Free Questions for SAA-C03 by certscare

Shared by Freeman on 24-05-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

A solutions architect is designing a user authentication solution for a company The solution must invoke two-factor authentication for users that log in from inconsistent geographical locations. IP addresses, or devices. The solution must also be able to scale up to accommodate millions of users.

Which solution will meet these requirements'?

Options:

A- Configure Amazon Cognito user pools for user authentication Enable the nsk-based adaptive authentication feature with multi-factor authentication (MFA)

B- Configure Amazon Cognito identity pools for user authentication Enable multi-factor authentication (MFA).

C- Configure AWS Identity and Access Management (1AM) users for user authentication Attach an 1AM policy that allows the AllowManageOwnUserMFA action

D- Configure AWS 1AM Identity Center (AWS Single Sign-On) authentication for user authentication Configure the permission sets to require multi-factor authentication
(MFA)

Answer:

A

Explanation:

Amazon Cognito user pools provide a secure and scalable user directory for user authentication and management. User pools support various authentication methods, such as username and password, email and password, phone number and password, and social identity providers. User pools also support multi-factor authentication (MFA), which adds an extra layer of security by requiring users to provide a verification code or a biometric factor in addition to their credentials. User pools can also enable risk-based adaptive authentication, which dynamically adjusts the authentication challenge based on the risk level of the sign-in attempt. For example, if a user tries to sign in from an unfamiliar device or location, the user pool can require a stronger authentication factor, such as SMS or email verification code. This feature helps to protect user accounts from unauthorized access and reduce the friction for legitimate users. User pools can scale up to millions of users and integrate with other AWS services, such as Amazon SNS, Amazon SES, AWS Lambda, and AWS KMS.

Amazon Cognito identity pools provide a way to federate identities from multiple identity providers, such as user pools, social identity providers, and corporate identity providers. Identity pools allow users to access AWS resources with temporary, limited-privilege credentials. Identity pools do not provide user authentication or management features, such as MFA or adaptive authentication. Therefore, option B is not correct.

AWS Identity and Access Management (IAM) is a service that helps to manage access to AWS resources. IAM users are entities that represent people or applications that need to interact with AWS. IAM users can be authenticated with a password or an access key. IAM users can also enable MFA for their own accounts, by using the AllowManageOwnUserMFA action in an IAM policy. However, IAM users are not suitable for user authentication for web or mobile applications, as they are intended for administrative purposes. IAM users also do not support adaptive authentication based on risk factors. Therefore, option C is not correct.

AWS IAM Identity Center (AWS Single Sign-On) is a service that enables users to sign in to multiple AWS accounts and applications with a single set of credentials. AWS SSO supports various identity sources, such as AWS SSO directory, AWS Managed Microsoft AD, and external identity providers. AWS SSO also supports MFA for user authentication, which can be configured in the permission sets that define the level of access for each user. However, AWS SSO does not support adaptive authentication based on risk factors. Therefore, option D is not correct.

Amazon Cognito User Pools

Adding Multi-Factor Authentication (MFA) to a User Pool

Risk-Based Adaptive Authentication

Amazon Cognito Identity Pools

IAM Users

Enabling MFA Devices

AWS Single Sign-On

How AWS SSO Works

Question 2

Question Type: MultipleChoice

A company wants to analyze and troubleshoot Access Denied errors and Unauthonzed errors that are related to 1AM permissions The company has AWS CloudTrail turned on Which solution will meet these requirements with the LEAST effort?

Options:

A- Use AWS Glue and write custom scripts to query CloudTrail logs for the errors

B- Use AWS Batch and write custom scripts to query CloudTrail logs for the errors

C- Search CloudTrail logs with Amazon Athena queries to identify the errors

D- Search CloudTrail logs with Amazon QuickSight. Create a dashboard to identify the errors.

Answer:

C

Explanation:

This solution meets the following requirements:

It is the least effort, as it does not require any additional AWS services, custom scripts, or data processing steps. Amazon Athena is a serverless interactive query service that allows you to analyze data in Amazon S3 using standard SQL. You can use Athena to query CloudTrail logs directly from the S3 bucket where they are stored, without any data loading or transformation. You can also use the AWS Management Console, the AWS CLI, or the Athena API to run and manage your queries.

It is effective, as it allows you to filter, aggregate, and join CloudTrail log data using SQL syntax. You can use various SQL functions and operators to specify the criteria for identifying Access Denied and Unauthorized errors, such as the error code, the user identity, the event source, the event name, the event time, and the resource ARN. You can also use subqueries, views, and common table expressions to simplify and optimize your queries.

It is flexible, as it allows you to customize and save your queries for future use. You can also export the query results to other formats, such as CSV or JSON, or integrate them with other AWS services, such as Amazon QuickSight, for further analysis and visualization.

Querying AWS CloudTrail Logs - Amazon Athena

Analyzing Data in S3 using Amazon Athena | AWS Big Data Blog

Troubleshoot IAM permisson access denied or unauthorized errors | AWS re:Post

Question 3

Question Type: MultipleChoice

A company's marketing data is uploaded from multiple sources to an Amazon S3 bucket A series ot data preparation jobs aggregate the data for reporting The data preparation jobs need to run at regular intervals in parallel A few jobs need to run in a specific order later

The company wants to remove the operational overhead of job error handling retry logic, and state management

Which solution will meet these requirements?

Options:

A- Use an AWS Lambda function to process the data as soon as the data is uploaded to the S3 bucket Invoke Other Lambda functions at regularly scheduled intervals

B- Use Amazon Athena to process the data Use Amazon EventBndge Scheduler to invoke Athena on a regular internal

C- Use AWS Glue DataBrew to process the data Use an AWS Step Functions state machine to run the DataBrew data preparation jobs

D- Use AWS Data Pipeline to process the data. Schedule Data Pipeline to process the data once at midnight.

Answer:

C

Explanation:

AWS Glue DataBrew is a visual data preparation tool that allows you to easily clean, normalize, and transform your data without writing any code. You can create and run data preparation jobs on your data stored in Amazon S3, Amazon Redshift, or other data sources. AWS Step Functions is a service that lets you coordinate multiple AWS services into serverless workflows. You can use Step Functions to orchestrate your DataBrew jobs, define the order and parallelism of execution, handle errors and retries, and monitor the state of your workflow. By using AWS Glue DataBrew and AWS Step Functions, you can meet the requirements of the company with minimal operational overhead, as you do not need to write any code, manage any servers, or deal with complex dependencies.

AWS Glue DataBrew

AWS Step Functions

Orchestrate AWS Glue DataBrew jobs using AWS Step Functions

Question 4

Question Type: MultipleChoice

A financial company needs to handle highly sensitive data The company will store the data in an Amazon S3 bucket The company needs to ensure that the data is encrypted in transit and at rest The company must manage the encryption keys outside the AWS Cloud

Which solution will meet these requirements?

Options:

A- Encrypt the data in the S3 bucket with server-side encryption (SSE) that uses an AWS Key Management Service (AWS KMS) customer managed key

B- Encrypt the data in the S3 bucket with server-side encryption (SSE) that uses an AWS Key Management Service (AWS KMS) AWS managed key

C- Encrypt the data in the S3 bucket with the default server-side encryption (SSE)

D- Encrypt the data at the company's data center before storing the data in the S3 bucket

Answer:

D

Explanation:

This option is the only solution that meets the requirements because it allows the company to encrypt the data with its own encryption keys and tools outside the AWS Cloud. By encrypting the data at the company's data center before storing the data in the S3 bucket, the company can ensure that the data is encrypted in transit and at rest, and that the company has full control over the encryption keys and processes. This option also avoids the need to use any AWS encryption services or features, which may not be compatible with the company's security policies or compliance standards.

A) Encrypt the data in the S3 bucket with server-side encryption (SSE) that uses an AWS Key Management Service (AWS KMS) customer managed key. This option does not meet the requirements because it does not allow the company to manage the encryption keys outside the AWS Cloud. Although the company can create and use its own customer managed key in AWS KMS, the key is still stored and managed by AWS KMS, which is a service within the AWS Cloud. Moreover, the company still needs to use the AWS encryption features and APIs to encrypt and decrypt the data in the S3 bucket, which may not be compatible with the company's security policies or compliance standards.

B) Encrypt the data in the S3 bucket with server-side encryption (SSE) that uses an AWS Key Management Service (AWS KMS) AWS managed key. This option does not meet the requirements because it does not allow the company to manage the encryption keys outside the AWS Cloud. In this option, the company uses the default AWS managed key in AWS KMS, which is created and managed by AWS on behalf of the company. The company has no control over the key rotation, deletion, or recovery policies. Moreover, the company still needs to use the AWS encryption features and APIs to encrypt and decrypt the data in the S3 bucket, which may not be compatible with the company's security policies or compliance standards.

C) Encrypt the data in the S3 bucket with the default server-side encryption (SSE). This option does not meet the requirements because it does not allow the company to manage the encryption keys outside the AWS Cloud. In this option, the company uses the default server-side encryption with Amazon S3 managed keys (SSE-S3), which is applied to every bucket in Amazon S3. The company has no visibility or control over the encryption keys, which are managed by Amazon S3. Moreover, the company still needs to use the AWS encryption features and APIs to encrypt and decrypt the data in the S3 bucket, which may not be compatible with the company's security policies or compliance standards.

1Protecting data with encryption - Amazon Simple Storage Service

2Protecting data with server-side encryption - Amazon Simple Storage Service

3Protecting data by using client-side encryption - Amazon Simple Storage Service

4AWS Key Management Service Concepts - AWS Key Management Service

Question 5

Question Type: MultipleChoice

A company maintains about 300 TB in Amazon S3 Standard storage month after month The S3 objects are each typically around 50 GB in size and are frequently replaced with multipart uploads by their global application The number and size of S3 objects remain constant but the company's S3 storage costs are increasing each month.

How should a solutions architect reduce costs in this situation?

Options:

A- Switch from multipart uploads to Amazon S3 Transfer Acceleration.

B- Enable an S3 Lifecycle policy that deletes incomplete multipart uploads.

C- Configure S3 inventory to prevent objects from being archived too quickly.

D- Configure Amazon CloudFront to reduce the number of objects stored in Amazon S3.

Answer:

B

Explanation:

This option is the most cost-effective way to reduce the S3 storage costs in this situation. Incomplete multipart uploads are parts of objects that are not completed or aborted by the application. They consume storage space and incur charges until they are deleted. By enabling an S3 Lifecycle policy that deletes incomplete multipart uploads, you can automatically remove them after a specified period of time (such as one day) and free up the storage space. This will reduce the S3 storage costs and also improve the performance of the application by avoiding unnecessary retries or errors.

Option A is not correct because switching from multipart uploads to Amazon S3 Transfer Acceleration will not reduce the S3 storage costs. Amazon S3 Transfer Acceleration is a feature that enables faster data transfers to and from S3 by using the AWS edge network. It is useful for improving the upload speed of large objects over long distances, but it does not affect the storage space or charges. In fact, it may increase the costs by adding a data transfer fee for using the feature.

Option C is not correct because configuring S3 inventory to prevent objects from being archived too quickly will not reduce the S3 storage costs. Amazon S3 Inventory is a feature that provides a report of the objects and their metadata in an S3 bucket. It is useful for managing and auditing the S3 objects, but it does not affect the storage space or charges. In fact, it may increase the costs by generating additional S3 objects for the inventory reports.

Option D is not correct because configuring Amazon CloudFront to reduce the number of objects stored in Amazon S3 will not reduce the S3 storage costs. Amazon CloudFront is a content delivery network (CDN) that distributes the S3 objects to edge locations for faster and lower latency access. It is useful for improving the download speed and availability of the S3 objects, but it does not affect the storage space or charges. In fact, it may increase the costs by adding a data transfer fee for using the service.Reference:

Managing your storage lifecycle

Using multipart upload

Amazon S3 Transfer Acceleration

Amazon S3 Inventory

What Is Amazon CloudFront?

Question 6

Question Type: **MultipleChoice**

A company has a business-critical application that runs on Amazon EC2 instances. The application stores data in an Amazon DynamoDB table. The company must be able to revert the table to any point within the last 24 hours.

Which solution meets these requirements with the LEAST operational overhead?

Options:

A- Configure point-in-time recovery for the table.

B- Use AWS Backup for the table.

C- Use an AWS Lambda function to make an on-demand backup of the table every hour.

D- Turn on streams on the table to capture a log of all changes to the table in the last 24 hours Store a copy of the stream in an Amazon S3 bucket.

Answer:

A

Explanation:

Point-in-time recovery (PITR) for DynamoDB is a feature that enables you to restore your table data to any point in time during the last 35 days. PITR helps protect your table from accidental write or delete operations, such as a test script writing to a production table or a user issuing a wrong command. PITR is easy to use, fully managed, fast, and scalable. You can enable PITR with a single click in the DynamoDB console or with a simple API call. You can restore a table to a new table using the console, the AWS CLI, or the DynamoDB API. PITR does not consume any provisioned table capacity and has no impact on the performance or availability of your production applications. PITR meets the requirements of the company with the least operational overhead, as it does not require any manual backup creation, scheduling, or maintenance. It also provides per-second granularity for restoring the table to any point within the last 24 hours.

Point-in-time recovery for DynamoDB - Amazon DynamoDB

Amazon DynamoDB point-in-time recovery (PITR)

Enable Point-in-Time Recovery (PITR) for Dynamodb global tables

Restoring a DynamoDB table to a point in time - Amazon DynamoDB

Point-in-time recovery: How it works - Amazon DynamoDB

Question 7

Question Type: **MultipleChoice**

A company has an organization in AWS Organizations that has all features enabled The company requires that all API calls and logins in any existing or new AWS account must be audited The company needs a managed solution to prevent additional work and to minimize costs The company also needs to know when any AWS account is not compliant with the AWS Foundational Security Best Practices (FSBP) standard.

Which solution will meet these requirements with the LEAST operational overhead?

Options:

A- Deploy an AWS Control Tower environment in the Organizations management account Enable AWS Security Hub and AWS Control Tower Account Factory in the environment.

B- Deploy an AWS Control Tower environment in a dedicated Organizations member account Enable AWS Security Hub and AWS Control Tower Account Factory in the environment.

C- Use AWS Managed Services (AMS) Accelerate to build a multi-account landing zone (MALZ) Submit an RFC to self-service provision Amazon GuardDuty in the MALZ.

D- Use AWS Managed Services (AMS) Accelerate to build a multi-account landing zone (MALZ) Submit an RFC to self-service provision AWS Security Hub in the MALZ.

Answer:

A

Explanation:

AWS Control Tower is a fully managed service that simplifies the setup and governance of a secure, compliant, multi-account AWS environment. It establishes a landing zone that is based on best-practices blueprints, and it enables governance using controls you can choose from a pre-packaged list. The landing zone is a well-architected, multi-account baseline that follows AWS best practices. Controls implement governance rules for security, compliance, and operations. AWS Security Hub is a service that provides a comprehensive view of your security posture across your AWS accounts. It aggregates, organizes, and prioritizes security alerts and findings from multiple AWS services, such as Amazon GuardDuty, Amazon Inspector, Amazon Macie, AWS Firewall Manager, and AWS IAM Access Analyzer, as well as from AWS Partner solutions. AWS Security Hub continuously monitors your environment using automated compliance checks based on the AWS best practices and industry standards, such as the AWS Foundational Security Best Practices (FSBP) standard. AWS Control Tower Account Factory is a feature that automates the provisioning of new AWS accounts that are preconfigured to meet your business, security, and compliance requirements. By deploying an AWS Control Tower environment in the Organizations management account, you can leverage the existing organization structure and policies, and enable AWS Security Hub and AWS Control Tower Account Factory in the environment. This way, you can audit all API calls and logins in any existing or new AWS account, monitor the compliance status of each account with the FSBP standard, and provision new accounts with ease and consistency. This solution meets the requirements with the least operational overhead, as you do not need to manage any infrastructure, perform any data migration, or submit any requests for changes.

AWS Control Tower

[AWS Security Hub]

[AWS Control Tower Account Factory]

Question 8

Question Type: **MultipleChoice**

An ecommerce company runs applications in AWS accounts that are part of an organization in AWS Organizations The applications run on Amazon Aurora PostgreSQL databases across all the accounts The company needs to prevent malicious activity and must identify abnormal failed and incomplete login attempts to the databases

Which solution will meet these requirements in the MOST operationally efficient way?

Options:

A- Attach service control policies (SCPs) to the root of the organization to identify the failed login attempts

B- Enable the Amazon RDS Protection feature in Amazon GuardDuty for the member accounts of the organization

C- Publish the Aurora general logs to a log group in Amazon CloudWatch Logs Export the log data to a central Amazon S3 bucket

D- Publish all the Aurora PostgreSQL database events in AWS CloudTrail to a central Amazon S3 bucket

Answer:

C

Explanation:

This option is the most operationally efficient way to meet the requirements because it allows the company to monitor and analyze the database login activity across all the accounts in the organization. By publishing the Aurora general logs to a log group in Amazon CloudWatch Logs, the company can enable the logging of the database connections, disconnections, and failed authentication attempts. By exporting the log data to a central Amazon S3 bucket, the company can store the log data in a durable and cost-effective way and use other AWS services or tools to perform further analysis or alerting on the log data. For example, the company can use Amazon Athena to query the log data in Amazon S3, or use Amazon SNS to send notifications based on the log data.

A) Attach service control policies (SCPs) to the root of the organization to identify the failed login attempts. This option is not effective because SCPs are not designed to identify the failed login attempts, but to restrict the actions that the users and roles can perform in the member accounts of the organization. SCPs are applied to the AWS API calls, not to the database login attempts. Moreover, SCPs do not provide any logging or analysis capabilities for the database activity.

B) Enable the Amazon RDS Protection feature in Amazon GuardDuty for the member accounts of the organization. This option is not optimal because the Amazon RDS Protection feature in Amazon GuardDuty is not available for Aurora PostgreSQL databases, but only for Amazon RDS for MySQL and Amazon RDS for MariaDB databases. Moreover, the Amazon RDS Protection feature does not monitor the database login attempts, but the network and API activity related to the RDS instances.

D) Publish all the Aurora PostgreSQL database events in AWS CloudTrail to a central Amazon S3 bucket. This option is not sufficient because AWS CloudTrail does not capture the database login attempts, but only the AWS API calls made by or on behalf of the Aurora PostgreSQL database. For example, AWS CloudTrail can record the events such as creating, modifying, or deleting the database instances, clusters, or snapshots, but not the events such as connecting, disconnecting, or failing to authenticate to the database.

1Working with Amazon Aurora PostgreSQL - Amazon Aurora

2Working with log groups and log streams - Amazon CloudWatch Logs

3Exporting Log Data to Amazon S3 - Amazon CloudWatch Logs

[4] Amazon GuardDuty FAQs

[5] Logging Amazon RDS API Calls with AWS CloudTrail - Amazon Relational Database Service

Question 9

Question Type: MultipleChoice

A company has a mobile game that reads most of its metadata from an Amazon RDS DB instance. As the game increased in popularity, developers noticed slowdowns related to the game's metadata load times Performance metrics indicate that simply scaling the database will not help A solutions architect must explore all options that include capabilities for snapshots, replication, and sub-millisecond response times

What should the solutions architect recommend to solve these issues'?

Options:

A- Migrate the database to Amazon Aurora with Aurora Replicas

B- Migrate the database to Amazon DynamoDB with global tables

C- Add an Amazon ElastiCache for Redis layer in front of the database.

D- Add an Amazon ElastiCache for Memcached layer in front of the database

Answer:

C

Explanation:

This option is the most suitable way to improve the game's metadata load times without migrating the database. Amazon ElastiCache for Redis is a fully managed, in-memory data store that provides sub-millisecond latency and high throughput for read-intensive workloads. You can use it as a caching layer in front of your RDS DB instance to store frequently accessed metadata and reduce the load on the database. You can also take advantage of Redis features such as snapshots, replication, and data persistence to ensure data durability and availability. ElastiCache for Redis scales automatically to meet your demand and integrates with other AWS services such as CloudFormation, CloudWatch, and IAM.

Option A is not optimal because migrating the database to Amazon Aurora with Aurora Replicas would incur additional costs and complexity. Amazon Aurora is a relational database service that provides high performance, availability, and compatibility with MySQL and PostgreSQL. Aurora Replicas are read-only copies of the primary database that can be used for scaling read capacity and enhancing availability. However, migrating the database to Aurora would require modifying the application code, testing the compatibility, and performing the data migration. Moreover, Aurora Replicas may not provide sub-millisecond response times as ElastiCache for Redis does.

Option B is not optimal because migrating the database to Amazon DynamoDB with global tables would incur additional costs and complexity. Amazon DynamoDB is a NoSQL database service that provides fast and flexible data access for any scale. Global tables are a feature of DynamoDB that enables you to replicate your data across multiple AWS Regions for high availability and performance. However, migrating the database to DynamoDB would require changing the data model, modifying the application code, and performing the data migration. Moreover, global tables may not be necessary for the game's metadata, as they are mainly used for cross-region data access and disaster recovery.

Option D is not optimal because adding an Amazon ElastiCache for Memcached layer in front of the database would not provide the same capabilities as ElastiCache for Redis. Amazon ElastiCache for Memcached is another fully managed, in-memory data store that provides high performance and scalability for caching workloads. However, Memcached does not support snapshots, replication, or data persistence, which means that the cached data may be lost in case of a node failure or a cache eviction. Moreover, Memcached does not integrate with other AWS services as well as Redis does. Therefore, ElastiCache for Redis is a better choice for this scenario.Reference:

What Is Amazon ElastiCache for Redis?

What Is Amazon Aurora?

What Is Amazon DynamoDB?

What Is Amazon ElastiCache for Memcached?