

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

Which approach meets these requirements?

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

Answer: C

Explanation:

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

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

NEW QUESTION 3

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 4

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 5

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 6

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 7

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 Region
- B. Create a new instance from the snapshot in the ap-southeast-1 Region
- 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 Region
- E. Reconfigure the ap-southeast-1 front-end dashboard to access this instance.
- F. Create a new RDS instance in the ap-southeast-1 Region
- G. Use AWS DMS and change data capture (CDC) to update the new instance in the ap-southeast-1 Region
- 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 Region
- 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 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 ride-hailing application uses an Amazon RDS for MySQL DB instance as persistent storage for bookings. This application is very popular and the company

expects a tenfold increase in the user base in next few months. The application experiences more traffic during the morning and evening hours. This application has two parts:

- > An in-house booking component that accepts online bookings that directly correspond to simultaneous requests from users.
 - > A third-party customer relationship management (CRM) component used by customer care representatives. The CRM uses queries to access booking data.
- A database specialist needs to design a cost-effective database solution to handle this workload. Which solution meets these requirements?

- A. Use Amazon ElastiCache for Redis to accept the 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: D

NEW QUESTION 10

A company maintains several databases using Amazon RDS for MySQL and PostgreSQL. Each RDS database generates log files with retention periods set to their default values. The company has now mandated that database logs be maintained for up to 90 days in a centralized repository to facilitate real-time and after-the-fact analyses.

What should a Database Specialist do to meet these requirements with minimal effort?

- A. Create an AWS Lambda function to pull logs from the RDS databases and consolidate the log files in an Amazon S3 bucket
- B. Set a lifecycle policy to expire the objects after 90 days.
- C. Modify the RDS databases to publish log to Amazon CloudWatch Log
- D. Change the log retention policy for each log group to expire the events after 90 days.
- E. Write a stored procedure in each RDS database to download the logs and consolidate the log files in an Amazon S3 bucket
- F. Set a lifecycle policy to expire the objects after 90 days.
- G. Create an AWS Lambda function to download the logs from the RDS databases and publish the logs to Amazon CloudWatch Log
- H. Change the log retention policy for the log group to expire the events after 90 days.

Answer: B

Explanation:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_LogAccess.html
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_LogAccess.Procedural.UploadtoCloudWat
<https://aws.amazon.com/premiumsupport/knowledge-center/rds-aurora-mysql-logs-cloudwatch/>
https://docs.aws.amazon.com/AmazonCloudWatchLogs/latest/APIReference/API_PutRetentionPolicy.html

NEW QUESTION 10

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 11

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 14

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 18

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 20

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 23

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 24

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

Which solution satisfies these criteria?

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

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

Answer: C

Explanation:

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

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

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 a SecretTargetAttachment resource with the SecretId and TargetId properties set to the Amazon Resource Names (ARNs) of the secret and the database
- 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 ARN
- F. Then, create the secret with a chosen user name and a randomly generated password set by the GenerateSecretString property
- 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 template
- I. Create a resource for the database and reference the secret string for the MasterUserName and MasterUserPassword properties
- J. Then, add a resource of type AWS::SecretsManager::SecretTargetAttachment 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 property
- L. Add a SecretTargetAttachment resource with the SecretId property set to the Amazon Resource Name (ARN) of the secret and the TargetId property set to a parameter value matching the desired database ARN
- M. Then, create a database with the MasterUserName and MasterUserPassword properties set to the previously created values in the secret.

Answer: C

NEW QUESTION 29

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

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

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

Answer: A

NEW QUESTION 31

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 Region
- 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 key
- 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 key
- 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 32

A business is operating an on-premises application that is divided into three tiers: web, application, and MySQL database. The database is predominantly accessed during business hours, with occasional bursts of activity throughout the day. As part of the company's shift to AWS, a database expert wants to increase the availability and minimize the cost of the MySQL database tier.

Which MySQL database choice satisfies these criteria?

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

Answer: B

Explanation:

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

NEW QUESTION 37

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

What should a database specialist do to prevent data loss during the stack upgrade?

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

Answer: D

Explanation:

To preserve your data, perform the following procedure:

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

NEW QUESTION 41

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

Which of the following will resolve this issue?

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

Answer: B

Explanation:

Max_connection is a formula in RDS parameter group:

$\text{GREATEST}(\{\log(\text{DBInstanceClassMemory}/805306368)*45\}, \{\log(\text{DBInstanceClassMemory}/8187281408)*100\})$

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Managing.Performance.htm> You can increase the maximum number of connections to your Aurora MySQL DB instance by scaling the instance up to a DB instance class with more memory, or by setting a larger value for the max_connections parameter in the DB parameter group for your instance, up to 16,000. You must change a larger value for the max_connections parameter in the DB parameter group, not edit my.cnf, it is not physical server hosting MySQL.

NEW QUESTION 44

A team of Database Specialists is currently investigating performance issues on an Amazon RDS for MySQL DB instance and is reviewing related metrics. The team wants to narrow the possibilities down to specific database wait events to better understand the situation.

How can the Database Specialists accomplish this?

- A. Enable the option to push all database logs to Amazon CloudWatch for advanced analysis
- B. Create appropriate Amazon CloudWatch dashboards to contain specific periods of time
- C. Enable Amazon RDS Performance Insights and review the appropriate dashboard
- D. Enable Enhanced Monitoring with the appropriate settings

Answer: C

Explanation:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_PerfInsights.Enabling.html <https://aws.amazon.com/rds/performance-insights/>
<https://aws.amazon.com/blogs/database/tuning-amazon-rds-for-mysql-with-performance-insights/>

NEW QUESTION 45

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 47

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 52

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 55

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 56

A company is using a Single-AZ Amazon RDS for MySQL DB instance for development. The DB instance is experiencing slow performance when queries are executed. Amazon CloudWatch metrics indicate that the instance requires more I/O capacity. Which actions can a database specialist perform to resolve this issue? (Choose two.)

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

Answer: BE

Explanation:

<https://aws.amazon.com/blogs/database/best-storage-practices-for-running-production-workloads-on-hosted-data> "If you find the pattern of IOPS usage consistently going beyond more than 16,000, you should modify the DB instance and change the storage type from gp2 to io1."

NEW QUESTION 58

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

How could a database professional approach these needs while using DynamoDB?

- A. Enable Amazon DynamoDB Streams on the tabl
- B. Create an AWS Lambda function triggered by the strea
- C. Write the log entries to an Amazon S3 object.
- D. Create a backup plan in AWS Backup to back up the DynamoDB table once a da
- E. Create an AWS Lambda function that restores the backup in another table and compares both tables for change
- F. Generate the log entries and write them to an Amazon S3 object.
- G. Enable AWS CloudTrail logs on the tabl
- H. Create an AWS Lambda function that reads the log files once an hour and filters DynamoDB API action
- I. Write the filtered log files to Amazon S3.
- J. Enable Amazon DynamoDB Streams on the tabl
- K. Create an AWS Lambda function triggered by the strea
- L. Write the log entries to an Amazon Kinesis Data Firehose delivery stream with buffering and Amazon S3 as the destination.

Answer: D

NEW QUESTION 61

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 instanc
- B. Create an encrypted copy of the unencrypted snapsho
- C. Restore the encrypted snapshot copy.
- D. Modify the RDS DB instanc
- 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 instanc
- 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 64

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 69

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

backups.

The program has corrupted the database logically, resulting in the application being unresponsive. The exact moment the corruption occurred has been determined, and it occurred within the backup retention period.

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

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

Answer: B

Explanation:

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

NEW QUESTION 74

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 75

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 77

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 82

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 83

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 84

Recently, a gaming firm purchased a popular iOS game that is especially popular during the Christmas season. The business has opted to include a leaderboard into the game, which will be powered by Amazon DynamoDB. The application's load is likely to increase significantly throughout the Christmas season.

Which solution satisfies these criteria at the lowest possible cost?

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

Answer: D

Explanation:

"On-demand is ideal for bursty, new, or unpredictable workloads whose traffic can spike in seconds or minutes"

vs.

"DynamoDB released auto scaling to make it easier for you to manage capacity efficiently, and auto scaling continues to help DynamoDB users lower the cost of workloads that have a predictable traffic pattern."

<https://aws.amazon.com/blogs/database/amazon-dynamodb-auto-scaling-performance-and-cost-optimization-at>

NEW QUESTION 86

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 91

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 95

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

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

Answer: A

NEW QUESTION 100

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

- > The networks and routes affected if a particular component fails.
- > The networks that have redundant routes between them.
- > The networks that do not have redundant routes between them.
- > The fastest path between two networks.

Which database engine meets these requirements?

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

Answer: B

NEW QUESTION 102

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 105

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 106

A database specialist needs to configure an Amazon RDS for MySQL DB instance to close non-interactive connections that are inactive after 900 seconds. What should the database specialist do to accomplish this task?

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

Answer: A

Explanation:

<https://aws.amazon.com/fr/blogs/database/best-practices-for-configuring-parameters-for-amazon-rds-for-mysql-> "You can set parameters globally using a parameter group. Alternatively, you can set them for a particular session using the SET command."

<https://aws.amazon.com/blogs/database/best-practices-for-configuring-parameters-for-amazon-rds-for-mysql-pa>

NEW QUESTION 109

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

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

Answer: A

Explanation:

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

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

Answer: C

NEW QUESTION 111

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 116

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

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

Answer: C

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.Reference.html> <https://blog.dbi-services.com/aws-aurora-xactsync-batch-commit/>

NEW QUESTION 121

On AWS, a business is developing a web application. The application needs that the database supports concurrent read and write activities in several AWS Regions. Additionally, the database must communicate data changes across Regions as they occur. The application must be highly available and have a latency of less than a few hundred milliseconds. Which solution satisfies these criteria?

- A. Amazon DynamoDB global tables
- B. Amazon DynamoDB streams with AWS Lambda to replicate the data
- C. An Amazon ElastiCache for Redis cluster with cluster mode enabled and multiple shards
- D. An Amazon Aurora global database

Answer: A

Explanation:

Aurora Global Databases provides a writer and a reader endpoints in the primary region but only a reader endpoints in other region. Although strongly consistent, it does not fulfill the requirements that "there are plenty of read / write activities" in all regions.

NEW QUESTION 126

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 129

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 Region
- 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 132

A business uses Amazon EC2 instances in VPC A to serve an internal file-sharing application. This application is supported by an Amazon ElastiCache cluster in VPC B that is peering with VPC A. The corporation migrates the instances of its applications from VPC A to VPC B. The file-sharing application is no longer able to connect to the ElastiCache cluster, as shown by the logs.

What is the best course of action for a database professional to take in order to remedy this issue?

- A. Create a second security group on the EC2 instance
- B. Add an outbound rule to allow traffic from the ElastiCache cluster security group.
- C. Delete the ElastiCache security group
- D. Add an interface VPC endpoint to enable the EC2 instances to connect to the ElastiCache cluster.
- E. Modify the ElastiCache security group by adding outbound rules that allow traffic to VPC CIDR blocks from the ElastiCache cluster.
- F. Modify the ElastiCache security group by adding an inbound rule that allows traffic from the EC2 instances security group to the ElastiCache cluster.

Answer: D

Explanation:

<https://docs.aws.amazon.com/vpc/latest/peering/vpc-peering-security-groups.html>

NEW QUESTION 133

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 137

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 140

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

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

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

Answer: C

NEW QUESTION 144

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 147

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

Which combination of actions should the Database Specialist take? (Choose three.)

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

Answer: BCF

Explanation:

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

NEW QUESTION 152

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 157

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-glob>

NEW QUESTION 158

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 163

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

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

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

Answer: A

Explanation:

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

NEW QUESTION 165

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 170

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

How should this issue be resolved?

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

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 172

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 174

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

"Amazon Invalid operation: S3ServiceException:Access Denied,Status 403,Error AccessDenied."

The developers need to load this data soon, so a database specialist must act quickly to solve this issue. What is the MOST secure solution?

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

Answer: B

Explanation:

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

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

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

NEW QUESTION 177

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 cant 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 181

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 auditin
- D. Enable Multi-AZ after auditing is enabled.
- E. Disable automated backup on the DB instance, and then enable auditin
- 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 186

A Database Specialist has migrated an on-premises Oracle database to Amazon Aurora PostgreSQL. The schema and the data have been migrated successfully.

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

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

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

Answer: C

Explanation:

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/events/Create-CloudWatch-Events-Scheduled-Rule.html> <https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/schedule-jobs-for-amazon-rds-and-aurora-postgresql.html> a job for data extraction or a job for data purging can easily be scheduled using cron. For these jobs, database credentials are typically either hard-coded or stored in a properties file. However, when you migrate to Amazon Relational Database Service (Amazon RDS) or Amazon Aurora PostgreSQL, you lose the ability to log in to the host instance to schedule cron jobs. This pattern describes how to use AWS Lambda and AWS Secrets Manager to schedule jobs for Amazon RDS and Aurora PostgreSQL databases after migration. <https://docs.aws.amazon.com/AmazonCloudWatch/latest/events/RunLambdaSchedule.html>

NEW QUESTION 187

A company uses the Amazon DynamoDB table contractDB in us-east-1 for its contract system with the following schema:
 orderID (primary key) timestamp (sort key) contract (map) createdBy (string) customerEmail (string)

After a problem in production, the operations team has asked a database specialist to provide an IAM policy to read items from the database to debug the application. In addition, the developer is not allowed to access the value of the customerEmail field to stay compliant.

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

A)

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

B)

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

C)

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

D)

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

- A. Option A
- B. Option B
- C. Option C
- D. Option C

Answer: A

NEW QUESTION 190

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 191

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

Occasionally, the PutItem request fails with a ResourceNotFoundException error, which causes the workflow to fail. The development team has confirmed that the same table name is used in the two API calls.

How should a database specialist fix this issue?

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

Answer: C

Explanation:

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

NEW QUESTION 196

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 200

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

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

Answer: C

Explanation:

"A VPC peering connection is a networking connection between two VPCs that enables you to route traffic between them using private IP addresses. Instances in either VPC can communicate with each other as if they are within the same network. You can create a VPC peering connection between your own VPCs, with a VPC in another AWS account, or with a VPC in a different AWS Region." https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_VPC.Scenarios.html

NEW QUESTION 205

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 209

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 instanc
- B. Write a CLI command that creates a backup of the DynamoDB tabl
- 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 tabl
- 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%2>
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/backuprestore_HowItWorks.html

NEW QUESTION 213

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 216

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 221

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 222

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 226

A bank intends to utilize Amazon RDS to host a MySQL database instance. The database should be able to handle high-volume read requests with extremely few repeated queries.

Which solution satisfies these criteria?

- A. Create an Amazon ElastiCache cluster
- B. Use a write-through strategy to populate the cache.
- C. Create an Amazon ElastiCache cluster
- D. Use a lazy loading strategy to populate the cache.
- E. Change the DB instance to Multi-AZ with a standby instance in another AWS Region.
- F. Create a read replica of the DB instance
- G. Use the read replica to distribute the read traffic.

Answer: D

NEW QUESTION 228

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 230

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 232

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

What is the most likely reason for this?

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

Answer: D

Explanation:

>Your source DB instance must have backup retention enabled.

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

NEW QUESTION 235

A large company is using an Amazon RDS for Oracle Multi-AZ DB instance with a Java application. As a part of its disaster recovery annual testing, the company would like to simulate an Availability Zone failure and record how the application reacts during the DB instance failover activity. The company does not want to make any code changes for this activity.

What should the company do to achieve this in the shortest amount of time?

- A. Use a blue-green deployment with a complete application-level failover test
- B. Use the RDS console to reboot the DB instance by choosing the option to reboot with failover
- C. Use RDS fault injection queries to simulate the primary node failure
- D. Add a rule to the NACL to deny all traffic on the subnets associated with a single Availability Zone

Answer: B

Explanation:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_RebootInstance.html <https://exain.wordpress.com/2017/07/12/amazon-rds-multi-az-setup-failover-simulation/>

"Rebooting with failover is beneficial when you want to simulate a failure of a DB instance for testing, or restore operations to the original AZ after a failover occurs."

NEW QUESTION 237

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

What is most secure solution to store the master password?

- A. Store the master password in a parameter file in each environment
- B. Reference the environment-specific parameter file in the CloudFormation template.
- C. Encrypt the master password using an AWS KMS key
- D. Store the encrypted master password in the CloudFormation template.

- E. Use the secretsmanager dynamic reference to retrieve the master password stored in AWS Secrets Manager and enable automatic rotation.
- F. Use the ssm dynamic reference to retrieve the master password stored in the AWS Systems Manager Parameter Store and enable automatic rotation.

Answer: C

Explanation:

"By using the secure string support in CloudFormation with dynamic references you can better maintain your infrastructure as code. You'll be able to avoid hard coding passwords into your templates and you can keep these runtime configuration parameters separated from your code. Moreover, when properly used, secure strings will help keep your development and production code as similar as possible, while continuing to make your infrastructure code suitable for continuous deployment pipelines."

<https://aws.amazon.com/blogs/mt/using-aws-systems-manager-parameter-store-secure-string-parameters-in-aws> <https://aws.amazon.com/blogs/security/how-to-use-aws-secrets-manager-rotate-credentials-amazon-rds-database>

NEW QUESTION 240

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 databases
- 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 241

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 242

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 243

.....

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](#)