



AWS-SAP

AWS Solutions Architect Professional

Exam Summary – Syllabus – Questions

Table of Contents

Introduction to AWS-SAP Exam on AWS Solutions Architect Professional.....	2
AWS AWS-SAP Certification Details:.....	2
AWS AWS-SAP Exam Syllabus:.....	3
AWS-SAP Sample Questions:.....	4
Answers to AWS-SAP Exam Questions:	7

Introduction to AWS-SAP Exam on AWS Solutions Architect Professional

The AWS-SAP exam preparation guide is designed to provide candidates with necessary information about the AWS-SAP exam. It includes exam summary, sample questions, practice test, objectives and ways to interpret the exam objectives to enable candidates to assess the types of questions-answers that may be asked during the AWS Certified Solutions Architect - Professional exam.

It is recommended for all the candidates to refer the AWS-SAP objectives and sample questions provided in this preparation guide. The AWS-SAP certification is mainly targeted to the candidates who want to build their career in Architecting domain and demonstrate their expertise. We suggest you to use practice exam listed in this cert guide to get used to with exam environment and identify the knowledge areas where you need more work prior to taking the actual AWS Solutions Architect Professional exam.

AWS AWS-SAP Certification Details:

Exam Name	AWS Solutions Architect Professional (AWS-SAP)
Exam Code	AWS-SAP
Exam Price	\$300 USD
Duration	170 minutes
Number of Questions	77
Passing Score	75%
Recommended Training / Books	Advanced Architecting on AWS
Schedule Exam	AWS Certification
Sample Questions	AWS-SAP Sample Questions
Recommended Practice	AWS Certified Solutions Architect - Professional Practice Test

AWS AWS-SAP Exam Syllabus:

Section	Objectives	Weight
High Availability and Business Continuity	<ul style="list-style-type: none"> - Demonstrate ability to architect the appropriate level of availability based on stakeholder requirements - Demonstrate ability to implement DR for systems based on RPO and RTO - Determine appropriate use of multi-Availability Zones vs. multi-Region architectures - Demonstrate ability to implement self-healing capabilities 	15%
Costing	<ul style="list-style-type: none"> - Demonstrate ability to make architectural decisions that minimize and optimize infrastructure cost - Apply the appropriate AWS account and billing set-up options based on scenario - Ability to compare and contrast the cost implications of different architectures 	6%
Deployment Management	<ul style="list-style-type: none"> - Ability to manage the lifecycle of an application on AWS - Demonstrate ability to implement the right architecture for development, testing, and staging environments - Position and select most appropriate AWS deployment mechanism based on scenario 	10%
Network Design	<ul style="list-style-type: none"> - Demonstrate ability to design and implement networking features of AWS - Demonstrate ability to design and implement connectivity features of AWS 	10%
Data Storage	<ul style="list-style-type: none"> - Demonstrate ability to make architectural trade off decisions involving storage options - Demonstrate ability to make architectural trade off decisions involving database options - Demonstrate ability to implement the most appropriate data storage architecture - Determine use of synchronous versus asynchronous replication 	15%
Security	<ul style="list-style-type: none"> - Design information security management systems and compliance controls - Design security controls with the AWS shared responsibility model and global infrastructure - Design identity and access management controls - Design protection of Data at Rest controls - Design protection of Data in Flight and Network Perimeter controls 	19%

Section	Objectives	Weight
Scalability & Elasticity	<ul style="list-style-type: none"> - Demonstrate the ability to design a loosely coupled system - Demonstrate ability to implement the most appropriate front-end scaling architecture - Demonstrate ability to implement the most appropriate middle-tier scaling architecture - Demonstrate ability to implement the most appropriate data storage scaling architecture - Determine trade-offs between vertical and horizontal scaling 	15%
Cloud Migration & Hybrid Architecture	<ul style="list-style-type: none"> - Plan and execute for applications migrations - Demonstrate ability to design hybrid cloud architectures 	10%

AWS-SAP Sample Questions:

01. You have an application running on an EC2 instance which will allow users to download files from a private S3 bucket using a pre-signed URL.

Before generating the URL, the application should verify the existence of the file in S3. How should the application use AWS credentials to access the S3 bucket securely?

- a) Create an IAM user for the application with permissions that allow list access to the S3 bucket; launch the instance as the IAM user, and retrieve the IAM user's credentials from the EC2 instance user data.
- b) Create an IAM role for EC2 that allows list access to objects in the S3 bucket; launch the instance with the role, and retrieve the role's credentials from the EC2 instance metadata.
- c) Create an IAM user for the application with permissions that allow list access to the S3 bucket; the application retrieves the IAM user credentials from a temporary directory with permissions that allow read access only to the application user.
- d) Use the AWS account access keys; the application retrieves the credentials from the source code of the application.

02. An ERP application is deployed across multiple AZs in a single region. In the event of failure, the Recovery Time Objective (RTO) must be less than 3 hours, and the Recovery Point Objective (RPO) must be 15 minutes.

The customer realizes that data corruption occurred roughly 1.5 hours ago. What DR strategy could be used to achieve this RTO and RPO in the event of this kind of failure?

- a) Take 15 minute DB backups stored in Glacier with transaction logs stored in S3 every 5 minutes.
- b) Use synchronous database master-slave replication between two availability zones.
- c) Take hourly DB backups to EC2 instance store volumes with transaction logs stored in S3 every 5 minutes.
- d) Take hourly DB backups to S3, with transaction logs stored in S3 every 5 minutes.

03. You would like to create a mirror image of your production environment in another region for disaster recovery purposes. Which of the following AWS resources do not need to be recreated in the second region?

(Choose 2 answers)

- a) Route53 Record Sets
- b) Launch Configurations
- c) EC2 Key Pairs
- d) Security Groups
- e) IAM Roles
- f) Elastic IP Addresses (EIP)

04. An administrator is using Amazon CloudFormation to deploy a three tier web application that consists of a web tier and application tier that will utilize Amazon DynamoDB for storage.

When creating the CloudFormation template which of the following would allow the application Instance access to the DynamoDB tables without exposing API credentials?

- a) Create an Identity and Access Management Role that has the required permissions to read and write from the .required DynamoDB table and associate the Role to the application instances by referencing an instance profile.
- b) Create an Identity and Access Management Role that has the required permissions to read and write from the required DynamoDB table and reference the Role in the instance profile property of the application instance.
- c) Use the Parameter section in the CloudFormation template to have the user input Access and Secret keys from an already created IAM user that has the permissions required to read and write from the required DynamoDB table.
- d) Create an Identity and Access Management user in the CloudFormation template that has permissions to read and write from the required DynamoDB table, use the GetAtt function to retrieve the Access and Secret keys and pass them to the application instance through user-data.

05. Your company policies require encryption of sensitive data at rest. You are considering the possible options for protecting data while storing it at rest on an EBS data volume, attached to an EC2 instance. Which of these options would allow you to encrypt your data at rest?

(Choose 3 answers)

- a) Implement third party volume encryption tools
- b) Implement SSL/TLS for all services running on the server
- c) Encrypt data inside your applications before storing it on EBS
- d) Encrypt data using native data encryption drivers at the file system level
- e) Do nothing as EBS volumes are encrypted by default

06. Your company plans to host a large donation website on Amazon Web Services (AWS). You anticipate a large and undetermined amount of traffic that will create many database writes.

To be certain that you do not drop any writes to a database hosted on AWS, which service should you use?

- a) Amazon Simple Queue Service (SQS) for capturing the writes and draining the queue to write to the database.
- b) Amazon DynamoDB with provisioned write throughput up to the anticipated peak write throughput.
- c) Amazon ElastiCache to store the writes until the writes are committed to the database.
- d) Amazon RDS with provisioned IOPS up to the anticipated peak write throughput.

07. Company B is launching a new game app for mobile devices. Users will log into the game using their existing social media account. To streamline data capture, Company B would like to directly save player data and scoring information from the mobile app to a DynamoDB table named ScoreData. When a user saves their game, the progress data will be stored to the GameState S3 bucket. What is the best approach for storing data to DynamoDB and S3?

- a) Use Login with Amazon allowing users to sign in with an Amazon account providing the mobile app with access to the ScoreData DynamoDB table and the GameState S3 bucket.
- b) Use temporary security credentials that assume a role providing access to the ScoreData DynamoDB table and the GameState S3 bucket using web identity federation
- c) Use an IAM user with access credentials assigned a role providing access to the ScoreData DynamoDB table and the GameState S3 bucket for distribution with the mobile app
- d) Use an EC2 instance that is launched with an EC2 role providing access to the ScoreData DynamoDB table and the GameState S3 bucket that communicates with the mobile app via web services

08. You require the ability to analyze a customer's clickstream data on a website, so they can do behavioral analysis. Your customer needs to know what sequence of pages and ads their customer clicked on.

This data will be used in real time to modify the page layouts as customers dick through the site, to increase stickiness and advertising click-through.

Which option meets the requirements for capturing and analyzing this data?

- a) Log dicks in weblogs by URL, store to Amazon S3, and then analyze with Elastic Map Reduce.
- b) Publish web clicks by session to an Amazon SQS queue; then periodically drain these events to Amazon RDS and analyze with SQL.
- c) Push web clicks by session to Amazon Kinesis, then analyze behavior using Kinesis workers.
- d) Write click events directly to Amazon Redshift, and then analyze with SQL.

09. A read only news reporting site with a combined web and application tier and a database tier that receives large and unpredictable traffic demands must be able to respond to these traffic fluctuations automatically. What AWS services should be used meet these requirements?

- a) Stateless instances for the web and application tier synchronized using ElastiCache Memcached in an autoscaling group monitored with CloudWatch, and RDS with read replicas
- b) Stateful instances for the web and application tier in an autoscaling group monitored with CloudWatch, and multi-AZ RDS
- c) Stateful instances for the web and application tier in an autoscaling group monitored with CloudWatch, and RDS with read replicas
- d) Stateless instances for the web and application tier synchronized using ElastiCache Memcached in an autoscaling group monitored with CloudWatch, and multi-AZ RDS

10. You require the ability to analyze a large amount of data which is stored on Amazon S3 using Amazon Elastic MapReduce. You are using the cc2.8xlarge instance type, whose CPUs are mostly idle during processing. Which of the below would be the most cost efficient way to reduce the runtime of the job?

- a) Create fewer, larger files m Amazon S3.
- b) Use smaller instances that have higher aggregate I/O performance.
- c) Create more, smaller files on Amazon S3.
- d) Add additional cc2.8xlarge instances by introducing a task group.

Answers to AWS-SAP Exam Questions:

Question: 01 Answer: b	Question: 02 Answer: c	Question: 03 Answer: a, e	Question: 04 Answer: b	Question: 05 Answer: b, c, d
Question: 06 Answer: d	Question: 07 Answer: b	Question: 08 Answer: c	Question: 09 Answer: a	Question: 10 Answer: b

Note: If you find any typo or data entry error in these sample questions, we request you to update us by commenting on this page or write an email on feedback@vmexam.com