



Amazon

Exam Questions AWS-Certified-Database-Specialty

AWS Certified Database - Specialty

About ExamBible

Your Partner of IT Exam

Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1

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.
- E. Then use the restore-table-from-backup command and modify the partition key.

Answer: A

Explanation:

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

NEW QUESTION 2

A database professional is tasked with the task of migrating 25 GB of data files from an on-premises storage system to an Amazon Neptune database. Which method of data loading is the FASTEST?

- A. Upload the data to Amazon S3 and use the Loader command to load the data from Amazon S3 into the Neptune database.
- B. Write a utility to read the data from the on-premises storage and run INSERT statements in a loop to load the data into the Neptune database.
- C. Use the AWS CLI to load the data directly from the on-premises storage into the Neptune database.
- D. Use AWS DataSync to load the data directly from the on-premises storage into the Neptune database.

Answer: A

Explanation:

- * 1. Copy the data files to an Amazon Simple Storage Service (Amazon S3) bucket.
- * 2. Create an IAM role with Read and List access to the bucket.
- * 3. Create an Amazon S3 VPC endpoint.
- * 4. Start the Neptune loader by sending a request via HTTP to the Neptune DB instance.
- * 5. The Neptune DB instance assumes the IAM role to load the data from the bucket.

NEW QUESTION 3

A Database Specialist is setting up a new Amazon Aurora DB cluster with one primary instance and three Aurora Replicas for a highly intensive, business-critical application. The Aurora DB cluster has one medium-sized primary instance, one large-sized replica, and two medium-sized replicas. The Database Specialist did not assign a promotion tier to the replicas.

In the event of a primary failure, what will occur?

- A. Aurora will promote an Aurora Replica that is of the same size as the primary instance
- B. Aurora will promote an arbitrary Aurora Replica
- C. Aurora will promote the largest-sized Aurora Replica
- D. Aurora will not promote an Aurora Replica

Answer: C

Explanation:

Priority: If you don't select a value, the default is tier-1. This priority determines the order in which Aurora

https://docs.amazonaws.cn/en_us/AmazonRDS/latest/AuroraUserGuide/aurora-replicas-adding.html

More than one Aurora Replica can share the same priority, resulting in promotion tiers. If two or more Aurora Replicas share the same priority, then Amazon RDS promotes the replica that is largest in size. If two or more Aurora Replicas share the same priority and size, then Amazon RDS promotes an arbitrary replica in the same promotion tier.

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Managing.Backups.html#Aurora.M> If two or more Aurora Replicas share the same priority, then Amazon RDS promotes the replica that is largest in size. If two or more Aurora Replicas share the same priority and size, then Amazon RDS promotes an arbitrary replica in the same promotion tier. <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Concepts.AuroraHighAvailability.html>

NEW QUESTION 4

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.
- B. Restore the snapshot to a new DB cluster using the AWS CLI.
- C. Create a dump file from the DB cluster.
- D. Load the dump file into a new DB cluster.
- E. Schedule a job to create a clone of the DB cluster at the end of the overnight batch process.
- F. 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.
- G. Set up a time for the AWS DMS stream to stop when the new cluster is current.

Answer: C

NEW QUESTION 5

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

Answer: D

Explanation:

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

NEW QUESTION 6

A business that specializes in internet advertising is developing an application that will show adverts to its customers. The program stores data in an Amazon DynamoDB database. Additionally, the application caches its reads using a DynamoDB Accelerator (DAX) cluster. The majority of reads come via the `GetItem` and `BatchGetItem` queries. The application does not need consistency of readings.

The application cache does not behave as intended after deployment. Specific extremely consistent queries to the DAX cluster are responding in several milliseconds rather than microseconds.

How can the business optimize cache behavior in order to boost application performance?

- A. Increase the size of the DAX cluster.
- B. Configure DAX to be an item cache with no query cache
- C. Use eventually consistent reads instead of strongly consistent reads.
- D. Create a new DAX cluster with a higher TTL for the item cache.

Answer: C

NEW QUESTION 7

An AWS CloudFormation stack that included an Amazon RDS DB instance was accidentally deleted and recent data was lost. A Database Specialist needs to add RDS settings to the CloudFormation template to reduce the chance of accidental instance data loss in the future.

Which settings will meet this requirement? (Choose three.)

- A. Set `DeletionProtection` to True
- B. Set `MultiAZ` to True
- C. Set `TerminationProtection` to True
- D. Set `DeleteAutomatedBackups` to False
- E. Set `DeletionPolicy` to Delete
- F. Set `DeletionPolicy` to Retain

Answer: ACF

NEW QUESTION 8

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

Answer: B

Explanation:

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

NEW QUESTION 9

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
- B. Reboot the read replica
- C. Connect to the read replica and create the tables.

- D. Change the read_only parameter to false (read_only=0) in the default parameter group of the read replic
- E. Perform a reboot without failove
- F. Connect to the read replica and create the tables using the local_only MySQL option.
- G. Change the read_only parameter to false (read_only=0) in the default parameter grou
- H. Reboot the read replic
- I. Connect to the read replica and create the tables.
- J. Create a new DB parameter grou
- K. Change the read_only parameter to false (read_only=0). Associate the read replica with the new grou
- L. Reboot the read replic
- M. Connect to the read replica and create the tables.

Answer: D

Explanation:

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

NEW QUESTION 10

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 value
- B. Then, create the secret with the user name and password set to the same default value
- C. Add aSecret Target Attachment resource with the SecretId and TargetId properties set to the Amazon Resource Names (ARNs) of the secret and the databas
- D. Finally, update the secret's password value with a randomly generated string set by the GenerateSecretString property.
- E. Add a Mapping property from the database Amazon Resource Name (ARN) to the secret AR
- F. Then, create the secret with a chosen user name and a randomly generated password set by the GenerateSecretString propert
- G. Add the database with the MasterUserName and MasterUserPassword properties set to the user name of the secret.
- H. Add a resource of type AWS::SecretsManager::Secret and specify the GenerateSecretString property. Then, define the database user name in the SecureStringTemplat
- I. Create a resource for the database and reference the secret string for the MasterUserName and MasterUserPassword propertie
- J. 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.
- K. Create the secret with a chosen user name and a randomly generated password set by the GenerateSecretString propert
- L. 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 AR
- M. Then, create a database with the MasterUserName and MasterUserPassword properties set to the previously created values in the secret.

Answer: C

NEW QUESTION 10

Recently, an ecommerce business transferred one of its SQL Server databases to an Amazon RDS for SQL Server Enterprise Edition database instance. The corporation anticipates an increase in read traffic as a result of an approaching sale. To accommodate the projected read load, a database professional must establish a read replica of the database instance.

Which procedures should the database professional do prior to establishing the read replica? (Select two.)

- A. Identify a potential downtime window and stop the application calls to the source DB instance.
- B. Ensure that automatic backups are enabled for the source DB instance.
- C. Ensure that the source DB instance is a Multi-AZ deployment with Always ON Availability Groups.
- D. Ensure that the source DB instance is a Multi-AZ deployment with SQL Server Database Mirroring(DBM).
- E. Modify the read replica parameter group setting and set the value to 1.

Answer: BC

Explanation:

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

NEW QUESTION 12

A company is hosting critical business data in an Amazon Redshift cluster. Due to the sensitive nature of the data, the cluster is encrypted at rest using AWS KMS. As a part of disaster recovery requirements, the company needs to copy the Amazon Redshift snapshots to another Region.

Which steps should be taken in the AWS Management Console to meet the disaster recovery requirements?

- A. Create a new KMS customer master key in the source Regio
- B. Switch to the destination Region, enable Amazon Redshift cross-Region snapshots, and use the KMS key of the source Region.
- C. Create a new IAM role with access to the KMS ke
- D. Enable Amazon Redshift cross-Region replication using the new IAM role, and use the KMS key of the source Region.
- E. Enable Amazon Redshift cross-Region snapshots in the source Region, and create a snapshot copy grant and use a KMS key in the destination Region.
- F. Create a new KMS customer master key in the destination Region and create a new IAM role with access to the new KMS ke
- G. Enable Amazon Redshift cross-Region replication in the source Region and use the KMS key of the destination Region.

Answer: C

Explanation:

If you want to enable cross-Region snapshot copy for an AWS KMS–encrypted cluster, you must configure a snapshot copy grant for a root key in the destination AWS Region Source-Region : configure a cross-Region snapshot for an AWS KMS–encrypted cluster In Destination AWS Region : choose the AWS Region to which to copy snapshots.

<https://docs.aws.amazon.com/redshift/latest/mgmt/managing-snapshots-console.html#xregioncopy-kms-encrypt>

NEW QUESTION 15

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 CM
- B. 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.
- C. Create an RDS for MySQL DB instance with an AWS managed CM
- D. 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.
- E. Create an RDS for MySQL DB instance with an AWS owned CM
- F. 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.
- G. Create an RDS for MySQL DB instance with an AWS CloudHSM ke
- H. 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.

Answer: A

Explanation:

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

NEW QUESTION 19

A company has a database monitoring solution that uses Amazon CloudWatch for its Amazon RDS for SQL Server environment. The cause of a recent spike in CPU utilization was not determined using the standard metrics that were collected. The CPU spike caused the application to perform poorly, impacting users. A Database Specialist needs to determine what caused the CPU spike.

Which combination of steps should be taken to provide more visibility into the processes and queries running during an increase in CPU load? (Choose two.)

- A. Enable Amazon CloudWatch Events and view the incoming T-SQL statements causing the CPU to spike.
- B. Enable Enhanced Monitoring metrics to view CPU utilization at the RDS SQL Server DB instance level.
- C. Implement a caching layer to help with repeated queries on the RDS SQL Server DB instance.
- D. Use Amazon QuickSight to view the SQL statement being run.
- E. Enable Amazon RDS Performance Insights to view the database load and filter the load by waits, SQL statements, hosts, or users.

Answer: BE

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/rds-instance-high-cpu/> "Several factors can cause an increase in CPU utilization. For example, user-initiated heavy workloads, analytic queries, prolonged deadlocks and lock waits, multiple concurrent transactions, long-running transactions, or other processes that utilize CPU resources. First, you can identify the source of the CPU usage by: Using Enhanced Monitoring Using Performance Insights"

NEW QUESTION 23

The Security team for a finance company was notified of an internal security breach that happened 3 weeks ago. A Database Specialist must start producing audit logs out of the production Amazon Aurora PostgreSQL cluster for the Security team to use for monitoring and alerting. The Security team is required to perform real-time alerting and monitoring outside the Aurora DB cluster and wants to have the cluster push encrypted files to the chosen solution.

Which approach will meet these requirements?

- A. Use pg_audit to generate audit logs and send the logs to the Security team.
- B. Use AWS CloudTrail to audit the DB cluster and the Security team will get data from Amazon S3.
- C. Set up database activity streams and connect the data stream from Amazon Kinesis to consumer applications.
- D. Turn on verbose logging and set up a schedule for the logs to be dumped out for the Security team.

Answer: C

Explanation:

<https://aws.amazon.com/about-aws/whats-new/2019/05/amazon-aurora-with-postgresql-compatibility-supports-> "Database Activity Streams for Amazon Aurora with PostgreSQL compatibility provides a near real-time data stream of the database activity in your relational database to help you monitor activity. When integrated with third party database activity monitoring tools, Database Activity Streams can monitor and audit database activity to provide safeguards for your database and help meet compliance and regulatory requirements."

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

NEW QUESTION 28

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.

Answer: A

Explanation:

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

NEW QUESTION 29

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

Answer: D

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 GB 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 over-provisioning and no need to manage disk utilization.

NEW QUESTION 33

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

Answer: D

Explanation:

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

NEW QUESTION 38

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.

Answer: D

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.

NEW QUESTION 42

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

Answer: D

Explanation:

<https://aws.amazon.com/blogs/database/amazon-dynamodb-gaming-use-cases-and-design-patterns/> "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."

NEW QUESTION 47

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
- B. Create an encrypted copy of the unencrypted snapshot
- C. Restore the encrypted snapshot copy.
- D. Modify the RDS DB instance
- E. Enable the AWS KMS encryption option that leverages the AWS CLI.
- F. Restore an unencrypted snapshot into a MySQL RDS DB instance that is encrypted.
- G. Create an encrypted read replica of the RDS DB instance
- H. Promote it the master.

Answer: A

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>

NEW QUESTION 49

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
- B. Once the EC2 instance and RDS DB instance are in sync, fail over from Amazon EC2 to Amazon RDS.
- C. Export the Oracle database from the EC2 instance using Oracle Data Pump and perform an import into Amazon RDS
- D. Stop the application for the entire process
- E. When the import is complete, change the database connection string and then restart the application.
- F. 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.
- G. Configure AWS DataSync with the EC2 instance as the source and the RDS DB instance as the destination
- H. Stop the application when the replication is in sync, change the database connection string, and then restart the application.

Answer: C

NEW QUESTION 50

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

Answer: A

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/global-database/?nc1=h_ls. Aurora Global Database lets you easily scale database reads across the world and place your applications close to your users.

NEW QUESTION 52

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

Answer: D

Explanation:

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

NEW QUESTION 55

The Development team recently executed a database script containing several data definition language (DDL) and data manipulation language (DML) statements on an Amazon Aurora MySQL DB cluster. The release accidentally deleted thousands of rows from an important table and broke some application functionality. This was discovered 4 hours after the release. Upon investigation, a Database Specialist tracked the issue to a DELETE command in the script with an incorrect WHERE clause filtering the wrong set of rows.

The Aurora DB cluster has Backtrack enabled with an 8-hour backtrack window. The Database Administrator also took a manual snapshot of the DB cluster before the release started. The database needs to be returned to the correct state as quickly as possible to resume full application functionality. Data loss must be minimal. How can the Database Specialist accomplish this?

- A. Quickly rewind the DB cluster to a point in time before the release using Backtrack.
- B. Perform a point-in-time recovery (PITR) of the DB cluster to a time before the release and copy the deleted rows from the restored database to the original database.
- C. Restore the DB cluster using the manual backup snapshot created before the release and change the application configuration settings to point to the new DB cluster.
- D. Create a clone of the DB cluster with Backtrack enable
- E. Rewind the cloned cluster to a point in time before the release
- F. Copy deleted rows from the clone to the original database.

Answer: A

NEW QUESTION 56

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

Answer: ACD

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-policy.html

NEW QUESTION 57

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.

Answer: B

Explanation:

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

NEW QUESTION 61

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
- C. Enable point-in-time recovery for both tables.
- D. Use a DynamoDB Accelerator table in another Region
- E. Enable point-in-time recovery for the table.
- F. Create an AWS Backup plan and assign the DynamoDB table as a resource.

Answer: C

NEW QUESTION 62

A company wants to migrate its existing on-premises Oracle database to Amazon Aurora PostgreSQL. The migration must be completed with minimal downtime using AWS DMS. A Database Specialist must validate that the data was migrated accurately from the source to the target before the cutover. The migration must have minimal impact on the performance of the source database.

Which approach will MOST effectively meet these requirements?

- A. Use the AWS Schema Conversion Tool (AWS SCT) to convert source Oracle database schemas to the target Aurora DB cluster
- B. Verify the datatype of the columns.
- C. Use the table metrics of the AWS DMS task created for migrating the data to verify the statistics for the tables being migrated and to verify that the data definition language (DDL) statements are completed.
- D. Enable the AWS Schema Conversion Tool (AWS SCT) premigration validation and review the premigration checklist to make sure there are no issues with the conversion.
- E. Enable AWS DMS data validation on the task so the AWS DMS task compares the source and target records, and reports any mismatches.

Answer: D

Explanation:

"To ensure that your data was migrated accurately from the source to the target, we highly recommend that you use data validation."
https://docs.aws.amazon.com/dms/latest/userguide/CHAP_BestPractices.html

NEW QUESTION 65

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

Answer: B

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.

NEW QUESTION 70

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

Answer: D

NEW QUESTION 72

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

Answer: D

Explanation:

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

NEW QUESTION 73

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

Answer: C

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/rds-cannot-connect/>

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

NEW QUESTION 77

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

Answer: B

Explanation:

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

NEW QUESTION 81

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
- D. Use SQL functions to extract sentiment analysis.
- E. Set up Aurora native integration with Amazon SageMaker
- F. Use SQL functions to extract sentiment analysis.

Answer: C

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.stackoverflow.com/2019/11/27/new-for-amazon-aurora-use-machine-learning-directly-from-your>

NEW QUESTION 86

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.

Answer: A

Explanation:

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

NEW QUESTION 88

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
- E. Replace the passwords by running CloudFormation using a sed command.

Answer: B

Explanation:

<https://aws.amazon.com/blogs/infrastructure-and-automation/securing-passwords-in-aws-quick-starts-using-aws>

NEW QUESTION 90

A Database Specialist is migrating a 2 TB Amazon RDS for Oracle DB instance to an RDS for PostgreSQL DB instance using AWS DMS. The source RDS Oracle DB instance is in a VPC in the us-east-1 Region. The target RDS for PostgreSQL DB instance is in a VPC in the us-west-2 Region.

Where should the AWS DMS replication instance be placed for the MOST optimal performance?

- A. In the same Region and VPC of the source DB instance
- B. In the same Region and VPC as the target DB instance
- C. In the same VPC and Availability Zone as the target DB instance
- D. In the same VPC and Availability Zone as the source DB instance

Answer: C

Explanation:

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_ReplicationInstance.VPC.html#CHAP_ReplicationIn In fact, all the configurations list on above url prefer the replication instance putting into target vpc region / subnet / az.
https://docs.aws.amazon.com/dms/latest/sbs/CHAP_SQLServer2Aurora.Steps.CreateReplicationInstance.html

NEW QUESTION 93

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
- B. Use the mysqldump utility to create an initial snapshot from the on-premises MySQL server, and copy it to an Amazon S3 bucket
- C. Import the snapshot into the DB instance utilizing the MySQL utilities running on an Amazon EC2 instance
- D. Immediately point the application to the DB instance.
- E. 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
- F. Use the mysqldump utility to create a snapshot of the on-premises MySQL server
- G. Copy the snapshot into the EC2 instance and restore it into the EC2 MySQL instance
- H. Use AWS DMS to migrate data into a new RDS for MySQL DB instance
- I. Point the application to the DB instance.
- J. 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
- K. Use the mysqldump utility to create a snapshot of the on-premises MySQL server
- L. 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
- M. Point the application to the DB instance.
- N. Deploy a new RDS for MySQL DB instance and configure it for access from the on-premises data center
- O. Use the mysqldump utility to create an initial snapshot from the on-premises MySQL server, and copy it to an Amazon S3 bucket
- P. Import the snapshot into the DB instance using the MySQL utilities running on an Amazon EC2 instance
- Q. Establish replication into the new DB instance using MySQL replication
- R. Stop application access to the on-premises MySQL server and let the remaining transactions replicate over
- S. Point the application to the DB instance.

Answer: D

Explanation:

<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>

NEW QUESTION 94

A company is deploying a solution in Amazon Aurora by migrating from an on-premises system. The IT department has established an AWS Direct Connect link from the company's data center. The company's Database Specialist has selected the option to require SSL/TLS for connectivity to prevent plaintext data from being sent over the network. The migration appears to be working successfully, and the data can be queried from a desktop machine. Two Data Analysts have been asked to query and validate the data in the new Aurora DB cluster. Both Analysts are unable to connect to Aurora. Their user names and passwords have been verified as valid and the Database Specialist can connect to the DB cluster using their accounts. The Database Specialist also verified that the security group configuration allows network from all corporate IP addresses. What should the Database Specialist do to correct the Data Analysts' inability to connect?

- A. Restart the DB cluster to apply the SSL change.
- B. Instruct the Data Analysts to download the root certificate and use the SSL certificate on the connection string to connect.
- C. Add explicit mappings between the Data Analysts' IP addresses and the instance in the security group assigned to the DB cluster.
- D. Modify the Data Analysts' local client firewall to allow network traffic to AWS.

Answer: B

Explanation:

- To connect using SSL:
- Provide the SSLTrust certificate (can be downloaded from AWS)
- Provide SSL options when connecting to database
- Not using SSL on a DB that enforces SSL would result in error <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/ssl-certificate-rotation-aurora-postgresql.html>

NEW QUESTION 99

A company is about to launch a new product, and test databases must be re-created from production data. The company runs its production databases on an Amazon Aurora MySQL DB cluster. A Database Specialist needs to deploy a solution to create these test databases as quickly as possible with the least amount of administrative effort. What should the Database Specialist do to meet these requirements?

- A. Restore a snapshot from the production cluster into test clusters
- B. Create logical dumps of the production cluster and restore them into new test clusters
- C. Use database cloning to create clones of the production cluster
- D. Add an additional read replica to the production cluster and use that node for testing

Answer: C

Explanation:

<https://aws.amazon.com/getting-started/hands-on/aurora-cloning-backtracking/>

"Cloning an Aurora cluster is extremely useful if you want to assess the impact of changes to your database, or if you need to perform workload-intensive operations—such as exporting data or running analytical queries, or simply if you want to use a copy of your production database in a development or testing environment. You can make multiple clones of your Aurora DB cluster. You can even create additional clones from other clones, with the constraint that the clone databases must be created in the same region as the source databases.

NEW QUESTION 104

An IT consulting company wants to reduce costs when operating its development environment databases. The company's workflow creates multiple Amazon Aurora MySQL DB clusters for each development group. The Aurora DB clusters are only used for 8 hours a day. The DB clusters can then be deleted at the end of the development cycle, which lasts 2 weeks.

Which of the following provides the MOST cost-effective solution?

- A. Use AWS CloudFormation template
- B. Deploy a stack with the DB cluster for each development group. Delete the stack at the end of the development cycle.
- C. Use the Aurora DB cloning feature
- D. Deploy a single development and test Aurora DB instance, and create clone instances for the development group
- E. Delete the clones at the end of the development cycle.
- F. Use Aurora Replica
- G. From the master automatic pause compute capacity option, create replicas for each development group, and promote each replica to master
- H. Delete the replicas at the end of the development cycle.
- I. Use Aurora Serverless
- J. Restore current Aurora snapshot and deploy to a serverless cluster for each development group
- K. Enable the option to pause the compute capacity on the cluster and set an appropriate timeout.

Answer: B

Explanation:

Aurora Serverless is not compatible to all Aurora provisioned engine version. However, you can do clone with most engine version. Meanwhile, I also consider the performance while restoring snapshot to Aurora Serverless.

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-serverless-how-it-works.html#aurora>

[https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-serverless.html#aurora-serverless.us](https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-serverless.html#aurora-serverless-us)

NEW QUESTION 106

A large e-commerce 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

Answer: D

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/bp-partition-key-design.html#bp-partition> "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."

NEW QUESTION 110

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
- B. Apply this setting immediately without waiting for the next scheduled RDS maintenance window.
- C. Stop the DB instance and create an encrypted snapshot
- D. Restore the encrypted snapshot to a new encrypted DB instance
- E. Delete the original DB instance, and update the applications to point to the new encrypted DB instance.
- F. Stop the DB instance and create a snapshot
- G. Copy the snapshot into another encrypted snapshot
- H. Restore the encrypted snapshot to a new encrypted DB instance
- I. Delete the original DB instance, and update the applications to point to the new encrypted DB instance.
- J. Create an encrypted read replica of the DB instance
- K. Promote the read replica to master
- L. Delete the original DB instance, and update the applications to point to the new encrypted DB instance.

Answer: C

NEW QUESTION 111

A single MySQL database was moved to Amazon Aurora by a business. The production data is stored in a database cluster in VPC PROD, whereas 12 testing environments are hosted in VPC TEST with the same AWS account. Testing has a negligible effect on the test data. The development team requires that each environment be updated nightly to ensure that each test database has daily production data.

Which migration strategy will be the quickest and least expensive to implement?

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

Answer: A

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Managing.Clone.html>

NEW QUESTION 114

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.

Answer: AC

Explanation:

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

NEW QUESTION 118

A database specialist is constructing an AWS CloudFormation stack using AWS CloudFormation. The database expert wishes to avoid the stack's Amazon RDS ProductionDatabase resource being accidentally deleted.

Which solution will satisfy this criterion?

- A. Create a stack policy to prevent update
- B. Include Effect : ProductionDatabase and Resource Deny in the policy.
- C. Create an AWS CloudFormation stack in XML form
- D. Set xAttribute as false.
- E. Create an RDS DB instance without the DeletionPolicy attribute
- F. Disable termination protection.
- G. Create a stack policy to prevent update
- H. Include Effect, Deny, and Resource :ProductionDatabase in the policy.

Answer: D

Explanation:

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/protect-stack-resources.html> "When you set a stack policy, all resources are protected by default. To allow updates on all resources, we add an Allow statement that allows all actions on all resources. Although the Allow statement specifies all resources, the explicit Deny statement overrides it for the resource with the ProductionDatabase logical ID. This Deny statement prevents all update actions, such as replacement or deletion, on the ProductionDatabase resource."

NEW QUESTION 120

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

Answer: C

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-cross-account>

NEW QUESTION 121

A company is using Amazon RDS for PostgreSQL. The Security team wants all database connection requests to be logged and retained for 180 days. The RDS for PostgreSQL DB instance is currently using the default parameter group. A Database Specialist has identified that setting the log_connections parameter to 1 will enable connections logging.

Which combination of steps should the Database Specialist take to meet the logging and retention requirements? (Choose two.)

- A. Update the log_connections parameter in the default parameter group
- B. Create a custom parameter group, update the log_connections parameter, and associate the parameter with the DB instance
- C. Enable publishing of database engine logs to Amazon CloudWatch Logs and set the event expiration to 180 days

- D. Enable publishing of database engine logs to an Amazon S3 bucket and set the lifecycle policy to 180 days
- E. Connect to the RDS PostgreSQL host and update the log_connections parameter in the postgresql.conf file

Answer: AE

NEW QUESTION 124

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)

Answer: A

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."

NEW QUESTION 129

A ride-hailing application stores bookings in a persistent Amazon RDS for MySQL DB instance. This program is very popular, and the corporation anticipates a tenfold rise in the application's user base over the next several months. The application receives a higher volume of traffic in the morning and evening.

This application is divided into two sections:

An internal booking component that takes online reservations in response to concurrent user queries. A component of a third-party customer relationship management (CRM) system that customer service professionals utilize. Booking data is accessed using queries in the CRM.

To manage this workload effectively, a database professional must create a cost-effective database system. Which solution satisfies these criteria?

- A. Use Amazon ElastiCache for Redis to accept the booking
- B. Associate an AWS Lambda function to capture changes and push the booking data to the RDS for MySQL DB instance used by the CRM.
- C. Use Amazon DynamoDB to accept the booking
- D. Enable DynamoDB Streams and associate an AWS Lambda function to capture changes and push the booking data to an Amazon SQS queue
- E. 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.
- F. Use Amazon ElastiCache for Redis to accept the booking
- G. Associate an AWS Lambda function to capture changes and push the booking data to an Amazon Redshift database used by the CRM.
- H. Use Amazon DynamoDB to accept the booking
- I. 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.

Answer: B

Explanation:

"AWS Lambda function to capture changes" capture changes to what? ElastiCache? The main use of ElastiCache is to cache frequently read data. Also "the company expects a tenfold increase in the user base" and "correspond to simultaneous requests from users"

NEW QUESTION 131

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

Answer: C

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.

NEW QUESTION 132

A manufacturing company's website uses an Amazon Aurora PostgreSQL DB cluster.

Which configurations will result in the LEAST application downtime during a failover? (Choose three.)

- A. Use the provided read and write Aurora endpoints to establish a connection to the Aurora DB cluster.
- B. Create an Amazon CloudWatch alert triggering a restore in another Availability Zone when the primary Aurora DB cluster is unreachable.
- C. Edit and enable Aurora DB cluster cache management in parameter groups.
- D. Set TCP keepalive parameters to a high value.

- E. Set JDBC connection string timeout variables to a low value.
- F. Set Java DNS caching timeouts to a high value.

Answer: ABC

NEW QUESTION 134

A retail company with its main office in New York and another office in Tokyo plans to build a database solution on AWS. The company's main workload consists of a mission-critical application that updates its application data in a data store. The team at the Tokyo office is building dashboards with complex analytical queries using the application data. The dashboards will be used to make buying decisions, so they need to have access to the application data in less than 1 second. Which solution meets these requirements?

- A. Use an Amazon RDS DB instance deployed in the us-east-1 Region with a read replica instance in the ap-northeast-1 Region
- B. Create an Amazon ElastiCache cluster in the ap-northeast-1 Region to cache application data from the replica to generate the dashboards.
- C. Use an Amazon DynamoDB global table in the us-east-1 Region with replication into the ap-northeast-1 Region
- D. Use Amazon QuickSight for displaying dashboard results.
- E. Use an Amazon RDS for MySQL DB instance deployed in the us-east-1 Region with a read replica instance in the ap-northeast-1 Region
- F. Have the dashboard application read from the read replica.
- G. Use an Amazon Aurora global database
- H. Deploy the writer instance in the us-east-1 Region and the replica in the ap-northeast-1 Region
- I. Have the dashboard application read from the replica in the ap-northeast-1 Region.

Answer: D

Explanation:

<https://aws.amazon.com/blogs/database/aurora-postgresql-disaster-recovery-solutions-using-amazon-aurora-global>

NEW QUESTION 138

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
- B. Watch these dashboards during the next slow period.
- C. 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.
- D. 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.
- E. Enable Amazon RDS Performance Insights on the PostgreSQL database
- F. Use the metrics to identify any queries that are related to spikes in the graph during the next slow period.

Answer: D

NEW QUESTION 142

A company is due for renewing its database license. The company wants to migrate its 80 TB transactional database system from on-premises to the AWS Cloud. The migration should incur the least possible downtime on the downstream database applications. The company's network infrastructure has limited network bandwidth that is shared with other applications.

Which solution should a database specialist use for a timely migration?

- A. Perform a full backup of the source database to AWS Snowball Edge appliances and ship them to be loaded to Amazon S3. Use AWS DMS to migrate change data capture (CDC) data from the source database to Amazon S3. Use a second AWS DMS task to migrate all the S3 data to the target database.
- B. Perform a full backup of the source database to AWS Snowball Edge appliances and ship them to be loaded to Amazon S3. Periodically perform incremental backups of the source database to be shipped in another Snowball Edge appliance to handle syncing change data capture (CDC) data from the source to the target database.
- C. Use AWS DMS to migrate the full load of the source database over a VPN tunnel using the internet for its primary connection
- D. Allow AWS DMS to handle syncing change data capture (CDC) data from the source to the target database.
- E. Use the AWS Schema Conversion Tool (AWS SCT) to migrate the full load of the source database over a VPN tunnel using the internet for its primary connection
- F. Allow AWS SCT to handle syncing change data capture (CDC) data from the source to the target database.

Answer: A

Explanation:

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Target.S3.html Using Amazon S3 as a target for AWS Database Migration Service

NEW QUESTION 147

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
- B. Use the mysqldump utility to create an initial snapshot from the on-premises MySQL server, and copy it to an Amazon S3 bucket
- C. Import the snapshot into the DB instance utilizing the MySQL utilities running on an Amazon EC2 instance
- D. Immediately point the application to the DB instance.
- E. 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
- F. Use the mysqldump utility to create a snapshot of the on-premises MySQL server
- G. Copy the snapshot into the EC2 instance and restore it into the EC2 MySQL instance
- H. Use AWS DMS to migrate data into a new RDS for MySQL DB instance
- I. Point the application to the DB instance.
- J. 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
- K. Use the mysqldump utility to create a snapshot of the on-premises MySQL server
- L. 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

instanc

M. Point the application to the DB instance.

N. Deploy a new RDS for MySQL DB instance and configure it for access from the on-premises data center

O. Use the mysqldump utility to create an initial snapshot from the on-premises MySQL server, and copy it to an Amazon S3 bucket

P. Import the snapshot into the DB instance using the MySQL utilities running on an Amazon EC2 instance

Q. Establish replication into the new DB instance using MySQL replication

R. Stop application access to the on-premises MySQL server and let the remaining transactions replicate over

S. Point the application to the DB instance.

Answer: B

NEW QUESTION 149

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

B. Launch an Amazon EC2 instance with a weekly cron job to execute the custom script.

C. 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

D. Schedule the Lambda function to run weekly using Amazon EventBridge (Amazon CloudWatch Events).

E. 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

F. Schedule both Lambda functions to run weekly using Amazon EventBridge (Amazon CloudWatch Events).

G. 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.

Answer: B

Explanation:

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

NEW QUESTION 151

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

B. 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.

C. 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.

D. 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.

E. Configure Amazon DynamoDB Streams for the DynamoDB table in the current account

F. 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.

Answer: C

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/dynamodb-cross-account-migration/> <https://aws.amazon.com/premiumsupport/knowledge-center/data-pipeline-account-access-dynamodb-s3/>

NEW QUESTION 153

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:WaitSync 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).

Answer: A

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>

NEW QUESTION 157

A company is planning to close for several days. A Database Specialist needs to stop all applications along with the DB instances to ensure employees do not have access to the systems during this time. All databases are running on Amazon RDS for MySQL.

The Database Specialist wrote and executed a script to stop all the DB instances. When reviewing the logs, the Database Specialist found that Amazon RDS DB instances with read replicas did not stop.

How should the Database Specialist edit the script to fix this issue?

- A. Stop the source instances before stopping their read replicas
- B. Delete each read replica before stopping its corresponding source instance
- C. Stop the read replicas before stopping their source instances
- D. Use the AWS CLI to stop each read replica and source instance at the same time

Answer: B

Explanation:

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

"The following are some limitations to stopping and starting a DB instance: You can't stop a DB instance that has a read replica, or that is a read replica." So if you can't stop a db with a read replica, you have to delete the read replica first to then stop it???

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_MySQL.Replication.ReadReplicas.html#U

NEW QUESTION 160

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
- D. Enable Multi-AZ after auditing is enabled.
- E. Disable automated backup on the DB instance, and then enable auditing
- F. Enable automated backup after auditing is enabled.
- G. 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.

Answer: AE

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

NEW QUESTION 162

A Database Specialist is creating a new Amazon Neptune DB cluster, and is attempting to load data from Amazon S3 into the Neptune DB cluster using the Neptune bulk loader API. The Database Specialist receives the following error:

"Unable to connect to s3 endpoint. Provided source = s3://mybucket/graphdata/ and region = us-east-1. Please verify your S3 configuration."

Which combination of actions should the Database Specialist take to troubleshoot the problem? (Choose two.)

- A. Check that Amazon S3 has an IAM role granting read access to Neptune
- B. Check that an Amazon S3 VPC endpoint exists
- C. Check that a Neptune VPC endpoint exists
- D. Check that Amazon EC2 has an IAM role granting read access to Amazon S3
- E. Check that Neptune has an IAM role granting read access to Amazon S3

Answer: BD

NEW QUESTION 166

A corporation wishes to move a 1 TB Oracle database from its current location to an Amazon Aurora PostgreSQL DB cluster. The database specialist at the firm noticed that the Oracle database stores 100 GB of large binary objects (LOBs) across many tables. The Oracle database supports LOBs up to 500 MB in size and an average of 350 MB. AWS DMS was picked by the Database Specialist to transfer the data with the most replication instances.

How should the database specialist improve the transfer of the database to AWS DMS?

- A. Create a single task using full LOB mode with a LOB chunk size of 500 MB to migrate the data and LOBs together
- B. Create two tasks: task1 with LOB tables using full LOB mode with a LOB chunk size of 500 MB and task2 without LOBs
- C. Create two tasks: task1 with LOB tables using limited LOB mode with a maximum LOB size of 500 MB and task 2 without LOBs
- D. Create a single task using limited LOB mode with a maximum LOB size of 500 MB to migrate data and LOBs together

Answer: C

Explanation:

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_BestPractices.html#CHAP_BestPractices.LOBS, "AWS DMS migrates LOB data in two phases: 1.

AWS DMS creates a new row in the target table and

populates the row with all data except the associated LOB value. 2. AWS DMS updates the row in the target table with the LOB data." This means that we would need two tasks, one per phase and use limited LOB mode for best performance.

NEW QUESTION 169

A Database Specialist must create a read replica to isolate read-only queries for an Amazon RDS for MySQL DB instance. Immediately after creating the read replica, users that query it report slow response times. What could be causing these slow response times?

- A. New volumes created from snapshots load lazily in the background
- B. Long-running statements on the master
- C. Insufficient resources on the master
- D. Overload of a single replication thread by excessive writes on the master

Answer: A

Explanation:

snapshot is lazy loaded If the volume is accessed where the data is not loaded, the application accessing the volume encounters a higher latency than normal while the data gets loaded

<https://aws.amazon.com/about-aws/whats-new/2019/11/amazon-efs-fast-snapshot-restore-eliminates-need-for-p>

NEW QUESTION 173

Amazon Neptune is being used by a corporation as the graph database for one of its products. During an ETL procedure, the company's data science team produced enormous volumes of temporary data by unintentionally. The Neptune DB cluster extended its storage capacity automatically to handle the added data, but the data science team erased the superfluous data.

What should a database professional do to prevent incurring extra expenditures for cluster volume space that is not being used?

- A. Take a snapshot of the cluster volume
- B. Restore the snapshot in another cluster with a smaller volume size.
- C. Use the AWS CLI to turn on automatic resizing of the cluster volume.
- D. Export the cluster data into a new Neptune DB cluster.
- E. Add a Neptune read replica to the cluster
- F. Promote this replica as a new primary DB instance
- G. Reset the storage space of the cluster.

Answer: C

Explanation:

The only way to shrink the storage space used by your DB cluster when you have a large amount of unused allocated space is to export all the data in your graph and then reload it into a new DB cluster. Creating and restoring a snapshot does not reduce the amount of storage allocated for your DB cluster, because a snapshot retains the original image of the cluster's underlying storage.

NEW QUESTION 174

A company is load testing its three-tier production web application deployed with an AWS CloudFormation template on AWS. The Application team is making changes to deploy additional Amazon EC2 and AWS Lambda resources to expand the load testing capacity. A Database Specialist wants to ensure that the changes made by the Application team will not change the Amazon RDS database resources already deployed.

Which combination of steps would allow the Database Specialist to accomplish this? (Choose two.)

- A. Review the stack drift before modifying the template
- B. Create and review a change set before applying it
- C. Export the database resources as stack outputs
- D. Define the database resources in a nested stack
- E. Set a stack policy for the database resources

Answer: BE

Explanation:

https://docs.amazonaws.cn/en_us/AWSCloudFormation/latest/UserGuide/best-practices.html#cfn-best-practices

NEW QUESTION 177

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

Answer: C

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/CreateBackup.html#:~:text=If%20you%20>
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/backuprestore_HowItWorks.html

NEW QUESTION 181

A company is looking to migrate a 1 TB Oracle database from on-premises to an Amazon Aurora PostgreSQL DB cluster. The company's Database Specialist discovered that the Oracle database is storing 100 GB of large binary objects (LOBs) across multiple tables. The Oracle database has a maximum LOB size of 500 MB with an average LOB size of 350 MB. The Database Specialist has chosen AWS DMS to migrate the data with the largest replication instances.

How should the Database Specialist optimize the database migration using AWS DMS?

- A. Create a single task using full LOB mode with a LOB chunk size of 500 MB to migrate the data and LOBs together
- B. Create two tasks: task1 with LOB tables using full LOB mode with a LOB chunk size of 500 MB and task2 without LOBs
- C. Create two tasks: task1 with LOB tables using limited LOB mode with a maximum LOB size of 500 MB and task 2 without LOBs
- D. Create a single task using limited LOB mode with a maximum LOB size of 500 MB to migrate data and LOBs together

Answer: C

NEW QUESTION 183

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.

Answer: D

Explanation:

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

NEW QUESTION 185

A company developed an AWS CloudFormation template used to create all new Amazon DynamoDB tables in its AWS account. The template configures provisioned throughput capacity using hard-coded values. The company wants to change the template so that the tables it creates in the future have independently configurable read and write capacity units assigned.

Which solution will enable this change?

- A. Add values for the rcuCount and wcuCount parameters to the Mappings section of the template. Configure DynamoDB to provision throughput capacity using the stack's mappings.
- B. Add values for two Number parameters, rcuCount and wcuCount, to the template.
- C. Replace the hard-coded values with calls to the Ref intrinsic function, referencing the new parameters.
- D. Add values for the rcuCount and wcuCount parameters as outputs of the template.
- E. Configure DynamoDB to provision throughput capacity using the stack outputs.
- F. Add values for the rcuCount and wcuCount parameters to the Mappings section of the template.
- G. Replace the hard-coded values with calls to the Ref intrinsic function, referencing the new parameters.

Answer: B

Explanation:

Input parameter and FindInMap You can use an input parameter with the Fn::FindInMap function to refer to a specific value in a map. For example, suppose you have a list of regions and environment types that map to a specific AMI ID. You can select the AMI ID that your stack uses by using an input parameter (EnvironmentType). To determine the region, use the AWS::Region pseudo parameter, which gets the AWS Region in which you create the stack.

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/parameters-section-structure.html>

NEW QUESTION 188

A Database Specialist is designing a disaster recovery strategy for a production Amazon DynamoDB table. The table uses provisioned read/write capacity mode, global secondary indexes, and time to live (TTL). The Database Specialist has restored the latest backup to a new table.

To prepare the new table with identical settings, which steps should be performed? (Choose two.)

- A. Re-create global secondary indexes in the new table
- B. Define IAM policies for access to the new table
- C. Define the TTL settings
- D. Encrypt the table from the AWS Management Console or use the update-table command
- E. Set the provisioned read and write capacity

Answer: BC

Explanation:

The following items need to be reconfigured after restoring the DynamoDB table.

- AutoScaling policy
- IAM policy
- CloudWatch settings
- Tags
- Stream settings
- TTL

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

NEW QUESTION 189

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.

Answer: D

Explanation:

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

NEW QUESTION 194

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

Answer: C

NEW QUESTION 199

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

Answer: AE

Explanation:

<https://docs.aws.amazon.com/AmazonElastiCache/latest/mem-ug/Strategies.html#Strategies.WriteThrough> <https://aws.amazon.com/products/databases/real-time-apps-elasticache-for-redis/>

NEW QUESTION 203

A business just transitioned from an on-premises Oracle database to Amazon Aurora PostgreSQL. Following the move, the organization observed that every day around 3:00 PM, the application's response time is substantially slower. The firm has determined that the problem is with the database, not the application. Which set of procedures should the Database Specialist do to locate the erroneous PostgreSQL query most efficiently?

- A. Create an Amazon CloudWatch dashboard to show the number of connections, CPU usage, and disk space consumption.
- B. Watch these dashboards during the next slow period.
- C. 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.
- D. 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.
- E. Enable Amazon RDS Performance Insights on the PostgreSQL database.
- F. Use the metrics to identify any queries that are related to spikes in the graph during the next slow period.

Answer: D

Explanation:

<https://aws.amazon.com/blogs/database/optimizing-and-tuning-queries-in-amazon-rds-postgresql-based-on-native> "AWS recently released a feature called Amazon RDS Performance Insights, which provides an easy-to-understand dashboard for detecting performance problems in terms of load." "AWS recently released a feature called Amazon RDS Performance Insights, which provides an easy-to-understand dashboard for detecting performance problems in terms of load."

NEW QUESTION 205

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

Answer: A

Explanation:

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

NEW QUESTION 207

A company has deployed an e-commerce web application in a new AWS account. An Amazon RDS for MySQL Multi-AZ DB instance is part of this deployment with a database-1.xxxxxxxxxx.us-east-1.rds.amazonaws.com endpoint listening on port 3306. The company's Database Specialist is able to log in to MySQL and run queries from the bastion host using these details.

When users try to utilize the application hosted in the AWS account, they are presented with a generic error message. The application servers are logging a "could not connect to server: Connection times out" error message to Amazon CloudWatch Logs. What is the cause of this error?

- A. The user name and password the application is using are incorrect.
- B. The security group assigned to the application servers does not have the necessary rules to allow inbound connections from the DB instance.
- C. The security group assigned to the DB instance does not have the necessary rules to allow inbound connections from the application servers.
- D. The user name and password are correct, but the user is not authorized to use the DB instance.

Answer: C

NEW QUESTION 209

In one AWS account, a business runs a two-tier e-commerce 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.

Answer: ACD

NEW QUESTION 211

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.
- B. Assign the DB instances to the backup plan.
- C. Modify the DB instances to enable the automated backup option.
- D. Select the required backup retention period.
- E. 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.
- F. Use AWS Lambda to schedule a daily manual snapshot of the DB instance.
- G. Delete snapshots that exceed the retention requirement.

Answer: A

Explanation:

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

NEW QUESTION 215

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

Answer: A

NEW QUESTION 216

A clothing company uses a custom ecommerce application and a PostgreSQL database to sell clothes to thousands of users from multiple countries. The company is migrating its application and database from its on-premises data center to the AWS Cloud. The company has selected Amazon EC2 for the application and Amazon RDS for PostgreSQL for the database. The company requires database passwords to be changed every 60 days. A Database Specialist needs to ensure that the credentials used by the web application to connect to the database are managed securely. Which approach should the Database Specialist take to securely manage the database credentials?

- A. Store the credentials in a text file in an Amazon S3 bucket.
- B. Restrict permissions on the bucket to the IAM role associated with the instance profile only.
- C. Modify the application to download the text file and retrieve the credentials on start up.
- D. Update the text file every 60 days.
- E. Configure IAM database authentication for the application to connect to the database.
- F. Create an IAM user and map it to a separate database user for each ecommerce user.
- G. Require users to update their passwords every 60 days.
- H. Store the credentials in AWS Secrets Manager.
- I. Restrict permissions on the secret to only the IAM role associated with the instance profile.
- J. Modify the application to retrieve the credentials from Secrets Manager on start up.
- K. Configure the rotation interval to 60 days.
- L. Store the credentials in an encrypted text file in the application AMI.
- M. Use AWS KMS to store the key for decrypting the text file.
- N. Modify the application to decrypt the text file and retrieve the credentials on start up.
- O. Update the text file and publish a new AMI every 60 days.

Answer: C

NEW QUESTION 221

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).

Answer: D

NEW QUESTION 226

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

Answer: A

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.

NEW QUESTION 228

.....

Relate Links

100% Pass Your AWS-Certified-Database-Specialty Exam with ExamBible Prep Materials

<https://www.exambible.com/AWS-Certified-Database-Specialty-exam/>

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>