny action and a StringNotEquals condition for the DevelopmentUnit resource tag that matches the development unit name Assign the SCP to the corresponding OU
- B.** Pass an attribute for DevelopmentUnit as an AWS Security Token Service (AWS STS) session tag during SAML federation Update the IAM policy for the developers' assumed IAM role with a deny action and a StringNotEquals condition for the DevelopmentUnit resource tag and aws PrincipalTag/DevelopmentUnit
- C.** Pass an attribute for DevelopmentUnit as an AWS Security Token Service (AWS STS) session tag during SAML federation Create an SCP with an allow action and a StrmgEquals condition for the DevelopmentUnit resource tag and aws Principal Tag 'DevelopmentUnit Assign the SCP to the root OU.
- D.** Create separate IAM policies for each development unit For every IAM policy add an allow action and a StringEquals condition for the DevelopmentUnit resource tag and the development unit name During SAML federation use AWS Security Token Service (AWS STS) to assign the IAM policy and match the development unit name to the assumed IAM role

Answer: A

NO.18 A company hosts a blog post application on AWS using Amazon API Gateway, Amazon DynamoDB, and AWS Lambda The application currently does not use API keys to authorize requests The API model is as follows:

- GET /posts/{postId} to get post details
- GET /users/{userId}. to get user details
- GET /comments/{commentId}: to get comments details

The company has noticed users are actively discussing topics in the comments section, and the company wants to increase user engagement by making the comments appear in real time Which design should be used to reduce comment latency and improve user experience?

- A.** Modify the blog application code to request GET/commentsV{commentId} every 10 seconds
- B.** Change the concurrency limit of the Lambda functions to lower the API response time.
- C.** Use edge-optimized API with Amazon CloudFront to cache API responses.
- D.** Use AWS AppSync and leverage WebSockets to deliver comments

Answer: D

NO.19 What should the solutions architect do to meet this requirement?

- A.** Create an AWS Lambda function to poll detailed metrics from the ECS cluster. When the number running Fargate tasks is greater than 80. invoke Amazon Simple Email Service (Amazon SES) to notify the development team
- B.** Use Amazon CloudWatch to monitor service quotas that are published under the AWS-'Usage metric namespace Set an alarm for when the math expression metricSERVICE_QUOTA(metric)"100 is greater than 80 Notify the development team by using Amazon Simple Notification Service (Amazon SNS).
- C.** Create an AWS Config rule to evaluate whether the Fargate SERVICE_QUOTA is greater than 80. Use Amazon Simple Email Service (Amazon SES) to notify the development team when the AWS

Config rule is not compliant.

D. / Use Amazon CloudWatch to monitor the Sample Count statistic for each service in the ECS cluster. Set an alarm for when the math expression `sample Notification SERVICE_QUOTA(service)"100` is greater than 80. Notify the development team by using Amazon Simple Notification Service (Amazon SNS)

Answer: B

NO.20 A company is configuring connectivity to a multi-account AWS environment to support application workloads that serve users in a single geographic region. The workloads depend on a highly available, on-premises legacy system deployed across two locations. It is critical for the AWS workloads to maintain connectivity to the legacy system, and a minimum of 5 Gbps of bandwidth is required. All application workloads within AWS must have connectivity with one another.

Which solution will meet these requirements?

A. Configure multiple AWS Direct Connect (DX) 10 Gbps dedicated connections from two DX partners for each on-premises location. Create a transit gateway and a DX gateway in a central network account. Create a transit virtual interface for each DX interface and associate them with the DX gateway. Create a gateway association between the DX gateway and the transit gateway.

B. Configure multiple AWS Direct Connect (DX) 10 Gbps dedicated connections from a DX partner for each on-premises location. Create private virtual interfaces on each connection for each AWS account VPC. Associate each private virtual interface with a virtual private gateway attached to each VPC.

C. Configure multiple AWS Direct Connect (DX) 10 Gbps dedicated connections from a DX partner for each on-premises location. Create and attach a virtual private gateway for each AWS account VPC. Create a transit gateway in a central network account and associate it with the virtual private gateways. Create a transit virtual interface on each DX connection and attach the interface to the transit gateway.

D. Configure multiple AWS Direct Connect (DX) 10 Gbps dedicated connections from two DX partners for each on-premises location. Create and attach a virtual private gateway for each AWS account VPC. Create a DX gateway in a central network account and associate it with the virtual private gateways. Create a public virtual interface on each DX connection and associate the interface with the DX gateway.

Answer: D

NO.21 A company has a project that is launching Amazon EC2 instances that are larger than required. The project's account cannot be part of the company's organization in AWS Organizations due to policy restrictions to keep this activity outside of corporate IT. The company wants to allow only the launch of t3.small EC2 instances by developers in the project's account. These EC2 instances must be restricted to the us-east-2 Region.

What should a solutions architect do to meet these requirements?

A. Create an SCP that denies the launch of all EC2 instances except t3.small EC2 instances in us-east-2. Attach the SCP to the project's account.

B. Create an IAM policy that allows the launch of only t3.small EC2 instances in us-east-2. Attach the policy to the roles and groups that the developers use in the project's account.

C. Create a new developer account. Move all EC2 instances, users, and assets into us-east-2. Add the account to the company's organization in AWS Organizations. Enforce a tagging policy that denotes

Region affinity.

D. Create and purchase a t3.small EC2 Reserved Instance for each developer in us-east-2. Assign each developer a specific EC2 instance with their name as the tag.

Answer: B

NO.22 A company is migrating its data centre from on premises to the AWS Cloud. The migration will take several months to complete. The company will use Amazon Route 53 for private DNS zones. During the migration, the company must Keep its AWS services pointed at the VPC's Route 53 Resolver for DNS. The company also must maintain the ability to resolve addresses from its on-premises DNS server A solutions architect must set up DNS so that Amazon EC2 instances can use native Route 53 endpoints to resolve on-premises DNS queries

Which configuration will meet these requirements?

A. Create a new outbound endpoint in Route 53. and attach me endpoint to the VP Ensure that the security groups that are attached to the endpoint can access the on-premises DNS server IP address on port 53 Create a new Route 53 Resolver rule that routes on-premises designated traffic to the on-premises DNS server.

B. Launch an EC2 instance that has DNS BIND installed and configured. Ensure that the security groups that are attached to the EC2 instance can access the on-premises DNS server IP address on port 53. Configure BIND to forward DNS queries to on-premises DNS server IP addresses Configure each migrated EC2 instances DNS settings to point to the BIND server IP address.

C. Configure VPC DHCP options set to point to on-premises DNS server IP addresses Ensure that security groups for EC2 instances allow outbound access to port 53 on those DNS server IP addresses.

D. Create a new private DNS zone in Route 53 with the same domain name as the on-premises domain. Create a single wildcard record with the on-premises DNS server IP address as the record's address.

Answer: A

NO.23 A financial services company receives a regular data feed from its credit card servicing partner Approximately 5.000 records are sent every 15 minutes in plaintext, delivered over HTTPS directly into an Amazon S3 bucket with server-side encryption. This feed contains sensitive credit card primary account number (PAN) data The company needs to automatically mask the PAN before sending the data to another S3 bucket for additional internal processing. The company also needs to remove and merge specific fields, and then transform the record into JSON format Additionally, extra feeds are likely to be added in the future, so any design needs to be easily expandable.

Which solutions will meet these requirements?

A. Create an AWS Glue crawler and custom classifier based upon the data feed formats and build a table definition to match. Perform an Amazon Athena query on file delivery to start an Amazon EMR ETL job to transform the entire record according to the processing and transformation requirements. Define the output format as JSON. Once complete, send the results to another S3 bucket for internal processing and scale down the EMR cluster.

B. Trigger an AWS Lambda function on file delivery that extracts each record and writes it to an Amazon SQS queue. Trigger another Lambda function when new messages arrive in the SQS queue to process the records, writing the results to a temporary location in Amazon S3. Trigger a final Lambda function once the SQS queue is empty to transform the records into JSON format and send the results to another S3 bucket for internal processing.

- C.** Trigger an AWS Lambda function on file delivery that extracts each record and writes it to an Amazon SQS queue. Configure an AWS Fargate container application to
- D.** automatically scale to a single instance when the SQS queue contains messages. Have the application process each record, and transform the record into JSON format. When the queue is empty, send the results to another S3 bucket for internal processing and scale down the AWS Fargate instance.
- E.** Create an AWS Glue crawler and custom classifier based on the data feed formats and build a table definition to match. Trigger an AWS Lambda function on file delivery to start an AWS Glue ETL job to transform the entire record according to the processing and transformation requirements. Define the output format as JSON. Once complete, have the ETL job send the results to another S3 bucket for internal processing.

Answer: D

Explanation:

You can use a Glue crawler to populate the AWS Glue Data Catalog with tables. The Lambda function can be triggered using S3 event notifications when object create events occur. The Lambda function will then trigger the Glue ETL job to transform the records masking the sensitive data and modifying the output format to JSON. This solution meets all requirements.

Create an AWS Glue crawler and custom classifier based on the data feed formats and build a table definition to match. Trigger an AWS Lambda function on file delivery to start an AWS Glue ETL job to transform the entire record according to the processing and transformation requirements. Define the output format as JSON. Once complete, have the ETL job send the results to another S3 bucket for internal processing.

<https://docs.aws.amazon.com/glue/latest/dg/trigger-job.html>

https://d1.awsstatic.com/Products/product-name/diagrams/product-page-diagram_Glue_Event-driven-ETL-Pipelines.e24d59bb79a9e24cdba7f43ffd234ec0482a60e2.png

NO.24 A solutions architect is designing a multi-account structure that has 10 existing accounts. The design must meet the following requirements

- * Consolidate all accounts into one organization
- * Allow full access to the Amazon EC2 service from the management account and the secondary accounts
- * Minimize the effort required to add additional secondary accounts

Which combination of steps should be included in the solution? (Select TWO)

- A.** Create an organization from the management account. Send invitations to the secondary accounts from the management account. Accept the invitations and create an OU
- B.** Create a full EC2 access policy and map the policy to a role in each account. Trust every other account to assume the role
- C.** Create a service control policy (SCP) that enables full EC2 access, and attach the policy to the OU
- D.** Create an organization from the management account. Send a join request to the management account from each secondary account. Accept the requests and create an OU
- E.** Create a VPC peering connection between the management account and the secondary accounts. Accept the request for the VPC peering connection

Answer: A,B

NO.25 A company is running an application in the AWS Cloud. The application runs on containers in

an Amazon Elastic Container Service (Amazon ECS) cluster. The ECS tasks use the Fargate launch type. The application's data is relational and is stored in Amazon Aurora MySQL. To meet regulatory requirements, the application must be able to recover to a separate AWS Region in the event of an application failure. In case of a failure, no data can be lost. Which solution will meet these requirements with the LEAST amount of operational overhead?

- A.** Provision an Aurora Replica in a different Region.
- B.** Use Amazon Data Lifecycle Manager (Amazon DLM) to schedule a snapshot every 5 minutes.
- C.** Set up AWS DataSync for continuous replication of the data to a different Region.
- D.** Set up AWS Database Migration Service (AWS DMS) to perform a continuous replication of the data to a different Region.

Answer: C

NO.26 A company is using AWS Organizations to manage multiple accounts. Due to regulatory requirements, the company wants to restrict specific member accounts to certain AWS Regions, where they are permitted to deploy resources. The resources in the accounts must be tagged, enforced based on a group standard, and centrally managed with minimal configuration. What should a solutions architect do to meet these requirements?

- A.** Associate the specific member accounts with a new OU. Apply a tag policy and an SCP using conditions to limit Regions.
- B.** From the AWS Billing and Cost Management console, in the master account, disable Regions for the specific member accounts and apply a tag policy on the root.
- C.** Associate the specific member accounts with the root. Apply a tag policy and an SCP using conditions to limit Regions.
- D.** Create an AWS Config rule in the specific member accounts to limit Regions and apply a tag policy.

Answer: A

NO.27 An auction website enables users to bid on collectible items. The auction rules require that each bid is processed only once and in the order it was received. The current implementation is based on a fleet of Amazon EC2 web servers that write bid records into Amazon Kinesis Data Streams. A single large instance has a cron job that runs the bid processor, which reads incoming bids from Kinesis Data Streams and processes each bid. The auction site is growing in popularity, but users are complaining that some bids are not registering.

Troubleshooting indicates that the bid processor is too slow during peak demand hours, sometimes crashes while processing, and occasionally loses track of which record is being processed.

What changes should be made to make the bid processing more reliable?

- A.** Refactor the web application to post each incoming bid to an Amazon SQS FIFO queue in place of Kinesis Data Streams. Refactor the bid processor to continuously consume the SQS queue. Place the bid processing EC2 instance in an Auto Scaling group with a minimum and a maximum size of 1.
- B.** Refactor the web application to use the Amazon Kinesis Producer Library (KPL) when posting bids to Kinesis Data Streams. Refactor the bid processor to flag each record in Kinesis Data Streams as being unread, processed, and processed. At the start of each bid processing run; scan Kinesis Data Streams for unprocessed records.
- C.** Refactor the web application to post each incoming bid to an Amazon SNS topic in place of Kinesis Data Streams. Configure the SNS topic to trigger an AWS Lambda function that processes each bid as

soon as a user submits it

D. Switch the EC2 instance type from t2 large to a larger general compute instance type Put the bid processor EC2 instances in an Auto Scaling group that scales out the number of EC2 instances running the bid processor based on the incomingRecords metric in Kinesis Data Streams

Answer: A

Explanation:

<https://aws.amazon.com/sqs/faqs/#:~:text=A%20single%20Amazon%20SQS%20message,20%2C000%20for%20a%20FIFO%20queue.>

NO.28 A company is running a three-tier web application in an on-premises data center. The frontend is served by an Apache web server, the middle tier is a monolithic Java application, and the storage tier is a PostgreSQL database.

During a recent marketing promotion, customers could not place orders through the application because the application crashed An analysis showed that all three tiers were overloaded. The application became unresponsive, and the database reached its capacity limit because of read operations. The company already has several similar promotions scheduled in the near future. A solutions architect must develop a plan for migration to AWS to resolve these issues. The solution must maximize scalability and must minimize operational effort.

Which combination of steps will meet these requirements? (Select THREE.)

- A.** Refactor the frontend so that static assets can be hosted on Amazon S3. Use Amazon CloudFront to serve the frontend to customers. Connect the frontend to the Java application.
- B.** Rehost the Apache web server of the frontend on Amazon EC2 instances that are in an Auto Scaling group. Use a load balancer in front of the Auto Scaling group. Use Amazon Elastic File System (Amazon EFS) to host the static assets that the Apache web server needs.
- C.** Refactor the Java application. Develop a Docker container to run the Java application. Use AWS Fargate to host the container.
- D.** Use AWS Database Migration Service (AWS DMS) to replatform the PostgreSQL database to an Amazon Aurora PostgreSQL database. Use Aurora Auto Scaling for read replicas.
- E.** Rehost the Java application in an AWS Elastic Beanstalk environment that includes auto scaling.
- F.** Rehost the PostgreSQL database on an Amazon EC2 instance that has twice as much memory as the on-premises server.

Answer: B,E,F

NO.29 A company provides a centralized Amazon EC2 application hosted in a single shared VPC. The centralized application must be accessible from client applications running in the VPCs of other business units. The centralized application front end is configured with a Network Load Balancer (NLB) for scalability.

Up to 10 business unit VPCs will need to be connected to the shared VPC. Some of the business unit VPC CIDR blocks overlap with the shared VPC. and some overlap with each other. Network connectivity to the centralized application in the shared VPC should be allowed from authorized business unit VPCs only.

Which network configuration should a solutions architect use to provide connectivity from the client applications in the business unit VPCs to the centralized application in the shared VPC?

- A.** Create a VPC peering connection from each business unit VPC to the shared VPC. Accept the VPC peering connections from the shared VPC console. Configure VPC routing tables to send traffic to the

VPC peering connection.

B. Create a VPC endpoint service using the centralized application NLB and enable (the option to require endpoint acceptance). Create a VPC endpoint in each of the business unit VPCs using the service name of the endpoint service. Accept authorized endpoint requests from the endpoint service console.

C. Create an AWS Transit Gateway. Attach the shared VPC and the authorized business unit VPCs to the transit gateway. Create a single transit gateway route table and associate it with all of the attached VPCs. Allow automatic propagation of routes from the attachments into the route table. Configure VPC routing tables to send traffic to the transit gateway.

D. Configure a virtual private gateway for the shared VPC and create customer gateways for each of the authorized business unit VPCs. Establish a Site-to-Site VPN connection from the business unit VPCs to the shared VPC. Configure VPC routing tables to send traffic to the VPN connection.

Answer: B

Explanation:

Amazon Transit Gateway doesn't support routing between Amazon VPCs with overlapping CIDRs. If you attach a new Amazon VPC that has a CIDR which overlaps with an already attached Amazon VPC, Amazon Transit Gateway will not propagate the new Amazon VPC route into the Amazon Transit Gateway route table.

<https://docs.aws.amazon.com/elasticloadbalancing/latest/network/load-balancer-target-groups.html#client-ip-preservation>

NO.30 A company is running a critical application that uses an Amazon RDS for MySQL database to store data

a. The RDS DB instance is deployed in Multi-AZ mode.

A recent RDS database failover test caused a 40-second outage to the application. A solutions architect needs to design a solution to reduce the outage time to less than 20 seconds.

Which combination of steps should the solutions architect take to meet these requirements? (Select THREE.)

A. Use Amazon ElastiCache for Memcached in front of the database

B. Create an Amazon Aurora Replica

C. Create an RDS for MySQL read replica

D. Use Amazon ElastiCache for Redis in front of the database.

E. Use RDS Proxy in front of the database

F. Migrate the database to Amazon Aurora MySQL

Answer: A,C,D