

Amazon

Exam Questions AWS-Certified-Database-Specialty

AWS Certified Database - Specialty



NEW QUESTION 1

A database specialist manages a critical Amazon RDS for MySQL DB instance for a company. The data stored daily could vary from .01% to 10% of the current database size. The database specialist needs to ensure that the DB instance storage grows as needed.

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

- A. Configure RDS Storage Auto Scaling.
- B. Configure RDS instance Auto Scaling.
- C. Modify the DB instance allocated storage to meet the forecasted requirements.
- D. Monitor the Amazon CloudWatch FreeStorageSpace metric daily and add storage as required.

Answer: A

Explanation:

If your workload is unpredictable, you can enable storage autoscaling for an Amazon RDS DB instance. With storage autoscaling enabled, when Amazon RDS detects that you are running out of free database space it automatically scales up your storage.

<https://aws.amazon.com/about-aws/whats-new/2019/06/rds-storage-auto-scaling/>

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

NEW QUESTION 2

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 3

A gaming company is designing a mobile gaming app that will be accessed by many users across the globe. The company wants to have replication and full support for multi-master writes. The company also wants to ensure low latency and consistent performance for app users.

Which solution meets these requirements?

- A. Use Amazon DynamoDB global tables for storage and enable DynamoDB automatic scaling
- B. Use Amazon Aurora for storage and enable cross-Region Aurora Replicas
- C. Use Amazon Aurora for storage and cache the user content with Amazon ElastiCache
- D. Use Amazon Neptune for storage

Answer: A

NEW QUESTION 4

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 5

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

- A. Create multiple AWS DMS tasks to migrate the large table.
- B. Configure the AWS DMS replication instance with Multi-AZ.
- C. Increase the capacity of the AWS DMS replication server.
- D. Establish an AWS Direct Connect connection between the on-premises data center and AWS.
- E. Enable an Amazon RDS Multi-AZ configuration.

F. Enable full large binary object (LOB) mode to migrate all LOB data for all large tables.

Answer: ACD

Explanation:

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

NEW QUESTION 6

A company with branch offices in Portland, New York, and Singapore has a three-tier web application that leverages a shared database. The database runs on Amazon RDS for MySQL and is hosted in the us-west-2 Region. The application has a distributed front end deployed in the us-west-2, ap-southeast-1, and us-east-2 Regions.

This front end is used as a dashboard for Sales Managers in each branch office to see current sales statistics. There are complaints that the dashboard performs more slowly in the Singapore location than it does in Portland or New York. A solution is needed to provide consistent performance for all users in each location. Which set of actions will meet these requirements?

- A. Take a snapshot of the instance in the us-west-2 Regio
- B. Create a new instance from the snapshot in the ap-southeast-1 Regio
- C. Reconfigure the ap-southeast-1 front-end dashboard to access this instance.
- D. Create an RDS read replica in the ap-southeast-1 Region from the primary RDS DB instance in the us- west-2 Regio
- E. Reconfigure the ap-southeast-1 front-end dashboard to access this instance.
- F. Create a new RDS instance in the ap-southeast-1 Regio
- G. Use AWS DMS and change data capture (CDC) to update the new instance in the ap-southeast-1 Regio
- H. Reconfigure the ap-southeast-1 front-end dashboard to access this instance.
- I. Create an RDS read replica in the us-west-2 Region where the primary instance reside
- J. Create a read replica in the ap-southeast-1 Region from the read replica located on the us-west-2 Regio
- K. Reconfigure the ap-southeast-1 front-end dashboard to access this instance.

Answer: B

Explanation:

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

"Amazon RDS Read Replicas provide enhanced performance and durability for RDS database (DB) instances.

They make it easy to elastically scale out beyond the capacity constraints of a single DB instance for read-heavy database workloads. You can create one or more replicas of a given source DB Instance and serve high-volume application read traffic from multiple copies of your data, thereby increasing aggregate read throughput. "

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

NEW QUESTION 7

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

Which database specialist solution should be recommended?

- A. Enable Aurora Database Activity Streams on the database in synchronous mod
- B. Connect the Amazon Kinesis data stream to Kinesis Data Firehos
- C. Set the Kinesis Data Firehose destination to an Amazon S3 bucket.
- D. Create an AWS CloudTrail trail in the Region where the database run
- E. Associate the database activity logs with the trail.
- F. Enable Aurora Database Activity Streams on the database in asynchronous mod
- G. Connect the Amazon Kinesis data stream to Kinesis Data Firehos
- H. Set the Firehose destination to an Amazon S3 bucket.
- I. Allow connections to the DB cluster through a bastion host onl
- J. Restrict database access to the bastion host and application server
- K. Push the bastion host logs to Amazon CloudWatch Logs using the CloudWatch Logs agent.

Answer: C

Explanation:

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

NEW QUESTION 8

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 9

A global digital advertising company captures browsing metadata to contextually display relevant images, pages, and links to targeted users. A single page load can generate multiple events that need to be stored individually. The maximum size of an event is 200 KB and the average size is 10 KB. Each page load must query the user's browsing history to provide targeting recommendations. The advertising company expects over 1 billion page visits per day from users in the United States, Europe, Hong Kong, and India. The structure of the metadata varies depending on the event. Additionally, the browsing metadata must be written and read with very low latency to ensure a good viewing experience for the users.

Which database solution meets these requirements?

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

Answer: C

NEW QUESTION 10

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

Which solution satisfies these criteria?

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

Answer: D

Explanation:

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

NEW QUESTION 10

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 13

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 18

A company is using 5 TB Amazon RDS DB instances and needs to maintain 5 years of monthly database backups for compliance purposes. A Database Administrator must provide Auditors with data within 24 hours. Which solution will meet these requirements and is the MOST operationally efficient?

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

Answer: A

Explanation:

Unlike automated backups, manual snapshots aren't subject to the backup retention period. Snapshots don't expire. For very long-term backups of MariaDB, MySQL, and PostgreSQL data, we recommend exporting snapshot data to Amazon S3. If the major version of your DB engine is no longer supported, you can't restore to that version from a snapshot. https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_CreateSnapshot.html

NEW QUESTION 20

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

How should the company identify the source of the problem?

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

Answer: B

Explanation:

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

NEW QUESTION 25

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 SecureStringTemplate templat
- 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 27

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

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

Answer: AD

Explanation:

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

NEW QUESTION 28

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

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

What should the database administrator do to meet these requirements?

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

Answer: C

Explanation:

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

NEW QUESTION 32

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 36

A media company is using Amazon RDS for PostgreSQL to store user data. The RDS DB instance currently has a publicly accessible setting enabled and is hosted in a public subnet. Following a recent AWS Well- Architected Framework review, a Database Specialist was given new security requirements. Only certain on-premises corporate network IPs should connect to the DB instance. Connectivity is allowed from the corporate network only.

Which combination of steps does the Database Specialist need to take to meet these new requirements? (Choose three.)

- A. Modify the pg_hba.conf file
- B. Add the required corporate network IPs and remove the unwanted IPs.
- C. Modify the associated security group
- D. Add the required corporate network IPs and remove the unwanted IPs.
- E. Move the DB instance to a private subnet using AWS DMS.
- F. Enable VPC peering between the application host running on the corporate network and the VPC associated with the DB instance.
- G. Disable the publicly accessible setting.
- H. Connect to the DB instance using private IPs and a VPN.

Answer: BEF

Explanation:

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

NEW QUESTION 41

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.html>. Unlike for other DB engines, for Aurora DB instances this metric reports the amount of storage available to each DB instance. This value depends on the DB instance class (for pricing information, see the Amazon RDS product page). You can increase the amount of free storage space for an instance by choosing a larger DB instance class for your instance."

NEW QUESTION 42

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 45

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

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

Answer: D

Explanation:

<https://aws.amazon.com/about-aws/whats-new/2018/11/amazon-aurora-simplifies-workload-management-with-c> You can now create custom endpoints for Amazon Aurora databases. This allows you to distribute and load balance workloads across different sets of database instances in your Aurora cluster. For example, you may provision a set of Aurora Replicas to use an instance type with higher memory capacity in order to run an analytics workload. A custom endpoint can then help you route the analytics workload to these appropriately-configured instances, while keeping other instances in your cluster isolated from this workload. As you add or remove instances from the custom endpoint to match your workload, the endpoint helps spread the load around.

NEW QUESTION 49

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 54

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 55

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

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

Answer: C

Explanation:

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

NEW QUESTION 57

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 62

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 66

A company is moving its fraud detection application from on premises to the AWS Cloud and is using Amazon Neptune for data storage. The company has set up a 1 Gbps AWS Direct Connect connection to migrate 25 TB of fraud detection data from the on-premises data center to a Neptune DB instance. The company already has an Amazon S3 bucket and an S3 VPC endpoint, and 80% of the company's network bandwidth is available.

How should the company perform this data load?

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

Answer: C

Explanation:

"AWS DataSync is an online data transfer service that simplifies, automates, and accelerates moving data between on-premises storage systems and AWS storage services, and also between AWS storage services."

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

NEW QUESTION 71

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 enabled.
- 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 74

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

Which action will allow AWS DMS to perform the replication?

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

Answer: A

Explanation:

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

NEW QUESTION 75

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 79

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

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

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

Answer: CDE

NEW QUESTION 84

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](https://www.amazonaws.com/en/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 88

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

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

Answer: D

Explanation:

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

NEW QUESTION 95

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 set 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 97

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 99

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 featur
- 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 maste
- H. Delete the replicas at the end of the development cycle.
- I. Use Aurora Serverles
- J. Restore current Aurora snapshot and deploy to a serverless cluster for each development grou
- 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>

NEW QUESTION 102

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

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

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

Answer: D

Explanation:

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

NEW QUESTION 104

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 Serverles
- 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 105

A Database Specialist is troubleshooting an application connection failure on an Amazon Aurora DB cluster with multiple Aurora Replicas that had been running with no issues for the past 2 months. The connection failure lasted for 5 minutes and corrected itself after that. The Database Specialist reviewed the Amazon RDS events and determined a failover event occurred at that time. The failover process took around 15 seconds to complete.

What is the MOST likely cause of the 5-minute connection outage?

- A. After a database crash, Aurora needed to replay the redo log from the last database checkpoint
- B. The client-side application is caching the DNS data and its TTL is set too high
- C. After failover, the Aurora DB cluster needs time to warm up before accepting client connections
- D. There were no active Aurora Replicas in the Aurora DB cluster

Answer: B

Explanation:

When your application tries to establish a connection after a failover, the new Aurora PostgreSQL writer will be a previous reader, which can be found using the Aurora read only endpoint before DNS updates have fully propagated. Setting the java DNS TTL to a low value helps cycle between reader nodes on subsequent connection attempts.

Amazon Aurora is designed to recover from a crash almost instantaneously and continue to serve your application data. Unlike other databases, after a crash Amazon Aurora does not need to replay the redo log from the last database checkpoint before making the database available for operations. Amazon Aurora performs crash recovery asynchronously on parallel threads, so your database is open and available immediately after a crash. Because the storage is organized in many small segments, each with its own redo log, the underlying storage can replay redo records on demand in parallel and asynchronously as part of a disk read after a crash. This approach reduces database restart times to less than 60 seconds in most cases

NEW QUESTION 109

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

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

Answer: BEF

Explanation:

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

NEW QUESTION 112

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 format
- 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 117

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 120

A company is using an Amazon Aurora PostgreSQL DB cluster with an xlarge primary instance master and two large Aurora Replicas for high availability and read-only workload scaling. A failover event occurs and application performance is poor for several minutes. During this time, application servers in all Availability Zones are healthy and responding normally.

What should the company do to eliminate this application performance issue?

- A. Configure both of the Aurora Replicas to the same instance class as the primary DB instance
- B. Enable cache coherence on the DB cluster, set the primary DB instance failover priority to tier-0, and assign a failover priority of tier-1 to the replicas.
- C. Deploy an AWS Lambda function that calls the DescribeDBInstances action to establish which instance has failed, and then use the PromoteReadReplica operation to promote one Aurora Replica to be the primary DB instance
- D. Configure an Amazon RDS event subscription to send a notification to an Amazon SNS topic to which the Lambda function is subscribed.
- E. Configure one Aurora Replica to have the same instance class as the primary DB instance
- F. Implement Aurora PostgreSQL DB cluster cache management
- G. Set the failover priority to tier-0 for the primary DB instance and one replica with the same instance class
- H. Set the failover priority to tier-1 for the other replicas.
- I. Configure both Aurora Replicas to have the same instance class as the primary DB instance
- J. Implement Aurora PostgreSQL DB cluster cache management
- K. Set the failover priority to tier-0 for the primary DB instance and to tier-1 for the replicas.

Answer: C

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.cluster-cache-mgmt.htm>

<https://aws.amazon.com/blogs/database/introduction-to-aurora-postgresql-cluster-cache-management/>

"You can customize the order in which your Aurora Replicas are promoted to the primary instance after a failure by assigning each replica a priority. Priorities range from 0 for the first priority to 15 for the last priority. If the primary instance fails, Amazon RDS promotes the Aurora Replica with the better priority to the new primary instance. You can modify the priority of an Aurora Replica at any time. Modifying the priority doesn't trigger a failover. 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. "

Amazon Aurora with PostgreSQL compatibility now supports cluster cache management, providing a faster path to full performance if there's a failover. With cluster cache management, you designate a specific reader DB instance in your Aurora PostgreSQL cluster as the failover target. Cluster cache management keeps the data in the designated reader's cache synchronized with the data in the read-write instance's cache. If a failover occurs, the designated reader is promoted to be the new read-write instance, and workloads benefit immediately from the data in its cache.

NEW QUESTION 123

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 128

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 132

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

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

An administrator observes the following upon review:

- No role with the dynamodb:CreateGlobalTable permission exists in the account.
- An empty table with the same name exists in the new Region where replication is desired.
- A global secondary index with the same partition key but a different sort key exists in the new Region where replication is desired.

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

- A. A global secondary index with the same partition key but a different sort key exists in the new Region where replication is desired.
- B. An empty table with the same name exists in the Region where replication is desired.
- C. No role with the dynamodb:CreateGlobalTable permission exists in the account.
- D. DynamoDB Streams is enabled for the table.
- E. The table is encrypted using a KMS customer managed key.

Answer: AB

NEW QUESTION 135

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 136

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 137

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 138

A Database Specialist is designing a new database infrastructure for a ride hailing application. The application data includes a ride tracking system that stores GPS coordinates for all rides. Real-time statistics and metadata lookups must be performed with high throughput and microsecond latency. The database should be fault tolerant with minimal operational overhead and development effort.

Which solution meets these requirements in the MOST efficient way?

- A. Use Amazon RDS for MySQL as the database and use Amazon ElastiCache
- B. Use Amazon DynamoDB as the database and use DynamoDB Accelerator
- C. Use Amazon Aurora MySQL as the database and use Aurora's buffer cache
- D. Use Amazon DynamoDB as the database and use Amazon API Gateway

Answer: B

Explanation:

[https://aws.amazon.com/dynamodb/dax/#:~:text=Amazon%20DynamoDB%20Accelerator%20\(DAX\)%20is,mil](https://aws.amazon.com/dynamodb/dax/#:~:text=Amazon%20DynamoDB%20Accelerator%20(DAX)%20is,mil) "Amazon DynamoDB Accelerator (DAX) is a fully managed, highly available, in-memory cache for DynamoDB that delivers up to a 10x performance improvement – from milliseconds to microseconds – even at millions of requests per second."

NEW QUESTION 143

An ecommerce company is using Amazon DynamoDB as the backend for its order-processing application.

The steady increase in the number of orders is resulting in increased DynamoDB costs. Order verification and reporting perform many repeated GetItem functions that pull similar datasets, and this read activity is contributing to the increased costs. The company wants to control these costs without significant development efforts.

How should a Database Specialist address these requirements?

- A. Use AWS DMS to migrate data from DynamoDB to Amazon DocumentDB
- B. Use Amazon DynamoDB Streams and Amazon Kinesis Data Firehose to push the data into Amazon Redshift
- C. Use an Amazon ElastiCache for Redis in front of DynamoDB to boost read performance
- D. Use DynamoDB Accelerator to offload the reads

Answer: D

Explanation:

https://docs.amazonaws.cn/en_us/amazondynamodb/latest/developerguide/DAX.html

"Applications that are read-intensive, but are also cost-sensitive. With DynamoDB, you provision the number of reads per second that your application requires. If read activity increases, you can increase your tables' provisioned read throughput (at an additional cost). Or, you can offload the activity from your application to a DAX cluster, and reduce the number of read capacity units that you need to purchase otherwise."

NEW QUESTION 146

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 147

A company is running a finance application on an Amazon RDS for MySQL DB instance. The application is governed by multiple financial regulatory agencies. The RDS DB instance is set up with security groups to allow access to certain Amazon EC2 servers only. AWS KMS is used for encryption at rest. Which step will provide additional security?

- A. Set up NACLs that allow the entire EC2 subnet to access the DB instance
- B. Disable the master user account
- C. Set up a security group that blocks SSH to the DB instance
- D. Set up RDS to use SSL for data in transit

Answer: D

NEW QUESTION 150

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 155

A company has an application that uses an Amazon DynamoDB table to store user data. Every morning, a single-threaded process calls the DynamoDB API Scan operation to scan the entire table and generate a critical start-of-day report for management. A successful marketing campaign recently doubled the number of items in the table, and now the process takes too long to run and the report is not generated in time.

A database specialist needs to improve the performance of the process. The database specialist notes that, when the process is running, 15% of the table's provisioned read capacity units (RCUs) are being used.

What should the database specialist do?

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

Answer: B

Explanation:

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

NEW QUESTION 157

An online shopping company has a large inflow of shopping requests daily. As a result, there is a consistent load on the company's Amazon RDS database. A database specialist needs to ensure the database is up and running at all times. The database specialist wants an automatic notification system for issues that may cause database downtime or for configuration changes made to the database.

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

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

Answer: BD

NEW QUESTION 159

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

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

Which solution will satisfy these criteria?

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

Answer: D

Explanation:

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

NEW QUESTION 162

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 volum
- 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 cluste
- F. Promote this replica as a new primary DB instanc
- 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 165

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

Dashboards with advanced analytical queries are used to present data in the eu-west-2 office. Because the corporation will use these dashboards to make purchasing choices, they must have less than a second to obtain application data.

Which solution satisfies these criteria and gives the MOST CURRENT dashboard?

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

Answer: C

Explanation:

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

NEW QUESTION 167

A Database Specialist needs to speed up any failover that might occur on an Amazon Aurora PostgreSQL DB cluster. The Aurora DB cluster currently includes the primary instance and three Aurora Replicas.
How can the Database Specialist ensure that failovers occur with the least amount of downtime for the application?

- A. Set the TCP keepalive parameters low
- B. Call the AWS CLI failover-db-cluster command
- C. Enable Enhanced Monitoring on the DB cluster
- D. Start a database activity stream on the DB cluster

Answer: A

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.BestPractices.html#Aur>

NEW QUESTION 172

A worldwide digital advertising corporation collects browser information in order to provide targeted visitors with contextually relevant pictures, websites, and connections. A single page load may create many events, each of which must be kept separately. A single event may have a maximum size of 200 KB and an average size of 10 KB. Each page load requires a query of the user's browsing history in order to deliver suggestions for targeted advertising. The advertising corporation anticipates daily page views of more than 1 billion from people in the United States, Europe, Hong Kong, and India. The information structure differs according to the event. Additionally, browsing information must be written and read with a very low latency to guarantee that consumers have a positive viewing experience.
Which database solution satisfies these criteria?

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

Answer: C

NEW QUESTION 175

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 178

A Database Specialist is constructing a new Amazon Neptune DB cluster and tries to load data from Amazon S3 using the Neptune bulk loader API. The Database Specialist is confronted with the following error message:

€Unable to establish a connection to the s3 endpoint. The source URL is s3:/mybucket/graphdata/ and the region code is us-east-1. Kindly confirm your Configuration S3.

Which of the following activities should the Database Specialist take to resolve the issue? (Select 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: BE

Explanation:

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

“An IAM role for the Neptune DB instance to assume that has an IAM policy that allows access to the data files in the S3 bucket. The policy must grant Read and List permissions.” “An Amazon S3 VPC endpoint. For more information, see the Creating an Amazon S3 VPC Endpoint section.”

NEW QUESTION 179

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 184

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 188

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 192

In one AWS account, a business runs a two-tier ecommerce application. An Amazon RDS for MySQL

Multi-AZ database instance serves as the application's backend. A developer removed the database instance in the production environment by accident. Although the organization recovers the database, the incident results in hours of outage and financial loss.

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

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

Answer: ACD

NEW QUESTION 193

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

AWS-Certified-Database-Specialty Practice Exam Features:

- * AWS-Certified-Database-Specialty Questions and Answers Updated Frequently
- * AWS-Certified-Database-Specialty Practice Questions Verified by Expert Senior Certified Staff
- * AWS-Certified-Database-Specialty Most Realistic Questions that Guarantee you a Pass on Your First Try
- * AWS-Certified-Database-Specialty Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The AWS-Certified-Database-Specialty Practice Test Here](#)