

FREE



Amazon-Web-Services

SAP-C02 Dumps

AWS Certified Solutions Architect - Professional Dumps

20%
OFF
All
AWS
Exams



A Comprehensive Guide to the AWS Certified Solutions Architect - Professional (SAP-C02) Exam

If you're aspiring to advance your career in cloud computing and demonstrate your expertise in designing distributed systems on Amazon Web Services (AWS), the AWS Certified Solutions Architect - Professional (SAP-C02) exam is an essential milestone. In this article, we'll delve into the details of the SAP-C02 exam, [SAP-C02 Dumps PDF](#), its key Features, and address two important FAQs. We'll also discuss a valuable money-back guarantee service if you don't pass the exam.

SAP-C02 Exam Overview

The AWS Certified Solutions Architect - Professional (SAP-C02) exam is a certification that validates your ability to design advanced solutions using AWS technologies. Here are some essential details about this certification:

Exam Format:

Exam Duration: 180 minutes (3 hours)

Number of Questions: 75 multiple-choice and multiple-response questions

Passing Score: 750 out of 1000

Exam Cost: \$300 USD

Exam Delivery: In-person at a testing center or online via Pearson VUE

Exam Domains:

The SAP-C02 exam covers various domains, including:

Design for Organizational Complexity: Focuses on designing for large organizations with complex requirements.

Design for New Solutions: Addresses designing greenfield projects on AWS.

Migration Planning: Covers strategies for migrating existing workloads to AWS.

Cost Control: Examines optimizing AWS costs.

Continuous Improvement for Existing Solutions: Focuses on enhancing existing AWS solutions.

Define Operationally Excellent Architectures: Covers operational best practices.

Security: Addresses security best practices in AWS.

Key Features of SAP-C02 Exam

Scenario-Based Questions: The SAP-C02 exam includes scenario-based questions that test your ability to apply AWS best practices to real-world situations.

Hands-on Experience: Practical experience with AWS services is crucial. You'll need to understand how to implement solutions effectively.

Updated Content: AWS regularly updates its certification exams to align with the latest AWS services and best practices. Stay updated with AWS documentation and training resources.

Time Management: Time is crucial in the SAP-C02 exam. Manage your time efficiently to ensure you can answer all questions.

Official AWS Training: Consider enrolling in AWS training courses, which cover exam-specific content and provide hands-on labs.

FAQs

FAQ 1: What Is the Best Way to Prepare for the SAP-C02 Exam?

The best way to prepare for the SAP-C02 exam is to follow these steps:

Study Material: Utilize AWS documentation, whitepapers, and official AWS training resources.

Hands-On Experience: Practice using AWS services by building real solutions.

Online Courses: Enroll in online courses offered by AWS or accredited training providers.

Practice Exams: Take practice exams to familiarize yourself with the exam format.

FAQ 2: What Happens If I Don't Pass the SAP-C02 Exam?

While failure is always a possibility, it shouldn't deter you from pursuing your AWS certification. If you don't pass the SAP-C02 exam on your first attempt, you have the option to retake it. Keep in mind that AWS has a retake policy and a waiting period between attempts.

Money-Back Guarantee Service

Studying for and passing the SAP-C02 exam requires dedication and effort. If you're looking for added assurance, consider utilizing a money-back guarantee service offered by some AWS training providers. These services often provide a refund of your exam fees if you don't pass the SAP-C02 exam on your first try, provided you meet their terms and conditions.

Conclusion

In conclusion, the AWS Certified Solutions Architect - Professional (SAP-C02) exam is a challenging but rewarding certification that can significantly boost your career in cloud computing. To succeed, invest in [SAP-C02 Study Material](#), gain hands-on experience, and consider utilizing money-back guarantee services for added peace of mind.

NEW QUESTION 1

A startup company recently migrated a large ecommerce website to AWS. The website has experienced a 70% increase in sales. Software engineers are using a private GitHub repository to manage code. The DevOps team is using Jenkins for builds and unit testing. The engineers need to receive notifications for bad builds and zero downtime during deployments. The engineers also need to ensure any changes to production are seamless for users and can be rolled back in the event of a major issue.

The software engineers have decided to use AWS CodePipeline to manage their build and deployment process.

Which solution will meet these requirements?

- A. Use GitHub websockets to trigger the CodePipeline pipeline
- B. Use the Jenkins plugin for AWS CodeBuild to conduct unit testing
- C. Send alerts to an Amazon SNS topic for any bad build
- D. Deploy in an in-place
- E. all-at-once deployment configuration using AWS CodeDeploy.
- F. Use GitHub webhooks to trigger the CodePipeline pipeline
- G. Use the Jenkins plugin for AWS CodeBuild to conduct unit testing
- H. Send alerts to an Amazon SNS topic for any bad build
- I. Deploy in a blue/green deployment using AWS CodeDeploy.
- J. Use GitHub websockets to trigger the CodePipeline pipeline
- K. Use AWS X-Ray for unit testing and static code analysis
- L. Send alerts to an Amazon SNS topic for any bad build
- M. Deploy in a blue/green deployment using AWS CodeDeploy.
- N. Use GitHub webhooks to trigger the CodePipeline pipeline
- O. Use AWS X-Ray for unit testing and static code analysis
- P. Send alerts to an Amazon SNS topic for any bad build
- Q. Deploy in an in-place, all-at-once deployment configuration using AWS CodeDeploy.

Answer: B

NEW QUESTION 2

A company is running an application distributed over several Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer. The security team requires that all application access attempts be made available for analysis. Information about the client IP address, connection type, and user agent must be included.

Which solution will meet these requirements?

- A. Enable EC2 detailed monitoring, and include network logs. Send all logs through Amazon Kinesis Data Firehose to an Amazon Elasticsearch Service (Amazon ES) cluster that the security team uses for analysis.
- B. Enable VPC Flow Logs for all EC2 instance network interfaces. Publish VPC Flow Logs to an Amazon S3 bucket. Have the security team use Amazon Athena to query and analyze the logs.
- C. Enable access logs for the Application Load Balancer, and publish the logs to an Amazon S3 bucket. Have the security team use Amazon Athena to query and analyze the logs.
- D. Enable Traffic Mirroring and specify all EC2 instance network interfaces as the source.
- E. Send all traffic information through Amazon Kinesis Data Firehose to an Amazon Elasticsearch Service (Amazon ES) cluster that the security team uses for analysis.

Answer: C

NEW QUESTION 3

A company has a complex web application that leverages Amazon CloudFront for global scalability and performance. Over time, users report that the web application is slowing down.

The company's operations team reports that the CloudFront cache hit ratio has been dropping steadily. The cache metrics report indicates that query strings on some URLs are inconsistently ordered and are specified sometimes in mixed-case letters and sometimes in lowercase letters.

Which set of actions should the solutions architect take to increase the cache hit ratio as quickly as possible?

- A. Deploy a Lambda Edge function to sort parameters by name and force them to be lowercase.
- B. Select the CloudFront viewer request trigger to invoke the function.
- C. Update the CloudFront distribution to disable caching based on query string parameters.
- D. Deploy a reverse proxy after the load balancer to post-process the emitted URLs in the application to force the URL strings to be lowercase.
- E. Update the CloudFront distribution to specify casing-insensitive query string processing.

Answer: A

NEW QUESTION 4

A company has an Amazon VPC that is divided into a public subnet and a private subnet. A web application runs in Amazon VPC, and each subnet has its own NACL. The public subnet has a CIDR of 10.0.0.0/24. An Application Load Balancer is deployed to the public subnet. The private subnet has a CIDR of 10.0.1.0/24. Amazon EC2 instances that run a web server on port 80 are launched into the private subnet.

Only network traffic that is required for the Application Load Balancer to access the web application can be allowed to travel between the public and private subnets.

What collection of rules should be written to ensure that the private subnet's NACL meets the requirement? (Select TWO.)

- A. An inbound rule for port 80 from source 0.0.0.0/0
- B. An inbound rule for port 80 from source 10.0.0.0/24
- C. An outbound rule for port 80 to destination 0.0.0.0/0
- D. An outbound rule for port 80 to destination 10.0.0.0/24
- E. An outbound rule for ports 1024 through 65535 to destination 10.0.0.0/24

Answer: BE

NEW QUESTION 5

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

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

What should the solutions architect do to meet these requirements?

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

Answer: B

NEW QUESTION 6

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
- B. 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
- C. Use the new file system as the shared storage for the duration of the job
- D. Delete the file system when the job is complete.
- E. Migrate the data from the existing shared file system to a large Amazon Elastic Block Store (Amazon EBS) volume with Multi-Attach enabled
- F. Attach the EBS volume to each of the instances by using a user data script in the Auto Scaling group launch template
- G. Use the EBS volume as the shared storage for the duration of the job
- H. Detach the EBS volume when the job is complete.
- I. Migrate the data from the existing shared file system to an Amazon S3 bucket that uses the S3 Standard storage class
- J. 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
- K. Use the new file system as the shared storage for the duration of the job
- L. Delete the file system when the job is complete.
- M. Migrate the data from the existing shared file system to an Amazon S3 bucket
- N. 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
- O. Delete the file gateway when the job is complete.

Answer: B

NEW QUESTION 7

A company is developing and hosting several projects in the AWS Cloud. The projects are developed across multiple AWS accounts under the same organization in AWS Organizations. The company requires the cost for cloud infrastructure to be allocated to the owning project. The team responsible for all of the AWS accounts has discovered that several Amazon EC2 instances are lacking the Project tag used for cost allocation.

Which actions should a solutions architect take to resolve the problem and prevent it from happening in the future? (Select THREE.)

- A. Create an AWS Config rule in each account to find resources with missing tags.
- B. Create an SCP in the organization with a deny action for ec2:RunInstances if the Project tag is missing.
- C. Use Amazon Inspector in the organization to find resources with missing tags.
- D. Create an IAM policy in each account with a deny action for ec2:RunInstances if the Project tag is missing.
- E. Create an AWS Config aggregator for the organization to collect a list of EC2 instances with the missing Project tag.
- F. Use AWS Security Hub to aggregate a list of EC2 instances with the missing Project tag.

Answer: BDE

NEW QUESTION 8

- (Exam Topic 1)

A company built an ecommerce website on AWS using a three-tier web architecture. The application is Java-based and composed of an Amazon CloudFront distribution, an Apache web server layer of Amazon EC2 instances in an Auto Scaling group, and a backend Amazon Aurora MySQL database.

Last month, during a promotional sales event, users reported errors and timeouts while adding items to their shopping carts. The operations team recovered the logs created by the web servers and reviewed Aurora DB cluster performance metrics. Some of the web servers were terminated before logs could be collected and the Aurora metrics were not sufficient for query performance analysis.

Which combination of steps must the solutions architect take to improve application performance visibility during peak traffic events? (Select THREE.)

- A. Configure the Aurora MySQL DB cluster to publish slow query and error logs to Amazon CloudWatch Logs.
- B. Implement the AWS X-Ray SDK to trace incoming HTTP requests on the EC2 instances and implement tracing of SQL queries with the X-Ray SDK for Java.
- C. Configure the Aurora MySQL DB cluster to stream slow query and error logs to Amazon Kinesis.
- D. Install and configure an Amazon CloudWatch Logs agent on the EC2 instances to send the Apache logs to CloudWatch Logs.
- E. Enable and configure AWS CloudTrail to collect and analyze application activity from Amazon EC2 and Aurora.
- F. Enable Aurora MySQL DB cluster performance benchmarking and publish the stream to AWS X-Ray.

Answer: ABD

NEW QUESTION 9

To abide by industry regulations, a solutions architect must design a solution that will store a company's critical data in multiple public AWS Regions, including in the United States, where the company's headquarters is located. The solutions architect is required to provide access to the data stored in AWS to the company's global WAN network. The security team mandates that no traffic accessing this data should traverse the public internet. How should the solutions architect design a highly available solution that meets the requirements and is cost-effective?

- A. Establish AWS Direct Connect connections from the company headquarters to all AWS Regions in use. Use the company WAN to send traffic over to the headquarters and then to the respective DX connection to access the data.
- B. Establish two AWS Direct Connect connections from the company headquarters to an AWS Region. Use the company WAN to send traffic over a DX connection to access the data.
- C. Use inter-region VPC peering to access the data in other AWS Regions.
- D. Establish two AWS Direct Connect connections from the company headquarters to an AWS Region. Use the company WAN to send traffic over a DX connection to access the data.
- E. Use an AWS transit VPC solution to access data in other AWS Regions.
- F. Establish two AWS Direct Connect connections from the company headquarters to an AWS Region. Use the company WAN to send traffic over a DX connection to access the data.
- G. Use Direct Connect Gateway to access data in other AWS Regions.

Answer: D

NEW QUESTION 10

A solutions architect is responsible for redesigning a legacy Java application to improve its availability, data durability, and scalability. Currently, the application runs on a single high-memory Amazon EC2 instance. It accepts HTTP requests from upstream clients, adds them to an in-memory queue, and responds with a 200 status. A separate application thread reads items from the queue, processes them, and persists the results to an Amazon RDS MySQL instance. The processing time for each item takes 90 seconds on average, most of which is spent waiting on external service calls, but the application is written to process multiple items in parallel.

Traffic to this service is unpredictable. During periods of high load, items may sit in the internal queue for over an hour while the application processes the backlog. In addition, the current system has issues with availability and data loss if the single application node fails.

Clients that access this service cannot be modified. They expect to receive a response to each HTTP request they send within 10 seconds before they will time out and retry the request.

Which approach would improve the availability and durability of the system while decreasing the processing latency and minimizing costs?

- A. Create an Amazon API Gateway REST API that uses Lambda proxy integration to pass requests to an AWS Lambda function.
- B. Migrate the core processing code to a Lambda function and write a wrapper class that provides a handler method that converts the proxy events to the internal application data model and invokes the processing module.
- C. Create an Amazon API Gateway REST API that uses a service proxy to put items in an Amazon SQS queue.
- D. Extract the core processing code from the existing application and update it to pull items from Amazon SQS instead of an in-memory queue.
- E. Deploy the new processing application to smaller EC2 instances within an Auto Scaling group that scales dynamically based on the approximate number of messages in the Amazon SQS queue.
- F. Modify the application to use Amazon DynamoDB instead of Amazon RDS.
- G. Configure Auto Scaling for the DynamoDB table.
- H. Deploy the application within an Auto Scaling group with a scaling policy based on CPU utilization.
- I. Back the in-memory queue with a memory-mapped file to an instance store volume and periodically write that file to Amazon S3.
- J. Update the application to use a Redis task queue instead of the in-memory queue.
- K. Build a Docker container image for the application.
- L. Create an Amazon ECS task definition that includes the application container and a separate container to host Redis.
- M. Deploy the new task definition as an ECS service using AWS Fargate, and enable Auto Scaling.

Answer: B