SAP-C01 Q&As

AWS Certified Solutions Architect - Professional (SAP-C01)

Pass Amazon SAP-C01 Exam with 100% Guarantee

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

https://www.passapply.com/aws-solution-architect-professional.html

100% Passing Guarantee
100% Money Back Assurance

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

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

QUESTION 1

You are running a news website in the eu-west-1 region that updates every 15 minutes. The website has a world-wide audience. It uses an Auto Scaling group behind an Elastic Load Balancer and an Amazon RDS database. Static content resides on Amazon S3, and is distributed through Amazon CloudFront. Your Auto Scaling group is set to trigger a scale up event at 60% CPU utilization. You use an Amazon RDS extra large DB instance with 10.000 Provisioned IOPS, its CPU utilization is around 80%, while freeable memory is in the 2 GB range.

Web analytics reports show that the average load time of your web pages is around 1.5 to 2 seconds, but your SEO consultant wants to bring down the average load time to under 0.5 seconds.

How would you improve page load times for your users? (Choose three.)

A. Lower the scale up trigger of your Auto Scaling group to 30% so it scales more aggressively.

B. Add an Amazon ElastiCache caching layer to your application for storing sessions and frequent DB queries

C. Configure Amazon CloudFront dynamic content support to enable caching of re-usable content from your site

D. Switch the Amazon RDS database to the high memory extra large Instance type

E. Set up a second installation in another region, and use the Amazon Route 53 latency-based routing feature to select the right region.

Correct Answer: BCD

QUESTION 2

A company operates a group of imaging satellites. The satellites stream data to one of the company\\'s ground stations where processing creates about 5 GB of images per minute. This data is added to network-attached storage, where 2 PB of data are already stored.

The company runs a website that allows its customers to access and purchase the images over the Internet. This website is also running in the ground station. Usage analysis shows that customers are most likely to access images that have been captured in the last 24 hours.

The company would like to migrate the image storage and distribution system to AWS to reduce costs and increase the number of customers that can be served.

Which AWS architecture and migration strategy will meet these requirements?

A. Use multiple AWS Snowball appliances to migrate the existing imagery to Amazon S3. Create a 1-Gb AWS Direct Connect connection from the ground station to AWS, and upload new data to Amazon S3 through the Direct Connect connection. Migrate the data distribution website to Amazon EC2 instances. By using Amazon S3 as an origin, have this website serve the data through Amazon CloudFront by creating signed URLs.

B. Create a 1-Gb Direct Connect connection from the ground station to AWS. Use the AWS Command Line Interface to copy the existing data and upload new data to Amazon S3 over the Direct Connect connection. Migrate the data distribution website to EC2 instances. By using Amazon S3 as an origin, have this website serve the data through CloudFront by creating signed URLs.

C. Use multiple Snowball appliances to migrate the existing images to Amazon S3. Upload new data by regularly using Snowball appliances to upload data from the network-attached storage. Migrate the data distribution website to EC2

instances. By using Amazon S3 as an origin, have this website serve the data through CloudFront by creating signed URLs.

D. Use multiple Snowball appliances to migrate the existing images to an Amazon EFS file system. Create a 1-Gb Direct Connect connection from the ground station to AWS, and upload new data by mounting the EFS file system over the Direct Connect connection. Migrate the data distribution website to EC2 instances. By using webservers in EC2 that mount the EFS file system as the origin, have this website serve the data through CloudFront by creating signed URLs.

Correct Answer: B

QUESTION 3

Can you configure multiple Load Balancers with a single Auto Scaling group?

A. No

B. Yes, you can but only if it is configured with Amazon Redshift.

C. Yes, you can provide the ELB is configured with Amazon AppStream.

D. Yes

Correct Answer: D

Yes, you can configure more than one load balancer with an autoscaling group. Auto Scaling integrates with Elastic Load Balancing to enable you to attach one or more load balancers to an existing Auto Scaling group. After you attach the load balancer, it automatically registers the instances in the group and distributes incoming traffic across the instances.

Reference: http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/AS_Concepts.html

QUESTION 4

While assigning a tag to an instance, which of the below mentioned options is not a valid tag key/value pair?

A. Key : "aws" Value:"aws"

B. Key: "aws:name" Value: "instanceAnswer: Aws"

C. Key: "Name :aws" Value: "instanceAnswer: Aws"

D. Key : "nameAnswer: Aws" Value:"aws:instance"

Correct Answer: B

In Amazon Web Services, to help manage EC2 instances as well their usage in a better way, the user can tag the instances. The tags are metadata assigned by the user which consists of a key and value. The tag key cannot have a prefix as "aws:", although it can have only "aws".

Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using_Tags.html

QUESTION 5

A solutions architect is designing a disaster recovery strategy for a three-tier application. The application has an RTO of 30 minutes and an RPO of 5 minutes for the data tier. The application and web tiers are stateless and leverage a fleet of Amazon EC2 instances. The data tier consists of a 50 TB Amazon Aurora database.

Which combination of steps satisfies the RTO and RPO requirements while optimizing costs? (Choose two.)

A. Create daily snapshots of the EC2 instances and replicate the snapshots to another Region.

B. Deploy a hot standby of the application to another Region.

C. Create snapshots of the Aurora database every 5 minutes.

D. Create a cross-Region Aurora Replica of the database.

E. Create an AWS Backup job to replicate data to another Region.

Correct Answer: AD

QUESTION 6

A media company is hosting a high-traffic news website on AWS. The website\\'s front end is based solely on HTML and JavaScript. The company loads all dynamic content by using dynamic asynchronous JavaScript requests to a dedicated backend infrastructure.

The front end runs on four Amazon EC2 instances as web servers. The dynamic backend runs in containers on an Amazon Elastic Container Service (Amazon ECS) cluster that uses an Auto Scaling group of EC2 instances. The ECS tasks are behind an Application Load Balancer (ALB).

Which solutions should a solutions architect recommend to optimize costs? (Choose two.)

A. Migrate the front end of the website to an Amazon S3 bucket Deploy an Amazon CloudFront distribution. Set the S3 bucket as the distribution\\'s origin.

B. Deploy an Amazon CloudFront distribution. Configure the distribution to use the ALB endpoint as the origin.

C. Migrate the front-end services to the ECS cluster. Increase the minimum number of nodes in the Auto Scaling group.

D. Turn on Auto Scaling for the front-end EC2 instances. Configure a new listener rule on the ALB to serve the front end.

E. Migrate the backend of the website to an Amazon S3 bucket. Deploy an Amazon CloudFront distribution. Set the S3 bucket as the distribution\\'s origin.

Correct Answer: BD

Reference: https://aws.amazon.com/blogs/security/how-to-enhance-amazon-cloudfront-origin-securitywith-aws-waf-and-aws-secrets-manager/ https://digitalcloud.training/certification-training/aws-solutionsarchitect-associate/compute/elastic-load-balancing/

QUESTION 7

A user is running a batch process on EBS backed EC2 instances. The batch process launches few EC2 instances to process Hadoop Map reduce jobs which can run between 50 ?600 minutes or sometimes for even more time. The user wants a configuration that can terminate the instance only when the process is completed.

How can the user configure this with CloudWatch?

A. Configure a job which terminates all instances after 600 minutes

B. It is not possible to terminate instances automatically

C. Configure the CloudWatch action to terminate the instance when the CPU utilization falls below 5%

D. Set up the CloudWatch with Auto Scaling to terminate all the instances

Correct Answer: C

Amazon CloudWatch alarm watches a single metric over a time period that the user specifies and performs one or more actions based on the value of the metric relative to a given threshold over a number of time periods. The user can setup an action which terminates the instances when their CPU utilization is below a certain threshold for a certain period of time. The EC2 action can either terminate or stop the instance as part of the EC2 action. Reference: http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/UsingAlarmActions.html

QUESTION 8

An Auto Scaling group is running at the desired capacity of 5 instances and receives a trigger from the Cloudwatch Alarm to increase the capacity by 1. The cool down period is 5 minutes. Cloudwatch sends another trigger after 2 minutes to decrease the desired capacity by 1.

What will be the count of instances at the end of 4 minutes?

A. 4

B. 5

C. 6

D. 7

Correct Answer: C

The cool down period is the time difference between the end of one scaling activity (can be start or terminate) and the start of another one (can be start or terminate). During the cool down period, Auto Scaling does not allow the desired capacity of the Auto Scaling group to be changed by any other CloudWatch alarm. Thus, in this case the trigger from the second alarm will have no effect.

Reference: http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/AS_Concepts.html#healthcheck

QUESTION 9

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.

Correct Answer: B

Reference: https://docs.aws.amazon.com/fsx/latest/WindowsGuide/high-availability-multiAZ.html

QUESTION 10

You have recently joined a startup company building sensors to measure street noise and air quality in urban areas. The company has been running a pilot deployment of around 100 sensors for 3 months each sensor uploads 1KB of sensor data every minute to a backend hosted on AWS. During the pilot, you measured a peak or 10 IOPS on the database, and you stored an average of 3GB of sensor data per month in the database. The current deployment consists of a load-balanced auto scaled Ingestion layer using EC2 instances and a PostgreSQL RDS database with 500GB standard storage. The pilot is considered a success and your CEO has managed to get the attention or some potential investors. The business plan requires a deployment of at least 100K sensors which needs to be supported by the backend. You also need to store sensor data for at least two years to be able to compare year over year Improvements. To secure funding, you have to make sure that the platform meets these requirements and leaves room for further scaling.

Which setup win meet the requirements?

A. Add an SQS queue to the ingestion layer to buffer writes to the RDS instance

B. Ingest data into a DynamoDB table and move old data to a Redshift cluster

C. Replace the RDS instance with a 6 node Redshift cluster with 96TB of storage

D. Keep the current architecture but upgrade RDS storage to 3TB and 10K provisioned IOPS

Correct Answer: C

The POC solution is being scaled up by 1000, which means it will require 72TB of Storage to retain 24

months\\' worth of data. This rules out RDS as a possible DB solution which leaves you with Redshift.

I believe DynamoDB is a more cost effective and scales better for ingest rather than using EC2 in an auto

scaling group.

Also, this example solution from AWS is somewhat similar for reference.

QUESTION 11

A user authenticating with Amazon Cognito will go through a multi-step process to bootstrap their

credentials.

Amazon Cognito has two different flows for authentication with public providers.

Which of the following are the two flows?

A. Authenticated and non-authenticated

B. Public and private

C. Enhanced and basic

D. Single step and multistep

Correct Answer: C

A user authenticating with Amazon Cognito will go through a multi-step process to bootstrap their credentials. Amazon Cognito has two different flows for authentication with public providers: enhanced and basic.

Reference: http://docs.aws.amazon.com/cognito/devguide/identity/concepts/authentication-flow/

QUESTION 12

A company is migrating its applications to AWS. The applications will be deployed to AWS accounts owned by business units. The company has several teams of developers who are responsible for the development and maintenance of all applications. The company is expecting rapid growth in the number of users.

The company\\'s chief technology officer has the following requirements:

1.

Developers must launch the AWS infrastructure using AWS CloudFormation.

2.

Developers must not be able to create resources outside of CloudFormation.

3.

The solution must be able to scale to hundreds of AWS accounts.

Which of the following would meet these requirements? (Choose two.)

A. Using CloudFormation, create an IAM role that can be assumed by CloudFormation that has permissions to create all the resources the company needs. Use CloudFormation StackSets to deploy this template to each AWS account.

B. In a central account, create an IAM role that can be assumed by developers, and attach a policy that allows interaction with CloudFormation. Modify the AssumeRolePolicyDocument action to allow the IAM role to be passed to CloudFormation.

C. Using CloudFormation, create an IAM role that can be assumed by developers, and attach policies that allow interaction with and passing a role to CloudFormation. Attach an inline policy to deny access to all other AWS services. Use CloudFormation StackSets to deploy this template to each AWS account.

D. Using CloudFormation, create an IAM role for each developer, and attach policies that allow interaction with CloudFormation. Use CloudFormation StackSets to deploy this template to each AWS account.

E. In a central AWS account, create an IAM role that can be assumed by CloudFormation that has permissions to create the resources the company requires. Create a CloudFormation stack policy that allows the IAM role to manage resources. Use CloudFormation StackSets to deploy the CloudFormation stack policy to each AWS account.

Correct Answer: AB

Reference: https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_boundaries.html

QUESTION 13

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 m the same subnets of a VPC with distinct security groups that are configured.

The password (or the database is stored m 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.

Correct Answer: ADE

QUESTION 14

A company is planning to host a web application on AWS and works to load balance the traffic across a group of Amazon EC2 instances. One of the security requirements is to enable end-to-end encryption in transit between the client and the web server.

Which solution will meet this requirement?

A. Place the EC2 instances behind an Application Load Balancer (ALB) Provision an SSL certificate using AWS Certificate Manager (ACM), and associate the SSL certificate with the ALB. Export the SSL certificate and install it on each EC2 instance. Configure the ALB to listen on port 443 and to forward traffic to port 443 on the instances.

B. Associate the EC2 instances with a target group. Provision an SSL certificate using AWS Certificate Manager (ACM). Create an Amazon CloudFront distribution and configure It to use the SSL certificate. Set CloudFront to use the target group as the origin server

C. Place the EC2 instances behind an Application Load Balancer (ALB). Provision an SSL certificate using AWS Certificate Manager (ACM), and associate the SSL certificate with the ALB. Provision a third-party SSL certificate and install it on each EC2 instance. Configure the ALB to listen on port 443 and to forward traffic to port 443 on the instances.

D. Place the EC2 instances behind a Network Load Balancer (NLB). Provision a third-party SSL certificate and install it on the NLB and on each EC2 instance. Configure the NLB to listen on port 443 and to forward traffic to port 443 on the instances.

Correct Answer: C

QUESTION 15

A company has an application that runs a web service on Amazon EC2 instances and stores .jpg images in Amazon S3. The web traffic has a predictable baseline, but often demand spikes unpredictably for short periods of time. The application is loosely coupled and stateless. The .jpg images stored in Amazon S3 are accessed frequently for the first 15 to 20 days, they are seldom accessed thereafter but always need to be immediately available. The CIO has asked to find ways to reduce costs.

Which of the following options will reduce costs? (Choose two.)

A. Purchase Reserved instances for baseline capacity requirements and use On-Demand instances for the demand spikes.

B. Configure a lifecycle policy to move the .jpg images on Amazon S3 to S3 IA after 30 days.

C. Use On-Demand instances for baseline capacity requirements and use Spot Fleet instances for the demand spikes.

D. Configure a lifecycle policy to move the .jpg images on Amazon S3 to Amazon Glacier after 30 days.

E. Create a script that checks the load on all web servers and terminates unnecessary On-Demand instances.

Correct Answer: AB