

Amazon-Web-Services

Exam Questions DBS-C01

AWS Certified Database - Specialty



NEW QUESTION 1

A company's applications store data in Amazon Aurora MySQL DB clusters. The company has separate AWS accounts for its production, test, and development environments. To test new functionality in the test environment, the company's development team requires a copy of the production database four times a day. Which solution meets this requirement with the MOST operational efficiency?

- A. Take a manual snapshot in the production account
- B. Share the snapshot with the test account
- C. Restore the database from the snapshot.
- D. Take a manual snapshot in the production account
- E. Export the snapshot to Amazon S3. Copy the snapshot to an S3 bucket in the test account
- F. Restore the database from the snapshot.
- G. Share the Aurora DB cluster with the test account
- H. Create a snapshot of the production database in the test account
- I. Restore the database from the snapshot.
- J. Share the Aurora DB cluster with the test account
- K. Create a clone of the production database in the test account.

Answer: D

Explanation:

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

NEW QUESTION 2

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

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

Answer: C

Explanation:

For details about using Aurora and Amazon Comprehend together, see [Using Amazon Comprehend for sentiment detection](https://www.stackovercloud.com/2019/11/27/new-for-amazon-aurora-use-machine-learning-directly-from-your). Aurora machine learning uses a highly optimized integration between the Aurora database and the AWS machine learning (ML) services SageMaker and Amazon Comprehend.
<https://www.stackovercloud.com/2019/11/27/new-for-amazon-aurora-use-machine-learning-directly-from-your>

NEW QUESTION 3

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 4

A database specialist has been entrusted by an ecommerce firm with designing a reporting dashboard that visualizes crucial business KPIs derived from the company's primary production database running on Amazon Aurora. The dashboard should be able to read data within 100 milliseconds after an update. The Database Specialist must conduct an audit of the Aurora DB cluster's present setup and provide a cost-effective alternative. The solution must support the unexpected read demand generated by the reporting dashboard without impairing the DB cluster's write availability and performance. Which solution satisfies these criteria?

- 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: D

NEW QUESTION 5

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

ERROR: cloud not write block 7507718 of temporary file: No space left on device

What is the cause of this error and what should the Database Specialist do to resolve this issue?

- A. The scaling of Aurora storage cannot catch up with the data loadin
- B. The Database Specialist needs to modify the workload to load the data slowly.
- C. The scaling of Aurora storage cannot catch up with the data loadin
- D. The Database Specialist needs to enable Aurora storage scaling.
- E. The local storage used to store temporary tables is ful
- F. The Database Specialist needs to scale up the instance.
- G. The local storage used to store temporary tables is ful
- H. The Database Specialist needs to enable local storage scaling.

Answer: C

NEW QUESTION 6

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 7

A database specialist is designing an enterprise application for a large company. The application uses Amazon DynamoDB with DynamoDB Accelerator (DAX). The database specialist observes that most of the queries are not found in the DAX cache and that they still require DynamoDB table reads.

What should the database specialist review first to improve the utility of DAX?

- A. The DynamoDB ConsumedReadCapacityUnits metric
- B. The trust relationship to perform the DynamoDB API calls
- C. The DAX cluster's TTL setting
- D. The validity of customer-specified AWS Key Management Service (AWS KMS) keys for DAX encryption at rest

Answer: C

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/DAX.cluster-management.html#DAX.clu>

NEW QUESTION 8

An online gaming company is using an Amazon DynamoDB table in on-demand mode to store game scores. After an intensive advertisement campaign in South America, the average number of concurrent users rapidly increases from 100,000 to 500,000 in less than 10 minutes every day around 5 PM.

The on-call software reliability engineer has observed that the application logs contain a high number of DynamoDB throttling exceptions caused by game score insertions around 5 PM. Customer service has also reported that several users are complaining about their scores not being registered.

How should the database administrator remediate this issue at the lowest cost?

- A. Enable auto scaling and set the target usage rate to 90%.
- B. Switch the table to provisioned mode and enable auto scaling.
- C. Switch the table to provisioned mode and set the throughput to the peak value.
- D. Create a DynamoDB Accelerator cluster and use it to access the DynamoDB table.

Answer: B

NEW QUESTION 9

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 authenticatio
- 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 agai
- 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 10

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

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

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

Which choice of keys is recommended for the DynamoDB table?

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

Answer: D

Explanation:

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

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

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

NEW QUESTION 10

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

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

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

Answer: B

Explanation:

<https://aws.amazon.com/blogs/database/new-aws-dms-and-aws-snowball-integration-enables-mass-database-mi>

NEW QUESTION 14

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 18

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

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

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

- L. Copy the snapshot into an Amazon S3 bucket and import the snapshot into a new RDS for MySQL DB instance using the MySQL utilities running on an EC2 instance
- M. Point the application to the DB instance.
- N. Deploy a new RDS for MySQL DB instance and configure it for access from the on-premises data center
- O. Use the mysqldump utility to create an initial snapshot from the on-premises MySQL server, and copy it to an Amazon S3 bucket
- P. Import the snapshot into the DB instance using the MySQL utilities running on an Amazon EC2 instance
- Q. Establish replication into the new DB instance using MySQL replication
- R. Stop application access to the on-premises MySQL server and let the remaining transactions replicate over
- S. Point the application to the DB instance.

Answer: D

Explanation:

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

NEW QUESTION 23

A company is looking to move an on-premises IBM Db2 database running AIX on an IBM POWER7 server. Due to escalating support and maintenance costs, the company is exploring the option of moving the workload to an Amazon Aurora PostgreSQL DB cluster. What is the quickest way for the company to gather data on the migration compatibility?

- A. Perform a logical dump from the Db2 database and restore it to an Aurora DB cluster
- B. Identify the gaps and compatibility of the objects migrated by comparing row counts from source and target tables.
- C. Run AWS DMS from the Db2 database to an Aurora DB cluster
- D. Identify the gaps and compatibility of the objects migrated by comparing the row counts from source and target tables.
- E. Run native PostgreSQL logical replication from the Db2 database to an Aurora DB cluster to evaluate the migration compatibility.
- F. Run the AWS Schema Conversion Tool (AWS SCT) from the Db2 database to an Aurora DB cluster. Create a migration assessment report to evaluate the migration compatibility.

Answer: D

NEW QUESTION 25

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 29

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

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

Answer: C

Explanation:

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

NEW QUESTION 30

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 35

A vehicle insurance company needs to choose a highly available database to track vehicle owners and their insurance details. The persisted data should be immutable in the database, including the complete and sequenced history of changes over time with all the owners and insurance transfer details for a vehicle. The data should be easily verifiable for the data lineage of an insurance claim. Which approach meets these requirements with MINIMAL effort?

- A. Create a blockchain to store the insurance detail
- B. Validate the data using a hash function to verify the data lineage of an insurance claim.
- C. Create an Amazon DynamoDB table to store the insurance detail
- D. Validate the data using AWS DMS validation by moving the data to Amazon S3 to verify the data lineage of an insurance claim.
- E. Create an Amazon QLDB ledger to store the insurance detail
- F. Validate the data by choosing the ledger name in the digest request to verify the data lineage of an insurance claim.
- G. Create an Amazon Aurora database to store the insurance detail
- H. Validate the data using AWS DMS validation by moving the data to Amazon S3 to verify the data lineage of an insurance claim.

Answer: C

NEW QUESTION 40

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

- A. Enable Amazon DynamoDB Streams on the table
- B. Create an AWS Lambda function triggered by the stream
- 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 day
- 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 table
- 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 table
- K. Create an AWS Lambda function triggered by the stream
- 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 45

A large IT hardware manufacturing company wants to deploy a MySQL database solution in the AWS Cloud. The solution should quickly create copies of the company's production databases for test purposes. The solution must deploy the test databases in minutes, and the test data should match the latest production data as closely as possible. Developers must also be able to make changes in the test database and delete the instances afterward. Which solution meets these requirements?

- A. Leverage Amazon RDS for MySQL with write-enabled replicas running on Amazon EC2. Create the test copies using a mysqldump backup from the RDS for MySQL DB instances and importing them into the new EC2 instances.
- B. Leverage Amazon Aurora MySQL
- C. Use database cloning to create multiple test copies of the production DB clusters.
- D. Leverage Amazon Aurora MySQL
- E. Restore previous production DB instance snapshots into new test copies of Aurora MySQL DB clusters to allow them to make changes.
- F. Leverage Amazon RDS for MySQL
- G. Use database cloning to create multiple developer copies of the production DB instance.

Answer: B

NEW QUESTION 48

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 53

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

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

Answer: A

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.AuroraMySQL.Monitoring.Metrics>. Unlike for other DB engines, for Aurora DB instances this metric reports the amount of storage available to each DB instance. This value depends on the DB instance class (for pricing information, see the Amazon RDS product page). You can increase the amount of free storage space for an instance by choosing a larger DB instance class for your instance."

NEW QUESTION 55

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

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

Answer: C

NEW QUESTION 56

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

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

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

Answer: C

Explanation:

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

[https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/HowItWorks.ReadWriteCapacityMode.h](https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/HowItWorks.ReadWriteCapacityMode.html)

Amazon DynamoDB auto scaling uses the AWS Application Auto Scaling service to dynamically adjust provisioned throughput capacity on your behalf

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

NEW QUESTION 60

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 63

A database specialist is working on an Amazon RDS for PostgreSQL DB instance that is experiencing application performance issues due to the addition of new workloads. The database has 5 TB of storage space with Provisioned IOPS. Amazon CloudWatch metrics show that the average disk queue depth is greater than 200 and that the disk I/O response time is significantly higher than usual.

What should the database specialist do to improve the performance of the application immediately?

- A. Increase the Provisioned IOPS rate on the storage.
- B. Increase the available storage space.
- C. Use General Purpose SSD (gp2) storage with burst credits.
- D. Create a read replica to offload Read IOPS from the DB instance.

Answer: A

NEW QUESTION 66

A company has an ecommerce website that runs on AWS. The website uses an Amazon RDS for MySQL database. A database specialist wants to enforce the use of temporary credentials to access the database.

Which solution will meet this requirement?

- A. Use MySQL native database authentication.
- B. Use AWS Secrets Manager to rotate the credentials.
- C. Use AWS Identity and Access Management (IAM) database authentication.
- D. Use AWS Systems Manager Parameter Store for authentication.

Answer: C

NEW QUESTION 71

A company hosts a 2 TB Oracle database in its on-premises data center. A database specialist is migrating the database from on premises to an Amazon Aurora PostgreSQL database on AWS.

The database specialist identifies a problem that relates to compatibility Oracle stores metadata in its data dictionary in uppercase, but PostgreSQL stores the metadata in lowercase. The database specialist must resolve this problem to complete the migration.

What is the MOST operationally efficient solution that meets these requirements?

- A. Override the default uppercase format of Oracle schema by encasing object names in quotation marks during creation.
- B. Use AWS Database Migration Service (AWS DMS) mapping rules with rule-action as convert-lowercase.
- C. Use the AWS Schema Conversion Tool conversion agent to convert the metadata from uppercase to lowercase.
- D. Use an AWS Glue job that is attached to an AWS Database Migration Service (AWS DMS) replication task to convert the metadata from uppercase to lowercase.

Answer: B

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/dms-mapping-oracle-postgresql/>

NEW QUESTION 74

A development team asks a database specialist to create a copy of a production Amazon RDS for MySQL DB instance every morning. The development team will use the copied DB instance as a testing environment for development. The original DB instance and the copy will be hosted in different VPCs of the same AWS account. The development team wants the copy to be available by 6 AM each day and wants to use the same endpoint address each day.

Which combination of steps should the database specialist take to meet these requirements MOST cost-effectively? (Choose three.)

- A. Create a snapshot of the production database each day before the 6 AM deadline.
- B. Create an RDS for MySQL DB instance from the snapshot.
- C. Select the desired DB instance size.
- D. Update a defined Amazon Route 53 CNAME record to point to the copied DB instance.
- E. Set up an AWS Database Migration Service (AWS DMS) migration task to copy the snapshot to the copied DB instance.
- F. Use the CopySnapshot action on the production DB instance to create a snapshot before 6 AM.
- G. Update a defined Amazon Route 53 alias record to point to the copied DB instance.

Answer: ABC

NEW QUESTION 76

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 77

The website of a manufacturing firm makes use of an Amazon Aurora PostgreSQL database cluster. Which settings will result in the LEAST amount of downtime for the application during failover? (Select

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: ACE

Explanation:

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

NEW QUESTION 80

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 81

A company is launching a new Amazon RDS for MySQL Multi-AZ DB instance to be used as a data store for a custom-built application. After a series of tests with point-in-time recovery disabled, the company decides that it must have point-in-time recovery reenabled before using the DB instance to store production data. What should a database specialist do so that point-in-time recovery can be successful?

- A. Enable binary logging in the DB parameter group used by the DB instance.
- B. Modify the DB instance and enable audit logs to be pushed to Amazon CloudWatch Logs.
- C. Modify the DB instance and configure a backup retention period.
- D. Set up a scheduled job to create manual DB instance snapshots.

Answer: C

Explanation:

You can restore a DB instance to a specific point in time (PITR), creating a new DB instance. To support PITR, your DB instances must have backup retention set to a nonzero value. <https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/custom-backup-sqlserver.html>
<https://aws.amazon.com/blogs/database/setting-up-a-binlog-server-for-amazon-rds-mysql-and-mariadb-using-m> "After you run the command, it's okay to enable backup retention on the RDS instance by using the AWS CLI or the console. Enabling backup retention also enables binary logging."
<https://aws.amazon.com/blogs/storage/point-in-time-recovery-and-continuous-backup-for-amazon-rds-with-aws>

NEW QUESTION 82

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 85

A company is running Amazon RDS for MySQL for its workloads. There is downtime when AWS operating system patches are applied during the Amazon RDS-specified maintenance window.

What is the MOST cost-effective action that should be taken to avoid downtime?

- A. Migrate the workloads from Amazon RDS for MySQL to Amazon DynamoDB.
- B. Enable cross-Region read replicas and direct read traffic to them when Amazon RDS is down.

- C. Enable a read replicas and direct read traffic to it when Amazon RDS is down
- D. Enable an Amazon RDS for MySQL Multi-AZ configuration

Answer: D

Explanation:

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

To minimize downtime, modify the Amazon RDS DB instance to a Multi-AZ deployment. For Multi-AZ deployments, OS maintenance is applied to the secondary instance first, then the instance fails over, and then the primary instance is updated. The downtime is during failover. For more information, see Maintenance for Multi-AZ Deployments. <https://aws.amazon.com/rds/faqs/> The availability benefits of Multi-AZ also extend to planned maintenance. For example, with automated backups, I/O activity is no longer suspended on your primary during your preferred backup window, since backups are taken from the standby. In the case of patching or DB instance class scaling, these operations occur first on the standby, prior to automatic fail over. As a result, your availability impact is limited to the time required for automatic failover to complete.

NEW QUESTION 88

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

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

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

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

Answer: ACD

NEW QUESTION 89

A gaming company has recently acquired a successful iOS game, which is particularly popular during the holiday season. The company has decided to add a leaderboard to the game that uses Amazon DynamoDB. The application load is expected to ramp up over the holiday season.

Which solution will meet these requirements at the lowest cost?

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

Answer: C

NEW QUESTION 93

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 95

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

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

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

Answer: BE

Explanation:

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

NEW QUESTION 97

A company hosts an on-premises Microsoft SQL Server Enterprise edition database with Transparent Data Encryption (TDE) enabled. The database is 20 TB in size and includes sparse tables. The company needs to migrate the database to Amazon RDS for SQL Server during a maintenance window that is scheduled for an upcoming weekend. Data-at-rest encryption must be enabled for the target DB instance.

Which combination of steps should the company take to migrate the database to AWS in the MOST operationally efficient manner? (Choose two.)

- A. Use AWS Database Migration Service (AWS DMS) to migrate from the on-premises source database to the RDS for SQL Server target database.
- B. Disable TD
- C. Create a database backup without encryptio
- D. Copy the backup to Amazon S3.
- E. Restore the backup to the RDS for SQL Server DB instanc
- F. Enable TDE for the RDS for SQL Server DB instance.
- G. Set up an AWS Snowball Edge devic
- H. Copy the database backup to the devic
- I. Send the device to AW
- J. Restore the database from Amazon S3.
- K. Encrypt the data with client-side encryption before transferring the data to Amazon RDS.

Answer: BC

Explanation:

<https://aws.amazon.com/blogs/database/migrate-tde-enabled-sql-server-databases-to-amazon-rds-for-sql-server/>

NEW QUESTION 99

A small startup company is looking to migrate a 4 TB on-premises MySQL database to AWS using an Amazon RDS for MySQL DB instance. Which strategy would allow for a successful migration with the LEAST amount of downtime?

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

Answer: B

NEW QUESTION 101

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 106

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

- A. Enable Aurora Database Activity Streams on the database in synchronous mod
- B. Connect the Amazon Kinesis data stream to Kinesis Data Firehos
- C. Set the Kinesis Data Firehose destination to an Amazon S3 bucket.
- D. Create an AWS CloudTrail trail in the Region where the database run

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

Answer: C

Explanation:

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

NEW QUESTION 111

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

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

Answer: BD

NEW QUESTION 116

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 121

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 123

A gaming company is evaluating Amazon ElastiCache as a solution to manage player leaderboards. Millions of players around the world will complete in annual tournaments. The company wants to implement an architecture that is highly available. The company also wants to ensure that maintenance activities have minimal impact on the availability of the gaming platform.

Which combination of steps should the company take to meet these requirements? (Choose two.)

- A. Deploy an ElastiCache for Redis cluster with read replicas and Multi-AZ enabled.
- B. Deploy an ElastiCache for Memcached global datastore.
- C. Deploy a single-node ElastiCache for Redis cluster with automatic backups enable
- D. In the event of a failure, create a new cluster and restore data from the most recent backup.
- E. Use the default maintenance window to apply any required system changes and mandatory updates as soon as they are available.
- F. Choose a preferred maintenance window at the time of lowest usage to apply any required changes and mandatory updates.

Answer: AE

Explanation:

<https://aws.amazon.com/blogs/database/configuring-amazon-elasticache-for-redis-for-higher-availability/>

NEW QUESTION 127

A company uses Amazon Aurora MySQL as the primary database engine for many of its applications. A database specialist must create a dashboard to provide the company with information about user connections to databases. According to compliance requirements, the company must retain all connection logs for at least 7 years.

Which solution will meet these requirements MOST cost-effectively?

- A. Enable advanced auditing on the Aurora cluster to log CONNECT event
- B. Export audit logs from Amazon CloudWatch to Amazon S3 by using an AWS Lambda function that is invoked by an Amazon EventBridge (Amazon CloudWatch Events) scheduled event
- C. Build a dashboard by using Amazon QuickSight.
- D. Capture connection attempts to the Aurora cluster with AWS Cloud Trail by using the DescribeEvents API operation
- E. Create a CloudTrail trail to export connection logs to Amazon S3. Build a dashboard by using Amazon QuickSight.
- F. Start a database activity stream for the Aurora cluster
- G. Push the activity records to an Amazon Kinesis data stream
- H. Build a dynamic dashboard by using AWS Lambda.
- I. Publish the DatabaseConnections metric for the Aurora DB instances to Amazon CloudWatch
- J. Build a dashboard by using CloudWatch dashboards.

Answer: A

Explanation:

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

NEW QUESTION 129

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 132

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 136

A pharmaceutical company's drug search API is using an Amazon Neptune DB cluster. A bulk uploader process automatically updates the information in the database a few times each week. A few weeks ago during a bulk upload, a database specialist noticed that the database started to respond frequently with a ThrottlingException error. The problem also occurred with subsequent uploads.

The database specialist must create a solution to prevent ThrottlingException errors for the database. The solution must minimize the downtime of the cluster. Which solution meets these requirements?

- A. Create a read replica that uses a larger instance size than the primary DB instance
- B. Fail over the primary DB instance to the read replica.
- C. Add a read replica to each Availability Zone
- D. Use an instance for the read replica that is the same size as the primary DB instance
- E. Keep the traffic between the API and the database within the Availability Zone.
- F. Create a read replica that uses a larger instance size than the primary DB instance
- G. Offload the reads from the primary DB instance.
- H. Take the latest backup, and restore it in a DB cluster of a larger size
- I. Point the application to the newly created DB cluster.

Answer: C

Explanation:

<https://docs.aws.amazon.com/neptune/latest/userguide/manage-console-add-replicas.html>

Neptune replicas connect to the same storage volume as the primary DB instance and support only read operations. Neptune replicas can offload read workloads from the primary DB instance.

NEW QUESTION 139

A startup company in the travel industry wants to create an application that includes a personal travel assistant to display information for nearby airports based on user location. The application will use Amazon DynamoDB and must be able to access and display attributes such as airline names, arrival times, and flight numbers. However, the application must not be able to access or display pilot names or passenger counts.

Which solution will meet these requirements MOST cost-effectively?

- A. Use a proxy tier between the application and DynamoDB to regulate access to specific tables, items, and attributes.
- B. Use IAM policies with a combination of IAM conditions and actions to implement fine-grained access control.
- C. Use DynamoDB resource policies to regulate access to specific tables, items, and attributes.
- D. Configure an AWS Lambda function to extract only allowed attributes from tables based on user profiles.

Answer: B

Explanation:

<https://aws.amazon.com/blogs/aws/fine-grained-access-control-for-amazon-dynamodb/>

NEW QUESTION 143

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 148

A finance company migrated its 3 TB on-premises PostgreSQL database to an Amazon Aurora PostgreSQL DB cluster. During a review after the migration, a database specialist discovers that the database is not encrypted at rest. The database must be encrypted at rest as soon as possible to meet security requirements. The database specialist must enable encryption for the DB cluster with minimal downtime.

Which solution will meet these requirements?

- A. Modify the unencrypted DB cluster using the AWS Management Console
- B. Enable encryption and choose to apply the change immediately.
- C. Take a snapshot of the unencrypted DB cluster and restore it to a new DB cluster with encryption enabled
- D. Update any database connection strings to reference the new DB cluster endpoint, and then delete the unencrypted DB cluster.
- E. Create an encrypted Aurora Replica of the unencrypted DB cluster
- F. Promote the Aurora Replica as the new master.
- G. Create a new DB cluster with encryption enabled and use the `pg_dump` and `pg_restore` utilities to load data to the new DB cluster
- H. Update any database connection strings to reference the new DB cluster endpoint, and then delete the unencrypted DB cluster.

Answer: B

Explanation:

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

NEW QUESTION 150

A Database Specialist modified an existing parameter group currently associated with a production Amazon RDS for SQL Server Multi-AZ DB instance. The change is associated with a static parameter type, which controls the number of user connections allowed on the most critical RDS SQL Server DB instance for the company. This change has been approved for a specific maintenance window to help minimize the impact on users.

How should the Database Specialist apply the parameter group change for the DB instance?

- A. Select the option to apply the change immediately
- B. Allow the preconfigured RDS maintenance window for the given DB instance to control when the change is applied
- C. Apply the change manually by rebooting the DB instance during the approved maintenance window
- D. Reboot the secondary Multi-AZ DB instance

Answer: C

Explanation:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_WorkingWithParamGroups.html#USER_W

NEW QUESTION 153

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 154

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 157

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

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

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

Answer: B

Explanation:

This link seems to indicate that more information is required to determine if the Enterprise instance is a candidate for downgrading to Standard.

<https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/determine-whether-your-microsoft-sql-server>

<https://calculator.aws/#/createCalculator/RDSSQLServer>

NEW QUESTION 161

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 164

.....

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

DBS-C01 Practice Exam Features:

- * DBS-C01 Questions and Answers Updated Frequently
- * DBS-C01 Practice Questions Verified by Expert Senior Certified Staff
- * DBS-C01 Most Realistic Questions that Guarantee you a Pass on Your First Try
- * DBS-C01 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The DBS-C01 Practice Test Here](#)