

Exam Code: DBS-C01

Exam Name: AWS Certified Database - Specialty

Exam A

QUESTION 1

A company is using Amazon with Aurora Replicas for read-only workload scaling. A Database Specialist needs to split up two read-only applications so each application always connects to a dedicated replica. The Database Specialist wants to implement load balancing and high availability for the read-only applications.

Which solution meets these requirements?

- A. Use a specific instance endpoint for each replica and add the instance endpoint to each read-only application connection string.
- B. Use reader endpoints for both the read-only workload applications.
- C. Use a reader endpoint for one read-only application and use an instance endpoint for the other read-only application.
- D. Use custom endpoints for the two read-only applications.

Correct Answer: D

Section:

Explanation:

<https://aws.amazon.com/about-aws/whats-new/2018/11/amazon-aurora-simplifies-workloadmanagement-with-custom-endpoints/>

QUESTION 2

An online gaming company is planning to launch a new game with Amazon DynamoDB as its data store. The database should be designated to support the following use cases:

Update scores in real time whenever a player is playing the game. Retrieve a player's score details for a specific game session.

A Database Specialist decides to implement a DynamoDB table. Each player has a unique user_id and each game has a unique game_id.

Which choice of keys is recommended for the DynamoDB table?

- A. Create a global secondary index with game_id as the partition key
- B. Create a global secondary index with user_id as the partition key
- C. Create a composite primary key with game_id as the partition key and user_id as the sort key
- D. Create a composite primary key with user_id as the partition key and game_id as the sort key

Correct Answer: D

Section:

Explanation:

<https://aws.amazon.com/blogs/database/amazon-dynamodb-gaming-use-cases-and-designpatterns/>"EA uses the user ID as the partition key and primary key (a 1:1 modeling pattern)."

<https://aws.amazon.com/blogs/database/choosing-the-right-dynamodb-partition-key/>

"Partition key and sort key: Referred to as a composite primary key, this type of key is composed of two attributes. The first attribute is the partition key, and the second attribute is the sort key."

QUESTION 3

A Database Specialist migrated an existing production MySQL database from on-premises to an Amazon RDS for MySQL DB instance. However, after the migration, the database needed to be encrypted at rest using AWS KMS. Due to the size of the database, reloading, the data into an encrypted database would be too time-consuming, so it is not an option.

How should the Database Specialist satisfy this new requirement?

- A. Create a snapshot of the unencrypted RDS DB instance. Create an encrypted copy of the unencrypted snapshot. Restore the encrypted snapshot copy.
- B. Modify the RDS DB instance. Enable the AWS KMS encryption option that leverages the AWS CLI.

- C. Restore an unencrypted snapshot into a MySQL RDS DB instance that is encrypted.
- D. Create an encrypted read replica of the RDS DB instance. Promote it the master.

Correct Answer: A

Section:

Explanation:

"However, because you can encrypt a copy of an unencrypted DB snapshot, you can effectively add encryption to an unencrypted DB instance. That is, you can create a snapshot of your DB instance, and then create an encrypted copy of that snapshot. You can then restore a DB instance from the encrypted snapshot, and thus you have an encrypted copy of your original DB instance. For more information, see Copying a Snapshot."

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.Encryption.html>

QUESTION 4

A Database Specialist is planning to create a read replica of an existing Amazon RDS for MySQL Multi-AZ DB instance. When using the AWS Management Console to conduct this task, the Database Specialist discovers that the source RDS DB instance does not appear in the read replica source selection box, so the read replica cannot be created.

What is the most likely reason for this?

- A. The source DB instance has to be converted to Single-AZ first to create a read replica from it.
- B. Enhanced Monitoring is not enabled on the source DB instance.
- C. The minor MySQL version in the source DB instance does not support read replicas.
- D. Automated backups are not enabled on the source DB instance.

Correct Answer: D

Section:

Explanation:

>Your source DB instance must have backup retention enabled.

https://docs.aws.amazon.com/AmazonRDS/latest/APIReference/API_CreateDBInstanceReadReplica.html

Reference: <https://aws.amazon.com/rds/features/read-replicas/>

QUESTION 5

A Database Specialist has migrated an on-premises Oracle database to Amazon Aurora PostgreSQL.

The schema and the data have been migrated successfully. The on-premises database server was also being used to run database maintenance cron jobs written in Python to perform tasks including data purging and generating data exports. The logs for these jobs show that, most of the time, the jobs completed within 5 minutes, but a few jobs took up to 10 minutes to complete. These maintenance jobs need to be set up for Aurora PostgreSQL.

How can the Database Specialist schedule these jobs so the setup requires minimal maintenance and provides high availability?

- A. Create cron jobs on an Amazon EC2 instance to run the maintenance jobs following the required schedule.
- B. Connect to the Aurora host and create cron jobs to run the maintenance jobs following the required schedule.
- C. Create AWS Lambda functions to run the maintenance jobs and schedule them with Amazon CloudWatch Events.
- D. Create the maintenance job using the Amazon CloudWatch job scheduling plugin.

Correct Answer: C

Section:

Explanation:

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/events/Create-CloudWatch-Events-Scheduled-Rule.html>

www.VCEplus.io

<https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/schedule-jobs-for-amazon-rdsand-aurora-postgresql-using-lambda-and-secrets-manager.html>

a job for data extraction or a job for data purging can easily be scheduled using cron. For these jobs, database credentials are typically either hard-coded or stored in a properties file. However, when you migrate to Amazon Relational Database Service (Amazon RDS) or Amazon Aurora PostgreSQL, you lose the ability to log in to the host instance to schedule cron jobs. This pattern describes how to use AWS Lambda and AWS Secrets Manager to schedule jobs for Amazon RDS and Aurora PostgreSQL databases after migration.

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/events/RunLambdaSchedule.html>

QUESTION 6

A company has an Amazon RDS Multi-AZ DB instances that is 200 GB in size with an RPO of 6 hours.

To meet the company's disaster recovery policies, the database backup needs to be copied into another Region. The company requires the solution to be cost-effective and operationally efficient.

What should a Database Specialist do to copy the database backup into a different Region?

- A. Use Amazon RDS automated snapshots and use AWS Lambda to copy the snapshot into another Region
- B. Use Amazon RDS automated snapshots every 6 hours and use Amazon S3 cross-Region replication to copy the snapshot into another Region
- C. Create an AWS Lambda function to take an Amazon RDS snapshot every 6 hours and use a second Lambda function to copy the snapshot into another Region
- D. Create a cross-Region read replica for Amazon RDS in another Region and take an automated snapshot of the read replica

Correct Answer: C

Section:

Explanation:

System snapshot can't fulfill 6 hours requirement. You need to control it by script

<https://aws.amazon.com/blogs/database/%C2%AD%C2%AD%C2%ADautomating-cross-region-crossaccount-snapshot-copies-with-the-snapshot-tool-for-amazon-aurora/>

QUESTION 7

An Amazon RDS EBS-optimized instance with Provisioned IOPS (PIOPS) storage is using less than half of its allocated IOPS over the course of several hours under constant load. The RDS instance exhibits multi-second read and write latency, and uses all of its maximum bandwidth for read throughput, yet the instance uses less than half of its CPU and RAM resources.

What should a Database Specialist do in this situation to increase performance and return latency to sub-second levels?

- A. Increase the size of the DB instance storage
- B. Change the underlying EBS storage type to General Purpose SSD (gp2)
- C. Disable EBS optimization on the DB instance
- D. Change the DB instance to an instance class with a higher maximum bandwidth

Correct Answer: D

Section:

Explanation:

https://docs.amazonaws.cn/en_us/AmazonRDS/latest/UserGuide/CHAP_BestPractices.html

QUESTION 8

After restoring an Amazon RDS snapshot from 3 days ago, a company's Development team cannot connect to the restored RDS DB instance. What is the likely cause of this problem?

- A. The restored DB instance does not have Enhanced Monitoring enabled
- B. The production DB instance is using a custom parameter group
- C. The restored DB instance is using the default security group
- D. The production DB instance is using a custom option group

Correct Answer: C

Section:

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/rds-cannot-connect/>
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_RestoreFromSnapshot.html

QUESTION 9

A gaming company has implemented a leaderboard in AWS using a Sorted Set data structure within Amazon ElastiCache for Redis. The ElastiCache cluster has been deployed with cluster mode disabled and has a replication group deployed with two additional replicas. The company is planning for a worldwide gaming event and is anticipating a higher write load than what the current cluster can handle. Which method should a Database Specialist use to scale the ElastiCache cluster ahead of the upcoming event?

- A. Enable cluster mode on the existing ElastiCache cluster and configure separate shards for the Sorted Set across all nodes in the cluster.
- B. Increase the size of the ElastiCache cluster nodes to a larger instance size.
- C. Create an additional ElastiCache cluster and load-balance traffic between the two clusters.
- D. Use the EXPIRE command and set a higher time to live (TTL) after each call to increment a given key.

Correct Answer: B

Section:

Explanation:

QUESTION 10

An ecommerce company has tasked a Database Specialist with creating a reporting dashboard that visualizes critical business metrics that will be pulled from the core production database running on Amazon Aurora. Data that is read by the dashboard should be available within 100 milliseconds of an update. The Database Specialist needs to review the current configuration of the Aurora DB cluster and develop a cost-effective solution. The solution needs to accommodate the unpredictable read workload from the reporting dashboard without any impact on the write availability and performance of the DB cluster. Which solution meets these requirements?

- A. Turn on the serverless option in the DB cluster so it can automatically scale based on demand.
- B. Provision a clone of the existing DB cluster for the new Application team.
- C. Create a separate DB cluster for the new workload, refresh from the source DB cluster, and set up ongoing replication using AWS DMS change data capture (CDC).
- D. Add an automatic scaling policy to the DB cluster to add Aurora Replicas to the cluster based on CPU consumption.

Correct Answer: A

Section:

Explanation:

QUESTION 11

A retail company is about to migrate its online and mobile store to AWS. The company's CEO has strategic plans to grow the brand globally. A Database Specialist has been challenged to provide predictable read and write database performance with minimal operational overhead. What should the Database Specialist do to meet these requirements?

- A. Use Amazon DynamoDB global tables to synchronize transactions
- B. Use Amazon EMR to copy the orders table data across Regions

- C. Use Amazon Aurora Global Database to synchronize all transactions
- D. Use Amazon DynamoDB Streams to replicate all DynamoDB transactions and sync them

Correct Answer: A

Section:

Explanation:

<https://aws.amazon.com/dynamodb/features/>

With global tables, your globally distributed applications can access data locally in the selected regions to get single-digit millisecond read and write performance.

Not Aurora Global Database, as per this link: https://aws.amazon.com/rds/aurora/globaldatabase/?nc1=h_ls.

Aurora Global Database lets you easily scale database reads across the world and place your applications close to your users.

QUESTION 12

A company is closing one of its remote data centers. This site runs a 100 TB on-premises data warehouse solution. The company plans to use the AWS Schema Conversion Tool (AWS SCT) and AWS DMS for the migration to AWS. The site network bandwidth is 500 Mbps. A Database Specialist wants to migrate the on-premises data using Amazon S3 as the data lake and Amazon Redshift as the data warehouse. This move must take place during a 2-week period when source systems are shut down for maintenance. The data should stay encrypted at rest and in transit.

Which approach has the least risk and the highest likelihood of a successful data transfer?

- A. Set up a VPN tunnel for encrypting data over the network from the data center to AWS. Leverage AWS SCT and apply the converted schema to Amazon Redshift. Once complete, start an AWS DMS task to move the data from the source to Amazon S3. Use AWS Glue to load the data from Amazon S3 to Amazon Redshift.
- B. Leverage AWS SCT and apply the converted schema to Amazon Redshift. Start an AWS DMS task with two AWS Snowball Edge devices to copy data from on-premises to Amazon S3 with AWS KMS encryption. Use AWS DMS to finish copying data to Amazon Redshift.
- C. Leverage AWS SCT and apply the converted schema to Amazon Redshift. Once complete, use a fleet of 10 TB dedicated encrypted drives using the AWS Import/Export feature to copy data from on-premises to Amazon S3 with AWS KMS encryption. Use AWS Glue to load the data to Amazon Redshift.
- D. Set up a VPN tunnel for encrypting data over the network from the data center to AWS. Leverage a native database export feature to export the data and compress the files. Use the aws S3 cp multipart upload command to upload these files to Amazon S3 with AWS KMS encryption. Once complete, load the data to Amazon Redshift using AWS Glue.

Correct Answer: B

Section:

Explanation:

<https://aws.amazon.com/blogs/database/new-aws-dms-and-aws-snowball-integration-enables-mass-database-migrations-and-migrations-of-large-databases/>

QUESTION 13

A company has multiple applications serving data from a secure on-premises database. The company is migrating all applications and databases to the AWS Cloud. The IT Risk and Compliance department requires that auditing be enabled on all secure databases to capture all log ins, log outs, failed logins, permission changes, and database schema changes. A Database Specialist has recommended Amazon Aurora MySQL as the migration target, and leveraging the Advanced Auditing feature in Aurora.

Which events need to be specified in the Advanced Auditing configuration to satisfy the minimum auditing requirements? (Choose three.)

- A. CONNECT
- B. QUERY_DCL
- C. QUERY_DDL
- D. QUERY_DML
- E. TABLE
- F. QUERY

Correct Answer: A, B, C

Section:

Explanation:

Connect - logins / DCL - authorizations (grant, revoke), DDL - schema updates

QUESTION 14

A gaming company has recently acquired a successful iOS game, which is particularly popular during the holiday season. The company has decided to add a leaderboard to the game that uses Amazon DynamoDB. The application load is expected to ramp up over the holiday season. Which solution will meet these requirements at the lowest cost?

- A. DynamoDB Streams
- B. DynamoDB with DynamoDB Accelerator
- C. DynamoDB with on-demand capacity mode
- D. DynamoDB with provisioned capacity mode with Auto Scaling

Correct Answer: C

Section:

Explanation:

Reference: https://aws.amazon.com/blogs/database/running-spiky-workloads-and-optimizing-costs-by-more-than-90-using-amazon-dynamodb-on-demand-capacity-mode/?nc1=b_rp

QUESTION 15

A company's Security department established new requirements that state internal users must connect to an existing Amazon RDS for SQL Server DB instance using their corporate Active Directory (AD) credentials. A Database Specialist must make the modifications needed to fulfill this requirement. Which combination of actions should the Database Specialist take? (Choose three.)

- A. Disable Transparent Data Encryption (TDE) on the RDS SQL Server DB instance.
- B. Modify the RDS SQL Server DB instance to use the directory for Windows authentication. Create appropriate new logins.
- C. Use the AWS Management Console to create an AWS Managed Microsoft AD. Create a trust relationship with the corporate AD.
- D. Stop the RDS SQL Server DB instance, modify it to use the directory for Windows authentication, and start it again. Create appropriate new logins.
- E. Use the AWS Management Console to create an AD Connector. Create a trust relationship with the corporate AD.
- F. Configure the AWS Managed Microsoft AD domain controller Security Group.

Correct Answer: B, C, F

Section:

Explanation:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_SQLServerWinAuth.html

QUESTION 16

A Database Specialist is performing a proof of concept with Amazon Aurora using a small instance to confirm a simple database behavior. When loading a large dataset and creating the index, the Database Specialist encounters the following error message from Aurora:

ERROR: cloud not write block 7507718 of temporary file: No space left on device What is the cause of this error and what should the Database Specialist do to resolve this issue?

- A. The scaling of Aurora storage cannot catch up with the data loading. The Database Specialist needs to modify the workload to load the data slowly.

- B. The scaling of Aurora storage cannot catch up with the data loading. The Database Specialist needs to enable Aurora storage scaling.
- C. The local storage used to store temporary tables is full. The Database Specialist needs to scale up the instance.
- D. The local storage used to store temporary tables is full. The Database Specialist needs to enable local storage scaling.

Correct Answer: C

Section:

Explanation:

Reference: <https://serverfault.com/QUESTION NO:s/109828/how-can-i-tune-postgres-to-avoid-thiserror>

QUESTION 17

A financial company wants to store sensitive user data in an Amazon Aurora PostgreSQL DB cluster.

The database will be accessed by multiple applications across the company. The company has mandated that all communications to the database be encrypted and the server identity must be validated. Any non-SSL-based connections should be disallowed access to the database.

Which solution addresses these requirements?

- A. Set the `rds.force_ssl=0` parameter in DB parameter groups. Download and use the Amazon RDS certificate bundle and configure the PostgreSQL connection string with `sslmode=allow`.
- B. Set the `rds.force_ssl=1` parameter in DB parameter groups. Download and use the Amazon RDS certificate bundle and configure the PostgreSQL connection string with `sslmode=disable`.
- C. Set the `rds.force_ssl=0` parameter in DB parameter groups. Download and use the Amazon RDS certificate bundle and configure the PostgreSQL connection string with `sslmode=verify-ca`.
- D. Set the `rds.force_ssl=1` parameter in DB parameter groups. Download and use the Amazon RDS certificate bundle and configure the PostgreSQL connection string with `sslmode=verify-full`.

Correct Answer: D

Section:

Explanation:

PostgreSQL: `sslrootcert=rds-cert.pem sslmode=[verify-ca | verify-full]`

Reference: <https://forums.aws.amazon.com/message.jspa?messageID=734076>

QUESTION 18

A company is using 5 TB Amazon RDS DB instances and needs to maintain 5 years of monthly database backups for compliance purposes. A Database Administrator must provide Auditors with data within 24 hours.

Which solution will meet these requirements and is the MOST operationally efficient?

- A. Create an AWS Lambda function to run on the first day of every month to take a manual RDS snapshot. Move the snapshot to the company's Amazon S3 bucket.
- B. Create an AWS Lambda function to run on the first day of every month to take a manual RDS snapshot.
- C. Create an RDS snapshot schedule from the AWS Management Console to take a snapshot every 30 days.
- D. Create an AWS Lambda function to run on the first day of every month to create an automated RDS snapshot.

Correct Answer: A

Section:

Explanation:

Unlike automated backups, manual snapshots aren't subject to the backup retention period.

Snapshots don't expire. For very long-term backups of MariaDB, MySQL, and PostgreSQL data, we recommend exporting snapshot data to Amazon S3. If the major version of your DB engine is no longer supported, you can't restore to that version from a snapshot.

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_CreateSnapshot.html

QUESTION 19

A company wants to automate the creation of secure test databases with random credentials to be stored safely for later use. The credentials should have sufficient information about each test database to initiate a connection and perform automated credential rotations. The credentials should not be logged or stored anywhere in an unencrypted form. Which steps should a Database Specialist take to meet these requirements using an AWS CloudFormation template?

- A. Create the database with the MasterUserName and MasterUserPassword properties set to the default values. Then, create the secret with the user name and password set to the same default values. Add a Secret Target Attachment resource with the SecretId and TargetId properties set to the Amazon Resource Names (ARNs) of the secret and the database. Finally, update the secret's password value with a randomly generated string set by the GenerateSecretString property.
- B. Add a Mapping property from the database Amazon Resource Name (ARN) to the secret ARN. Then, create the secret with a chosen user name and a randomly generated password set by the GenerateSecretString property. Add the database with the MasterUserName and MasterUserPassword properties set to the user name of the secret.
- C. Add a resource of type AWS::SecretsManager::Secret and specify the GenerateSecretString property. Then, define the database user name in the SecureStringTemplate template. Create a resource for the database and reference the secret string for the MasterUserName and MasterUserPassword properties. Then, add a resource of type AWS::SecretsManagerSecretTargetAttachment with the SecretId and TargetId properties set to the Amazon Resource Names (ARNs) of the secret and the database.
- D. Create the secret with a chosen user name and a randomly generated password set by the GenerateSecretString property. Add an SecretTargetAttachment resource with the SecretId property set to the Amazon Resource Name (ARN) of the secret and the TargetId property set to a parameter value matching the desired database ARN. Then, create a database with the MasterUserName and MasterUserPassword properties set to the previously created values in the secret.

Correct Answer: C

Section:

Explanation:

Reference: <https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-secretsmanager-secrettargetattachment.html>

QUESTION 20

A company is going to use an Amazon Aurora PostgreSQL DB cluster for an application backend. The DB cluster contains some tables with sensitive data. A Database Specialist needs to control the access privileges at the table level.

How can the Database Specialist meet these requirements?

- A. Use AWS IAM database authentication and restrict access to the tables using an IAM policy.
- B. Configure the rules in a NACL to restrict outbound traffic from the Aurora DB cluster.
- C. Execute GRANT and REVOKE commands that restrict access to the tables containing sensitive data.
- D. Define access privileges to the tables containing sensitive data in the pg_hba.conf file.

Correct Answer: C

Section:

Explanation:

Reference: <https://aws.amazon.com/blogs/database/managing-postgresql-users-and-roles/>

QUESTION 21

A Database Specialist is working with a company to launch a new website built on Amazon Aurora with several Aurora Replicas. This new website will replace an on-premises website connected to a legacy relational database. Due to stability issues in the legacy database, the company would like to test the resiliency of Aurora.

Which action can the Database Specialist take to test the resiliency of the Aurora DB cluster?

- A. Stop the DB cluster and analyze how the website responds
- B. Use Aurora fault injection to crash the master DB instance

- C. Remove the DB cluster endpoint to simulate a master DB instance failure
- D. Use Aurora Backtrack to crash the DB cluster

Correct Answer: B

Section:

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Managing.FaultInjectionQueries.html>

"You can test the fault tolerance of your Amazon Aurora DB cluster by using fault injection queries.

Fault injection queries are issued as SQL commands to an Amazon Aurora instance and they enable you to schedule a simulated occurrence of one of the following events: A crash of a writer or reader DB instance A failure of an Aurora Replica A disk failure Disk congestion When a fault injection query specifies a crash, it forces a crash of the Aurora DB instance. The other fault injection queries result in simulations of failure events, but don't cause the event to occur.

When you submit a fault injection query, you also specify an amount of time for the failure event simulation to occur for."

QUESTION 22

A company just migrated to Amazon Aurora PostgreSQL from an on-premises Oracle database. After the migration, the company discovered there is a period of time every day around 3:00 PM where the response time of the application is noticeably slower. The company has narrowed down the cause of this issue to the database and not the application.

Which set of steps should the Database Specialist take to most efficiently find the problematic PostgreSQL query?

- A. Create an Amazon CloudWatch dashboard to show the number of connections, CPU usage, and disk space consumption. Watch these dashboards during the next slow period.
- B. Launch an Amazon EC2 instance, and install and configure an open-source PostgreSQL monitoring tool that will run reports based on the output error logs.
- C. Modify the logging database parameter to log all the queries related to locking in the database and then check the logs after the next slow period for this information.
- D. Enable Amazon RDS Performance Insights on the PostgreSQL database. Use the metrics to identify any queries that are related to spikes in the graph during the next slow period.

Correct Answer: D

Section:

Explanation:

www.VCEplus.io

QUESTION 23

A company has a web-based survey application that uses Amazon DynamoDB. During peak usage, when survey responses are being collected, a Database Specialist sees the ProvisionedThroughputExceededException error.

What can the Database Specialist do to resolve this error? (Choose two.)

- A. Change the table to use Amazon DynamoDB Streams
- B. Purchase DynamoDB reserved capacity in the affected Region
- C. Increase the write capacity units for the specific table
- D. Change the table capacity mode to on-demand
- E. Change the table type to throughput optimized

Correct Answer: C, D

Section:

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/switching.capacitymode.html>

QUESTION 24

A company is running a two-tier ecommerce application in one AWS account. The web server is deployed using an Amazon RDS for MySQL Multi-AZ DB instance. A Developer mistakenly deleted the database in the production environment. The database has been restored, but this resulted in hours of downtime and lost revenue.

Which combination of changes in existing IAM policies should a Database Specialist make to prevent an error like this from happening in the future? (Choose three.)

- A. Grant least privilege to groups, users, and roles
- B. Allow all users to restore a database from a backup that will reduce the overall downtime to restore the database
- C. Enable multi-factor authentication for sensitive operations to access sensitive resources and API operations
- D. Use policy conditions to restrict access to selective IP addresses
- E. Use AccessList Controls policy type to restrict users for database instance deletion
- F. Enable AWS CloudTrail logging and Enhanced Monitoring

Correct Answer: A, C, D

Section:

Explanation:

<https://aws.amazon.com/blogs/database/using-iam-multifactor-authentication-with-amazon-rds/>

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/security_iam_id-based-policyhtml

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/DataDurability.html>

QUESTION 25

A company is building a new web platform where user requests trigger an AWS Lambda function that performs an insert into an Amazon Aurora MySQL DB cluster. Initial tests with less than 10 users on the new platform yielded successful execution and fast response times. However, upon more extensive tests with the actual target of 3,000 concurrent users, Lambda functions are unable to connect to the DB cluster and receive too many connections errors.

Which of the following will resolve this issue?

- A. Edit the my.cnf file for the DB cluster to increase max_connections
- B. Increase the instance size of the DB cluster
- C. Change the DB cluster to Multi-AZ
- D. Increase the number of Aurora Replicas

Correct Answer: B

Section:

Explanation:

Max_connection is a formula in RDS parameter group:

```
GREATEST({log(DBInstanceClassMemory/805306368)*45},{log(DBInstanceClassMemory/818728140 8)*1000})
```

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Managing.Performance.html>

You can increase the maximum number of connections to your Aurora MySQL DB instance by scaling the instance up to a DB instance class with more memory, or by setting a larger value for the max_connections parameter in the DB parameter group for your instance, up to 16,000. You must change a larger value for the max_connections parameter in the DB parameter group, not edit my.cnf, it is not physical server hosting MySQL.

QUESTION 26

A company is developing a multi-tier web application hosted on AWS using Amazon Aurora as the database. The application needs to be deployed to production and other non-production environments. A Database Specialist needs to specify different MasterUsername and MasterUserPassword properties in the AWS CloudFormation templates used for automated deployment. The CloudFormation templates are version controlled in the company's code repository.

The company also needs to meet compliance requirement by routinely rotating its database master password for production.

What is most secure solution to store the master password?

- A. Store the master password in a parameter file in each environment. Reference the environment specific parameter file in the CloudFormation template.
- B. Encrypt the master password using an AWS KMS key. Store the encrypted master password in the CloudFormation template.
- C. Use the secretsmanager dynamic reference to retrieve the master password stored in AWS Secrets Manager and enable automatic rotation.
- D. Use the ssm dynamic reference to retrieve the master password stored in the AWS Systems Manager Parameter Store and enable automatic rotation.

Correct Answer: C

Section:

Explanation:

"By using the secure string support in CloudFormation with dynamic references you can better maintain your infrastructure as code. You'll be able to avoid hard coding passwords into your templates and you can keep these runtime configuration parameters separated from your code.

Moreover, when properly used, secure strings will help keep your development and production code as similar as possible, while continuing to make your infrastructure code suitable for continuous deployment pipelines." <https://aws.amazon.com/blogs/mt/using-aws-systems-manager-parameterstore-secure-string-parameters-in-aws-cloudformation-templates/>

<https://aws.amazon.com/blogs/security/how-to-use-aws-secrets-manager-rotate-credentialsamazon-rds-database-types-oracle/>

QUESTION 27

A company is writing a new survey application to be used with a weekly televised game show. The application will be available for 2 hours each week. The company expects to receive over 500,000 entries every week, with each survey asking 2-3 multiple choice questions of each user. A Database Specialist needs to select a platform that is highly scalable for a large number of concurrent writes to handle the anticipated volume.

Which AWS services should the Database Specialist consider? (Choose two.)

- A. Amazon DynamoDB
- B. Amazon Redshift
- C. Amazon Neptune
- D. Amazon Elasticsearch Service
- E. Amazon ElastiCache

www.VCEplus.io

Correct Answer: A, E

Section:

Explanation:

<https://docs.aws.amazon.com/AmazonElastiCache/latest/memug/Strategies.html#Strategies.WriteThrough>

<https://aws.amazon.com/products/databases/real-time-apps-elasticache-for-redis/>

QUESTION 28

A company has migrated a single MySQL database to Amazon Aurora. The production data is hosted in a DB cluster in VPC_PROD, and 12 testing environments are hosted in VPC_TEST using the same AWS account. Testing results in minimal changes to the test data.

The Development team wants each environment refreshed nightly so each test database contains fresh production data every day.

Which migration approach will be the fastest and most cost-effective to implement?

- A. Run the master in Amazon Aurora MySQL. Create 12 clones in VPC_TEST, and script the clones to be deleted and re-created nightly.
- B. Run the master in Amazon Aurora MySQL. Take a nightly snapshot, and restore it into 12 databases in VPC_TEST using Aurora Serverless.
- C. Run the master in Amazon Aurora MySQL. Create 12 Aurora Replicas in VPC_TEST, and script the replicas to be deleted and re-created nightly.
- D. Run the master in Amazon Aurora MySQL using Aurora Serverless. Create 12 clones in VPC_TEST, and script the clones to be deleted and re-created nightly.

Correct Answer: A

Section:

Explanation:

QUESTION 29

A large ecommerce company uses Amazon DynamoDB to handle the transactions on its web portal.

Traffic patterns throughout the year are usually stable; however, a large event is planned. The company knows that traffic will increase by up to 10 times the normal load over the 3-day event.

When sale prices are published during the event, traffic will spike rapidly.

How should a Database Specialist ensure DynamoDB can handle the increased traffic?

- A. Ensure the table is always provisioned to meet peak needs
- B. Allow burst capacity to handle the additional load
- C. Set an AWS Application Auto Scaling policy for the table to handle the increase in traffic
- D. Preprovision additional capacity for the known peaks and then reduce the capacity after the event

Correct Answer: D

Section:

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/bp-partition-keydesign.html#bp-partition-key-throughput-bursting>

"DynamoDB provides some flexibility in your per-partition throughput provisioning by providing burst capacity. Whenever you're not fully using a partition's throughput, DynamoDB reserves a portion of that unused capacity for later bursts of throughput to handle usage spikes. DynamoDB currently retains up to 5 minutes (300 seconds) of unused read and write capacity. During an occasional burst of read or write activity, these extra capacity units can be consumed quickly— even faster than the per-second provisioned throughput capacity that you've defined for your table.

DynamoDB can also consume burst capacity for background maintenance and other tasks without prior notice. Note that these burst capacity details might change in the future."

QUESTION 30

A Database Specialist is migrating an on-premises Microsoft SQL Server application database to Amazon RDS for PostgreSQL using AWS DMS. The application requires minimal downtime when the RDS DB instance goes live.

What change should the Database Specialist make to enable the migration?

- A. Configure the on-premises application database to act as a source for an AWS DMS full load with ongoing change data capture (CDC)
- B. Configure the AWS DMS replication instance to allow both full load and ongoing change data capture (CDC)
- C. Configure the AWS DMS task to generate full logs to allow for ongoing change data capture (CDC)
- D. Configure the AWS DMS connections to allow two-way communication to allow for ongoing change data capture (CDC)

Correct Answer: A

Section:

Explanation:

"requires minimal downtime when the RDS DB instance goes live" in order to do CDC: "you must first ensure that ARCHIVELOG MODE is on to provide information to LogMiner. AWS DMS uses LogMiner to read information from the archive logs so that AWS DMS can capture changes"

<https://docs.aws.amazon.com/dms/latest/sbs/chap-oracle2postgresql.steps.configureoracle.html>

"If you want to capture and apply changes (CDC), then you also need the following privileges."

QUESTION 31

A financial company has allocated an Amazon RDS MariaDB DB instance with large storage capacity to accommodate migration efforts. Post-migration, the company purged unwanted data from the instance. The company now want to downsize storage to save money. The solution must have the least impact on production and near-zero downtime.

Which solution would meet these requirements?

- A. Create a snapshot of the old databases and restore the snapshot with the required storage
- B. Create a new RDS DB instance with the required storage and move the databases from the old instances to the new instance using AWS DMS
- C. Create a new database using native backup and restore
- D. Create a new read replica and make it the primary by terminating the existing primary

Correct Answer: B

Section:

Explanation:

[https://aws.amazon.com/premiumsupport/knowledge-center/rds-db-storage-size/Use AWS Database Migration Service \(AWS DMS\) for minimal downtime.](https://aws.amazon.com/premiumsupport/knowledge-center/rds-db-storage-size/Use-AWS-Database-Migration-Service-(AWS-DMS)-for-minimal-downtime)

QUESTION 32

A large financial services company requires that all data be encrypted in transit. A Developer is attempting to connect to an Amazon RDS DB instance using the company VPC for the first time with credentials provided by a Database Specialist. Other members of the Development team can connect, but this user is consistently receiving an error indicating a communications link failure. The Developer asked the Database Specialist to reset the password a number of times, but the error persists.

Which step should be taken to troubleshoot this issue?

- A. Ensure that the database option group for the RDS DB instance allows ingress from the Developer machine's IP address
- B. Ensure that the RDS DB instance's subnet group includes a public subnet to allow the Developer to connect
- C. Ensure that the RDS DB instance has not reached its maximum connections limit
- D. Ensure that the connection is using SSL and is addressing the port where the RDS DB instance is listening for encrypted connections

Correct Answer: D

Section:

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/SQLServer.Concepts.General.SSL.Using.html>

QUESTION 33

A database specialist is responsible for an Amazon RDS for MySQL DB instance with one read replica. The DB instance and the read replica are assigned to the default parameter group. The database team currently runs test queries against a read replica. The database team wants to create additional tables in the read replica that will only be accessible from the read replica to benefit the tests.

Which should the database specialist do to allow the database team to create the test tables?

- A. Contact AWS Support to disable read-only mode on the read replica. Reboot the read replica. Connect to the read replica and create the tables.
- B. Change the read_only parameter to false (read_only=0) in the default parameter group of the read replica. Perform a reboot without failover. Connect to the read replica and create the tables using the local_only MySQL option.
- C. Change the read_only parameter to false (read_only=0) in the default parameter group. Reboot the read replica. Connect to the read replica and create the tables.
- D. Create a new DB parameter group. Change the read_only parameter to false (read_only=0). Associate the read replica with the new group. Reboot the read replica. Connect to the read replica and create the tables.

Correct Answer: D

Section:

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/rds-read-replica/>

QUESTION 34

A company has a heterogeneous six-node production Amazon Aurora DB cluster that handles online transaction processing (OLTP) for the core business and OLAP reports for the human resources department. To match compute resources to the use case, the company has decided to have the reporting workload for the human resources department be directed to two small nodes in the Aurora DB cluster, while every other workload goes to four large nodes in the same DB cluster. Which option would ensure that the correct nodes are always available for the appropriate workload while meeting these requirements?

- A. Use the writer endpoint for OLTP and the reader endpoint for the OLAP reporting workload.
- B. Use automatic scaling for the Aurora Replica to have the appropriate number of replicas for the desired workload.
- C. Create additional readers to cater to the different scenarios.
- D. Use custom endpoints to satisfy the different workloads.

Correct Answer: D

Section:

Explanation:

<https://aws.amazon.com/about-aws/whats-new/2018/11/amazon-aurora-simplifies-workload-management-with-custom-endpoints/>

You can now create custom endpoints for Amazon Aurora databases. This allows you to distribute and load balance workloads across different sets of database instances in your Aurora cluster. For example, you may provision a set of Aurora Replicas to use an instance type with higher memory capacity in order to run an analytics workload. A custom endpoint can then help you route the analytics workload to these appropriately-configured instances, while keeping other instances in your cluster isolated from this workload. As you add or remove instances from the custom endpoint to match your workload, the endpoint helps spread the load around.

QUESTION 35

Developers have requested a new Amazon Redshift cluster so they can load new third-party marketing data. The new cluster is ready and the user credentials are given to the developers. The developers indicate that their copy jobs fail with the following error message:

"Amazon Invalid operation: S3ServiceException:Access Denied,Status 403,Error AccessDenied." The developers need to load this data soon, so a database specialist must act quickly to solve this issue.

What is the MOST secure solution?

- A. Create a new IAM role with the same user name as the Amazon Redshift developer user ID. Provide the IAM role with read-only access to Amazon S3 with the assume role action.
- B. Create a new IAM role with read-only access to the Amazon S3 bucket and include the assume role action. Modify the Amazon Redshift cluster to add the IAM role.
- C. Create a new IAM role with read-only access to the Amazon S3 bucket with the assume role action. Add this role to the developer IAM user ID used for the copy job that ended with an error message.
- D. Create a new IAM user with access keys and a new role with read-only access to the Amazon S3 bucket. Add this role to the Amazon Redshift cluster. Change the copy job to use the access keys created.

Correct Answer: B

Section:

Explanation:

<https://docs.aws.amazon.com/redshift/latest/gsg/rs-gsg-create-an-iam-role.html>

"Now that you have created the new role, your next step is to attach it to your cluster. You can attach the role when you launch a new cluster or you can attach it to an existing cluster. In the next step, you attach the role to a new cluster."

https://docs.aws.amazon.com/redshift/latest/dg/copy-usage_notes-access-permissions.html

QUESTION 36

A database specialist at a large multi-national financial company is in charge of designing the disaster recovery strategy for a highly available application that is in development. The application uses an Amazon DynamoDB table as its data store. The application requires a recovery time objective (RTO) of 1 minute and a recovery point objective (RPO) of 2 minutes.

Which operationally efficient disaster recovery strategy should the database specialist recommend for the DynamoDB table?

- A. Create a DynamoDB stream that is processed by an AWS Lambda function that copies the data to a DynamoDB table in another Region.
- B. Use a DynamoDB global table replica in another Region. Enable point-in-time recovery for both tables.
- C. Use a DynamoDB Accelerator table in another Region. Enable point-in-time recovery for the table.

D. Create an AWS Backup plan and assign the DynamoDB table as a resource.

Correct Answer: C

Section:

Explanation:

QUESTION 37

A small startup company is looking to migrate a 4 TB on-premises MySQL database to AWS using an Amazon RDS for MySQL DB instance.

Which strategy would allow for a successful migration with the LEAST amount of downtime?

- A. Deploy a new RDS for MySQL DB instance and configure it for access from the on-premises data center. Use the mysqldump utility to create an initial snapshot from the on-premises MySQL server, and copy it to an Amazon S3 bucket. Import the snapshot into the DB instance utilizing the MySQL utilities running on an Amazon EC2 instance. Immediately point the application to the DB instance.
- B. Deploy a new Amazon EC2 instance, install the MySQL software on the EC2 instance, and configure networking for access from the on-premises data center. Use the mysqldump utility to create a snapshot of the on-premises MySQL server. Copy the snapshot into the EC2 instance and restore it into the EC2 MySQL instance. Use AWS DMS to migrate data into a new RDS for MySQL DB instance. Point the application to the DB instance.
- C. Deploy a new Amazon EC2 instance, install the MySQL software on the EC2 instance, and configure networking for access from the on-premises data center. Use the mysqldump utility to create a snapshot of the on-premises MySQL server. Copy the snapshot into an Amazon S3 bucket and import the snapshot into a new RDS for MySQL DB instance using the MySQL utilities running on an EC2 instance. Point the application to the DB instance.
- D. Deploy a new RDS for MySQL DB instance and configure it for access from the on-premises data center. Use the mysqldump utility to create an initial snapshot from the on-premises MySQL server, and copy it to an Amazon S3 bucket. Import the snapshot into the DB instance using the MySQL utilities running on an Amazon EC2 instance. Establish replication into the new DB instance using MySQL replication. Stop application access to the on-premises MySQL server and let the remaining transactions replicate over. Point the application to the DB instance.

Correct Answer: B

Section:

Explanation:

www.VCEplus.io

QUESTION 38

A software development company is using Amazon Aurora MySQL DB clusters for several use cases, including development and reporting. These use cases place unpredictable and varying demands on the Aurora DB clusters, and can cause momentary spikes in latency. System users run ad-hoc queries sporadically throughout the week. Cost is a primary concern for the company, and a solution that does not require significant rework is needed.

Which solution meets these requirements?

- A. Create new Aurora Serverless DB clusters for development and reporting, then migrate to these new DB clusters.
- B. Upgrade one of the DB clusters to a larger size, and consolidate development and reporting activities on this larger DB cluster.
- C. Use existing DB clusters and stop/start the databases on a routine basis using scheduling tools.
- D. Change the DB clusters to the burstable instance family.

Correct Answer: A

Section:

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Concepts.DBInstanceClass.html>

QUESTION 39

A database specialist is building a system that uses a static vendor dataset of postal codes and related territory information that is less than 1 GB in size. The dataset is loaded into the application's cache at start up.

The company needs to store this data in a way that provides the lowest cost with a low application startup time. Which approach will meet these requirements?

- A. Use an Amazon RDS DB instance. Shut down the instance once the data has been read.
- B. Use Amazon Aurora Serverless. Allow the service to spin resources up and down, as needed.
- C. Use Amazon DynamoDB in on-demand capacity mode.
- D. Use Amazon S3 and load the data from flat files.

Correct Answer: D

Section:

Explanation:

<https://www.sumologic.com/insight/s3-cost-optimization/>

For example, for 1 GB file stored on S3 with 1 TB of storage provisioned, you are billed for 1 TB only.

In a lot of other services such as Amazon EC2, Amazon Elastic Block Storage (Amazon EBS) and Amazon DynamoDB you pay for provisioned capacity. For example, in the case of Amazon EBS disk you pay for the size of 1 TB of disk even if you just save 1 GB file. This makes managing S3 cost easier than many other services including Amazon EBS and Amazon EC2. On S3 there is no risk of overprovisioning and no need to manage disk utilization.

QUESTION 40

A database specialist needs to review and optimize an Amazon DynamoDB table that is experiencing performance issues. A thorough investigation by the database specialist reveals that the partition key is causing hot partitions, so a new partition key is created. The database specialist must effectively apply this new partition key to all existing and new data.

How can this solution be implemented?

- A. Use Amazon EMR to export the data from the current DynamoDB table to Amazon S3. Then use Amazon EMR again to import the data from Amazon S3 into a new DynamoDB table with the new partition key.
- B. Use AWS DMS to copy the data from the current DynamoDB table to Amazon S3. Then import the DynamoDB table to create a new DynamoDB table with the new partition key.
- C. Use the AWS CLI to update the DynamoDB table and modify the partition key.
- D. Use the AWS CLI to back up the DynamoDB table. Then use the restore-table-from-backup command and modify the partition key.

Correct Answer: A

Section:

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/back-up-dynamodb-s3/>

QUESTION 41

A company is going through a security audit. The audit team has identified cleartext master user password in the AWS CloudFormation templates for Amazon RDS for MySQL DB instances. The audit team has flagged this as a security risk to the database team.

What should a database specialist do to mitigate this risk?

- A. Change all the databases to use AWS IAM for authentication and remove all the cleartext passwords in CloudFormation templates.
- B. Use an AWS Secrets Manager resource to generate a random password and reference the secret in the CloudFormation template.
- C. Remove the passwords from the CloudFormation templates so Amazon RDS prompts for the password when the database is being created.
- D. Remove the passwords from the CloudFormation template and store them in a separate file. Replace the passwords by running CloudFormation using a sed command.

Correct Answer: B

Section:

Explanation:

<https://aws.amazon.com/blogs/infrastructure-and-automation/securing-passwords-in-aws-quickstarts-using-aws-secrets-manager/>

QUESTION 42

A company's database specialist disabled TLS on an Amazon DocumentDB cluster to perform benchmarking tests. A few days after this change was implemented, a database specialist trainee accidentally deleted multiple tables. The database specialist restored the database from available snapshots. An hour after restoring the cluster, the database specialist is still unable to connect to the new cluster endpoint. What should the database specialist do to connect to the new, restored Amazon DocumentDB cluster?

- A. Change the restored cluster's parameter group to the original cluster's custom parameter group.
- B. Change the restored cluster's parameter group to the Amazon DocumentDB default parameter group.
- C. Configure the interface VPC endpoint and associate the new Amazon DocumentDB cluster.
- D. Run the syncInstances command in AWS DataSync.

Correct Answer: A

Section:

Explanation:

You can't modify the parameter settings of the default parameter groups. You can use a DB parameter group to act as a container for engine configuration values that are applied to one or more DB instances. If you create a DB instance without specifying a DB parameter group, the DB instance uses a default DB parameter group. Each default DB parameter group contains database engine defaults and Amazon RDS system defaults. You can't modify the parameter settings of a default parameter group. Instead, you create your own parameter group where you choose your own parameter settings. Not all DB engine parameters can be changed in a parameter group that you create.

QUESTION 43

A company runs a customer relationship management (CRM) system that is hosted on-premises with a MySQL database as the backend. A custom stored procedure is used to send email notifications to another system when data is inserted into a table. The company has noticed that the performance of the CRM system has decreased due to database reporting applications used by various teams. The company requires an AWS solution that would reduce maintenance, improve performance, and accommodate the email notification feature.

Which AWS solution meets these requirements?

- A. Use MySQL running on an Amazon EC2 instance with Auto Scaling to accommodate the reporting applications. Configure a stored procedure and an AWS Lambda function that uses Amazon SES to send email notifications to the other system.
- B. Use Amazon Aurora MySQL in a multi-master cluster to accommodate the reporting applications. Configure Amazon RDS event subscriptions to publish a message to an Amazon SNS topic and subscribe the other system's email address to the topic.
- C. Use MySQL running on an Amazon EC2 instance with a read replica to accommodate the reporting applications. Configure Amazon SES integration to send email notifications to the other system.
- D. Use Amazon Aurora MySQL with a read replica for the reporting applications. Configure a stored procedure and an AWS Lambda function to publish a message to an Amazon SNS topic. Subscribe the other system's email address to the topic.

Correct Answer: D

Section:

Explanation:

RDS event subscriptions do not cover "data is inserted into a table" - see

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/USER_Events.Messages.html

We can use stored procedure to invoke Lambda function -

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Integrating.Lambda.html>

QUESTION 44

A company needs to migrate Oracle Database Standard Edition running on an Amazon EC2 instance to an Amazon RDS for Oracle DB instance with Multi-AZ. The database supports an ecommerce website that runs

continuously. The company can only provide a maintenance window of up to 5 minutes.
Which solution will meet these requirements?

- A. Configure Oracle Real Application Clusters (RAC) on the EC2 instance and the RDS DB instance. Update the connection string to point to the RAC cluster. Once the EC2 instance and RDS DB instance are in sync, fail over from Amazon EC2 to Amazon RDS.
- B. Export the Oracle database from the EC2 instance using Oracle Data Pump and perform an import into Amazon RDS. Stop the application for the entire process. When the import is complete, change the database connection string and then restart the application.
- C. Configure AWS DMS with the EC2 instance as the source and the RDS DB instance as the destination. Stop the application when the replication is in sync, change the database connection string, and then restart the application.
- D. Configure AWS DataSync with the EC2 instance as the source and the RDS DB instance as the destination. Stop the application when the replication is in sync, change the database connection string, and then restart the application.

Correct Answer: C

Section:

Explanation:

Reference: https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_UpgradeDBInstance.Oracle.html

QUESTION 45

A company is using Amazon Aurora PostgreSQL for the backend of its application. The system users are complaining that the responses are slow. A database specialist has determined that the queries to Aurora take longer during peak times. With the Amazon RDS Performance Insights dashboard, the load in the chart for average active sessions is often above the line that denotes maximum CPU usage and the wait state shows that most wait events are IO:XactSync.

What should the company do to resolve these performance issues?

- A. Add an Aurora Replica to scale the read traffic.
- B. Scale up the DB instance class.
- C. Modify applications to commit transactions in batches.
- D. Modify applications to avoid conflicts by taking locks.

Correct Answer: C

Section:

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.Reference.html>

<https://blog.dbi-services.com/aws-aurora-xactsync-batch-commit/>

QUESTION 46

A database specialist deployed an Amazon RDS DB instance in Dev-VPC1 used by their development team. Dev-VPC1 has a peering connection with Dev-VPC2 that belongs to a different development team in the same department. The networking team confirmed that the routing between VPCs is correct; however, the database engineers in Dev-VPC2 are getting a timeout connections error when trying to connect to the database in Dev- VPC1.

What is likely causing the timeouts?

- A. The database is deployed in a VPC that is in a different Region.
- B. The database is deployed in a VPC that is in a different Availability Zone.
- C. The database is deployed with misconfigured security groups.
- D. The database is deployed with the wrong client connect timeout configuration.

Correct Answer: C

Section:

Explanation:

"A VPC peering connection is a networking connection between two VPCs that enables you to route traffic between them using private IP addresses. Instances in either VPC can communicate with each other as if they are within the same network. You can create a VPC peering connection between your own VPCs, with a VPC in another AWS account, or with a VPC in a different AWS Region."

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_VPC.Scenarios.html

QUESTION 47

A company has a production environment running on Amazon RDS for SQL Server with an in-house web application as the front end. During the last application maintenance window, new functionality was added to the web application to enhance the reporting capabilities for management. Since the update, the application is slow to respond to some reporting queries.

How should the company identify the source of the problem?

- A. Install and configure Amazon CloudWatch Application Insights for Microsoft .NET and Microsoft SQL Server. Use a CloudWatch dashboard to identify the root cause.
- B. Enable RDS Performance Insights and determine which query is creating the problem. Request changes to the query to address the problem.
- C. Use AWS X-Ray deployed with Amazon RDS to track query system traces.
- D. Create a support request and work with AWS Support to identify the source of the issue.

Correct Answer: B

Section:

Explanation:

Amazon RDS Performance Insights is a database performance tuning and monitoring feature that helps you quickly assess the load on your database, and determine when and where to take action. Performance Insights allows non-experts to detect performance problems with an easy-to-understand dashboard that visualizes database load. <https://aws.amazon.com/rds/performanceinsights/>

QUESTION 48

An electric utility company wants to store power plant sensor data in an Amazon DynamoDB table.

The utility company has over 100 power plants and each power plant has over 200 sensors that send data every 2 seconds. The sensor data includes time with milliseconds precision, a value, and a fault attribute if the sensor is malfunctioning. Power plants are identified by a globally unique identifier.

Sensors are identified by a unique identifier within each power plant. A database specialist needs to design the table to support an efficient method of finding all faulty sensors within a given power plant.

Which schema should the database specialist use when creating the DynamoDB table to achieve the fastest query time when looking for faulty sensors?

- A. Use the plant identifier as the partition key and the measurement time as the sort key. Create a global secondary index (GSI) with the plant identifier as the partition key and the fault attribute as the sort key.
- B. Create a composite of the plant identifier and sensor identifier as the partition key. Use the measurement time as the sort key. Create a local secondary index (LSI) on the fault attribute.
- C. Create a composite of the plant identifier and sensor identifier as the partition key. Use the measurement time as the sort key. Create a global secondary index (GSI) with the plant identifier as the partition key and the fault attribute as the sort key.
- D. Use the plant identifier as the partition key and the sensor identifier as the sort key. Create a local secondary index (LSI) on the fault attribute.

Correct Answer: D

Section:

Explanation:

Plant id as partition key and Sensor id as a sort key. Fault can be identified quickly using the local secondary index and associated plant and sensor can be identified easily.

QUESTION 49

A company is releasing a new mobile game featuring a team play mode. As a group of mobile device users play together, an item containing their statuses is updated in an Amazon DynamoDB table.

Periodically, the other users' devices read the latest statuses of their teammates from the table using the BatchGetItem operation.

Prior to launch, some testers submitted bug reports claiming that the status data they were seeing in the game was not up-to-date. The developers are unable to replicate this issue and have asked a database specialist for a recommendation.

Which recommendation would resolve this issue?

- A. Ensure the DynamoDB table is configured to be always consistent.
- B. Ensure the BatchGetItem operation is called with the ConsistentRead parameter set to false.
- C. Enable a stream on the DynamoDB table and subscribe each device to the stream to ensure all devices receive up-to-date status information.
- D. Ensure the BatchGetItem operation is called with the ConsistentRead parameter set to true.

Correct Answer: D

Section:

Explanation:

https://docs.aws.amazon.com/ja_jp/amazondynamodb/latest/developerguide/API_BatchGetItem_v20111205.html

By default, BatchGetItem performs eventually consistent reads on every table in the request. If you want strongly consistent reads instead, you can set ConsistentRead to true for any or all tables.

QUESTION 50

A company is running an Amazon RDS for MySQL Multi-AZ DB instance for a business-critical workload. RDS encryption for the DB instance is disabled. A recent security audit concluded that all business-critical applications must encrypt data at rest. The company has asked its database specialist to formulate a plan to accomplish this for the DB instance.

Which process should the database specialist recommend?

- A. Create an encrypted snapshot of the unencrypted DB instance. Copy the encrypted snapshot to Amazon S3. Restore the DB instance from the encrypted snapshot using Amazon S3.
- B. Create a new RDS for MySQL DB instance with encryption enabled. Restore the unencrypted snapshot to this DB instance.
- C. Create a snapshot of the unencrypted DB instance. Create an encrypted copy of the snapshot. Restore the DB instance from the encrypted snapshot.
- D. Temporarily shut down the unencrypted DB instance. Enable AWS KMS encryption in the AWS Management Console using an AWS managed CMK. Restart the DB instance in an encrypted state.

Correct Answer: C

Section:

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.Encryption.html#Overview.Encryption.Limitations>

QUESTION 51

A company is migrating its on-premises database workloads to the AWS Cloud. A database specialist performing the move has chosen AWS DMS to migrate an Oracle database with a large table to Amazon RDS. The database specialist notices that AWS DMS is taking significant time to migrate the data.

Which actions would improve the data migration speed? (Choose three.)

- A. Create multiple AWS DMS tasks to migrate the large table.
- B. Configure the AWS DMS replication instance with Multi-AZ.
- C. Increase the capacity of the AWS DMS replication server.
- D. Establish an AWS Direct Connect connection between the on-premises data center and AWS.
- E. Enable an Amazon RDS Multi-AZ configuration.
- F. Enable full large binary object (LOB) mode to migrate all LOB data for all large tables.

Correct Answer: C, D, E

Section:

Explanation:

QUESTION 52

A company is migrating a mission-critical 2-TB Oracle database from on premises to Amazon Aurora. The cost for the database migration must be kept to a minimum, and both the on-premises Oracle database and the Aurora DB cluster must remain open for write traffic until the company is ready to completely cut over to Aurora.

Which combination of actions should a database specialist take to accomplish this migration as quickly as possible? (Choose two.)

- A. Use the AWS Schema Conversion Tool (AWS SCT) to convert the source database schema. Then restore the converted schema to the target Aurora DB cluster.
- B. Use Oracle's Data Pump tool to export a copy of the source database schema and manually edit the schema in a text editor to make it compatible with Aurora.
- C. Create an AWS DMS task to migrate data from the Oracle database to the Aurora DB cluster. Select the migration type to replicate ongoing changes to keep the source and target databases in sync until the company is ready to move all user traffic to the Aurora DB cluster.
- D. Create an AWS DMS task to migrate data from the Oracle database to the Aurora DB cluster. Once the initial load is complete, create an AWS Kinesis Data Firehose stream to perform change data capture (CDC) until the company is ready to move all user traffic to the Aurora DB cluster.
- E. Create an AWS Glue job and related resources to migrate data from the Oracle database to the Aurora DB cluster. Once the initial load is complete, create an AWS DMS task to perform change data capture (CDC) until the company is ready to move all user traffic to the Aurora DB cluster.

Correct Answer: A, C

Section:

Explanation:

QUESTION 53

A company has a 20 TB production Amazon Aurora DB cluster. The company runs a large batch job overnight to load data into the Aurora DB cluster. To ensure the company's development team has the most up-to-date data for testing, a copy of the DB cluster must be available in the shortest possible time after the batch job completes.

How should this be accomplished?

- A. Use the AWS CLI to schedule a manual snapshot of the DB cluster. Restore the snapshot to a new DB cluster using the AWS CLI.
- B. Create a dump file from the DB cluster. Load the dump file into a new DB cluster.
- C. Schedule a job to create a clone of the DB cluster at the end of the overnight batch process.
- D. Set up a new daily AWS DMS task that will use cloning and change data capture (CDC) on the DB cluster to copy the data to a new DB cluster. Set up a time for the AWS DMS stream to stop when the new cluster is current.

Correct Answer: C

Section:

Explanation:

QUESTION 54

A company has two separate AWS accounts: one for the business unit and another for corporate analytics. The company wants to replicate the business unit data stored in Amazon RDS for MySQL in us-east-1 to its corporate analytics Amazon Redshift environment in us-west-1. The company wants to use AWS DMS with Amazon RDS as the source endpoint and Amazon Redshift as the target endpoint.

Which action will allow AWS DMS to perform the replication?

- A. Configure the AWS DMS replication instance in the same account and Region as Amazon Redshift.
- B. Configure the AWS DMS replication instance in the same account as Amazon Redshift and in the same Region as Amazon RDS.
- C. Configure the AWS DMS replication instance in its own account and in the same Region as Amazon Redshift.

D. Configure the AWS DMS replication instance in the same account and Region as Amazon RDS.

Correct Answer: A

Section:

Explanation:

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Target.Redshift.html

QUESTION 55

A database specialist is managing an application in the us-west-1 Region and wants to set up disaster recovery in the us-east-1 Region. The Amazon Aurora MySQL DB cluster needs an RPO of 1 minute and an RTO of 2 minutes.

Which approach meets these requirements with no negative performance impact?

- A. Enable synchronous replication.
- B. Enable asynchronous binlog replication.
- C. Create an Aurora Global Database.
- D. Copy Aurora incremental snapshots to the us-east-1 Region.

Correct Answer: C

Section:

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-global-databasedisaster-recovery.html>

QUESTION 56

A gaming company is developing a new mobile game and decides to store the data for each user in Amazon DynamoDB. To make the registration process as easy as possible, users can log in with their existing Facebook or Amazon accounts. The company expects more than 10,000 users.

How should a database specialist implement access control with the LEAST operational effort?

- A. Use web identity federation on the mobile app and AWS STS with an attached IAM role to get temporary credentials to access DynamoDB.
- B. Use web identity federation on the mobile app and create individual IAM users with credentials to access DynamoDB.
- C. Use a self-developed user management system on the mobile app that lets users access the data from DynamoDB through an API.
- D. Use a single IAM user on the mobile app to access DynamoDB.

Correct Answer: A

Section:

Explanation:

Reference: <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/WIF.html>

QUESTION 57

A large retail company recently migrated its three-tier ecommerce applications to AWS. The company's backend database is hosted on Amazon Aurora PostgreSQL. During peak times, users complain about longer page load times. A database specialist reviewed Amazon RDS Performance Insights and found a spike in IO:XactSync wait events. The SQL attached to the wait events are all single INSERT statements.

How should this issue be resolved?

- A. Modify the application to commit transactions in batches

- B. Add a new Aurora Replica to the Aurora DB cluster.
- C. Add an Amazon ElastiCache for Redis cluster and change the application to write through.
- D. Change the Aurora DB cluster storage to Provisioned IOPS (PIOPS).

Correct Answer: A

Section:

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.Reference.html>

"This wait most often arises when there is a very high rate of commit activity on the system. You can sometimes alleviate this wait by modifying applications to commit transactions in batches. "

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/apg-waits.xactsync.html>

QUESTION 58

A company uses Amazon DynamoDB as the data store for its ecommerce website. The website receives little to no traffic at night, and the majority of the traffic occurs during the day. The traffic growth during peak hours is gradual and predictable on a daily basis, but it can be orders of magnitude higher than during off-peak hours.

The company initially provisioned capacity based on its average volume during the day without accounting for the variability in traffic patterns. However, the website is experiencing a significant amount of throttling during peak hours. The company wants to reduce the amount of throttling while minimizing costs.

What should a database specialist do to meet these requirements?

- A. Use reserved capacity. Set it to the capacity levels required for peak daytime throughput.
- B. Use provisioned capacity. Set it to the capacity levels required for peak daytime throughput.
- C. Use provisioned capacity. Create an AWS Application Auto Scaling policy to update capacity based on consumption.
- D. Use on-demand capacity.

Correct Answer: C

Section:

Explanation:

On-demand mode is a good option if any of the following are true: You create new tables with unknown workloads. You have unpredictable application traffic. You prefer the ease of paying for only what you use.

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/HowItWorks.ReadWriteCapacityMode.html#HowItWorks.ProvisionedThroughput.Manual> Amazon DynamoDB auto scaling uses the AWS Application Auto Scaling service to dynamically adjust provisioned throughput capacity on your behalf

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/AutoScaling.html>

QUESTION 59

A company uses an Amazon RDS for PostgreSQL DB instance for its customer relationship management (CRM) system. New compliance requirements specify that the database must be encrypted at rest.

Which action will meet these requirements?

- A. Create an encrypted copy of manual snapshot of the DB instance. Restore a new DB instance from the encrypted snapshot.
- B. Modify the DB instance and enable encryption.
- C. Restore a DB instance from the most recent automated snapshot and enable encryption.
- D. Create an encrypted read replica of the DB instance. Promote the read replica to a standalone instance.

Correct Answer: A

Section:

Explanation:

<https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/encrypt-an-existing-amazonrds-for-postgresql-db-instance.html>

You can enable encryption for an Amazon RDS DB instance when you create it, but not after it's created. However, you can add encryption to an unencrypted DB instance by creating a snapshot of your DB instance, and then creating an encrypted copy of that snapshot. You can then restore a DB instance from the encrypted snapshot to get an encrypted copy of your original DB instance. The pattern uses AWS Database Migration Service (AWS DMS) to migrate data and AWS Key Management Service (AWS KMS) for encryption.

QUESTION 60

A database specialist was alerted that a production Amazon RDS MariaDB instance with 100 GB of storage was out of space. In response, the database specialist modified the DB instance and added 50 GB of storage capacity. Three hours later, a new alert is generated due to a lack of free space on the same DB instance. The database specialist decides to modify the instance immediately to increase its storage capacity by 20 GB.

What will happen when the modification is submitted?

- A. The request will fail because this storage capacity is too large.
- B. The request will succeed only if the primary instance is in active status.
- C. The request will succeed only if CPU utilization is less than 10%.
- D. The request will fail as the most recent modification was too soon.

Correct Answer: D

Section:

Explanation:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_PIOPS.StorageTypes.html

QUESTION 61

A company uses Amazon Aurora for secure financial transactions. The data must always be encrypted at rest and in transit to meet compliance requirements.

Which combination of actions should a database specialist take to meet these requirements?

(Choose two.)

- A. Create an Aurora Replica with encryption enabled using AWS Key Management Service (AWS KMS). Then promote the replica to master.
- B. Use SSL/TLS to secure the in-transit connection between the financial application and the Aurora DB cluster.
- C. Modify the existing Aurora DB cluster and enable encryption using an AWS Key Management Service (AWS KMS) encryption key. Apply the changes immediately.
- D. Take a snapshot of the Aurora DB cluster and encrypt the snapshot using an AWS Key Management Service (AWS KMS) encryption key. Restore the snapshot to a new DB cluster and update the financial application database endpoints.
- E. Use AWS Key Management Service (AWS KMS) to secure the in-transit connection between the financial application and the Aurora DB cluster.

Correct Answer: A, B

Section:

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-replicas-adding.html>

QUESTION 62

A company is running a website on Amazon EC2 instances deployed in multiple Availability Zones (AZs). The site performs a high number of repetitive reads and writes each second on an Amazon RDS for MySQL Multi-AZ DB instance with General Purpose SSD (gp2) storage. After comprehensive testing and analysis, a database specialist discovers that there is high read latency and high CPU utilization on the DB instance. Which approach should the database specialist take to resolve this issue without changing the application?

- A. Implementing sharding to distribute the load to multiple RDS for MySQL databases.
- B. Use the same RDS for MySQL instance class with Provisioned IOPS (PIOPS) storage.

- C. Add an RDS for MySQL read replica.
- D. Modify the RDS for MySQL database class to a bigger size and implement Provisioned IOPS (PIOPS).

Correct Answer: D

Section:

Explanation:

QUESTION 63

A banking company recently launched an Amazon RDS for MySQL DB instance as part of a proof-of-concept project. A database specialist has configured automated database snapshots. As a part of routine testing, the database specialist noticed one day that the automated database snapshot was not created. Which of the following are possible reasons why the snapshot was not created? (Choose two.)

- A. A copy of the RDS automated snapshot for this DB instance is in progress within the same AWS Region.
- B. A copy of the RDS automated snapshot for this DB instance is in progress in a different AWS Region.
- C. The RDS maintenance window is not configured.
- D. The RDS DB instance is in the STORAGE_FULL state.
- E. RDS event notifications have not been enabled.

Correct Answer: A, D

Section:

Explanation:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_WorkingWithAutomatedBackups.html

QUESTION 64

An online shopping company has a large inflow of shopping requests daily. As a result, there is a consistent load on the company's Amazon RDS database. A database specialist needs to ensure the database is up and running at all times.

The database specialist wants an automatic notification system for issues that may cause database downtime or for configuration changes made to the database.

What should the database specialist do to achieve this? (Choose two.)

- A. Create an Amazon CloudWatch Events event to send a notification using Amazon SNS on every API call logged in AWS CloudTrail.
- B. Subscribe to an RDS event subscription and configure it to use an Amazon SNS topic to send notifications.
- C. Use Amazon SES to send notifications based on configured Amazon CloudWatch Events events.
- D. Configure Amazon CloudWatch alarms on various metrics, such as FreeStorageSpace for the RDS instance.
- E. Enable email notifications for AWS Trusted Advisor.

Correct Answer: B, D

Section:

Explanation:

QUESTION 65

A large company has a variety of Amazon DB clusters. Each of these clusters has various configurations that adhere to various requirements. Depending on the team and use case, these configurations can be organized into broader categories.

A database administrator wants to make the process of storing and modifying these parameters more systematic. The database administrator also wants to ensure that changes to individual categories of

configurations are automatically applied to all instances when required.
Which AWS service or feature will help automate and achieve this objective?

- A. AWS Systems Manager Parameter Store
- B. DB parameter group
- C. AWS Config
- D. AWS Secrets Manager

Correct Answer: B

Section:

Explanation:

Reference: https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/USER_WorkingWithParamGroups.html

QUESTION 66

A company is developing a new web application. An AWS CloudFormation template was created as a part of the build process. Recently, a change was made to an AWS::RDS::DBInstance resource in the template. The CharacterSetName property was changed to allow the application to process international text. A change set was generated using the new template, which indicated that the existing DB instance should be replaced during an upgrade. What should a database specialist do to prevent data loss during the stack upgrade?

- A. Create a snapshot of the DB instance. Modify the template to add the DBSnapshotIdentifier property with the ID of the DB snapshot. Update the stack.
- B. Modify the stack policy using the aws cloudformation update-stack command and the set-stackpolicy command, then make the DB resource protected.
- C. Create a snapshot of the DB instance. Update the stack. Restore the database to a new instance.
- D. Deactivate any applications that are using the DB instance. Create a snapshot of the DB instance. Modify the template to add the DBSnapshotIdentifier property with the ID of the DB snapshot. Update the stack and reactivate the applications.

Correct Answer: D

Section:

Explanation:

To preserve your data, perform the following procedure:

1. Deactivate any applications that are using the DB instance so that there's no activity on the DB instance.
2. Create a snapshot of the DB instance. For more information about creating DB snapshots
3. If you want to restore your instance using a DB snapshot, modify the updated template with your DB instance changes and add the DBSnapshotIdentifier property with the ID of the DB snapshot that you want to use
4. Update the stack.

Reference: <https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-properties-rds-database-instance.html>

QUESTION 67

A company recently acquired a new business. A database specialist must migrate an unencrypted 12 TB Amazon RDS for MySQL DB instance to a new AWS account. The database specialist needs to minimize the amount of time required to migrate the database.

Which solution meets these requirements?

- A. Create a snapshot of the source DB instance in the source account. Share the snapshot with the destination account. In the target account, create a DB instance from the snapshot.
- B. Use AWS Resource Access Manager to share the source DB instance with the destination account. Create a DB instance in the destination account using the shared resource.
- C. Create a read replica of the DB instance. Give the destination account access to the read replica. In the destination account, create a snapshot of the shared read replica and provision a new RDS for MySQL DB

instance.

D. Use mysqldump to back up the source database. Create an RDS for MySQL DB instance in the destination account. Use the mysql command to restore the backup in the destination database.

Correct Answer: A

Section:

Explanation:

Sharing an unencrypted manual DB snapshot enables authorized AWS accounts to directly restore a DB instance from the snapshot instead of taking a copy of it and restoring from that.

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ShareSnapshot.html However Resource Access Manager could not share non-Aurora cluster.

<https://docs.aws.amazon.com/ram/latest/userguide/shareable.html>

QUESTION 68

A company has applications running on Amazon EC2 instances in a private subnet with no internet connectivity. The company deployed a new application that uses Amazon DynamoDB, but the application cannot connect to the DynamoDB tables. A developer already checked that all permissions are set correctly.

What should a database specialist do to resolve this issue while minimizing access to external resources?

- A. Add a route to an internet gateway in the subnet's route table.
- B. Add a route to a NAT gateway in the subnet's route table.
- C. Assign a new security group to the EC2 instances with an outbound rule to ports 80 and 443.
- D. Create a VPC endpoint for DynamoDB and add a route to the endpoint in the subnet's route table.

Correct Answer: D

Section:

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/vpc-endpointsdynamodb.html>

www.VCEplus.io

QUESTION 69

The Amazon CloudWatch metric for FreeLocalStorage on an Amazon Aurora MySQL DB instance shows that the amount of local storage is below 10 MB. A database engineer must increase the local storage available in the Aurora DB instance.

How should the database engineer meet this requirement?

- A. Modify the DB instance to use an instance class that provides more local SSD storage.
- B. Modify the Aurora DB cluster to enable automatic volume resizing.
- C. Increase the local storage by upgrading the database engine version.
- D. Modify the DB instance and configure the required storage volume in the configuration section.

Correct Answer: A

Section:

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.AuroraMySQL.Monitoring.Metrics.html>

Unlike for other DB engines, for Aurora DB instances this metric reports the amount of storage available to each DB instance. This value depends on the DB instance class (for pricing information, see the Amazon RDS product page).

You can increase the amount of free storage space for an instance by choosing a larger DB instance class for your instance."

QUESTION 70

A company has an ecommerce web application with an Amazon RDS for MySQL DB instance. The marketing team has noticed some unexpected updates to the product and pricing information on the website,

which is impacting sales targets. The marketing team wants a database specialist to audit future database activity to help identify how and when the changes are being made. What should the database specialist do to meet these requirements? (Choose two.)

- A. Create an RDS event subscription to the audit event type.
- B. Enable auditing of CONNECT and QUERY_DML events.
- C. SSH to the DB instance and review the database logs.
- D. Publish the database logs to Amazon CloudWatch Logs.
- E. Enable Enhanced Monitoring on the DB instance.

Correct Answer: B, D

Section:

Explanation:

<https://aws.amazon.com/blogs/database/configuring-an-audit-log-to-capture-database-activitiesfor-amazon-rds-for-mysql-and-amazon-aurora-with-mysql-compatibility/>

QUESTION 71

A large gaming company is creating a centralized solution to store player session state for multiple online games. The workload required key-value storage with low latency and will be an equal mix of reads and writes. Data should be written into the AWS Region closest to the user across the games' geographically distributed user base. The architecture should minimize the amount of overhead required to manage the replication of data between Regions.

Which solution meets these requirements?

- A. Amazon RDS for MySQL with multi-Region read replicas
- B. Amazon Aurora global database
- C. Amazon RDS for Oracle with GoldenGate
- D. Amazon DynamoDB global tables

www.VCEplus.io

Correct Answer: A

Section:

Explanation:

QUESTION 72

A company is running an on-premises application comprised of a web tier, an application tier, and a MySQL database tier. The database is used primarily during business hours with random activity peaks throughout the day. A database specialist needs to improve the availability and reduce the cost of the MySQL database tier as part of the company's migration to AWS. Which MySQL database option would meet these requirements?

- A. Amazon RDS for MySQL with Multi-AZ
- B. Amazon Aurora Serverless MySQL cluster
- C. Amazon Aurora MySQL cluster
- D. Amazon RDS for MySQL with read replica

Correct Answer: C

Section:

Explanation:

QUESTION 73

A company wants to migrate its Microsoft SQL Server Enterprise Edition database instance from onpremises to AWS. A deep review is performed and the AWS Schema Conversion Tool (AWS SCT) provides options for running this workload on Amazon RDS for SQL Server Enterprise Edition, Amazon RDS for SQL Server Standard Edition, Amazon Aurora MySQL, and Amazon Aurora PostgreSQL. The company does not want to use its own SQL server license and does not want to change from Microsoft SQL Server.

What is the MOST cost-effective and operationally efficient solution?

- A. Run SQL Server Enterprise Edition on Amazon EC2.
- B. Run SQL Server Standard Edition on Amazon RDS.
- C. Run SQL Server Enterprise Edition on Amazon RDS.
- D. Run Amazon Aurora MySQL leveraging SQL Server on Linux compatibility libraries.

Correct Answer: B

Section:

Explanation:

This link seems to indicate that more information is required to determine if the Enterprise instance is a candidate for downgrading to Standard. <https://docs.aws.amazon.com/prescriptiveguidance/latest/patterns/determine-whether-your-microsoft-sql-server-database-can-be-downgraded-from-enterprise-to-standard-edition.html>
<https://calculator.aws/#/createCalculator/RDSSQLServer>

QUESTION 74

A company's ecommerce website uses Amazon DynamoDB for purchase orders. Each order is made up of a Customer ID and an Order ID. The DynamoDB table uses the Customer ID as the partition key and the Order ID as the sort key.

To meet a new requirement, the company also wants the ability to query the table by using a third attribute named Invoice ID. Queries using the Invoice ID must be strongly consistent. A database specialist must provide this capability with optimal performance and minimal overhead.

What should the database administrator do to meet these requirements?

- A. Add a global secondary index on Invoice ID to the existing table.
- B. Add a local secondary index on Invoice ID to the existing table.
- C. Recreate the table by using the latest snapshot while adding a local secondary index on Invoice ID.
- D. Use the partition key and a FilterExpression parameter with a filter on Invoice ID for all queries.

Correct Answer: C

Section:

Explanation:

as Local secondary index can only be created while creating the Dynamodb table. and query needs to use third attribute on top of primary and sort key, so Local Secondary index has primary and sort key as well as the third attribute. Global secondary index can be created without primary and sort key

QUESTION 75

A company wants to migrate its on-premises MySQL databases to Amazon RDS for MySQL. To comply with the company's security policy, all databases must be encrypted at rest. RDS DB instance snapshots must also be shared across various accounts to provision testing and staging environments.

Which solution meets these requirements?

- A. Create an RDS for MySQL DB instance with an AWS Key Management Service (AWS KMS) customer managed CMK. Update the key policy to include the Amazon Resource Name (ARN) of the other AWS accounts as a principal, and then allow the kms:CreateGrant action.
- B. Create an RDS for MySQL DB instance with an AWS managed CMK. Create a new key policy to include the Amazon Resource Name (ARN) of the other AWS accounts as a principal, and then allow the kms:CreateGrant action.

- C. Create an RDS for MySQL DB instance with an AWS owned CMK. Create a new key policy to include the administrator user name of the other AWS accounts as a principal, and then allow the kms:CreateGrant action.
- D. Create an RDS for MySQL DB instance with an AWS CloudHSM key. Update the key policy to include the Amazon Resource Name (ARN) of the other AWS accounts as a principal, and then allow the kms:CreateGrant action.

Correct Answer: A

Section:

Explanation:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ShareSnapshot.html

QUESTION 76

A retail company manages a web application that stores data in an Amazon DynamoDB table. The company is undergoing account consolidation efforts. A database engineer needs to migrate the DynamoDB table from the current AWS account to a new AWS account.

Which strategy meets these requirements with the LEAST amount of administrative work?

- A. Use AWS Glue to crawl the data in the DynamoDB table. Create a job using an available blueprint to export the data to Amazon S3. Import the data from the S3 file to a DynamoDB table in the new account.
- B. Create an AWS Lambda function to scan the items of the DynamoDB table in the current account and write to a file in Amazon S3. Create another Lambda function to read the S3 file and restore the items of a DynamoDB table in the new account.
- C. Use AWS Data Pipeline in the current account to export the data from the DynamoDB table to a file in Amazon S3. Use Data Pipeline to import the data from the S3 file to a DynamoDB table in the new account.
- D. Configure Amazon DynamoDB Streams for the DynamoDB table in the current account. Create an AWS Lambda function to read from the stream and write to a file in Amazon S3. Create another Lambda function to read the S3 file and restore the items to a DynamoDB table in the new account.

Correct Answer: C

Section:

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/dynamodb-cross-account-migration/>

<https://aws.amazon.com/premiumsupport/knowledge-center/data-pipeline-account-accessdynamodb-s3/>

QUESTION 77

A company uses the Amazon DynamoDB table contractDB in us-east-1 for its contract system with the following schema: orderID (primary key) timestamp (sort key) contract (map) createdBy (string) customerEmail (string) After a problem in production, the operations team has asked a database specialist to provide an IAM policy to read items from the database to debug the application. In addition, the developer is not allowed to access the value of the customerEmail field to stay compliant.

Which IAM policy should the database specialist use to achieve these requirements?

A.

www.VCEplus.io

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "IAMPolicy",
      "Effect": "Allow",
      "Action": [
        "dynamodb: Query"
      ],
      "Resource": [
        "arn:aws:dynamodb:us-east-1:123456789012:table/contractDB"
      ],
      "Condition": {
        "ForAllValues:StringLike": {
          "dynamodb:Attributes": [
            "orderId",
            "timestamp",
            "contract",
            "createdBy"
          ]
        },
        "StringEquals": {
          "dynamodb:Select": "SPECIFIC_ATTRIBUTES"
        }
      }
    }
  ]
}
```

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "IAMPolicy",
      "Effect": "Allow",
      "Action": [
        "dynamodb: Query"
      ],
      "Resource": [
        "arn:aws:dynamodb:us-east-1:123456789012:table/contractDB"
      ],
      "Condition": {
        "ForAllValues:StringLike": {
          "dynamodb:Attributes": [
            "orderId",
            "timestamp",
            "contract",
            "createdBy"
          ]
        },
        "StringEquals": {
          "dynamodb:Select": "SPECIFIC_ATTRIBUTES"
        }
      }
    }
  ]
}
```

www.VCEplus.io

B.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "IAMPolicy",
      "Effect": "Allow",
      "Action": [
        "dynamodb: Query"
      ],
      "Resource": [
        "arn:aws:dynamodb:us-east-1:123456789012:table/contractDB"
      ],
      "Condition": {
        "ForAllValues:StringLike": {
          "dynamodb:Attributes": [
            "customerEmail"
          ]
        },
        "StringEquals": {
          "dynamodb:Select": "SPECIFIC_ATTRIBUTES"
        }
      }
    }
  ]
}
```

C.

www.VCEplus.io

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "IAMPolicy",
      "Effect": "Deny",
      "Action": [
        "dynamodb:Query"
      ],
      "Resource": [
        "arn:aws:dynamodb:us-east-1:123456789012:table/contractDB"
      ],
      "Condition": {
        "ForAllValues:StringLike": {
          "dynamodb:Attributes": [
            "customerEmail"
          ]
        },
        "StringEquals": {
          "dynamodb:Select": "SPECIFIC_ATTRIBUTES"
        }
      }
    }
  ]
}
```

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "IAMPolicy",
      "Effect": "Deny",
      "Action": [
        "dynamodb:Query"
      ],
      "Resource": [
        "arn:aws:dynamodb:us-east-1:123456789012:table/contractDB"
      ],
      "Condition": {
        "ForAllValues:StringLike": {
          "dynamodb:Attributes": [
            "customerEmail"
          ]
        },
        "StringEquals": {
          "dynamodb:Select": "SPECIFIC_ATTRIBUTES"
        }
      }
    }
  ]
}
```

www.VCEplus.io

D.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "IAMPolicy",
      "Effect": "Deny",
      "Action": [
        "dynamodb:Query"
      ],
      "Resource": [
        "arn:aws:dynamodb:us-east-1:123456789012:table/contractDB"
      ],
      "Condition": {
        "ForAllValues:StringLike": {
          "dynamodb:Attributes": [
            "orderId",
            "timestamp",
            "contract",
            "createdBy"
          ]
        },
        "StringEquals": {
          "dynamodb:Select": "SPECIFIC_ATTRIBUTES"
        }
      }
    }
  ]
}

```

www.VCEplus.io

Correct Answer: A

Section:

Explanation:

QUESTION 78

A company has an application that uses an Amazon DynamoDB table to store user data. Every morning, a single-threaded process calls the DynamoDB API Scan operation to scan the entire table and generate a critical start-of-day report for management. A successful marketing campaign recently doubled the number of items in the table, and now the process takes too long to run and the report is not generated in time. A database specialist needs to improve the performance of the process. The database specialist notes that, when the process is running, 15% of the table's provisioned read capacity units (RCUs) are being used. What should the database specialist do?

- A. Enable auto scaling for the DynamoDB table.
- B. Use four threads and parallel DynamoDB API Scan operations.
- C. Double the table's provisioned RCUs.
- D. Set the Limit and Offset parameters before every call to the API.

Correct Answer: B

Section:

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Scan.html#Scan.ParallelScan>

QUESTION 79

A company is building a software as a service application. As part of the new user sign-on workflow, a Python script invokes the CreateTable operation using the Amazon DynamoDB API. After the call returns, the script attempts to call PutItem.

Occasionally, the PutItem request fails with a ResourceNotFoundException error, which causes the workflow to fail. The development team has confirmed that the same table name is used in the two API calls. How should a database specialist fix this issue?

- A. Add an allow statement for the dynamodb:PutItem action in a policy attached to the role used by the application creating the table.
- B. Set the StreamEnabled property of the StreamSpecification parameter to true, then call PutItem.
- C. Change the application to call DescribeTable periodically until the TableStatus is ACTIVE, then call PutItem.
- D. Add a ConditionExpression parameter in the PutItem request.

Correct Answer: C

Section:

Explanation:

https://docs.aws.amazon.com/amazondynamodb/latest/APIReference/API_DescribeTable.html

QUESTION 80

To meet new data compliance requirements, a company needs to keep critical data durably stored and readily accessible for 7 years. Data that is more than 1 year old is considered archival data and must automatically be moved out of the Amazon Aurora MySQL DB cluster every week. On average, around 10 GB of new data is added to the database every month. A database specialist must choose the most operationally efficient solution to migrate the archival data to Amazon S3.

Which solution meets these requirements?

- A. Create a custom script that exports archival data from the DB cluster to Amazon S3 using a SQL view, then deletes the archival data from the DB cluster. Launch an Amazon EC2 instance with a weekly cron job to execute the custom script.
- B. Configure an AWS Lambda function that exports archival data from the DB cluster to Amazon S3 using a SELECT INTO OUTFILE S3 statement, then deletes the archival data from the DB cluster. Schedule the Lambda function to run weekly using Amazon EventBridge (Amazon CloudWatch Events).
- C. Configure two AWS Lambda functions: one that exports archival data from the DB cluster to Amazon S3 using the mysqldump utility, and another that deletes the archival data from the DB cluster. Schedule both Lambda functions to run weekly using Amazon EventBridge (Amazon CloudWatch Events).
- D. Use AWS Database Migration Service (AWS DMS) to continually export the archival data from the DB cluster to Amazon S3. Configure an AWS Data Pipeline process to run weekly that executes a custom SQL script to delete the archival data from the DB cluster.

Correct Answer: B

Section:

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Integrating.SaveIntoS3.html>

QUESTION 81

A company developed a new application that is deployed on Amazon EC2 instances behind an Application Load Balancer. The EC2 instances use the security group named sg-application-servers. The company needs a database to store the data from the application and decides to use an Amazon RDS for MySQL DB instance. The DB instance is deployed in private DB subnet.

What is the MOST restrictive configuration for the DB instance security group?

- A. Only allow incoming traffic from the sg-application-servers security group on port 3306.

- B. Only allow incoming traffic from the sg-application-servers security group on port 443.
- C. Only allow incoming traffic from the subnet of the application servers on port 3306.
- D. Only allow incoming traffic from the subnet of the application servers on port 443.

Correct Answer: A

Section:

Explanation:

most restrictive approach is to allow only incoming connections from SG of EC2 instance on port 3306

QUESTION 82

A company is moving its fraud detection application from on premises to the AWS Cloud and is using Amazon Neptune for data storage. The company has set up a 1 Gbps AWS Direct Connect connection to migrate 25 TB of fraud detection data from the on-premises data center to a Neptune DB instance.

The company already has an Amazon S3 bucket and an S3 VPC endpoint, and 80% of the company's network bandwidth is available.

How should the company perform this data load?

- A. Use an AWS SDK with a multipart upload to transfer the data from on premises to the S3 bucket. Use the Copy command for Neptune to move the data in bulk from the S3 bucket to the Neptune DB instance.
- B. Use AWS Database Migration Service (AWS DMS) to transfer the data from on premises to the S3 bucket. Use the Loader command for Neptune to move the data in bulk from the S3 bucket to the Neptune DB instance.
- C. Use AWS DataSync to transfer the data from on premises to the S3 bucket. Use the Loader command for Neptune to move the data in bulk from the S3 bucket to the Neptune DB instance.
- D. Use the AWS CLI to transfer the data from on premises to the S3 bucket. Use the Copy command for Neptune to move the data in bulk from the S3 bucket to the Neptune DB instance.

Correct Answer: C

Section:

Explanation:

"AWS DataSync is an online data transfer service that simplifies, automates, and accelerates moving data between on-premises storage systems and AWS storage services, and also between AWS storage services."
<https://docs.aws.amazon.com/neptune/latest/userguide/bulk-load.html>

QUESTION 83

A company migrated one of its business-critical database workloads to an Amazon Aurora Multi-AZ DB cluster. The company requires a very low RTO and needs to improve the application recovery time after database failovers.

Which approach meets these requirements?

- A. Set the max_connections parameter to 16,000 in the instance-level parameter group.
- B. Modify the client connection timeout to 300 seconds.
- C. Create an Amazon RDS Proxy database proxy and update client connections to point to the proxy endpoint.
- D. Enable the query cache at the instance level.

Correct Answer: C

Section:

Explanation:

Amazon RDS Proxy allows applications to pool and share connections established with the database, improving database efficiency and application scalability. With RDS Proxy, failover times for Aurora and RDS databases are reduced by up to 66% and database credentials, authentication, and access can be managed through integration with AWS Secrets Manager and AWS Identity and Access Management (IAM).

<https://aws.amazon.com/rds/proxy/>

QUESTION 84

A company is using an Amazon RDS for MySQL DB instance for its internal applications. A security audit shows that the DB instance is not encrypted at rest. The company's application team needs to encrypt the DB instance.

What should the team do to meet this requirement?

- A. Stop the DB instance and modify it to enable encryption. Apply this setting immediately without waiting for the next scheduled RDS maintenance window.
- B. Stop the DB instance and create an encrypted snapshot. Restore the encrypted snapshot to a new encrypted DB instance. Delete the original DB instance, and update the applications to point to the new encrypted DB instance.
- C. Stop the DB instance and create a snapshot. Copy the snapshot into another encrypted snapshot. Restore the encrypted snapshot to a new encrypted DB instance. Delete the original DB instance, and update the applications to point to the new encrypted DB instance.
- D. Create an encrypted read replica of the DB instance. Promote the read replica to master. Delete the original DB instance, and update the applications to point to the new encrypted DB instance.

Correct Answer: C

Section:

Explanation:

QUESTION 85

A database specialist must create nightly backups of an Amazon DynamoDB table in a mission-critical workload as part of a disaster recovery strategy.

Which backup methodology should the database specialist use to MINIMIZE management overhead?

- A. Install the AWS CLI on an Amazon EC2 instance. Write a CLI command that creates a backup of the DynamoDB table. Create a scheduled job or task that executes the command on a nightly basis.
- B. Create an AWS Lambda function that creates a backup of the DynamoDB table. Create an Amazon CloudWatch Events rule that executes the Lambda function on a nightly basis.
- C. Create a backup plan using AWS Backup, specify a backup frequency of every 24 hours, and give the plan a nightly backup window.
- D. Configure DynamoDB backup and restore for an on-demand backup frequency of every 24 hours.

Correct Answer: C

Section:

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/CreateBackup.html#:~:text=If%20you%20don%27t%20want%20to%20create%20scheduling%20scripts%20and%20cleanup%20jobs%2C%20you%20can%20use%20AWS%20Backup%20to%20create%20backup%20plans%20with%20schedules%20and%20retention%20policies%20for%20your%20DynamoDB%20tables.%20AWS%20Backup%20runs%20the%20backups%20and%20deletes%20them%20when%20they%20expire.%20For%20more%20information%2C%20see%20the%20AWS%20Backup%20Developer%20Guide.>

https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/backuprestore_HowItWorks.html

QUESTION 86

A company is using a Single-AZ Amazon RDS for MySQL DB instance for development. The DB instance is experiencing slow performance when queries are executed. Amazon CloudWatch metrics indicate that the instance requires more I/O capacity.

Which actions can a database specialist perform to resolve this issue? (Choose two.)

- A. Restart the application tool used to execute queries.
- B. Change to a database instance class with higher throughput.
- C. Convert from Single-AZ to Multi-AZ.
- D. Increase the I/O parameter in Amazon RDS Enhanced Monitoring.
- E. Convert from General Purpose to Provisioned IOPS (PIOPS).

Correct Answer: B, E

Section:

Explanation:

<https://aws.amazon.com/blogs/database/best-storage-practices-for-running-production-workloadson-hosted-databases-with-amazon-rds-or-amazon-ec2/>

"If you find the pattern of IOPS usage consistently going beyond more than 16,000, you should modify the DB instance and change the storage type from gp2 to io1."

QUESTION 87

A company has an AWS CloudFormation template written in JSON that is used to launch new Amazon RDS for MySQL DB instances. The security team has asked a database specialist to ensure that the master password is automatically rotated every 30 days for all new DB instances that are launched using the template.

What is the MOST operationally efficient solution to meet these requirements?

- A. Save the password in an Amazon S3 object. Encrypt the S3 object with an AWS KMS key. Set the KMS key to be rotated every 30 days by setting the EnableKeyRotation property to true. Use a CloudFormation custom resource to read the S3 object to extract the password.
- B. Create an AWS Lambda function to rotate the secret. Modify the CloudFormation template to add an AWS::SecretsManager::RotationSchedule resource. Configure the RotationLambdaARN value and, for the RotationRules property, set the AutomaticallyAfterDays parameter to 30.
- C. Modify the CloudFormation template to use the AWS KMS key as the database password. Configure an Amazon EventBridge rule to invoke the KMS API to rotate the key every 30 days by setting the ScheduleExpression parameter to `*/30*`.
- D. Integrate the Amazon RDS for MySQL DB instances with AWS IAM and centrally manage the master database user password.

Correct Answer: B

Section:

Explanation:

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-secretsmanager-rotationschedule.html>

QUESTION 88

A startup company is building a new application to allow users to visualize their on-premises and cloud networking components. The company expects billions of components to be stored and requires responses in milliseconds. The application should be able to identify:

The networks and routes affected if a particular component fails.

The networks that have redundant routes between them.

The networks that do not have redundant routes between them. The fastest path between two networks.

Which database engine meets these requirements?

- A. Amazon Aurora MySQL
- B. Amazon Neptune
- C. Amazon ElastiCache for Redis
- D. Amazon DynamoDB

Correct Answer: B

Section:

Explanation:

QUESTION 89

An online retail company is planning a multi-day flash sale that must support processing of up to 5,000 orders per second. The number of orders and exact schedule for the sale will vary each day.

During the sale, approximately 10,000 concurrent users will look at the deals before buying items.

Outside of the sale, the traffic volume is very low. The acceptable performance for read/write queries should be under 25 ms. Order items are about 2 KB in size and have a unique identifier. The company requires the most cost-effective solution that will automatically scale and is highly available.

Which solution meets these requirements?

- A. Amazon DynamoDB with on-demand capacity mode
- B. Amazon Aurora with one writer node and an Aurora Replica with the parallel query feature enabled
- C. Amazon DynamoDB with provisioned capacity mode with 5,000 write capacity units (WCUs) and 10,000 read capacity units (RCUs)
- D. Amazon Aurora with one writer node and two cross-Region Aurora Replicas

Correct Answer: A

Section:

Explanation:

The number of orders and exact schedule for the sale will vary each day. During the sale, approximately 10,000 concurrent users will look at the deals before buying items. Outside of the sale, the traffic volume is very low ==> Setting provisioning DynamoDB fix read 5000/write 10000 with will waste the resource when the traffic is low. It is not cost-effective.

QUESTION 90

A ride-hailing application uses an Amazon RDS for MySQL DB instance as persistent storage for bookings. This application is very popular and the company expects a tenfold increase in the user base in next few months. The application experiences more traffic during the morning and evening hours.

This application has two parts:

An in-house booking component that accepts online bookings that directly correspond to simultaneous requests from users.

A third-party customer relationship management (CRM) component used by customer care representatives. The CRM uses queries to access booking data.

A database specialist needs to design a cost-effective database solution to handle this workload.

Which solution meets these requirements?

- A. Use Amazon ElastiCache for Redis to accept the bookings. Associate an AWS Lambda function to capture changes and push the booking data to the RDS for MySQL DB instance used by the CRM.
- B. Use Amazon DynamoDB to accept the bookings. Enable DynamoDB Streams and associate an AWS Lambda function to capture changes and push the booking data to an Amazon SQS queue. This triggers another Lambda function that pulls data from Amazon SQS and writes it to the RDS for MySQL DB instance used by the CRM.
- C. Use Amazon ElastiCache for Redis to accept the bookings. Associate an AWS Lambda function to capture changes and push the booking data to an Amazon Redshift database used by the CRM.
- D. Use Amazon DynamoDB to accept the bookings. Enable DynamoDB Streams and associate an AWS Lambda function to capture changes and push the booking data to Amazon Athena, which is used by the CRM.

Correct Answer: D

Section:

Explanation:

QUESTION 91

Amazon Aurora MySQL is being used by an ecommerce business to migrate its main application database. The firm is now doing OLTP stress testing using concurrent database connections. A database professional detected sluggish performance for several particular write operations during the first round of testing.

Examining the Amazon CloudWatch stats for the Aurora DB cluster revealed a CPU usage of 90%.

Which actions should the database professional take to determine the main cause of excessive CPU use and sluggish performance most effectively? (Select two.)

- A. Enable Enhanced Monitoring at less than 30 seconds of granularity to review the operating system metrics before the next round of tests.
- B. Review the VolumeBytesUsed metric in CloudWatch to see if there is a spike in write I/O.
- C. Review Amazon RDS Performance Insights to identify the top SQL statements and wait events.

- D. Review Amazon RDS API calls in AWS CloudTrail to identify long-running queries.
- E. Enable Advance Auditing to log QUERY events in Amazon CloudWatch before the next round of tests.

Correct Answer: A, C

Section:

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/rds-instance-high-cpu/>
<https://aws.amazon.com/premiumsupport/knowledge-center/rds-mysql-slow-query/>

QUESTION 92

A financial organization must ensure that the most current 90 days of MySQL database backups are accessible. Amazon RDS for MySQL DB instances are used to host all MySQL databases. A database expert must create a solution that satisfies the criteria for backup retention with the least amount of development work feasible.

Which strategy should the database administrator take?

- A. Use AWS Backup to build a backup plan for the required retention period. Assign the DB instances to the backup plan.
- B. Modify the DB instances to enable the automated backup option. Select the required backup retention period.
- C. Automate a daily cron job on an Amazon EC2 instance to create MySQL dumps, transfer to Amazon S3, and implement an S3 Lifecycle policy to meet the retention requirement.
- D. Use AWS Lambda to schedule a daily manual snapshot of the DB instances. Delete snapshots that exceed the retention requirement.

Correct Answer: A

Section:

Explanation:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_WorkingWithAutomatedBackups.html

QUESTION 93

A business is operating an on-premises application that is divided into three tiers: web, application, and MySQL database. The database is predominantly accessed during business hours, with occasional bursts of activity throughout the day.

As part of the company's shift to AWS, a database expert wants to increase the availability and minimize the cost of the MySQL database tier.

Which MySQL database choice satisfies these criteria?

- A. Amazon RDS for MySQL with Multi-AZ
- B. Amazon Aurora Serverless MySQL cluster
- C. Amazon Aurora MySQL cluster
- D. Amazon RDS for MySQL with read replica

Correct Answer: B

Section:

Explanation:

Amazon Aurora Serverless v1 is a simple, cost-effective option for infrequent, intermittent, or unpredictable workloads. <https://aws.amazon.com/rds/aurora/serverless/>

QUESTION 94

A stock market analysis firm maintains two locations: one in the us-east-1 Region and another in the eu-west-2 Region. The business want to build an AWS database solution capable of providing rapid and accurate updates.

Dashboards with advanced analytical queries are used to present data in the eu-west-2 office.

Because the corporation will use these dashboards to make purchasing choices, they must have less than a second to obtain application data. Which solution satisfies these criteria and gives the MOST CURRENT dashboard?

- A. Deploy an Amazon RDS DB instance in us-east-1 with a read replica instance in eu-west-2. Create an Amazon ElastiCache cluster in eu-west-2 to cache data from the read replica to generate the dashboards.
- B. Use an Amazon DynamoDB global table in us-east-1 with replication into eu-west-2. Use multiactive replication to ensure that updates are quickly propagated to eu-west-2.
- C. Use an Amazon Aurora global database. Deploy the primary DB cluster in us-east-1. Deploy the secondary DB cluster in eu-west-2. Configure the dashboard application to read from the secondary cluster.
- D. Deploy an Amazon RDS for MySQL DB instance in us-east-1 with a read replica instance in eu-west-2. Configure the dashboard application to read from the read replica.

Correct Answer: C

Section:

Explanation:

Amazon Aurora global databases span multiple AWS Regions, enabling low latency global reads and providing fast recovery from the rare outage that might affect an entire AWS Region. An Aurora global database has a primary DB cluster in one Region, and up to five secondary DB clusters in different Regions. <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/auroraglobal-database.html>

QUESTION 95

A significant automotive manufacturer is switching a mission-critical finance application's database to Amazon DynamoDB. According to the company's risk and compliance policy, any update to the database must be documented as a log entry for auditing purposes. Each minute, the system anticipates about 500,000 log entries. Log entries should be kept in Apache Parquet files in batches of at least 100,000 records per file. How could a database professional approach these needs while using DynamoDB?

- A. Enable Amazon DynamoDB Streams on the table. Create an AWS Lambda function triggered by the stream. Write the log entries to an Amazon S3 object.
- B. Create a backup plan in AWS Backup to back up the DynamoDB table once a day. Create an AWS Lambda function that restores the backup in another table and compares both tables for changes. Generate the log entries and write them to an Amazon S3 object.
- C. Enable AWS CloudTrail logs on the table. Create an AWS Lambda function that reads the log files once an hour and filters DynamoDB API actions. Write the filtered log files to Amazon S3.
- D. Enable Amazon DynamoDB Streams on the table. Create an AWS Lambda function triggered by the stream. Write the log entries to an Amazon Kinesis Data Firehose delivery stream with buffering and Amazon S3 as the destination.

Correct Answer: D

Section:

Explanation:

QUESTION 96

A major organization maintains a number of Amazon DB clusters. Each of these clusters is configured differently to meet certain needs. These configurations may be classified into wider groups based on the team and use case.

A database administrator wishes to streamline the process of storing and updating these settings.

Additionally, the database administrator wants to guarantee that changes to certain configuration categories are automatically implemented to all instances as necessary.

Which AWS service or functionality will assist in automating and achieving this goal?

- A. AWS Systems Manager Parameter Store
- B. DB parameter group
- C. AWS Config
- D. AWS Secrets Manager

Correct Answer: B

Section:

Explanation:

Database parameters specify how the database is configured. For example, database parameters can specify the amount of resources, such as memory, to allocate to a database.

QUESTION 97

In one AWS account, a business runs a two-tier ecommerce application. An Amazon RDS for MySQL Multi-AZ database instance serves as the application's backend. A developer removed the database instance in the production environment by accident. Although the organization recovers the database, the incident results in hours of outage and financial loss.

Which combination of adjustments would reduce the likelihood that this error will occur again in the future? (Select three.)

- A. Grant least privilege to groups, IAM users, and roles.
- B. Allow all users to restore a database from a backup.
- C. Enable deletion protection on existing production DB instances.
- D. Use an ACL policy to restrict users from DB instance deletion.
- E. Enable AWS CloudTrail logging and Enhanced Monitoring.

Correct Answer: A, C, D

Section:

Explanation:

QUESTION 98

A database specialist needs to configure an Amazon RDS for MySQL DB instance to close noninteractive connections that are inactive after 900 seconds.

What should the database specialist do to accomplish this task?

- A. Create a custom DB parameter group and set the `wait_timeout` parameter value to 900. Associate the DB instance with the custom parameter group.
- B. Connect to the MySQL database and run the `SET SESSION wait_timeout=900` command.
- C. Edit the `my.cnf` file and set the `wait_timeout` parameter value to 900. Restart the DB instance.
- D. Modify the default DB parameter group and set the `wait_timeout` parameter value to 900.

Correct Answer: A

Section:

Explanation:

<https://aws.amazon.com/fr/blogs/database/best-practices-for-configuring-parameters-for-amazonrds-for-mysql-part-3-parameters-related-to-security-operational-manageability-and-connectivity-timeout/>
"You can set parameters globally using a parameter group. Alternatively, you can set them for a particular session using the SET command." <https://aws.amazon.com/blogs/database/best-practicesfor-configuring-parameters-for-amazon-rds-for-mysql-part-1-parameters-related-to-performance/>

QUESTION 99

A database professional maintains a fleet of Amazon RDS database instances that are configured to utilize the default database parameter group. A database expert must connect a custom parameter group with certain database instances.

When will the instances be allocated to this new parameter group once the database specialist performs this change?

- A. Instantaneously after the change is made to the parameter group
- B. In the next scheduled maintenance window of the DB instances
- C. After the DB instances are manually rebooted
- D. Within 24 hours after the change is made to the parameter group

Correct Answer: C

Section:

Explanation:

When you associate a new DB parameter group with a DB instance, the modified static and dynamic parameters are applied only after the DB instance is rebooted.

QUESTION 100

A small startup firm wishes to move a 4 TB MySQL database from on-premises to AWS through an Amazon RDS for MySQL DB instance.

Which migration approach would result in the LEAST amount of downtime?

- A. Deploy a new RDS for MySQL DB instance and configure it for access from the on-premises data center. Use the mysqldump utility to create an initial snapshot from the on-premises MySQL server, and copy it to an Amazon S3 bucket. Import the snapshot into the DB instance utilizing the MySQL utilities running on an Amazon EC2 instance. Immediately point the application to the DB instance.
- B. Deploy a new Amazon EC2 instance, install the MySQL software on the EC2 instance, and configure networking for access from the on-premises data center. Use the mysqldump utility to create a snapshot of the on-premises MySQL server. Copy the snapshot into the EC2 instance and restore it into the EC2 MySQL instance. Use AWS DMS to migrate data into a new RDS for MySQL DB instance. Point the application to the DB instance.
- C. Deploy a new Amazon EC2 instance, install the MySQL software on the EC2 instance, and configure networking for access from the on-premises data center. Use the mysqldump utility to create a snapshot of the on-premises MySQL server. Copy the snapshot into an Amazon S3 bucket and import the snapshot into a new RDS for MySQL DB instance using the MySQL utilities running on an EC2 instance. Point the application to the DB instance.
- D. Deploy a new RDS for MySQL DB instance and configure it for access from the on-premises data center. Use the mysqldump utility to create an initial snapshot from the on-premises MySQL server, and copy it to an Amazon S3 bucket. Import the snapshot into the DB instance using the MySQL utilities running on an Amazon EC2 instance. Establish replication into the new DB instance using MySQL replication. Stop application access to the on-premises MySQL server and let the remaining transactions replicate over. Point the application to the DB instance.

Correct Answer: D

Section:

Explanation:

www.VCEplus.io

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/MySQL.Procedural.Importing.NonRDSRepl.html>

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/MySQL.Procedural.Importing.External.Repl.html>

Repl.html

QUESTION 101

Amazon RDS for Oracle with Transparent Data Encryption is used by a financial services organization (TDE). At all times, the organization is obligated to encrypt its data at rest. The decryption key must be widely distributed, and access to the key must be restricted. The organization must be able to rotate the encryption key on demand to comply with regulatory requirements. If any possible security vulnerabilities are discovered, the organization must be able to disable the key. Additionally, the company's overhead must be kept to a minimal.

What method should the database administrator use to configure the encryption to fulfill these specifications?

- A. AWS CloudHSM
- B. AWS Key Management Service (AWS KMS) with an AWS managed key
- C. AWS Key Management Service (AWS KMS) with server-side encryption
- D. AWS Key Management Service (AWS KMS) CMK with customer-provided material

Correct Answer: D

Section:

Explanation:

<https://docs.aws.amazon.com/whitepapers/latest/kms-best-practices/aws-managed-and-customermanaged-cmks.html>

QUESTION 102

A business is launching a new Amazon RDS for SQL Server database instance. The organization wishes to allow auditing of the SQL Server database. Which measures should a database professional perform in combination to achieve this requirement? (Select two.)

- A. Create a service-linked role for Amazon RDS that grants permissions for Amazon RDS to store audit logs on Amazon S3.
- B. Set up a parameter group to configure an IAM role and an Amazon S3 bucket for audit log storage. Associate the parameter group with the DB instance.
- C. Disable Multi-AZ on the DB instance, and then enable auditing. Enable Multi-AZ after auditing is enabled.
- D. Disable automated backup on the DB instance, and then enable auditing. Enable automated backup after auditing is enabled.
- E. Set up an options group to configure an IAM role and an Amazon S3 bucket for audit log storage. Associate the options group with the DB instance.

Correct Answer: A, E

Section:

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Appendix.SQLServer.Options.Audit.html>

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/security_iam_service-with-iam.html

QUESTION 103

A business maintains a SQL Server database on-premises. Active Directory authentication is used to provide users access to the database. The organization transferred their database successfully to Amazon RDS for SQL Server. The organization, however, has reservations regarding user authentication in the AWS Cloud environment. Which authentication solution should a database professional provide?

- A. Deploy Active Directory Federation Services (AD FS) on premises and configure it with an on-premises Active Directory. Set up delegation between the on-premises AD FS and AWS Security Token Service (AWS STS) to map user identities to a role using the AmazonRDSDirectoryServiceAccess managed IAM policy.
- B. Establish a forest trust between the on-premises Active Directory and AWS Directory Service for Microsoft Active Directory. Use AWS SSO to configure an Active Directory user delegated to access the databases in RDS for SQL Server.
- C. Use Active Directory Connector to redirect directory requests to the company's on-premises Active Directory without caching any information in the cloud. Use the RDS master user credentials to connect to the DB instance and configure SQL Server logins and users from the Active Directory users and groups.
- D. Establish a forest trust between the on-premises Active Directory and AWS Directory Service for Microsoft Active Directory. Ensure RDS for SQL Server is using mixed mode authentication. Use the RDS master user credentials to connect to the DB instance and configure SQL Server logins and users from the Active Directory users and groups.

Correct Answer: D

Section:

Explanation:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_SQLServerWinAuth.html

QUESTION 104

A database expert is responsible for building a highly available online transaction processing (OLTP) solution that makes use of Amazon RDS for MySQL production databases. Disaster recovery criteria include a cross-regional deployment and an RPO and RTO of 5 and 30 minutes, respectively.

What should the database professional do to ensure that the database meets the criteria for high availability and disaster recovery?

- A. Use a Multi-AZ deployment in each Region.
- B. Use read replica deployments in all Availability Zones of the secondary Region.
- C. Use Multi-AZ and read replica deployments within a Region.
- D. Use Multi-AZ and deploy a read replica in a secondary Region.

Correct Answer: D

Section:

QUESTION 105

A financial services organization employs an Amazon Aurora PostgreSQL DB cluster to host an application on AWS. No log files detailing database administrator activity were discovered during a recent examination. A database professional must suggest a solution that enables access to the database and maintains activity logs. The solution should be simple to implement and have a negligible effect on performance. Which database specialist solution should be recommended?

- A. Enable Aurora Database Activity Streams on the database in synchronous mode. Connect the Amazon Kinesis data stream to Kinesis Data Firehose. Set the Kinesis Data Firehose destination to an Amazon S3 bucket.
- B. Create an AWS CloudTrail trail in the Region where the database runs. Associate the database activity logs with the trail.
- C. Enable Aurora Database Activity Streams on the database in asynchronous mode. Connect the Amazon Kinesis data stream to Kinesis Data Firehose. Set the Firehose destination to an Amazon S3 bucket.
- D. Allow connections to the DB cluster through a bastion host only. Restrict database access to the bastion host and application servers. Push the bastion host logs to Amazon CloudWatch Logs using the CloudWatch Logs agent.

Correct Answer: C

Section:

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/DBActivityStreams.Overview.html>

QUESTION 106

A company is running its customer feedback application on Amazon Aurora MySQL. The company runs a report every day to extract customer feedback, and a team reads the feedback to determine if the customer comments are positive or negative. It sometimes takes days before the company can contact unhappy customers and take corrective measures. The company wants to use machine learning to automate this workflow.

Which solution meets this requirement with the LEAST amount of effort?

- A. Export the Aurora MySQL database to Amazon S3 by using AWS Database Migration Service (AWS DMS). Use Amazon Comprehend to run sentiment analysis on the exported files.
- B. Export the Aurora MySQL database to Amazon S3 by using AWS Database Migration Service (AWS DMS). Use Amazon SageMaker to run sentiment analysis on the exported files.
- C. Set up Aurora native integration with Amazon Comprehend. Use SQL functions to extract sentiment analysis.
- D. Set up Aurora native integration with Amazon SageMaker. Use SQL functions to extract sentiment analysis.

Correct Answer: C

Section:

Explanation:

For details about using Aurora and Amazon Comprehend together, see Using Amazon Comprehend for sentiment detection. Aurora machine learning uses a highly optimized integration between the Aurora database and the AWS machine learning (ML) services SageMaker and Amazon Comprehend.

<https://www.stackovercloud.com/2019/11/27/new-for-amazon-aurora-use-machine-learningdirectly-from-your-databases/>

QUESTION 107

Recently, a financial institution created a portfolio management service. The application's backend is powered by Amazon Aurora, which supports MySQL.

The firm demands a response time of five minutes and a response time of five minutes. A database professional must create a disaster recovery system that is both efficient and has a low replication latency. How should the database professional tackle these requirements?

- A. Configure AWS Database Migration Service (AWS DMS) and create a replica in a different AWS Region.

- B. Configure an Amazon Aurora global database and add a different AWS Region.
- C. Configure a binlog and create a replica in a different AWS Region.
- D. Configure a cross-Region read replica.

Correct Answer: B

Section:

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-global-database-disaster-recovery.html>
<https://aws.amazon.com/blogs/database/how-to-choose-the-best-disaster-recovery-option-for-your-amazon-aurora-mysql-cluster/>
<https://aws.amazon.com/about-aws/whats-new/2019/11/aurora-supports-in-place-conversion-to-global-database/>

QUESTION 108

For the first time, a database professional is establishing a test graph database on Amazon Neptune. The database expert must input millions of rows of test observations from an Amazon S3.csv file. The database professional uploaded the data to the Neptune DB instance through a series of API calls. Which sequence of actions enables the database professional to upload the data most quickly? (Select three.)

- A. Ensure Amazon Cognito returns the proper AWS STS tokens to authenticate the Neptune DB instance to the S3 bucket hosting the CSV file.
- B. Ensure the vertices and edges are specified in different .csv files with proper header column formatting.
- C. Use AWS DMS to move data from Amazon S3 to the Neptune Loader.
- D. Curl the S3 URI while inside the Neptune DB instance and then run the addVertex or addEdge commands.
- E. Ensure an IAM role for the Neptune DB instance is configured with the appropriate permissions to allow access to the file in the S3 bucket.
- F. Create an S3 VPC endpoint and issue an HTTP POST to the database's loader endpoint.

Correct Answer: B, E, F

Section:

Explanation:

<https://docs.aws.amazon.com/neptune/latest/userguide/bulk-load-optimize.html>

QUESTION 109

A financial services company has an application deployed on AWS that uses an Amazon Aurora PostgreSQL DB cluster. A recent audit showed that no log files contained database administrator activity. A database specialist needs to recommend a solution to provide database access and activity logs. The solution should use the least amount of effort and have a minimal impact on performance. Which solution should the database specialist recommend?

- A. Enable Aurora Database Activity Streams on the database in synchronous mode. Connect the Amazon Kinesis data stream to Kinesis Data Firehose. Set the Kinesis Data Firehose destination to an Amazon S3 bucket.
- B. Create an AWS CloudTrail trail in the Region where the database runs. Associate the database activity logs with the trail.
- C. Enable Aurora Database Activity Streams on the database in asynchronous mode. Connect the Amazon Kinesis data stream to Kinesis Data Firehose. Set the Firehose destination to an Amazon S3 bucket.
- D. Allow connections to the DB cluster through a bastion host only. Restrict database access to the bastion host and application servers. Push the bastion host logs to Amazon CloudWatch Logs using the CloudWatch Logs agent.

Correct Answer: C

Section:

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/DBActivityStreams.Overview.html>

QUESTION 110

A corporation is transitioning from an IBM Informix database to an Amazon RDS for SQL Server Multi-AZ implementation with Always On Availability Groups (AGs). SQL Server Agent tasks are scheduled to execute at 5-minute intervals on the Always On AG listener to synchronize data between the Informix and SQL Server databases. After a successful failover to the backup node with minimum delay, users endure hours of stale data.

How can a database professional guarantee that consumers view the most current data after a failover?

- A. Set TTL to less than 30 seconds for cached DNS values on the Always On AG listener.
- B. Break up large transactions into multiple smaller transactions that complete in less than 5 minutes.
- C. Set the databases on the secondary node to read-only mode.
- D. Create the SQL Server Agent jobs on the secondary node from a script when the secondary node takes over after a failure.

Correct Answer: D

Section:

Explanation:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_SQLServerMultiAZ.html

If you have SQL Server Agent jobs, recreate them on the secondary. You do so because these jobs are stored in the msdb database, and you can't replicate this database by using Database Mirroring (DBM) or Always On Availability Groups (AGs). Create the jobs first in the original primary, then fail over, and create the same jobs in the new primary.

QUESTION 111

A business is transferring a database from one AWS Region to another using an Amazon RDS for SQL Server DB instance. The organization wishes to keep database downtime to a minimum throughout the transfer. Which migration strategy should the organization use for this cross-regional move?

- A. Back up the source database using native backup to an Amazon S3 bucket in the same Region. Then restore the backup in the target Region.
- B. Back up the source database using native backup to an Amazon S3 bucket in the same Region. Use Amazon S3 Cross-Region Replication to copy the backup to an S3 bucket in the target Region. Then restore the backup in the target Region.
- C. Configure AWS Database Migration Service (AWS DMS) to replicate data between the source and the target databases. Once the replication is in sync, terminate the DMS task.
- D. Add an RDS for SQL Server cross-Region read replica in the target Region. Once the replication is in sync, promote the read replica to master.

Correct Answer: C

Section:

Explanation:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.XRgn.html

With Amazon RDS, you can create a MariaDB, MySQL, Oracle, or PostgreSQL read replica in a different AWS Region from the source DB instance. Creating a cross-Region read replica isn't supported for SQL Server on Amazon RDS.

QUESTION 112

A financial institution uses AWS to host its online application. Amazon RDS for MySQL is used to host the application's database, which includes automatic backups.

The program has corrupted the database logically, resulting in the application being unresponsive.

The exact moment the corruption occurred has been determined, and it occurred within the backup retention period.

How should a database professional restore a database to its previous state prior to corruption?

- A. Use the point-in-time restore capability to restore the DB instance to the specified time. No changes to the application connection string are required.
- B. Use the point-in-time restore capability to restore the DB instance to the specified time. Change the application connection string to the new, restored DB instance.
- C. Restore using the latest automated backup. Change the application connection string to the new, restored DB instance.
- D. Restore using the appropriate automated backup. No changes to the application connection string are required.

Correct Answer: B

Section:

Explanation:

When you perform a restore operation to a point in time or from a DB Snapshot, a new DB Instance is created with a new endpoint (the old DB Instance can be deleted if so desired). This is done to enable you to create multiple DB Instances from a specific DB Snapshot or point in time."

QUESTION 113

An worldwide gaming company's development team is experimenting with using Amazon DynamoDB to store in-game events for three mobile titles. Maximum concurrent users for the most popular game is 500,000, while the least popular game is 10,000. The typical event is 20 KB in size, while the average user session generates one event each second. Each event is assigned a millisecond time stamp and a globally unique identification.

The lead developer generated a single DynamoDB database with the following structure for the events:

Partition key: game name Sort key: event identifier Local secondary index: player identifier Event time In a small-scale development setting, the tests were successful. When the application was deployed to production, however, new events were not being added to the database, and the logs indicated DynamoDB failures with the `ItemCollectionSizeLimitExceededException` issue code.

Which design modification should a database professional offer to the development team?

- A. Use the player identifier as the partition key. Use the event time as the sort key. Add a global secondary index with the game name as the partition key and the event time as the sort key.
- B. Create two tables. Use the game name as the partition key in both tables. Use the event time as the sort key for the first table. Use the player identifier as the sort key for the second table.
- C. Replace the sort key with a compound value consisting of the player identifier collated with the event time, separated by a dash. Add a local secondary index with the player identifier as the sort key.
- D. Create one table for each game. Use the player identifier as the partition key. Use the event time as the sort key.

Correct Answer: D

Section:

Explanation:

www.VCEplus.io

QUESTION 114

A business's mission-critical production workload is being operated on a 500 GB Amazon Aurora MySQL DB cluster. A database engineer must migrate the workload without causing data loss to a new Amazon Aurora Serverless MySQL DB cluster.

Which approach will result in the LEAST amount of downtime and the LEAST amount of application impact?

- A. Modify the existing DB cluster and update the Aurora configuration to Serverless.
- B. Create a snapshot of the existing DB cluster and restore it to a new Aurora Serverless DB cluster.
- C. Create an Aurora Serverless replica from the existing DB cluster and promote it to primary when the replica lag is minimal.
- D. Replicate the data between the existing DB cluster and a new Aurora Serverless DB cluster by using AWS Database Migration Service (AWS DMS) with change data capture (CDC) enabled.

Correct Answer: D

Section:

Explanation:

<https://medium.com/@souri29/how-to-migrate-from-amazon-rds-aurora-or-mysql-to-amazonaurora-serverless-55f9a4a74078>

QUESTION 115

A database professional is developing an application that will respond to single-instance requests.

The program will query large amounts of client data and offer end users with results.

These reports may include a variety of fields. The database specialist want to enable users to query the database using any of the fields offered.

During peak periods, the database's traffic volume will be significant yet changeable. However, the database will see little activity over the rest of the day.

Which approach will be the most cost-effective in meeting these requirements?

- A. Amazon DynamoDB with provisioned capacity mode and auto scaling
- B. Amazon DynamoDB with on-demand capacity mode
- C. Amazon Aurora with auto scaling enabled
- D. Amazon Aurora in a serverless mode

Correct Answer: D

Section:

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Limits.html#limits-items>

QUESTION 116

A huge gaming firm is developing a centralized method for storing the status of various online games' user sessions. The workload requires low-latency key-value storage and will consist of an equal number of reads and writes. Across the games' geographically dispersed user base, data should be written to the AWS Region nearest to the user. The design should reduce the burden associated with managing data replication across Regions.

Which solution satisfies these criteria?

- A. Amazon RDS for MySQL with multi-Region read replicas
- B. Amazon Aurora global database
- C. Amazon RDS for Oracle with GoldenGate
- D. Amazon DynamoDB global tables

Correct Answer: D

Section:

Explanation:

https://aws.amazon.com/dynamodb/?nc1=h_ls

QUESTION 117

A business's production databases are housed on a 3 TB Amazon Aurora MySQL DB cluster. The database cluster is installed in the region us-east-1. For disaster recovery (DR) requirements, the company's database expert needs to fast deploy the DB cluster in another AWS Region to handle the production load with an RTO of less than two hours.

Which approach is the MOST OPERATIONALLY EFFECTIVE in meeting these requirements?

- A. Implement an AWS Lambda function to take a snapshot of the production DB cluster every 2 hours, and copy that snapshot to an Amazon S3 bucket in the DR Region. Restore the snapshot to an appropriately sized DB cluster in the DR Region.
- B. Add a cross-Region read replica in the DR Region with the same instance type as the current primary instance. If the read replica in the DR Region needs to be used for production, promote the read replica to become a standalone DB cluster.
- C. Create a smaller DB cluster in the DR Region. Configure an AWS Database Migration Service (AWS DMS) task with change data capture (CDC) enabled to replicate data from the current production DB cluster to the DB cluster in the DR Region.
- D. Create an Aurora global database that spans two Regions. Use AWS Database Migration Service (AWS DMS) to migrate the existing database to the new global database.

Correct Answer: B

Section:

Explanation:

www.VCEplus.io

RTO is 2 hours. With 3 TB database, cross-region replica is a better option

QUESTION 118

A business's production database is hosted on a single-node Amazon RDS for MySQL DB instance. The database instance is hosted in a United States AWS Region.

A week before a significant sales event, a fresh database maintenance update is released. The maintenance update has been designated as necessary. The firm wants to minimize the database instance's downtime and requests that a database expert make the database instance highly accessible until the sales event concludes.

Which solution will satisfy these criteria?

- A. Defer the maintenance update until the sales event is over.
- B. Create a read replica with the latest update. Initiate a failover before the sales event.
- C. Create a read replica with the latest update. Transfer all read-only traffic to the read replica during the sales event.
- D. Convert the DB instance into a Multi-AZ deployment. Apply the maintenance update.

Correct Answer: D

Section:

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/rds-required-maintenance/>

QUESTION 119

A corporation intends to migrate a 500-GB Oracle database to Amazon Aurora PostgreSQL utilizing the AWS Schema Conversion Tool (AWS SCT) and AWS Data Management Service (AWS DMS). The database does not have any stored procedures, but does contain several huge or partitioned tables.

Because the program is vital to the company, it is preferable to migrate with little downtime.

Which measures should a database professional perform in combination to expedite the transfer process? (Select three.)

- A. Use the AWS SCT data extraction agent to migrate the schema from Oracle to Aurora PostgreSQL.
- B. For the large tables, change the setting for the maximum number of tables to load in parallel and perform a full load using AWS DMS.
- C. For the large tables, create a table settings rule with a parallel load option in AWS DMS, then perform a full load using DMS.
- D. Use AWS DMS to set up change data capture (CDC) for continuous replication until the cutover date.
- E. Use AWS SCT to convert the schema from Oracle to Aurora PostgreSQL.
- F. Use AWS DMS to convert the schema from Oracle to Aurora PostgreSQL and for continuous replication.

Correct Answer: C, D, E

Section:

Explanation:

QUESTION 120

A business needs a data warehouse system that stores data consistently and in a highly organized fashion. The organization demands rapid response times for end-user inquiries including current-year data, and users must have access to the whole 15-year dataset when necessary. Additionally, this solution must be able to manage a variable volume of incoming inquiries. Costs associated with storing the 100 TB of data must be maintained to a minimum.

Which solution satisfies these criteria?

- A. Leverage an Amazon Redshift data warehouse solution using a dense storage instance type while keeping all the data on local Amazon Redshift storage. Provision enough instances to support high demand.
- B. Leverage an Amazon Redshift data warehouse solution using a dense storage instance to store the most recent data. Keep historical data on Amazon S3 and access it using the Amazon Redshift Spectrum layer.

Provision enough instances to support high demand.

- C. Leverage an Amazon Redshift data warehouse solution using a dense storage instance to store the most recent data. Keep historical data on Amazon S3 and access it using the Amazon Redshift Spectrum layer. Enable Amazon Redshift Concurrency Scaling.
- D. Leverage an Amazon Redshift data warehouse solution using a dense storage instance to store the most recent data. Keep historical data on Amazon S3 and access it using the Amazon Redshift Spectrum layer. Leverage Amazon Redshift elastic resize.

Correct Answer: C

Section:

Explanation:

<https://docs.aws.amazon.com/redshift/latest/dg/concurrency-scaling.html>

"With the Concurrency Scaling feature, you can support virtually unlimited concurrent users and concurrent queries, with consistently fast query performance. When concurrency scaling is enabled, Amazon Redshift automatically adds additional cluster capacity when you need it to process an increase in concurrent read queries. Write operations continue as normal on your main cluster. Users always see the most current data, whether the queries run on the main cluster or on a concurrency scaling cluster. You're charged for concurrency scaling clusters only for the time they're in use. For more information about pricing, see Amazon Redshift pricing. You manage which queries are sent to the concurrency scaling cluster by configuring WLM queues. When you enable concurrency scaling for a queue, eligible queries are sent to the concurrency scaling cluster instead of waiting in line."

QUESTION 121

A business is transferring its on-premises database workloads to the Amazon Web Services (AWS) Cloud. A database professional migrating an Oracle database with a huge table to Amazon RDS has picked AWS DMS. The database professional observes that AWS DMS is consuming considerable time migrating the data. Which activities would increase the pace of data migration? (Select three.)

- A. Create multiple AWS DMS tasks to migrate the large table.
- B. Configure the AWS DMS replication instance with Multi-AZ.
- C. Increase the capacity of the AWS DMS replication server.
- D. Establish an AWS Direct Connect connection between the on-premises data center and AWS.
- E. Enable an Amazon RDS Multi-AZ configuration.
- F. Enable full large binary object (LOB) mode to migrate all LOB data for all large tables.

Correct Answer: A, C, D

Section:

Explanation:

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Tasks.LOBSupport.html

QUESTION 122

An internet advertising firm stores its data in an Amazon DynamoDb table. Amazon DynamoDB Streams are enabled on the table, and one of the keys has a global secondary index. The table is encrypted using a customer-managed AWS Key Management Service (AWS KMS) key.

The firm has chosen to grow worldwide and want to duplicate the database using DynamoDB global tables in a new AWS Region.

An administrator observes the following upon review:

No role with the dynamodb: CreateGlobalTable permission exists in the account.

An empty table with the same name exists in the new Region where replication is desired.

A global secondary index with the same partition key but a different sort key exists in the new Region where replication is desired.

Which settings will prevent you from creating a global table or replica in the new Region? (Select two.)

- A. A global secondary index with the same partition key but a different sort key exists in the new Region where replication is desired.
- B. An empty table with the same name exists in the Region where replication is desired.

- C. No role with the dynamodb:CreateGlobalTable permission exists in the account.
- D. DynamoDB Streams is enabled for the table.
- E. The table is encrypted using a KMS customer managed key.

Correct Answer: A, B

Section:

Explanation:

QUESTION 123

In North America, a business launched a mobile game that swiftly expanded to 10 million daily active players. The game's backend is hosted on AWS and makes considerable use of a TTL-configured Amazon DynamoDB table.

When an item is added or changed, its TTL is set to 600 seconds plus the current epoch time. The game logic is reliant on the purging of outdated data in order to compute rewards points properly. At times, items from the table are read that are many hours beyond their TTL expiration.

How should a database administrator resolve this issue?

- A. Use a client library that supports the TTL functionality for DynamoDB.
- B. Include a query filter expression to ignore items with an expired TTL.
- C. Set the ConsistentRead parameter to true when querying the table.
- D. Create a local secondary index on the TTL attribute.

Correct Answer: B

Section:

Explanation:

[www.VCEplus.io](https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/howitworks-ttl.html)

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/howitworks-ttl.html>

QUESTION 124

On a single Amazon RDS DB instance, a business hosts a MySQL database for its ecommerce application. Automatically saving application purchases to the database results in high-volume writes. Employees routinely create purchase reports for the company. The organization wants to boost database performance and minimize downtime associated with upgrade patching.

Which technique will satisfy these criteria with the LEAST amount of operational overhead?

- A. Enable a Multi-AZ deployment of the RDS for MySQL DB instance, and enable Memcached in the MySQL option group.
- B. Enable a Multi-AZ deployment of the RDS for MySQL DB instance, and set up replication to a MySQL DB instance running on Amazon EC2.
- C. Enable a Multi-AZ deployment of the RDS for MySQL DB instance, and add a read replica.
- D. Add a read replica and promote it to an Amazon Aurora MySQL DB cluster master. Then enable Amazon Aurora Serverless.

Correct Answer: C

Section:

Explanation:

QUESTION 125

A worldwide digital advertising corporation collects browser information in order to provide targeted visitors with contextually relevant pictures, websites, and connections. A single page load may create many events, each of which must be kept separately. A single event may have a maximum size of 200 KB and an average size of 10 KB. Each page load requires a query of the user's browsing history in order to deliver suggestions for targeted advertising. The advertising corporation anticipates daily page views of more than 1 billion from people in the United States, Europe, Hong Kong, and India. The information structure

differs according to the event. Additionally, browsing information must be written and read with a very low latency to guarantee that consumers have a positive viewing experience. Which database solution satisfies these criteria?

- A. Amazon DocumentDB
- B. Amazon RDS Multi-AZ deployment
- C. Amazon DynamoDB global table
- D. Amazon Aurora Global Database

Correct Answer: C

Section:

Explanation:

QUESTION 126

A company migrated an on-premises Oracle database to Amazon RDS for Oracle. A database specialist needs to monitor the latency of the database. Which solution will meet this requirement with the LEAST operational overhead?

- A. Publish RDS Performance insights metrics to Amazon CloudWatch. Add AWS CloudTrail filters to monitor database performance
- B. Install Oracle Statspack. Enable the performance statistics feature to collect, store, and display performance data to monitor database performance.
- C. Enable RDS Performance Insights to visualize the database load. Enable Enhanced Monitoring to view how different threads use the CPU
- D. Create a new DB parameter group that includes the AllocatedStorage, DBInstanceClassMemory, and DBInstanceVCPU variables. Enable RDS Performance Insights

Correct Answer: C

Section:

Explanation:

Explanation from Amazon documents: Amazon RDS for Oracle is a fully managed relational database service that supports Oracle Database. Amazon RDS for Oracle provides several features to monitor the performance and health of your database, such as RDS Performance Insights, Enhanced Monitoring, Amazon CloudWatch, and AWS CloudTrail. RDS Performance Insights is a feature that helps you quickly assess the load on your database and determine when and where to take action. RDS Performance Insights displays a dashboard that shows the database load in terms of average active sessions (AAS), which is the average number of sessions that are actively running SQL statements at any given time. RDS Performance Insights also shows the top SQL statements, waits, hosts, and users that are contributing to the database load. Enhanced Monitoring is a feature that provides metrics in real time for the operating system (OS) that your DB instance runs on. Enhanced Monitoring metrics include CPU utilization, memory, file system, disk I/O, network I/O, process list, and thread count. Enhanced Monitoring allows you to view how different threads use the CPU and how much memory each thread consumes. By enabling RDS Performance Insights and Enhanced Monitoring for the RDS for Oracle DB instance, the database specialist can monitor the latency of the database with the least operational overhead. This solution will allow the database specialist to use the RDS console or API to enable these features and view the metrics and dashboards without installing any additional software or tools. This solution will also provide comprehensive and granular information about the database load and resource utilization. Therefore, option C is the correct solution to meet the requirement. Option A is not optimal because publishing RDS Performance Insights metrics to Amazon CloudWatch and adding AWS CloudTrail filters to monitor database performance will incur additional operational overhead and cost. Amazon CloudWatch is a service that collects monitoring and operational data in the form of logs, metrics, and events. AWS CloudTrail is a service that records AWS API calls for your account and delivers log files to you. These services are useful for monitoring performance trends and auditing activities, but they are not necessary for monitoring latency in real time. Option B is not optimal because installing Oracle Statspack and enabling the performance statistics feature will require manual intervention and configuration on the RDS for Oracle DB instance. Oracle Statspack is a tool that collects, stores, and displays performance data for Oracle Database. The performance statistics feature is an option that enables Statspack to collect additional statistics such as wait events, latches, SQL statements, segments, rollback segments, etc. These tools are useful for performance tuning and troubleshooting, but they are not as easy to use as RDS Performance Insights and Enhanced Monitoring. Option D is not relevant because creating a new DB parameter group that includes the AllocatedStorage, DBInstanceClassMemory, and DBInstanceVCPU variables will not help monitor the latency of the database. A DB parameter group is a collection of DB engine configuration values that define how a DB instance operates. The AllocatedStorage parameter specifies the allocated storage size in gibibytes (GiB). The DBInstanceClassMemory parameter specifies the amount of memory available to an instance class in bytes. The DBInstanceVCPU parameter specifies the number of virtual CPUs available to an instance class. These parameters are used to configure the capacity and performance of a DB instance, but they do not provide any monitoring or metrics information. Enabling RDS Performance Insights alone will not provide enough information about the OS-level metrics such as CPU utilization or memory usage.

QUESTION 127

A database specialist needs to enable IAM authentication on an existing Amazon Aurora PostgreSQL DB cluster. The database specialist already has modified the DB cluster settings, has created IAM and database

credentials, and has distributed the credentials to the appropriate users.

What should the database specialist do next to establish the credentials for the users to use to log in to the DB cluster?

- A. Add the users' IAM credentials to the Aurora cluster parameter group.
- B. Run the generate-db-auth-token command with the user names to generate a temporary password for the users.
- C. Add the users' IAM credentials to the default credential profile, Use the AWS Management Console to access the DB cluster.
- D. Use an AWS Security Token Service (AWS STS) token by sending the IAM access key and secret key as headers to the DB cluster API endpoint.

Correct Answer: B

Section:

Explanation:

Explanation from Amazon documents: Amazon Aurora PostgreSQL supports IAM authentication, which is a method of using AWS Identity and Access Management (IAM) to manage database access. IAM authentication allows you to use IAM users and roles to control who can access your Aurora PostgreSQL DB cluster, instead of using a traditional database username and password. IAM authentication also provides more security by using temporary credentials that are automatically rotated. To enable IAM authentication on an existing Aurora PostgreSQL DB cluster, the database specialist needs to do the following: Modify the DB cluster settings to enable IAM database authentication. This can be done using the AWS Management Console, the AWS CLI, or the RDS API. Create IAM and database credentials for each user who needs access to the DB cluster. The IAM credentials consist of an access key ID and a secret access key. The database credentials consist of a database username and an optional password. The IAM credentials and the database username must match. Distribute the IAM and database credentials to the appropriate users. The users must keep their credentials secure and not share them with anyone else. Run the generate-db-auth-token command with the user names to generate a temporary password for the users. This command is part of the AWS CLI and it generates an authentication token that is valid for 15 minutes. The authentication token is a string that has the same format as a password. The users can use this token as their password when they connect to the DB cluster using a SQL client. Therefore, option B is the correct solution to establish the credentials for the users to use to log in to the DB cluster. Option A is incorrect because adding the users' IAM credentials to the Aurora cluster parameter group is not necessary or possible. A cluster parameter group is a collection of DB engine configuration values that define how a DB cluster operates. Option C is incorrect because adding the users' IAM credentials to the default credential profile and using the AWS Management Console to access the DB cluster is not supported or secure. The default credential profile is a file that stores your AWS credentials for use by AWS CLI or SDKs. The AWS Management Console does not allow you to connect to an Aurora PostgreSQL DB cluster using IAM authentication. Option D is incorrect because using an AWS Security Token Service (AWS STS) token by sending the IAM access key and secret key as headers to the DB cluster API endpoint is not supported or secure. AWS STS is a service that enables you to request temporary, limited-privilege credentials for IAM users or federated users. The DB cluster API endpoint is an endpoint that allows you to perform administrative actions on your DB cluster using RDS API calls.

QUESTION 128

A large financial services company uses Amazon ElastiCache for Redis for its new application that has a global user base. A database administrator must develop a caching solution that will be available across AWS Regions and include low-latency replication and failover capabilities for disaster recovery (DR). The company's security team requires the encryption of cross-Region data transfers.

Which solution meets these requirements with the LEAST amount of operational effort?

- A. Enable cluster mode in ElastiCache for Redis. Then create multiple clusters across Regions and replicate the cache data by using AWS Database Migration Service (AWS DMS). Promote a cluster in the failover Region to handle production traffic when DR is required.
- B. Create a global data store in ElastiCache for Redis. Then create replica clusters in two other Regions. Promote one of the replica clusters as primary when DR is required.
- C. Disable cluster mode in ElastiCache for Redis. Then create multiple replication groups across Regions and replicate the cache data by using AWS Database Migration Service (AWS DMS). Promote a replication group in the failover Region to primary when DR is required.
- D. Create a snapshot of ElastiCache for Redis in the primary Region and copy it to the failover Region. Use the snapshot to restore the cluster from the failover Region when DR is required.

Correct Answer: B

Section:

Explanation:

Explanation from Amazon documents: Amazon ElastiCache for Redis is a fully managed in-memory data store that supports Redis, an open source, key-value database. Amazon ElastiCache for Redis provides several features to enhance the performance, availability, scalability, and security of your Redis data, such as cluster mode, global data store, replication groups, snapshots, and encryption. A global data store is a feature that allows you to create a cross-Region read replica of your ElastiCache for Redis cluster. A global data store consists of a primary cluster that is replicated across up to two other Regions as secondary clusters. A global data store provides low-latency reads and high availability for your Redis data across Regions. A global data store also supports encryption of cross-Region data transfers using AWS Key Management Service (AWS KMS). To create a global data store in ElastiCache for Redis, you need to do the following: Create a primary cluster in one Region. You can use an existing cluster or create a new one. The cluster must have cluster mode enabled and use Redis engine version 5.0.6 or later. Create a global data store and add the primary cluster to it. You can use the AWS Management Console, the AWS CLI, or the ElastiCache API to

create a global datastore. Create one or two secondary clusters in other Regions and add them to the global datastore. The secondary clusters must have the same specifications as the primary cluster, such as node type, number of shards, and number of replicas per shard. Enable encryption in transit and at rest for the primary and secondary clusters. Specify a customer master key (CMK) from AWS KMS for each cluster. By creating a global datastore in ElastiCache for Redis and creating replica clusters in two other Regions, the database administrator can develop a caching solution that will be available across Regions and include low-latency replication and failover capabilities for DR. This solution will also meet the security requirement of encrypting cross-Region data transfers using AWS KMS. This solution will also require the least amount of operational effort, as it does not involve any data migration or manual intervention. Therefore, option B is the correct solution to meet the requirements. Option A is not optimal because enabling cluster mode in ElastiCache for Redis and creating multiple clusters across Regions will not provide cross-Region replication or failover capabilities. Using AWS DMS to replicate the cache data by using AWS DMS will incur additional time and cost, and may not support encryption of cross-Region data transfers. Option C is not optimal because disabling cluster mode in ElastiCache for Redis and creating multiple replication groups across Regions will not provide cross-Region replication or failover capabilities. Using AWS DMS to replicate the cache data will incur additional time and cost, and may not support encryption of cross-Region data transfers. Option D is not optimal because creating a snapshot of ElastiCache for Redis in the primary Region and copying it to the failover Region will not provide low-latency replication or high availability for the Redis data across Regions. Using the snapshot to restore the cluster from the failover Region when DR is required will involve manual intervention and downtime.

QUESTION 129

A company uses an Amazon Redshift cluster to run its analytical workloads. Corporate policy requires that the company's data be encrypted at rest with customer managed keys. The company's disaster recovery plan requires that backups of the cluster be copied into another AWS Region on a regular basis.

How should a database specialist automate the process of backing up the cluster data in compliance with these policies?

- A. Copy the AWS Key Management Service (AWS KMS) customer managed key from the source Region to the destination Region. Set up an AWS Glue job in the source Region to copy the latest snapshot of the Amazon Redshift cluster from the source Region to the destination Region. Use a time-based schedule in AWS Glue to run the job on a daily basis.
- B. Create a new AWS Key Management Service (AWS KMS) customer managed key in the destination Region. Create a snapshot copy grant in the destination Region specifying the new key. In the source Region, configure cross-Region snapshots for the Amazon Redshift cluster specifying the destination Region, the snapshot copy grant, and retention periods for the snapshot.
- C. Copy the AWS Key Management Service (AWS KMS) customer-managed key from the source Region to the destination Region. Create Amazon S3 buckets in each Region using the keys from their respective Regions. Use Amazon EventBridge (Amazon CloudWatch Events) to schedule an AWS Lambda function in the source Region to copy the latest snapshot to the S3 bucket in that Region. Configure S3 Cross-Region Replication to copy the snapshots to the destination Region, specifying the source and destination KMS key IDs in the replication configuration.
- D. Use the same customer-supplied key materials to create a CMK with the same private key in the destination Region. Configure cross-Region snapshots in the source Region targeting the destination Region. Specify the corresponding CMK in the destination Region to encrypt the snapshot.

Correct Answer: B

Section:

Explanation:

Create a new AWS Key Management Service (AWS KMS) Explanation from Amazon documents: Amazon Redshift supports encryption at rest using AWS Key Management Service (AWS KMS) customer master keys (CMKs). To copy encrypted snapshots across Regions, you need to create a snapshot copy grant in the destination Region and specify a CMK in that Region. You also need to configure cross-Region snapshots in the source Region and provide the destination Region, the snapshot copy grant, and retention periods for the snapshots. This way, you can automate the process of backing up the cluster data in compliance with the corporate policies. Option A is incorrect because you cannot copy a CMK from one Region to another. You can only import key material from an external source into a CMK in a specific Region. Option C is incorrect because it involves unnecessary steps of copying snapshots to S3 buckets and using S3 Cross-Region Replication. Option D is incorrect because it is not possible to create a CMK with the same private key as another CMK in a different Region. You can only use customer-supplied key material to create a CMK with a specific key ID in a specific Region.