SAP-C02 Q&As

AWS Certified Solutions Architect - Professional

Pass Amazon SAP-C02 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.pass4itsure.com/sap-c02.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Amazon Official Exam Center

- **Instant Download** After Purchase
- **100% Money Back** Guarantee
- **365 Days** Free Update
- **800,000+** Satisfied Customers

QUESTION 1

A company is developing a new serverless API by using Amazon API Gateway and AWS Lambda. The company integrated the Lambda functions with API Gateway to use several shared libraries and custom classes.

A solutions architect needs to simplify the deployment of the solution and optimize for code reuse.

Which solution will meet these requirements?

A. Deploy the shared libraries and custom classes into a Docker image. Store the image in an S3 bucket. Create a Lambda layer that uses the Docker image as the source. Deploy the API\\'s Lambda functions as Zip packages. Configure the packages to use the Lambda layer.

B. Deploy the shared libraries and custom classes to a Docker image. Upload the image to Amazon Elastic Container Registry (Amazon ECR). Create a Lambda layer that uses the Docker image as the source. Deploy the API\\'s Lambda functions as Zip packages. Configure the packages to use the Lambda layer.

C. Deploy the shared libraries and custom classes to a Docker container in Amazon Elastic Container Service (Amazon ECS) by using the AWS Fargate launch type. Deploy the API\\'s Lambda functions as Zip packages. Configure the packages to use the deployed container as a Lambda layer.

D. Deploy the shared libraries, custom classes, and code for the API\\'s Lambda functions to a Docker image. Upload the image to Amazon Elastic Container Registry (Amazon ECR). Configure the API\\'s Lambda functions to use the Docker image as the deployment package.

Correct Answer: B

Deploying the shared libraries and custom classes to a Docker image and uploading the image to Amazon Elastic Container Registry (Amazon ECR) and creating a Lambda layer that uses the Docker image as the source. Then, deploying

the API\\'s Lambda functions as Zip packages and configuring the packages to use the Lambda layer would meet the requirements for simplifying the deployment and optimizing for code reuse. A Lambda layer is a distribution mechanism for

libraries, custom runtimes, and other function dependencies. It allows you to manage your in-development function code separately from your dependencies, this way you can easily update your dependencies without having to update your

entire function code. By deploying the shared libraries and custom classes to a Docker image and uploading the image to Amazon Elastic Container Registry (ECR), it makes it easy to manage and version the dependencies. This way, the

company can use the same version of the dependencies across different Lambda functions.

By creating a Lambda layer that uses the Docker image as the source, the company can configure the API\\'s Lambda functions to use the layer, reducing the need to include the dependencies in each function package, and making it easy to

update the dependencies across all functions at once.

Reference:

AWS Lambda Layers documentation:

https://docs.aws.amazon.com/lambda/latest/dg/configuration-layers.html

AWS Elastic Container Registry (ECR) documentation:

https://aws.amazon.com/ecr/

Building Lambda Layers with Docker documentation:

https://aws.amazon.com/blogs/compute/building-lambda-layers-with-docker/

QUESTION 2

A company has VPC flow logs enabled for its NAT gateway. The company is seeing Action = ACCEPT for inbound traffic that comes from public IP address

198.51.100.2 destined for a private Amazon EC2 instance.

A solutions architect must determine whether the traffic represents unsolicited inbound connections from the internet. The first two octets of the VPC CIDR block are 203.0.

Which set of steps should the solutions architect take to meet these requirements?

A. Open the AWS CloudTrail console. Select the log group that contains the NAT gateway\\'s elastic network interface and the private instance\\'s elastic network interface. Run a query to filter with the destination address set as "like 203.0" and the source address set as "like 198.51.100.2". Run the stats command to filter the sum of bytes transferred by the source address and the destination address.

B. Open the Amazon CloudWatch console. Select the log group that contains the NAT gateway\\'s elastic network interface and the private instance\\'s elastic network interface. Run a query to filter with the destination address set as "like 203.0" and the source address set as "like 198.51.100.2". Run the stats command to filter the sum of bytes transferred by the source address and the destination address.

C. Open the AWS CloudTrail console. Select the log group that contains the NAT gateway\\'s elastic network interface and the private instance\\'s elastic network interface. Run a query to filter with the destination address set as "like 198.51.100.2" and the source address set as "like 203.0". Run the stats command to filter the sum of bytes transferred by the source address and the destination address.

D. Open the Amazon CloudWatch console. Select the log group that contains the NAT gateway\\'s elastic network interface and the private instance\\'s elastic network interface. Run a query to filter with the destination address set as "like 198.51.100.2" and the source address set as "like 203.0". Run the stats command to filter the sum of bytes transferred by the source address and the destination address.

Correct Answer: D

https://aws.amazon.com/premiumsupport/knowledge-center/vpc-analyze-inbound-traffic-nat-gateway/ by Cloudxie says "select appropriate log"

QUESTION 3

A company is deploying AWS Lambda functions that access an Amazon RDS for PostgreSQL database. The company needs to launch the Lambda functions in a QA environment and in a production environment.

The company must not expose credentials within application code and must rotate passwords automatically. Which solution will meet these requirements?

A. Store the database credentials for both environments in AWS Systems Manager Parameter Store. Encrypt the credentials by using an AWS Key Management Service (AWS KMS) key. Within the application code of the Lambda functions, pull the credentials from the Parameter Store parameter by using the AWS SDK for Python (Bot03). Add a role to the Lambda functions to provide access to the Parameter Store parameter.

B. Store the database credentials for both environments in AWS Secrets Manager with distinct key entry for the QA environment and the production environment. Turn on rotation. Provide a reference to the Secrets Manager key as an environment variable for the Lambda functions.

C. Store the database credentials for both environments in AWS Key Management Service (AWS KMS). Turn on rotation. Provide a reference to the credentials that are stored in AWS KMS as an environment variable for the Lambda functions.

D. Create separate S3 buckets for the QA environment and the production environment. Turn on server-side encryption with AWS KMS keys (SSE-KMS) for the S3 buckets. Use an object naming pattern that gives each Lambda function\\'s application code the ability to pull the correct credentials for the function\\'s corresponding environment. Grant each Lambda function\\'s execution role access to Amazon S3.

Correct Answer: B

The best solution is to store the database credentials for both environments in AWS Secrets Manager with distinct key entry for the QA environment and the production environment. AWS Secrets Manager is a web service that can securely store, manage, and retrieve secrets, such as database credentials. AWS Secrets Manager also supports automatic rotation of secrets by using Lambda functions or built-in rotation templates. By storing the database credentials for both environments in AWS Secrets Manager, the company can avoid exposing credentials within application code and rotate passwords automatically. By providing a reference to the Secrets Manager key as an environment variable for the Lambda functions, the company can easily access the credentials from the code by using the AWS SDK. This solution meets all the requirements of the company. References: AWS Secrets Manager Documentation, Using AWS Lambda with AWS Secrets Manager, Using environment variables - AWS Lambda

---

QUESTION 4

A solutions architect needs to improve an application that is hosted in the AWS Cloud. The application uses an Amazon Aurora MySQL DB instance that is experiencing overloaded connections. Most of the application\\'s operations insert records into the database. The application currently stores credentials in a text-based configuration file.

The solutions architect needs to implement a solution so that the application can handle the current connection load. The solution must keep the credentials secure and must provide the ability to rotate the credentials automatically on a regular basis.

Which solution will meet these requirements?

A. Deploy an Amazon RDS Proxy layer in front of the DB instance. Store the connection credentials as a secret in AWS Secrets Manager.

B. Deploy an Amazon RDS Proxy layer in front of the DB instance. Store the connection credentials in AWS Systems Manager Parameter Store.

C. Create an Aurora Replica. Store the connection credentials as a secret in AWS Secrets Manager.

D. Create an Aurora Replica. Store the connection credentials in AWS Systems Manager Parameter Store.

Correct Answer: A

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/rds-proxy.html

QUESTION 5

A company has applications in an AWS account that is named Source. The account is in an organization in AWS Organizations. One of the applications uses AWS Lambda functions and stores inventory data in an Amazon Aurora database. The application deploys the Lambda functions by using a deployment package. The company has configured automated backups for Aurora.

The company wants to migrate the Lambda functions and the Aurora database to a new AWS account that is named Target. The application processes critical data, so the company must minimize downtime.

Which solution will meet these requirements?

A. Download the Lambda function deployment package from the Source account. Use the deployment package and create new Lambda functions in the Target account. Share the automated Aurora DB cluster snapshot with the Target account.

B. Download the Lambda function deployment package from the Source account. Use the deployment package and create new Lambda functions in the Target account Share the Aurora DB cluster with the Target account by using AWS Resource Access Manager {AWS RAM). Grant the Target account permission to clone the Aurora DB cluster.

C. Use AWS Resource Access Manager (AWS RAM) to share the Lambda functions and the Aurora DB cluster with the Target account. Grant the Target account permission to clone the Aurora DB cluster.

D. Use AWS Resource Access Manager (AWS RAM) to share the Lambda functions with the Target account. Share the automated Aurora DB cluster snapshot with the Target account.

Correct Answer: C

This solution uses a combination of AWS Resource Access Manager (RAM) and automated backups to migrate the Lambda functions and the Aurora database to the Target account while minimizing downtime. In this solution, the Lambda function deployment package is downloaded from the Source account and used to create new Lambda functions in the Target account. The Aurora DB cluster is shared with the Target account using AWS RAM and the Target account is granted permission to clone the Aurora DB cluster, allowing for a new copy of the Aurora database to be created in the Target account. This approach allows for the data to be migrated to the Target account while minimizing downtime, as the Target account can use the cloned Aurora database while the original Aurora database continues to be used in the Source account.

QUESTION 6

A solutions architect needs to copy data from an Amazon S3 bucket m an AWS account to a new S3 bucket in a new AWS account. The solutions architect must implement a solution that uses the AWS CLI.

Which combination of steps will successfully copy the data? (Choose three.)

A. Create a bucket policy to allow the source bucket to list its contents and to put objects and set object ACLs in the destination bucket. Attach the bucket policy to the destination bucket.

B. Create a bucket policy to allow a user In the destination account to list the source bucket\\'s contents and read the source bucket\\'s objects. Attach the bucket policy to the source bucket.

C. Create an IAM policy in the source account. Configure the policy to allow a user In the source account to list contents and get objects In the source bucket, and to list contents, put objects, and set object ACLs in the destination bucket. Attach the policy to the user _

D. Create an IAM policy in the destination account. Configure the policy to allow a user In the destination account to list contents and get objects In the source bucket, and to list contents, put objects, and set objectACLs in the destination bucket. Attach the policy to the user.

E. Run the aws s3 sync command as a user in the source account. Specify\\' the source and destination buckets to copy the data.

F. Run the aws s3 sync command as a user in the destination account. Specify\\' the source and destination buckets to copy the data.

Correct Answer: BDF

Step B is necessary so that the user in the destination account has the necessary permissions to access the source bucket and list its contents, read its objects. Step D is needed so that the user in the destination account has the necessary permissions to access the destination bucket and list contents, put objects, and set object ACLs Step F is necessary because the aws s3 sync command needs to be run using the IAM user credentials from the destination account, so that the objects will have the appropriate permissions for the user in the destination account once they are copied.

QUESTION 7

A company is planning to migrate its business-critical applications from an on-premises data center to AWS. The company has an on-premises installation of a

Microsoft SQL Server Always On cluster. The company wants to migrate to an AWS managed database service. A solutions architect must design a heterogeneous database migration on AWS.

Which solution will meet these requirements?

A. Migrate the SQL Server databases to Amazon RDS for MySQL by using backup and restore utilities.

B. Use an AWS Snowball Edge Storage Optimized device to transfer data to Amazon S3. Set up Amazon RDS for MySQL. Use S3 integration with SQL Server features, such as BULK INSERT.

C. Use the AWS Schema Conversion Tool to translate the database schema to Amazon RDS for MeSQL. Then use AWS Database Migration Service (AWS DMS) to migrate the data from on-premises databases to Amazon RDS.

D. Use AWS DataSync to migrate data over the network between on-premises storage and Amazon S3. Set up Amazon RDS for MySQL. Use S3 integration with SQL Server features, such as BULK INSERT.

Correct Answer: C

https://aws.amazon.com/dms/schema-conversion-tool/

AWS Schema Conversion Tool (SCT) can automatically convert the database schema from Microsoft SQL Server to Amazon RDS for MySQL. This allows for a smooth transition of the database schema without any manual intervention. AWS

DMS can then be used to migrate the data from the on-premises databases to the newly created Amazon RDS for MySQL instance. This service can perform a one-time migration of the data or can set up ongoing replication of data changes

to keep the on-premises and AWS databases in sync.

QUESTION 8

A company wants to record key performance indicators (KPIs) from its application as part of a strategy to convert to a user-based licensing schema. The application is a multi-tier application with a web-based UI. The company saves all log files to Amazon CloudWatch by using the CloudWatch agent. All logins to the application are saved in a log file.

As part of the new license schema, the company needs to find out how many unique users each client has on a daily basis, weekly basis, and monthly basis.

Which solution will provide this information with the LEAST change to the application?

A. Configure an Amazon CloudWatch Logs metric filter that saves each successful login as a metric. Configure the user name and client name as dimensions for the metric.

B. Change the application logic to make each successful login generate a call to the AWS SDK to increment a custom metric that records user name and client name dimensions in CloudWatch.

C. Configure the CloudWatch agent to extract successful login metrics from the logs. Additionally, configure the CloudWatch agent to save the successful login metrics as a custom metric that uses the user name and client name as dimensions for the metric.

D. Configure an AWS Lambda function to consume an Amazon CloudWatch Logs stream of the application logs. Additionally, configure the Lambda function to increment a custom metric in CloudWatch that uses the user name and client name as dimensions for the metric.

Correct Answer: A

QUESTION 9

A solutions architect is reviewing an application\\'s resilience before launch. The application runs on an Amazon EC2 instance that is deployed in a private subnet of a VPC.

The EC2 instance is provisioned by an Auto Scaling group that has a minimum capacity of I and a maximum capacity of I. The application stores data on an Amazon RDS for MySQL DB instance. The VPC has subnets configured in three Availability Zones and is configured with a single NAT gateway.

The solutions architect needs to recommend a solution to ensure that the application will operate across multiple Availability Zones.

Which solution will meet this requirement?

A. Deploy an additional NAT gateway in the other Availability Zones. Update the route tables with appropriate routes. Modify the RDS for MySQL DB instance to a Multi-AZ configuration. Configure the Auto Scaling group to launch instances across Availability Zones. Set the minimum capacity and maximum capacity of the Auto Scaling group to 3.

B. Replace the NAT gateway with a virtual private gateway. Replace the RDS for MySQL DB instance with an Amazon Aurora MySQL DB cluster. Configure the Auto Scaling group to launch instances across all subnets in the VPC. Set the minimum capacity and maximum capacity of the Auto Scaling group to 3.

C. Replace the NAT gateway with a NAT instance. Migrate the RDS for MySQL DB instance to an RDS for PostgreSQL DB instance. Launch a new EC2 instance in the other Availability Zones.

D. Deploy an additional NAT gateway in the other Availability Zones. Update the route tables with appropriate routes. Modify the RDS for MySQL DB instance to turn on automatic backups and retain the backups for 7 days. Configure the Auto Scaling group to launch instances across all subnets in the VPC. Keep the minimum capacity and the maximum

capacity of the Auto Scaling group at 1.

Correct Answer: A

QUESTION 10

A solutions architect works for a government agency that has strict disaster recovery requirements All Amazon Elastic Block Store (Amazon EBS) snapshots are required to be saved in at least two additional AWS Regions. The agency also is required to maintain the lowest possible operational overhead.

Which solution meets these requirements?

A. Configure a policy in Amazon Data Lifecycle Manager (Amazon DLMJ to run once daily to copy the EBS snapshots to the additional Regions.

B. Use Amazon EventBridge (Amazon CloudWatch Events) to schedule an AWS Lambda function to copy the EBS snapshots to the additional Regions.

C. Set up AWS Backup to create the EBS snapshots. Configure Amazon S3 cross-Region replication to copy the EBS snapshots to the additional Regions.

D. Schedule Amazon EC2 Image Builder to run once daily to create an AMI and copy the AMI to the additional Regions.

Correct Answer: A

Amazon DLM features:

Automated snapshot and AMI creation:

Create a policy that automates the creation, retention, and deletion of EBS snapshots and EBS-backed AMIs.

Fast snapshot restore integration:

Automate the creation of snapshots that are enabled for fast snapshot restore. Fast snapshot restore enables you to restore volumes that are fully initialized at creation and instantly deliver all of their provisioned performance.

Built-in cross-Region copy:

Automatically copy snapshots that are created by a lifecycle policy to up to three AWS Regions.

Automated cross-account snapshot copy:

Use cross-account sharing in conjunction with a cross-account copy event policy to automatically share and copy snapshots created by a policy across accounts.

QUESTION 11

A company wants to use Amazon Workspaces in combination with thin client devices to replace aging desktops Employees use the desktops to access applications that work with clinical trial data Corporate security policy states that access to the applications must be restricted to only company branch office locations. The company is considering adding an additional branch office in the next 6 months.

Which solution meets these requirements with the MOST operational efficiency?

A. Create an IP access control group rule with the list of public addresses from the branch offices Associate the IP access control group with the Workspaces directory

B. Use AWS Firewall Manager to create a web ACL rule with an IPSet with the list of public addresses from the branch office locations Associate the web ACL with the Workspaces directory

C. Use AWS Certificate Manager (ACM) to issue trusted device certificates to the machines deployed in the branch office locations Enable restricted access on the Workspaces directory

D. Create a custom Workspace image with Windows Firewall configured to restrict access to the public addresses of the branch offices Use the image to deploy the Workspaces.

Correct Answer: A

QUESTION 12

A company has deployed an Amazon Connect contact center. Contact center agents are reporting large numbers of computer-generated calls. The company is concerned about the cost and productivity effects of these calls. The company wants a solution that will allow agents to flag the call as spam and automatically block the numbers from going to an agent in the future.

What is the MOST operationally efficient solution to meet these requirements?

A. Customize the Contact Control Panel (CCP) by adding a flag call button that will invoke an AWS Lambda function that calls the UpdateContactAttributes API. Use an Amazon DynamoDB table to store the spam numbers. Modify the contact flows to look for the updated attribute and to use a Lambda function to read and write to the DynamoDB table.

B. Use a Contact Lens for Amazon Connect rule that will look for spam calls. Use an Amazon DynamoDB table to store the spam numbers. Modify the contact flows to look for the rule and to invoke an AWS Lambda function to read and write to the DynamoDB table.

C. Use an Amazon DynamoDB table to store the spam numbers. Create a quick connect that the agents can transfer the spam call to from the Contact Control Panel (CCP). Modify the quick connect contact flow to invoke an AWS Lambda function to write to the DynamoDB table.

D. Modify the initial contact flow to ask for caller input. If the agent does not receive input, the agent should mark the caller as spam. Use an Amazon DynamoDB table to store the spam numbers. Use an AWS Lambda function to read and write to the DynamoDB table.

Correct Answer: A

QUESTION 13

A company uses an Amazon Aurora PostgreSQL DB cluster for applications in a single AWS Region. The company\\'s database team must monitor all data activity on all the databases.

Which solution will achieve this goal?

A. Set up an AWS Database Migration Service (AWS DMS) change data capture (CDC) task. Specify the Aurora DB

cluster as the source. Specify Amazon Kinesis Data Firehose as the target. Use Kinesis Data Firehose to upload the data into an Amazon OpenSearch Service cluster for further analysis.

B. Start a database activity stream on the Aurora DB cluster to capture the activity stream in Amazon EventBridge. Define an AWS Lambda function as a target for EventBridge. Program the Lambda function to decrypt the messages from EventBridge and to publish all database activity to Amazon S3 for further analysis.

C. Start a database activity stream on the Aurora DB cluster to push the activity stream to an Amazon Kinesis data stream. Configure Amazon Kinesis Data Firehose to consume the Kinesis data stream and to deliver the data to Amazon S3 for further analysis.

D. Set up an AWS Database Migration Service (AWS DMS) change data capture (CDC) task. Specify the Aurora DB cluster as the source. Specify Amazon Kinesis Data Firehose as the target. Use Kinesis Data Firehose to upload the data into an Amazon Redshift cluster. Run queries on the Amazon Redshift data to determine database activities on the Aurora database.

Correct Answer: C

QUESTION 14

A company has an application that generates reports and stores them in an Amazon S3 bucket. When a user accesses their report, the application generates a signed URL to allow the user to download the report. The company\\'s security team has discovered that the files are public and that anyone can download them without authentication. The company has suspended the generation of new reports until the problem is resolved.

Which set of actions will immediately remediate the security issue without impacting the application\\'s normal workflow?

A. Create an AWS Lambda function that applies a deny all policy for users who are not authenticated. Create a scheduled event to invoke the Lambda function.

B. Review the AWS Trusted Advisor bucket permissions check and implement the recommended actions.

C. Run a script that puts a private ACL on all of the objects in the bucket.

D. Use the Block Public Access feature in Amazon S3 to set the IgnorePublicAcls option to TRUE on the bucket.

Correct Answer: D

The S3 bucket is allowing public access and this must be immediately disabled. Setting the IgnorePublicAcls option to TRUE causes Amazon S3 to ignore all public ACLs on a bucket and any objects that it contains. The other settings you can configure with the Block Public Access Feature are:

1.

BlockPublicAcls - PUT bucket ACL and PUT objects requests are blocked if granting public access.

2.

BlockPublicPolicy - Rejects requests to PUT a bucket policy if granting public access.

3.

RestrictPublicBuckets - Restricts access to principles in the bucket owners\\' AWS account.
https://aws.amazon.com/s3/features/block-public-access/

QUESTION 15

A delivery company needs to migrate its third-party route planning application to AWS. The third party supplies a supported Docker image from a public registry. The image can run in as many containers as required to generate the route map.

The company has divided the delivery area into sections with supply hubs so that delivery drivers travel the shortest distance possible from the hubs to the customers. To reduce the time necessary to generate route maps, each section uses its own set of Docker containers with a custom configuration that processes orders only in the section\\'s area. The company needs the ability to allocate resources cost-effectively based on the number of running containers.

Which solution will meet these requirements with the LEAST operational overhead?

A. Create an Amazon Elastic Kubernetes Service (Amazon EKS) cluster on Amazon EC2. Use the Amazon EKS CLI to launch the planning application in pods by using the -tags option to assign a custom tag to the pod.

B. Create an Amazon Elastic Kubernetes Service (Amazon EKS) cluster on AWS Fargate. Use the Amazon EKS CLI to launch the planning application. Use the AWS CLI tag- resource API call to assign a custom tag to the pod.

C. Create an Amazon Elastic Container Service (Amazon ECS) cluster on Amazon EC2. Use the AWS CLI with run-tasks set to true to launch the planning application by using the - tags option to assign a custom tag to the task.

D. Create an Amazon Elastic Container Service (Amazon ECS) cluster on AWS Fargate. Use the AWS CLI run-task command and set enableECSManagedTags to true to launch the planning application. Use the --tags option to assign a custom tag to the task.

Correct Answer: D

Amazon Elastic Container Service (ECS) on AWS Fargate is a fully managed service that allows you to run containers without having to manage the underlying infrastructure. When you launch tasks on Fargate, resources are automatically allocated based on the number of tasks running, which reduces the operational overhead. Using ECS on Fargate allows you to assign custom tags to tasks using the -tags option in the run-task command, as described in the documentation:https://docs.aws.amazon.com/cli/latest/reference/ecs/run-task.html You can also set enableECSManagedTags to true, which allows the service to automatically add the cluster name and service name as t ags.https://docs.aws.amazon.com/AmazonECS/latest/developerguide/task-placementconstraints.html#tag-based-scheduling

SAP-C02 PDF Dumps SAP-C02 Study Guide SAP-C02 Exam Questions