

➤ **Vendor: Amazon**

➤ **Exam Code: AWS-Certified-Solutions-Architect-Professional**

➤ **Exam Name: AWS Certified Solutions Architect - Professional (SAP-C01)**

➤ **New Updated Questions from [Braindump2go](https://www.braindump2go.com) (Updated in [August/2021](#))**

Visit Braindump2go and Download Full Version SAP-C01 Exam Dumps

QUESTION 889

A company has a new application that needs to run on five Amazon EC2 instances in a single AWS Region. The application requires high-throughput, low-latency network connections between all of the EC2 instances where the application will run.

There is no requirement for the application to be fault tolerant.

Which solution will meet these requirements?

- A. Launch five new EC2 instances into a cluster placement group. Ensure that the EC2 instance type supports enhanced networking.
- B. Launch five new EC2 instances into an Auto Scaling group in the same Availability Zone. Attach an extra elastic network interface to each EC2 instance.
- C. Launch five new EC2 instances into a partition placement group. Ensure that the EC2 instance type supports enhanced networking.
- D. Launch five new EC2 instances into a spread placement group. Attach an extra elastic network interface to each EC2 instance.

Answer: D

QUESTION 890

A company is deploying a new cluster for big data analytics on AWS. The cluster will run across many Linux Amazon EC2 instances that are spread across multiple Availability Zones.

All of the nodes in the cluster must have read and write access to common underlying file storage. The file storage must be highly available, must be resilient, must be compatible with the Portable Operating System Interface (POSIX), and must accommodate high levels of throughput. Which storage solution will meet these requirements?

- A. Provision an AWS Storage Gateway file gateway NFS file share that is attached to an Amazon S3 bucket. Mount the NFS file share on each EC2 instance in the cluster.
- B. Provision a new Amazon Elastic File System (Amazon EFS) file system that uses General Purpose performance mode. Mount the EFS file system on each EC2 instance in the cluster.
- C. Provision a new Amazon Elastic Block Store (Amazon EBS) volume that uses the io2 volume type. Attach the EBS volume to all of the EC2 instances in the cluster.
- D. Provision a new Amazon Elastic File System (Amazon EFS) file system that uses Max I/O performance mode. Mount the EFS file system on each EC2 instance in the cluster.

Answer: D

QUESTION 891

[SAP-C01 Exam Dumps](#) [SAP-C01 Exam Questions](#) [SAP-C01 PDF Dumps](#) [SAP-C01 VCE Dumps](#)

<https://www.braindump2go.com/aws-certified-solutions-architect-professional.html>

A company is migrating applications from on premises to the AWS Cloud. These applications power the company's internal web forms.

These web forms collect data for specific events several times each quarter. The web forms use simple SQL statements to save the data to a local relational database.

Data collection occurs for each event, and the on-premises servers are idle most of the time. The company needs to minimize the amount of idle infrastructure that supports the web forms.

Which solution will meet these requirements?

- A. Use Amazon EC2 Image Builder to create AMIs for the legacy servers.
Use the AMIs to provision EC2 instances to recreate the applications in the AWS Cloud.
Place an Application Load Balancer (ALB) in front of the EC2 instances.
Use Amazon Route 53 to point the DNS names of the web forms to the ALB.
- B. Create one Amazon DynamoDB table to store data for all the data input.
Use the application form name as the table key to distinguish data items.
Create an Amazon Kinesis data stream to receive the data input and store the input in DynamoDB.
Use Amazon Route 53 to point the DNS names of the web forms to the Kinesis data stream's endpoint.
- C. Create Docker images for each server of the legacy web form applications.
Create an Amazon Elastic Container Service (Amazon ECS) cluster on AWS Fargate.
Place an Application Load Balancer in front of the ECS cluster.
Use Fargate task storage to store the web form data.
- D. Provision an Amazon Aurora Serverless cluster.
Build multiple schemas for each web form's data storage.
Use Amazon API Gateway and an AWS Lambda function to recreate the data input forms.
Use Amazon Route 53 to point the DNS names of the web forms to their corresponding API Gateway endpoint.

Answer: A

QUESTION 892

A company has developed a single-page web application in JavaScript. The source code is stored in a single Amazon S3 bucket in the us-east-1 Region.

The company serves the web application to a global user base through Amazon CloudFront.

The company wants to experiment with two versions of the website without informing application users. Each version of the website will reside in its own S3 bucket.

The company wants to determine which version is most successful in marketing a new product. The solution must send application users that are based in Europe to the new website design.

The solution must send application users that are based in the United States to the current website design.

However, some exceptions exist. The company needs to be able to redirect specific users to the new website design, regardless of the users' location.

Which solution meets these requirements?

- A. Configure two CloudFront distributions.
Configure a geolocation routing policy in Amazon Route 53 to route traffic to the appropriate CloudFront endpoint based on the location of clients.
- B. Configure a single CloudFront distribution.
Create a behavior with different paths for each version of the site.
Configure Lambda@Edge on the default path to generate redirects and send the client to the correct version of the website.
- C. Configure a single CloudFront distribution.
Configure an alternate domain name on the distribution.
Configure two behaviors to route users to the different S3 origins based on the domain name that the client uses in the HTTP request.
- D. Configure a single CloudFront distribution with Lambda@Edge.
Use Lambda@Edge to send user requests to different origins based on request attributes.

Answer: A

QUESTION 893

A company is running a data-intensive application on AWS. The application runs on a cluster of hundreds of Amazon EC2 instances. A shared file system also runs on several EC2 instances that store 200 TB of data. The application reads and modifies the data on the shared file system and generates a report. The job runs once monthly, reads a subset of the files from the shared file system, and takes about 72 hours to complete. The compute instances scale in an Auto Scaling group, but the instances that host the shared file system run continuously. The compute and storage instances are all in the same AWS Region.

A solutions architect needs to reduce costs by replacing the shared file system instances. The file system must provide high performance access to the needed data for the duration of the 72-hour run.

Which solution will provide the LARGEST overall cost reduction while meeting these requirements?

- A. Migrate the data from the existing shared file system to an Amazon S3 bucket that uses the S3 Intelligent-Tiering storage class.
Before the job runs each month, use Amazon FSx for Lustre to create a new file system with the data from Amazon S3 by using lazy loading.
Use the new file system as the shared storage for the duration of the job.
Delete the file system when the job is complete.
- B. Migrate the data from the existing shared file system to a large Amazon Elastic Block Store (Amazon EBS) volume with Multi-Attach enabled.
Attach the EBS volume to each of the instances by using a user data script in the Auto Scaling group launch template.
Use the EBS volume as the shared storage for the duration of the job.
Detach the EBS volume when the job is complete.
- C. Migrate the data from the existing shared file system to an Amazon S3 bucket that uses the S3 Standard storage class.
Before the job runs each month, use Amazon FSx for Lustre to create a new file system with the data from Amazon S3 by using batch loading.
Use the new file system as the shared storage for the duration of the job.
Delete the file system when the job is complete.
- D. Migrate the data from the existing shared file system to an Amazon S3 bucket.
Before the job runs each month, use AWS Storage Gateway to create a file gateway with the data from Amazon S3.
Use the file gateway as the shared storage for the job.
Delete the file gateway when the job is complete.

Answer: B

QUESTION 894

A company runs a popular web application in an on-premises data center. The application receives four million views weekly. The company expects traffic to increase by 200% because of an advertisement that will be published soon. The company needs to decrease the load on the origin before the increase of traffic occurs. The company does not have enough time to move the entire application to the AWS Cloud.

Which solution will meet these requirements?

- A. Create an Amazon CloudFront content delivery network (CDN).
Enable query forwarding to the origin.
Create a managed cache policy that includes query strings.
Use an on-premises load balancer as the origin.
Offload the DNS querying to AWS to handle CloudFront CDN traffic.
- B. Create an Amazon CloudFront content delivery network (CDN) that uses a Real Time Messaging Protocol (RTMP) distribution.
Enable query forwarding to the origin. Use an on-premises load balancer as the origin.
Offload the DNS querying to AWS to handle CloudFront CDN traffic.
- C. Create an accelerator in AWS Global Accelerator.
Add listeners for HTTP and HTTPS TCP ports.

Create an endpoint group. Create a Network Load Balancer (NLB), and attach it to the endpoint group.

Point the NLB to the on-premises servers.

Offload the DNS querying to AWS to handle AWS Global Accelerator traffic.

D. Create an accelerator in AWS Global Accelerator.

Add listeners for HTTP and HTTPS TCP ports.

Create an endpoint group.

Create an Application Load Balancer (ALB), and attach it to the endpoint group.

Point the ALB to the on-premises servers.

Offload the DNS querying to AWS to handle AWS Global Accelerator traffic.

Answer: C

QUESTION 895

A an Auto Scaling group behind an Application Load Balancer. The load on the application varies throughout the day, and EC2 instances are scaled in and out on a regular basis. Log files from the EC2 instances are copied to a central Amazon S3 bucket every 15 minutes. The security team discovers that log files are missing from some of the terminated EC2 instances.

Which set of actions will ensure that log files are copied to the central S3 bucket from the terminated EC2 instances?

A. Create a script to copy log files to Amazon S3, and store the script in a file on the EC2 instance. Create an Auto Scaling lifecycle hook and an Amazon EventBridge (Amazon CloudWatch Events) rule to detect lifecycle events from the Auto Scaling group.

Invoke an AWS Lambda function on the autoscaling:EC2_INSTANCE_TERMINATING transition to send ABANDON to the Auto Scaling group to prevent termination, run the script to copy the log files, and terminate the instance using the AWS SDK.

B. Create an AWS Systems Manager document with a script to copy log files to Amazon S3. Create an Auto Scaling lifecycle hook and an Amazon EventBridge (Amazon CloudWatch Events) rule to detect lifecycle events from the Auto Scaling group.

Invoke an AWS Lambda function on the autoscaling:EC2_INSTANCE_TERMINATING transition to call the AWS Systems Manager API SendCommand operation to run the document to copy the log files and send CONTINUE to the Auto Scaling group to terminate the instance.

C. Change the log delivery rate to every 5 minutes. Create a script to copy log files to Amazon S3, and add the script to EC2 instance user data.

Create an Amazon EventBridge (Amazon CloudWatch Events) rule to detect EC2 instance termination.

Invoke an AWS Lambda function from the EventBridge (CloudWatch Events) rule that uses the AWS CLI to run the user-data script to copy the log files and terminate the instance.

D. Create an AWS Systems Manager document with a script to copy log files to Amazon S3.

Create an Auto Scaling lifecycle hook that publishes a message to an Amazon Simple Notification Service (Amazon SNS) topic.

From the SNS notification, call the AWS Systems Manager API SendCommand operation to run the document to copy the log files and send ABANDON to the Auto Scaling group to terminate the instance.

Answer: A

QUESTION 896

A company has developed a web application. The company is hosting the application on a group of Amazon EC2 instances behind an Application Load Balancer. The company wants to improve the security posture of the application and plans to use AWS WAF web ACLs. The solution must not adversely affect legitimate traffic to the application. How should a solutions architect configure the web ACLs to meet these requirements?

A. Set the action of the web ACL rules to Count.

Enable AWS WAF logging Analyze the requests for false positives.

Modify the rules to avoid any false positive.

Over time change the action of the web ACL rules from Count to Block.

- B. Use only rate-based rules in the web ACLs.
And set the throttle limit as high as possible.
Temporarily block all requests that exceed the limit.
Define nested rules to narrow the scope of the rate tracking.
- C. Set the action of the web ACL rules to Block.
Use only AWS managed rule groups in the web ACLs.
Evaluate the rule groups by using Amazon CloudWatch metrics with AWS WAF sampled requests or AWS WAF logs.
- D. Use only custom rule groups in the web ACLs.
And set the action to Allow Enable AWS WAF logging.
Analyze the requests for false positives.
Modify the rules to avoid any false positive .
Over time, change the action of the web ACL rules from Allow to Block.

Answer: B

QUESTION 897

A company manages an on-premises JavaScript front-end web application. The application is hosted on two servers secured with a corporate Active Directory. The application calls a set of Java-based microservices on an application server and stores data in a clustered MySQL database. The application is heavily used during the day on weekdays. It is lightly used during the evenings and weekends. Daytime traffic to the application has increased rapidly, and reliability has diminished as a result. The company wants to migrate the application to AWS with a solution that eliminates the need for server maintenance, with an API to securely connect to the microservices. Which combination of actions will meet these requirements? (Select THREE.)

- A. Host the web application on Amazon S3.
Use Amazon Cognito identity pools (federated identities) with SAML for authentication and authorization.
- B. Host the web application on Amazon EC2 with Auto Scaling.
Use Amazon Cognito federation and Login with Amazon for authentication and authorization.
- C. Create an API layer with Amazon API Gateway.
Rehost the microservices on AWS Fargate containers.
- D. Create an API layer with Amazon API Gateway.
Rehost the microservices on Amazon Elastic Container Service (Amazon ECS) containers.
- E. Replatform the database to Amazon RDS for MySQL.
- F. Replatform the database to Amazon Aurora MySQL Serverless.

Answer: ACD

QUESTION 898

A company wants to host a new global website that consists of static content. A solutions architect is working on a solution that uses Amazon CloudFront with an origin access identity (OAI) to access website content that is stored in a private Amazon S3 bucket. During testing, the solutions architect receives 404 errors from the S3 bucket. Error messages appear only for attempts to access paths that end with a forward slash, such as example.com/path/. These requests should return the existing S3 object path/index.html. Any potential solution must not prevent CloudFront from caching the content.

What should the solutions architect do to resolve this problem?

- A. Change the CloudFront origin to an Amazon API Gateway proxy endpoint.
Rewrite the S3 request URL by using an AWS Lambda function.
- B. Change the CloudFront origin to an Amazon API Gateway endpoint.
Rewrite the S3 request URL in an AWS service integration.
- C. Change the CloudFront configuration to use an AWS Lambda@Edge function that is invoked by a viewer request event to rewrite the S3 request URL.
- D. Change the CloudFront configuration to use an AWS Lambda@Edge function that is invoked by an origin request event to rewrite the S3 request URL.

Answer: C

QUESTION 899

A company has developed an application that is running Windows Server on VMware vSphere VMs that the company hosts on-premises. The application data is stored in a proprietary format that must be read through the application. The company manually provisioned the servers and the application. As part of its disaster recovery plan, the company wants the ability to host its application on AWS temporarily if the company's on-premises environment becomes unavailable. The company wants the application to return to on-premises hosting after a disaster recovery event is complete within the RPO in 5 minutes.

Which solution meets these requirements with the LEAST amount of operational overhead?

- A. Configure AWS DataSync. Replicate the data to Amazon Elastic Block Store (Amazon EBS) volumes.
When the on-premises environment is unavailable, use AWS CloudFormation templates to provision Amazon EC2 instances and attach the EBS volumes.
- B. Configure CloudEndure Disaster Recovery. Replicate the data to replication Amazon EC2 instances that are attached to Amazon Elastic Block Store (Amazon EBS) volumes.
When the on-premises environment is unavailable, use CloudEndure to launch EC2 instances that use the replicated volumes.
- C. Provision an AWS Storage Gateway File gateway. Replicate the data to an Amazon S3 bucket.
When the on-premises environment is unavailable, use AWS Backup to restore the data to Amazon Elastic Block Store (Amazon EBS) volumes and launch Amazon EC2 instances from these EBS volumes.
- D. Provision an Amazon FSx for Windows File Server file system on AWS. Replicate the data to the data system.
When the on-premises environment is unavailable, use AWS CloudFormation templates to provision Amazon EC2 instances and use `AWS::CloudFormation::Init` commands to mount the Amazon FSx file shares.

Answer: D

QUESTION 900

A company plans to refactor a monolithic application into a modern application designed to be deployed on AWS. The CI/CD pipeline needs to be upgraded to support the modern design for the application with the following requirements:

• It should allow changes to be released several times every hour.

• It should be able to roll back the changes as quickly as possible. Which design will meet these requirements?

- A. Deploy a CI-CD pipeline that incorporates AMIs to contain the application and their configurations.
Deploy the application by replacing Amazon EC2 instances.
- B. Specify AWS Elastic Beanstalk to serve as a secondary environment as the deployment target for the CI/CD pipeline of the application.
To deploy, swap the staging and production environment URLs.
- C. Use AWS Systems Manager to re-provision the infrastructure for each deployment.
Update the Amazon EC2 user data to pull the latest code artifact from Amazon S3 and use Amazon Route 53 weighted routing to point to the new environment.
- D. Roll out application updates as part of an Auto Scaling event using prebuilt AMIs.
Use new versions of the AMIs to add instances, and phase out all instances that use the previous AMI version with the configured termination policy during a deployment event.

Answer: B

QUESTION 901

A data analytics company has an Amazon Redshift cluster that consists of several reserved nodes. The cluster is experiencing unexpected bursts of usage because a team of employees is compiling a deep audit analysis report.

The queries to generate the report are complex read queries and are CPU intensive. Business requirements dictate that the cluster must be able to service read and write queries at at) times. A solutions architect must devise a solution that accommodates the bursts of usage. Which solution meets these requirements MOST cost-effectively?

- A. Provision an Amazon EMR duster Offload the complex data processing tasks
- B. Deploy an AWS Lambda function to add capacity to the Amazon Redshift cluster by using a classic resize operation when the duster's CPU metrics in Amazon CloudWatch reach 80%.
- C. Deploy an AWS Lambda function to add capacity to the Amazon Redshift duster by using an elastic resize operation when the duster's CPU metrics in Amazon CloudWatch leach 80%.
- D. Turn on the Concurrency Scaling feature for the Amazon Redshift duster

Answer: D

QUESTION 902

A solutions architect needs to provide AWS Cost and Usage Report data from a company's AWS Organizations management account.

The company already has an Amazon S3 bucket to store the reports. The reports must be automatically ingested into a database that can be visualized with other toots.

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

- A. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that a new object creation in the S3 bucket will trigger
- B. Create an AWS Cost and Usage Report configuration to deliver the data into the S3 bucket
- C. Configure an AWS Glue crawler that a new object creation in the S3 bucket will trigger.
- D. Create an AWS Lambda function that a new object creation in the S3 bucket will trigger
- E. Create an AWS Glue crawler that me AWS Lambda function will trigger to crawl objects in me S3 bucket
- F. Create an AWS Glue crawler that the Amazon EventBridge (Amazon CloudWatCh Events) rule will trigger to crawl objects m the S3 bucket

Answer: BDF

QUESTION 903

A company used Amazon EC2 instances to deploy a web fleet to host a blog site. The EC2 instances are behind an Application Load Balancer (ALB) and are configured in an Auto ScaSng group.

The web application stores all blog content on an Amazon EFS volume. The company recently added a feature 'or Moggers to add video to their posts, attracting 10 times the previous user traffic. At peak times of day, users report buffering and timeout issues while attempting to reach the site or watch videos.

Which is the MOST cost-efficient and scalable deployment that win resolve the issues for users?

- A. Reconfigure Amazon EFS to enable maxmum I/O.
- B. Update the Nog site to use instance store volumes tor storage. Copy the site contents to the volumes at launch and to Amazon S3 al shutdown.
- C. Configure an Amazon CloudFront distribution. Point the distribution to an S3 bucket, and migrate the videos from EFS to Amazon S3.
- D. Set up an Amazon CloudFront distribution for all site contents, and point the distribution at the ALB.

Answer: C

QUESTION 904

A company is creating a sequel for a popular online game.

A large number of users from all over the world will play the game within the first week after launch.

Currently, the game consists of the following components deployed in a single AWS Region:

- Amazon S3 bucket that stores game assets

- Amazon DynamoDB table that stores player scores

A solutions architect needs to design a multi-Region solution that will reduce latency improve reliability, and require the least effort to implement.

What should the solutions architect do to meet these requirements?

- A. Create an Amazon CloudFront distribution to serve assets from the S3 bucket.
Configure S3 Cross-Region Replication.
Create a new DynamoDB table in a new Region Use the new table as a replica target for DynamoDB global tables.
- B. Create an Amazon CloudFront distribution to serve assets from the S3 bucket.
Configure S3 Same-Region Replication.
Create a new DynamoDB table in a new Region.
Configure asynchronous replication between the DynamoDB tables by using AWS Database Migration Service (AWS DMS) with change data capture (CDC)
- C. Create another S3 bucket in a new Region and configure S3 Cross-Region Replication between the buckets.
Create an Amazon CloudFront distribution and configure origin failover with two origins accessing the S3 buckets in each Region.
Configure DynamoDB global tables by enabling Amazon DynamoDB Streams, and add a replica table in a new Region.
- D. Create another S3 bucket in the same Region, and configure S3 Same-Region Replication between the buckets.
Create an Amazon CloudFront distribution and configure origin failover with two origin accessing the S3 buckets.
Create a new DynamoDB table in a new Region Use the new table as a replica target for DynamoDB global tables.

Answer: C

QUESTION 905

A company is creating a sequel for a popular online game.

A large number of users from all over the world will play the game within the first week after launch.

Currently, the game consists of the following components deployed in a single AWS Region:

- Amazon S3 bucket that stores game assets
- Amazon DynamoDB table that stores player scores

A solutions architect needs to design a Region solution that will reduce latency improve reliability, and require the least effort to implement

What should the solutions architect do to meet these requirements?

- A. Create an Amazon CloudFront distribution to serve assets from the S3 bucket.
Configure S3 Cross-Region Replication.
Create a new DynamoDB table in a new Region.
Use the new table as a replica target for DynamoDB global tables.
- B. Create an Amazon CloudFront distribution to serve assets from the S3 bucket.
Configure S3 Same-Region Replication.
Create a new DynamoDB table in a new Region.
Configure asynchronous replication between the DynamoDB tables by using AWS Database Migration Service (AWS DMS) with change data capture (CDC)
- C. Create another S3 bucket in a new Region and configure S3 Cross-Region Replication between the buckets.
Create an Amazon CloudFront distribution and configure origin failover with two origins accessing the S3 buckets in each Region.
Configure DynamoDB global tables by enabling Amazon DynamoDB Streams, and add a replica table in a new Region.
- D. Create another S3 bucket in the same Region, and configure S3 Same-Region Replication between the buckets.
Create an Amazon CloudFront distribution and configure origin failover with two origin accessing

the S3 buckets.

Create a new DynamoDB table in a new Region. Use the new table as a replica target for DynamoDB global tables.

Answer: B

QUESTION 906

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

Answer: A

QUESTION 907

A company's solution architect is designing a disaster recovery (DR) solution for an application that runs on AWS. The application uses PostgreSQL 11.7 as its database. The company has an RPO of 30 seconds. The solution architect must design a DR solution with the primary database in the us-east-1 Region and the secondary database in the us-west-2 Region.

What should the solution architect do to meet these requirements with minimum application change?

- A. Migrate the database to Amazon RDS for PostgreSQL in us-east-1. Set up a read replica in us-west-2. Set the managed RPO for the RDS database to 30 seconds.
- B. Migrate the database to Amazon Aurora PostgreSQL in us-east-1. Set up a standby replica in an Availability Zone in us-west-2. Set the managed RPO for the RDS database to 30 seconds.
- C. Migrate the database to an Amazon Aurora PostgreSQL global database with the primary Region as us-east-1 and the secondary Region as us-west-2. Set the managed RPO for the Aurora database to 30 seconds.
- D. Migrate the database to Amazon DynamoDB in us-east-1. Set up global tables with replica tables that are created in us-west-2.

Answer: A

QUESTION 908

A company's solution architect is designing a disaster recovery (DR) solution for an application that runs on AWS. The application uses PostgreSQL 11.7 as its database. The company has an RPO of 30 seconds. The solution architect

[SAP-C01 Exam Dumps](#) [SAP-C01 Exam Questions](#) [SAP-C01 PDF Dumps](#) [SAP-C01 VCE Dumps](#)

<https://www.braindump2go.com/aws-certified-solutions-architect-professional.html>

must design a DR solution with the primary database in the us-east-1 Region and the database in the us-west-2 Region.

What should the solution architect do to meet these requirements with minimum application change?

- A. Migrate the database to Amazon RDS for PostgreSQL in us-east-1.
Set up a read replica up a read replica in us-west-2.
Set the managed PRO for the RDS database to 30 seconds.
- B. Migrate the database to Amazon for PostgreSQL in us-east-1.
Set up a standby replica in an Availability Zone in us-west-2.
Set the managed PRO for the RDS database to 30 seconds.
- C. Migrate the database to an Amazon Aurora PostgreSQL global database with the primary Region as us-east-1 and the secondary Region as us-west-2.
Set the managed PRO for the Aurora database to 30 seconds.
- D. Migrate the database to Amazon DynamoDB in us-east-1.
Set up global tables with replica tables that are created in us-west-2.

Answer: A

QUESTION 909

An online magazine will launch its latest edition this month. This edition will be the first to be distributed globally. The magazine's dynamic website currently uses an Application Load Balancer in front of the web tier a fleet of Amazon EC2 instances for web and application servers, and Amazon Aurora MySQL. Portions of the website include static content and almost all traffic is read-only. The magazine is expecting a significant spike in internet traffic when the new edition is launched. Optimal performance is a top priority for the week following the launch.

Which combination of steps should a solutions architect take to reduce system response times for a global audience? (Select TWO)

- A. Use logical cross-Region replication to replicate the Aurora MySQL database to a secondary Region.
Replace the web servers with Amazon S3 Deploy S3 buckets in cross-Region replication mode
- B. Ensure the web and application tiers are each in Auto Scaling groups.
Introduce an AWS Direct Connect connection Deploy the web and application tiers in Regions across the world
- C. Migrate the database from Amazon Aurora to Amazon RDS for MySQL.
Ensure all three of the application tiers--web, application, and database--are in private subnets.
- D. Use an Aurora global database for physical cross-Region replication.
Use Amazon S3 with cross-Region replication for static content and resources.
Deploy the web and application tiers in Regions across the world
- E. Introduce Amazon Route 53 with latency-based routing and Amazon CloudFront distributions.
Ensure the web and application tiers are each in Auto Scaling groups

Answer: DE

QUESTION 910

A company is planning to migrate an application from on premises to the AWS Cloud. The company will begin the migration by moving the application's underlying data storage to AWS. The application data is stored on a shared file system on premises, and the application servers connect to the shared file system through SMB.

A solutions architect must implement a solution that uses an Amazon S3 bucket for shared storage. Until the application is fully migrated and code is rewritten to use native Amazon S3 APIs, the application must continue to have access to the data through SMB. The solutions architect must migrate the application data to AWS to its new location while still allowing the on-premises application to access the data.

Which solution will meet these requirements?

- A. Create a new Amazon FSx for Windows File Server file system.
Configure AWS DataSync with one location for the on-premises file share and one location for the new Amazon FSx file system.
Create a new DataSync task to copy the data from the on-premises file share location to the

[SAP-C01 Exam Dumps](#) [SAP-C01 Exam Questions](#) [SAP-C01 PDF Dumps](#) [SAP-C01 VCE Dumps](#)

<https://www.braindump2go.com/aws-certified-solutions-architect-professional.html>

Amazon FSx file system

- B. Create an S3 bucket for the application.
Copy the data from the on-premises storage to the S3 bucket
- C. Deploy an AWS Server Migration Service (AWS SMS) VM to the on-premises environment.
Use AWS SMS to migrate the file storage server from on premises to an Amazon EC2 instance
- D. Create an S3 bucket for the application.
Deploy a new AWS Storage Gateway Me gateway on an on-premises VM.
Create a new file share that stores data in the S3 bucket and is associated with the tie gateway.
Copy the data from the on-premises storage to the new file gateway endpoint.

Answer: A

QUESTION 911

A company is planning to migrate an application from on premises to AWS. The application currently uses an Oracle database and the company can tolerate a brief downtime of 1 hour when performing the switch to the new infrastructure As part of the migration. the database engine will be changed to MySQL.

A solutions architect needs to determine which AWS services can be used to perform the migration while minimizing the amount of work and time required.

Which of the following will meet the requirements?

- A. Use AWS SCT to generate the schema scripts and apply them on the target prior to migration.
Use AWS DMS to analyse the current schema and provide a recommendation for the optimal database engine.
Then, use AWS DMS to migrate to the recommended engine.
Use AWS SCT to identify what embedded SQL code in the application can be converted and what has to be done manually
- B. Use AWS SCT to generate the schema scripts and apply them on the target prior to migration.
Use AWS DMS to begin moving data from the on-premises database to AWS.
After the initial copy continue to use AWS DMS to keep the databases in sync until cutting over to the new database.
Use AWS SCT to identify what embedded SQL code in the application can be converted and what has to be done manually.
- C. Use AWS DMS lo help identify the best target deployment between installing the database engine on Amazon EC2 directly or moving to Amazon RDS.
Then, use AWS DMS to migrate to the platform.
Use AWS Application Discovery Service to identify what embedded SQL code in the application can be converted and what has to be done manually.
- D. Use AWS DMS to begin moving data from the on-premises database to AWS.
After the initial copy, continue to use AWS DMS to keep the databases in sync until cutting over to the new database use AWS Application Discovery Service to identify what embedded SQL code in the application can be converted and what has to be done manually

Answer: B

QUESTION 912

A company is collecting a large amount of data from a fleet of IoT devices. Data is stored as Optimized Row Columnar (ORC) files in the Hadoop Distributed File System (HDFS) on a persistent Amazon EMR cluster. The company's data analytics team queries the data by using SQL in Apache Presto deployed on the same EMR cluster Queries scan large amounts of data always run for less than 15 minutes, and run only between 5 PM and 10 PM. The company is concerned about the high cost associated with the current solution. A solutions architect must propose the most cost-effective solution that will allow SQL data queries.

Which solution will meet these requirements?

- A. Store data m Amazon S3. Use Amazon Redshift Spectrum to query data.
- B. Store data m Amazon S3. Use the AWS Glue Data Catalog and Amazon Athena to query data.
- C. Store data in EMR File System (EMRFS). Use Presto n Amazon EMR to query data.

D. Store data in Amazon Redshift. Use Amazon Redshift to query data

Answer: B

QUESTION 913

A company has deployed an application to multiple environments in AWS. Including production and testing the company has separate accounts for production and testing, and users are allowed to create additional application users for team members or services, as needed. The security team has asked the operations team for better isolation between production and testing with centralized controls on security credentials and improved management of permissions between environments. Which of the following options would MOST securely accomplish this goal?

- A. Create a new AWS account to hold user and service accounts, such as an identity account.
Create users and groups in the identity account.
Create roles with appropriate permissions in the production and testing accounts.
Add the identity account to the trust policies for the roles
- B. Modify permissions in the production and testing accounts to limit creating new IAM users to members of the operations team.
Set a strong IAM password policy on each account.
Create new IAM users and groups in each account to Limit developer access to just the services required to complete their job function.
- C. Create a script that runs on each account that checks user accounts.
For adherence to a security policy.
Disable any user or service accounts that do not comply.
- D. Create all user accounts in the production account.
Create roles for access in the production account and testing accounts.
Grant cross-account access from the production account to the testing account

Answer: A

QUESTION 914

A solutions architect needs to implement a client-side encryption mechanism for objects that will be stored in a new Amazon S3 bucket. The solutions architect created a CMK that is stored in AWS Key Management Service (AWS KMS) for this purpose. The solutions architect created the following IAM policy and attached it to an IAM role:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "GetObjectUpload",
      "Action": [
        "s3:GetObject",
        "s3:PutObject",
        "s3:DeleteObject",
        "s3:ListBucket",
        "s3:PutObjectAcl"
      ],
      "Effect": "Allow",
      "Resource": "*"
    },
    {
      "Sid": "KMSAccess",
      "Action": [
        "kms:Decrypt",
        "kms:Encrypt"
      ],
      "Effect": "Allow",
      "Resource": "*"
    }
  ]
}

```

During tests, the solutions architect was able to successfully get existing test objects in the S3 bucket. However, attempts to upload a new object resulted in an error message.

The error message stated that the action was forbidden.

Which action must the solutions architect add to the IAM policy to meet all the requirements?

- A. kms:GenerateDataKey
- B. kms:GetKeyPolicy
- C. kms:GetPublicKey
- D. kms:SKjn

Answer: A

QUESTION 915

A new application is running on Amazon Elastic Container Service (Amazon ECS) with AWS Fargate. The application uses an Amazon Aurora MySQL database. The application and the database run in the same subnets of a VPC with distinct security groups that are configured. The password of the database is stored in AWS Secrets Manager and is passed to the application through the D8_PASSWORD environment variable. The hostname of the database is passed to the application through the DB_HOST environment variable. The application is failing to access the database. Which combination of actions should a solutions architect take to resolve this error? (Select THREE)

- A. Ensure that the container has the environment variable with name "DB_PASSWORD" specified with a "ValueFrom" and the ARN of the secret
- B. Ensure that the container has the environment variable with name "D8_PASSWORD" specified with a "ValueFrom" and the secret name of the secret.
- C. Ensure that the Fargate service security group allows inbound network traffic from the Aurora MySQL database on the MySQL TCP port 3306.
- D. Ensure that the Aurora MySQL database security group allows inbound network traffic from the Fargate service on the MySQL TCP port 3306.
- E. Ensure that the container has the environment variable with name "D8_HOST" specified with the hostname of a DB instance endpoint.
- F. Ensure that the container has the environment variable with name "DB_HOST" specified with the hostname of the OB duster endpoint.

Answer: ADE

QUESTION 916

A company's site reliability engineer is performing a review of Amazon FSx for Windows File Server deployments within an account that the company acquired. Company policy states that all Amazon FSx file systems must be configured to be highly available across Availability Zones. During the review, the site reliability engineer discovers that one of the Amazon FSx file systems used a deployment type of Single-AZ 2. A solutions architect needs to minimize downtime while aligning this Amazon FSx file system with company policy. What should the solutions architect do to meet these requirements?

- A. Reconfigure the deployment type to Multi-AZ for this Amazon FSx file system
- B. Create a new Amazon FSx file system with a deployment type of Multi-AZ. Use AWS DataSync to transfer data to the new Amazon FSx file system. Point users to the new location
- C. Create a second Amazon FSx file system with a deployment type of Single-AZ 2. Use AWS DataSync to keep the data in sync. Switch users to the second Amazon FSx file system in the event of failure
- D. Use the AWS Management Console to take a backup of the Amazon FSx file system. Create a new Amazon FSx file system with a deployment type of Multi-AZ. Restore the backup to the new Amazon FSx file system. Point users to the new location.

Answer: B

QUESTION 917

A company is using multiple AWS accounts. The DNS records are stored in a private hosted zone for Amazon Route 53 in Account A. The company's applications and databases are running in Account B. A solutions architect will deploy a two-net application in a new VPC. To simplify the configuration, the db.example.com CNAME record set for the Amazon RDS endpoint was created in a private hosted zone for Amazon Route 53. During deployment, the application failed to start. Troubleshooting revealed that db.example.com is not resolvable on the Amazon EC2 instance. The solutions architect confirmed that the record set was created correctly in Route 53. Which combination of steps should the solutions architect take to resolve this issue? (Select TWO)

- A. Deploy the database on a separate EC2 instance in the new VPC. Create a record set for the instance's private IP in the private hosted zone
- B. Use SSH to connect to the application tier EC2 instance. Add an RDS endpoint IP address to the /etc/resolv.conf file
- C. Create an authorization lo associate the private hosted zone in Account A with the new VPC in Account B.
- D. Create a private hosted zone for the example.com domain in Account B. Configure Route 53 replication between AWS accounts
- E. Associate a new VPC in Account B with a hosted zone in Account A. Delete the association authorization In Account A.

Answer: BC