

➤ **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/2020](#))**

Visit Braindump2go and Download Full Version AWS-Certified-Solutions-Architect-Professional (SAP-C01) Exam Dumps

QUESTION 654

A manufacturing company is growing exponentially and has secured funding to improve its IT infrastructure and ecommerce presence. The company's ecommerce platform consists of:

- Static assets primarily comprised of product images stored in Amazon S3.
- Amazon DynamoDB tables that store product information, user information, and order information.
- Web servers containing the application's front-end behind Elastic Load Balancers.

The company wants to set up a disaster recovery site in a separate Region.

Which combination of actions should the solutions architect take to implement the new design while meeting all the requirements? (Choose three.)

- A. Enable Amazon Route 53 health checks to determine if the primary site is down, and route traffic to the disaster recovery site if there is an issue.
- B. Enable Amazon S3 cross-Region replication on the buckets that contain static assets.
- C. Enable multi-Region targets on the Elastic Load Balancer and target Amazon EC2 instances in both Regions.
- D. Enable DynamoDB global tables to achieve a multi-Region table replication.
- E. Enable Amazon CloudWatch and create CloudWatch alarms that route traffic to the disaster recovery site when application latency exceeds the desired threshold.
- F. Enable Amazon S3 versioning on the source and destination buckets containing static assets to ensure there is a rollback version available in the event of data corruption.

Answer: CEF

QUESTION 655

A company is developing a gene reporting device that will collect genomic information to assist researchers will collecting large samples of data from a diverse population. The device will push 8 KB of genomic data every second to a data platform that will need to process and analyze the data and provide information back to researchers. The data platform must meet the following requirements:

- Provide near-real-time analytics of the inbound genomic data
- Ensure the data is flexible, parallel, and durable
- Deliver results of processing to a data warehouse

Which strategy should a solutions architect use to meet these requirements?

[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>

Time!

- A. Use Amazon Kinesis Data Firehouse to collect the inbound sensor data, analyze the data with Kinesis clients, and save the results to an Amazon RDS instance.
- B. Use Amazon Kinesis Data Streams to collect the inbound sensor data, analyze the data with Kinesis clients, and save the results to an Amazon Redshift cluster using Amazon EMR.
- C. Use Amazon S3 to collect the inbound device data, analyze the data from Amazon SQS with Kinesis, and save the results to an Amazon Redshift cluster.
- D. Use an Amazon API Gateway to put requests into an Amazon SQS queue, analyze the data with an AWS Lambda function, and save the results to an Amazon Redshift cluster using Amazon EMR.

Answer: B

QUESTION 656

A company needs to move its on-premises resources to AWS. The current environment consists of 100 virtual machines (VMs) with a total of 40 TB of storage. Most of the VMs can be taken offline because they support functions during business hours only, however, some are mission critical, so downtime must be minimized.

The administrator of the on-premises network provisioned 10 Mbps of internet bandwidth for the migration. The on-premises network throughput has reached capacity and would be costly to increase. A solutions architect must design a migration solution that can be performed within the next 3 months.

Which method would fulfill these requirements?

- A. Set up a 1 Gbps AWS Direct Connect connection. Then, provision a private virtual interface, and use AWS Server Migration Service (SMS) to migrate the VMs into Amazon EC2.
- B. Use AWS Application Discovery Service to assess each application, and determine how to refactor and optimize each using AWS services or AWS Marketplace solutions.
- C. Export the VMs locally, beginning with the most mission-critical servers first. Use AWS Transfer for SFTP to securely upload each VM to Amazon S3 after they are exported. Use VM Import/Export to import the VMs into Amazon EC2.
- D. Migrate mission-critical VMs with AWS SMS. Export the other VMs locally and transfer them to Amazon S3 using AWS Snowball. Use VM Import/Export to import the VMs into Amazon EC2.

Answer: B

QUESTION 657

A company runs a popular public-facing ecommerce website. Its user base is growing quickly from a local market to a national market. The website is hosted in an on-premises data center with web servers and a MySQL database. The company wants to migrate its workload to AWS. A solutions architect needs to create a solution to:

- Improve security
- Improve reliability
- Improve availability
- Reduce latency
- Reduce maintenance

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

- A. Use Amazon EC2 instances in two Availability Zones for the web servers in an Auto Scaling group behind an Application Load Balancer.
- B. Migrate the database to a Multi-AZ Amazon Aurora MySQL DB cluster.
- C. Use Amazon EC2 instances in two Availability Zones to host a highly available MySQL database cluster.

[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>

Time!

- D. Host static website content in Amazon S3. Use S3 Transfer Acceleration to reduce latency while serving webpages. Use AWS WAF to improve website security.
- E. Host static website content in Amazon S3. Use Amazon CloudFront to reduce latency while serving webpages. Use AWS WAF to improve website security.
- F. Migrate the database to a single-AZ Amazon RDS for MySQL DB instance.

Answer: DEF

QUESTION 658

A company has an internal application running on AWS that is used to track and process shipments in the company's warehouse. Currently, after the system receives an order, it emails the staff the information needed to ship a package. Once the package is shipped, the staff replies to the email and the order is marked as shipped.

The company wants to stop using email in the application and move to a serverless application model.

Which architecture solution meets these requirements?

- A. Use AWS Batch to configure the different tasks required to ship a package.
Have AWS Batch trigger an AWS Lambda function that creates and prints a shipping label.
Once that label is scanned, as it leaves the warehouse, have another Lambda function move the process to the next step in the AWS Batch job.
- B. When a new order is created, store the order information in Amazon SQS.
Have AWS Lambda check the queue every 5 minutes and process any needed work.
When an order needs to be shipped, have Lambda print the label in the warehouse.
Once the label has been scanned, as it leaves the warehouse, have an Amazon EC2 instance update Amazon SQS.
- C. Update the application to store new order information in Amazon DynamoDB.
When a new order is created, trigger an AWS Step Functions workflow, mark the orders as "in progress", and print a package label to the warehouse.
Once the label has been scanned and fulfilled, the application will trigger an AWS Lambda function that will mark the order as shipped and complete the workflow.
- D. Store new order information in Amazon EFS.
Have instances pull the new information from the NFS and send that information to printers in the warehouse.
Once the label has been scanned, as it leaves the warehouse, have Amazon API Gateway call the instances to remove the order information from Amazon EFS.

Answer: A

QUESTION 659

A company has developed a mobile game. The backend for the game runs on several virtual machines located in an on-premises data center. The business logic is exposed using a REST API with multiple functions. Player session data is stored in central file storage. Backend services use different API keys for throttling and to distinguish between live and test traffic.

The load on the game backend varies throughout the day. During peak hours, the server capacity is not sufficient. There are also latency issues when fetching player session data. Management has asked a solutions architect to present a cloud architecture that can handle the game's varying load and provide low-latency data access. The API model should not be changed.

Which solution meets these requirements?

- A. Implement the REST API using a Network Load Balancer (NLB).
Run the business logic on an Amazon EC2 instance behind the NLB.
Store player session data in Amazon Aurora Serverless.

[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>

Time!

- B. Implement the REST API using an Application Load Balancer (ALB).
Run the business logic in AWS Lambda.
Store player session data in Amazon DynamoDB with on-demand capacity.
- C. Implement the REST API using Amazon API Gateway.
Run the business logic in AWS Lambda.
Store player session data in Amazon DynamoDB with on-demand capacity.
- D. Implement the REST API using AWS AppSync.
Run the business logic in AWS Lambda.
Store player session data in Amazon Aurora Serverless.

Answer: A

QUESTION 660

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 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 AWSPrivateMarketplaceAdminFullAccess managed policy.
- B. 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 AWSPrivateMarketplaceAdminFullAccess managed policy and attach it to all the developer roles.
- C. Create an IAM role named procurement-manager-role in all the shared services accounts in the organization.
Add the AWSPrivateMarketplaceAdminFullAccess 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.
- D. Create an IAM role named procurement-manager-role in all AWS accounts that will be used by developers.
Add the AWSPrivateMarketplaceAdminFullAccess managed policy to the role.
Create an SCP in 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 organization.

Answer: A

QUESTION 661

A solutions architect is designing the data storage and retrieval architecture for a new application that a company will be launching soon. The application is designed to ingest millions of small records per minute from devices all around the world. Each record is less than 4 KB in size and needs to be stored in a durable location where it can be retrieved with low latency. The data is ephemeral and the company is required to store the data for 120 days only, after which the

data can be deleted.

The solutions architect calculates that, during the course of year, the storage requirements would be about 10-15 TB.

Which storage strategy is the MOST cost-effective and meets the design requirements?

- A. Design the application to store each incoming record as a single .csv file in an Amazon S3 bucket to allow for indexed retrieval.
Configure a lifecycle policy to delete data older than 120 days.
- B. Design the application to store each incoming record in an Amazon DynamoDB table property configured for the scale.
Configure the DynamoDB Time to Live (TTL) feature to delete records older than 120 days.
- C. Design the application to store each incoming record in a single table in an Amazon RDS MySQL database.
Run a nightly cron job that executes a query to delete any records older than 120 days.
- D. Design the application to batch incoming records before writing them to an Amazon S3 bucket.
Update the metadata for the object to contain the list of records in the batch and use the Amazon S3 metadata search feature to retrieve the data.
Configure a lifecycle policy to delete the data after 120 days.

Answer: C

QUESTION 662

A company provides auction services for artwork and has users across North America and Europe. The company hosts its application in Amazon EC2 instances in the us-east-1 Region. Artists upload photos of their work as large-size, high-resolution image files from their mobile phones to a centralized Amazon S3 bucket created in the us-east-1 Region. The users in Europe are reporting slow performance for their image uploads.

How can a solutions architect improve the performance of the image upload process?

- A. Redeploy the application to use S3 multipart uploads.
- B. Create an Amazon CloudFront distribution and point to the application as a custom origin.
- C. Configure the buckets to use S3 Transfer Acceleration.
- D. Create an Auto Scaling group for the EC2 instances and create a scaling policy.

Answer: A

QUESTION 663

A company has developed a new release of a popular video game and wants to make it available for public download. The new release package is approximately 5 GB in size. The company provides downloads for existing releases from a Linux-based, publicly facing FTP site hosted in an on-premises data center. The company expects the new release will be downloaded by users worldwide. The company wants a solution that provides improved download performance and low transfer costs, regardless of a user's location.

Which solutions will meet these requirements?

- A. Store the game files on Amazon EBS volumes mounted on Amazon EC2 instances within an Auto Scaling group.
Configure an FTP service on the EC2 instances.
Use an Application Load Balancer in front of the Auto Scaling group.
Publish the game download URL for users to download the package.
- B. Store the game files on Amazon EFS volumes that are attached to Amazon EC2 instances within an Auto Scaling group.
Configure an FTP service on the EC2 instances.
Use an Application Load Balancer in front of the Auto Scaling group.

[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>

Time!

- Publish the game download URL for users to download the package.
- C. Configure Amazon Route 53 and an Amazon S3 bucket for website hosting.
Upload the game files to the S3 bucket.
Use Amazon CloudFront for the website.
Publish the game download URL for users to download the package.
- D. Configure Amazon Route 53 and an Amazon S3 bucket for website hosting.
Upload the game files to the S3 bucket.
Set Requester Pays for the S3 bucket.
Publish the game download URL for users to download the package.

Answer: B

QUESTION 664

A new startup is running a serverless application using AWS Lambda as the primary source of compute. New versions of the application must be made available to a subset of users before deploying changes to all users. Developers should also have the ability to abort the deployment and have access to an easy rollback mechanism. A solutions architect decides to use AWS CodeDeploy changes when a new version is available.

Which CodeDeploy configuration should the solutions architect use?

- A. A blue/green deployment
- B. A linear deployment
- C. A canary deployment
- D. An all-at-once deployment

Answer: D

QUESTION 665

A solutions architect is implementing federated access to AWS for users of the company's mobile application. Due to regulatory and security requirements, the application must use a custom-built solution for authenticating users and must use IAM roles for authorization.

Which of the following actions would enable authentication and authorization and satisfy the requirements? (Choose two.)

- A. Use a custom-built SAML-compatible solution for authentication and AWS SSO for authorization.
- B. Create a custom-built LDAP connector using Amazon API Gateway and AWS Lambda for authentication.
Store authorization tokens in Amazon DynamoDB, and validate authorization requests using another Lambda function that reads the credentials from DynamoDB.
- C. Use a custom-built OpenID Connect-compatible solution with AWS SSO for authentication and authorization.
- D. Use a custom-built SAML-compatible solution that uses LDAP for authentication and uses a SAML assertion to perform authorization to the IAM identity provider.
- E. Use a custom-built OpenID Connect-compatible solution for authentication and use Amazon Cognito for authorization.

Answer: BC

QUESTION 666

A company has developed a custom tool used in its workflow that runs within a Docker container. The company must perform manual steps each time the container code is updated to make the container image available to new workflow executions. The company wants to automate this process to eliminate manual effort and ensure a new container image

is generated every time the tool code is updated.

Which combination of actions should a solutions architect take to meet these requirements? (Choose three.)

- A. Configure an Amazon ECR repository for the tool. Configure an AWS CodeCommit repository containing code for the tool being deployed to the container image in Amazon ECR.
- B. Configure an AWS CodeDeploy application that triggers an application version update that pulls the latest tool container image from Amazon ECR, updates the container with code from the source AWS CodeCommit repository, and pushes the updated container image to Amazon ECR.
- C. Configuration an AWS CodeBuild project that pulls the latest tool container image from Amazon ECR, updates the container with code from the source AWS CodeCommit repository, and pushes the updated container image to Amazon ECR.
- D. Configure an AWS CodePipeline pipeline that sources the tool code from the AWS CodeCommit repository and initiates an AWS CodeDeploy application update.
- E. Configure an Amazon EventBridge rule that triggers on commits to the AWS CodeCommit repository for the tool. Configure the event to trigger an update to the tool container image in Amazon ECR. Push the updated container image to Amazon ECR.
- F. Configure an AWS CodePipeline pipeline that sources the tool code from the AWS CodeCommit repository and initiates an AWS CodeBuild build.

Answer: ACD

QUESTION 667

A company hosts an application on Amazon EC2 instance and needs to store files in Amazon S3. The files should never traverse the public internet, and only the application EC2 instances are granted access to a specific Amazon S3 bucket. A solutions architect has created a VPC endpoint for Amazon S3 and connected the endpoint to the application VPC.

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

- A. Assign an endpoint policy to the endpoint that restricts access to a specific S3 bucket.
Attach a bucket policy to the S3 bucket that grants access to the VPC endpoint.
Add the gateway prefix list to a NACL of the instances to limit access to the application EC2 instances only.
- B. Attach a bucket policy to the S3 bucket that grants access to application EC2 instances only using the aws:SourceIp condition.
Update the VPC route table so only the application EC2 instances can access the VPC endpoint.
- C. Assign an endpoint policy to the VPC endpoint that restricts access to a specific S3 bucket.
Attach a bucket policy to the S3 bucket that grants access to the VPC endpoint.
Assign an IAM role to the application EC2 instances and only allow access to this role in the S3 bucket's policy.
- D. Assign an endpoint policy to the VPC endpoint that restricts access to S3 in the current Region.
Attach a bucket policy to the S3 bucket that grants access to the VPC private subnets only.
Add the gateway prefix list to a NACL to limit access to the application EC2 instances only.

Answer: C

QUESTION 668

A financial services company has an on-premises environment that ingests market data feeds from stock exchanges, transforms the data, and sends the data to an internal Apache Kafka cluster. Management wants to leverage AWS services to build a scalable and near-real-time solution with consistent network performance to provide stock market data to a web application.

Which steps should a solutions architect take to build the solution? (Choose three.)

[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. Establish an AWS Direct Connect connection from the on-premises data center to AWS.
- B. Create an Amazon EC2 Auto Scaling group to pull the messages from the on-premises Kafka cluster and use the Amazon Consumer Library to put the data into an Amazon Kinesis data stream.
- C. Create an Amazon EC2 Auto Scaling group to pull the messages from the on-premises Kafka cluster and use the Amazon Kinesis Producer Library to put the data into a Kinesis data stream.
- D. Create a WebSocket API in Amazon API Gateway, create an AWS Lambda function to process an Amazon Kinesis data stream, and use the @connections command to send callback messages to connected clients.
- E. Create a GraphQL API in AWS AppSync, create an AWS Lambda function to process the Amazon Kinesis data stream, and use the @connections command to send callback messages to connected clients.
- F. Establish a Site-to-Site VPN from the on-premises data center to AWS.

Answer: BCD