

Exam Questions AWS-Certified-Database-Specialty

AWS Certified Database - Specialty

<https://www.2passeasy.com/dumps/AWS-Certified-Database-Specialty/>



NEW QUESTION 1

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

How can this solution be implemented?

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

Answer: A

Explanation:

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

NEW QUESTION 2

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

Which database specialist solution should be recommended?

- A. Enable Aurora Database Activity Streams on the database in synchronous mode.
- 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 3

A user has a non-relational key-value database. The user is looking for a fully managed AWS service that will offload the administrative burdens of operating and scaling distributed databases. The solution must be cost-effective and able to handle unpredictable application traffic.

What should a Database Specialist recommend for this user?

- A. Create an Amazon DynamoDB table with provisioned capacity mode.
- B. Create an Amazon DocumentDB cluster.
- C. Create an Amazon DynamoDB table with on-demand capacity mode.
- D. Create an Amazon Aurora Serverless DB cluster.

Answer: C

NEW QUESTION 4

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

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

- A. Set up a VPN tunnel for encrypting data over the network from the data center to AWS.
- B. Leverage AWS SCT and apply the converted schema to Amazon Redshift.
- 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 Redshift.
- E. Start an AWS DMS task with two AWS Snowball Edge devices to copy data from on-premises to Amazon S3 with AWS KMS encryption.
- F. Use AWS DMS to finish copying data to Amazon Redshift.
- G. Leverage AWS SCT and apply the converted schema to Amazon Redshift.
- 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 encryption.
- 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 AWS.
- 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 encryption.
- 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 5

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 6

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 7

A company has an ecommerce web application with an Amazon RDS for MySQL DB instance. The marketing team has noticed some unexpected updates to the product and pricing information on the website, which is impacting sales targets. The marketing team wants a database specialist to audit future database activity to help identify how and when the changes are being made.

What should the database specialist do to meet these requirements? (Choose two.)

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

Answer: BD

Explanation:

<https://aws.amazon.com/blogs/database/configuring-an-audit-log-to-capture-database-activities-for-amazon-rds>

NEW QUESTION 8

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

The firm demands a response time of five minutes and a response time of five minutes. A database professional must create a disaster recovery system that is both efficient and has a low replication latency.

How should the database professional tackle these requirements?

- A. Configure AWS Database Migration Service (AWS DMS) and create a replica in a different AWS Region.
- B. Configure an Amazon Aurora global database and add a different AWS Region.
- C. Configure a binlog and create a replica in a different AWS Region.
- D. Configure a cross-Region read replica.

Answer: B

Explanation:

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

NEW QUESTION 9

A gaming company has implemented a leaderboard in AWS using a Sorted Set data structure within Amazon ElastiCache for Redis. The ElastiCache cluster has been deployed with cluster mode disabled and has a replication group deployed with two additional replicas. The company is planning for a worldwide gaming event and is anticipating a higher write load than what the current cluster can handle.

Which method should a Database Specialist use to scale the ElastiCache cluster ahead of the upcoming event?

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

Answer: B

NEW QUESTION 10

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

Which database solution meets these requirements?

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

Answer: C

NEW QUESTION 10

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

Which solution satisfies these criteria?

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

Answer: D

Explanation:

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

NEW QUESTION 14

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 localstorage scaling.

Answer: C

NEW QUESTION 15

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

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

milliseconds rather than microseconds.

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

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

Answer: C

NEW QUESTION 18

A 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 22

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-rotationchedule>

NEW QUESTION 26

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 31

A banking company recently launched an Amazon RDS for MySQL DB instance as part of a proof-of-concept project. A database specialist has configured

automated database snapshots. As a part of routine testing, the database specialist noticed one day that the automated database snapshot was not created. Which of the following are possible reasons why the snapshot was not created? (Choose two.)

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

Answer: AD

Explanation:

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

NEW QUESTION 33

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 34

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 35

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 36

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 38

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 42

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

Which approach will meet these requirements?

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

Answer: C

Explanation:

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

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

NEW QUESTION 46

A company is concerned about the cost of a large-scale, transactional application using Amazon DynamoDB that only needs to store data for 2 days before it is deleted. In looking at the tables, a Database Specialist notices that much of the data is months old, and goes back to when the application was first deployed.

What can the Database Specialist do to reduce the overall cost?

- A. Create a new attribute in each table to track the expiration time and create an AWS Glue transformation to delete entries more than 2 days old.
- B. Create a new attribute in each table to track the expiration time and enable DynamoDB Streams on each table.
- C. Create a new attribute in each table to track the expiration time and enable time to live (TTL) on each table.
- D. Create an Amazon CloudWatch Events event to export the data to Amazon S3 daily using AWS Data Pipeline and then truncate the Amazon DynamoDB table.

Answer: C

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/TTL.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 49

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 51

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 55

A financial services company is developing a shared data service that supports different applications from throughout the company. A Database Specialist designed a solution to leverage Amazon ElastiCache for Redis with cluster mode enabled to enhance performance and scalability. The cluster is configured to listen on port 6379.

Which combination of steps should the Database Specialist take to secure the cache data and protect it from unauthorized access? (Choose three.)

- A. Enable in-transit and at-rest encryption on the ElastiCache cluster.
- B. Ensure that Amazon CloudWatch metrics are configured in the ElastiCache cluster.
- C. Ensure the security group for the ElastiCache cluster allows all inbound traffic from itself and inbound traffic on TCP port 6379 from trusted clients only.
- D. Create an IAM policy to allow the application service roles to access all ElastiCache API actions.
- E. Ensure the security group for the ElastiCache clients authorize inbound TCP port 6379 and port 22 traffic from the trusted ElastiCache cluster's security group.
- F. Ensure the cluster is created with the auth-token parameter and that the parameter is used in all subsequent commands.

Answer: ACF

Explanation:

<https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/encryption.html>

NEW QUESTION 58

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

- A. Use Amazon DynamoDB global tables to synchronize transactions
- B. Use Amazon EMR to copy the orders table data across Regions
- C. Use Amazon Aurora Global Database to synchronize all transactions
- D. Use Amazon DynamoDB Streams to replicate all DynamoDB transactions and sync them

Answer: A

Explanation:

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

With global tables, your globally distributed applications can access data locally in the selected regions to get single-digit millisecond read and write performance. Not Aurora Global Database, as per this link: https://aws.amazon.com/rds/aurora/global-database/?nc1=h_ls . Aurora Global Database lets you easily scale database reads across the world and place your applications close to your users.

NEW QUESTION 63

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 67

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 72

A company is releasing a new mobile game featuring a team play mode. As a group of mobile device users play together, an item containing their statuses is updated in an Amazon DynamoDB table. Periodically, the other users' devices read the latest statuses of their teammates from the table using the BatchGetItem operation.

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

Which recommendation would resolve this issue?

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

Answer: D

Explanation:

https://docs.aws.amazon.com/ja_jp/amazondynamodb/latest/developerguide/API_BatchGetItem_v20111205.htm By default, BatchGetItem performs eventually consistent reads on every table in the request. If you want strongly consistent reads instead, you can set ConsistentRead to true for any or all tables.

NEW QUESTION 75

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

Which approach will MOST effectively meet these requirements?

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

Answer: D

Explanation:

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

NEW QUESTION 78

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

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

Answer: C

Explanation:

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

NEW QUESTION 81

An online advertising website uses an Amazon DynamoDB table with on-demand capacity mode as its data store. The website also has a DynamoDB Accelerator (DAX) cluster in the same VPC as its web application server. The application needs to perform infrequent writes and many strongly consistent reads from the data store by querying the DAX cluster.

During a performance audit, a systems administrator notices that the application can look up items by using the DAX cluster. However, the QueryCacheHits metric for the DAX cluster consistently shows 0 while the QueryCacheMisses metric continuously keeps growing in Amazon CloudWatch. What is the MOST likely reason for this occurrence?

- A. A VPC endpoint was not added to access DynamoDB.
- B. Strongly consistent reads are always passed through DAX to DynamoDB.
- C. DynamoDB is scaling due to a burst in traffic, resulting in degraded performance.
- D. A VPC endpoint was not added to access CloudWatch.

Answer: B

Explanation:

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

"If the request specifies strongly consistent reads, DAX passes the request through to DynamoDB. The results from DynamoDB are not cached in DAX. Instead, they are simply returned to the application."

NEW QUESTION 82

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

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

Answer: B

Explanation:

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

NEW QUESTION 91

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

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

Answer: A

Explanation:

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

NEW QUESTION 92

A Database Specialist is creating Amazon DynamoDB tables, Amazon CloudWatch alarms, and associated infrastructure for an Application team using a development AWS account. The team wants a deployment method that will standardize the core solution components while managing environment-specific settings separately, and wants to minimize rework due to configuration errors.

Which process should the Database Specialist recommend to meet these requirements?

- A. Organize common and environmental-specific parameters hierarchically in the AWS Systems Manager Parameter Store, then reference the parameters dynamically from an AWS CloudFormation template
- B. Deploy the CloudFormation stack using the environment name as a parameter.
- C. Create a parameterized AWS CloudFormation template that builds the required object
- D. Keep separate environment parameter files in separate Amazon S3 bucket
- E. Provide an AWS CLI command that deploys the CloudFormation stack directly referencing the appropriate parameter bucket.
- F. Create a parameterized AWS CloudFormation template that builds the required object
- G. Import the template into the CloudFormation interface in the AWS Management Console
- H. Make the required changes to the parameters and deploy the CloudFormation stack.
- I. Create an AWS Lambda function that builds the required objects using an AWS SD
- J. Set the required parameter values in a test event in the Lambda console for each environment that the Application team can modify, as needed
- K. Deploy the infrastructure by triggering the test event in the console.

Answer: A

Explanation:

<https://aws.amazon.com/blogs/mt/integrating-aws-cloudformation-with-aws-systems-manager-parameter-store/>

NEW QUESTION 94

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 96

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 101

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 105

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

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 108

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 112

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 117

An electric utility company wants to store power plant sensor data in an Amazon DynamoDB table. The utility company has over 100 power plants and each power plant has over 200 sensors that send data every 2 seconds. The sensor data includes time with milliseconds precision, a value, and a fault attribute if the sensor is malfunctioning. Power plants are identified by a globally unique identifier. Sensors are identified by a unique identifier within each power plant. A database specialist needs to design the table to support an efficient method of finding all faulty sensors within a given power plant.

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

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

Answer: D

Explanation:

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

NEW QUESTION 118

For the first time, a database professional is establishing a test graph database on Amazon Neptune. The database expert must input millions of rows of test observations from an Amazon S3.csv file. The database professional uploaded the data to the Neptune DB instance through a series of API calls.

Which sequence of actions enables the database professional to upload the data most quickly? (Select three.)

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

Answer: BEF

Explanation:

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

NEW QUESTION 120

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

Which solution will satisfy this criterion?

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

Answer: D

Explanation:

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

NEW QUESTION 125

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 126

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 127

AWS CloudFormation stack including an Amazon RDS database instance was mistakenly removed, resulting in the loss of recent data. A Database Specialist must apply RDS parameters to the CloudFormation template in order to minimize the possibility of future inadvertent instance data loss.

Which settings will satisfy this criterion? (Select 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: ADF

Explanation:

A <https://aws.amazon.com/about-aws/whats-new/2018/09/amazon-rds-now-provides-database-deletion-protection/>

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

F - <https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute-deletionpolicy.html>

NEW QUESTION 129

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 132

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

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

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

Answer: A

Explanation:

"requires minimal downtime when the RDS DB instance goes live" in order to do CDC: "you must first ensure that ARCHIVELOG MODE is on to provide information to LogMiner. AWS DMS uses LogMiner to read information from the archive logs so that AWS DMS can capture changes"
<https://docs.aws.amazon.com/dms/latest/sbs/chap-oracle2postgresql.steps.configureoracle.html> "If you want to capture and apply changes (CDC), then you also need the following privileges."

NEW QUESTION 137

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

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

Answer: C

NEW QUESTION 138

An internet advertising firm stores its data in an Amazon DynamoDB table. Amazon DynamoDB Streams are enabled on the table, and one of the keys has a global secondary index. The table is encrypted using a customer-managed AWS Key Management Service (AWS KMS) key. The firm has chosen to grow worldwide and want to duplicate the database using DynamoDB global tables in a new AWS Region. An administrator observes the following upon review:

No role with the dynamodb: CreateGlobalTable permission exists in the account.
An empty table with the same name exists in the new Region where replication is desired.
A global secondary index with the same partition key but a different sort key exists in the new Region where replication is desired.
Which settings will prevent you from creating a global table or replica in the new Region? (Select two.)

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

Answer: AB

NEW QUESTION 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 143

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 148

A company has a production Amazon Aurora Db cluster that serves both online transaction processing (OLTP) transactions and compute-intensive reports. The reports run for 10% of the total cluster uptime while the OLTP transactions run all the time. The company has benchmarked its workload and determined that a six-node Aurora DB cluster is appropriate for the peak workload.

The company is now looking at cutting costs for this DB cluster, but needs to have a sufficient number of nodes in the cluster to support the workload at different times. The workload has not changed since the previous benchmarking exercise.

How can a Database Specialist address these requirements with minimal user involvement?

- A. Split up the DB cluster into two different clusters: one for OLTP and the other for reporting
- B. Monitor and set up replication between the two clusters to keep data consistent.
- C. Review and evaluate the peak combined workload
- D. Ensure that utilization of the DB cluster node is at an acceptable level
- E. Adjust the number of instances, if necessary.
- F. Use the stop cluster functionality to stop all the nodes of the DB cluster during times of minimal workload
- G. The cluster can be restarted again depending on the workload at the time.
- H. Set up automatic scaling on the DB cluster
- I. This will allow the number of reader nodes to adjust automatically to the reporting workload, when needed.

Answer: D

NEW QUESTION 149

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 151

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

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

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

Answer: D

NEW QUESTION 153

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

Which solution meets these requirements in the MOST efficient way?

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

D. Use Amazon DynamoDB as the database and use Amazon API Gateway

Answer: B

Explanation:

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

NEW QUESTION 156

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

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

NEW QUESTION 158

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 160

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

- A. Create a snapshot of the source DB instance in the source account
- B. Share the snapshot with the destination account
- C. In the target account, create a DB instance from the snapshot.
- D. Use AWS Resource Access Manager to share the source DB instance with the destination account. Create a DB instance in the destination account using the shared resource.
- E. Create a read replica of the DB instance
- F. Give the destination account access to the read replica
- G. In the destination account, create a snapshot of the shared read replica and provision a new RDS for MySQL DB instance.
- H. Use mysqldump to back up the source database
- I. Create an RDS for MySQL DB instance in the destination account
- J. Use the mysql command to restore the backup in the destination database.

Answer: A

Explanation:

Sharing an unencrypted manual DB snapshot enables authorized AWS accounts to directly restore a DB instance from the snapshot instead of taking a copy of it and restoring from that. https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ShareSnapshot.html However Resource Access Manager could not share non-Aurora cluster. <https://docs.aws.amazon.com/ram/latest/userguide/shareable.html>

NEW QUESTION 163

A Database Specialist needs to define a database migration strategy to migrate an on-premises Oracle database to an Amazon Aurora MySQL DB cluster. The company requires near-zero downtime for the data migration. The solution must also be cost-effective. Which approach should the Database Specialist take?

- A. Dump all the tables from the Oracle database into an Amazon S3 bucket using datapump (expdp). Run data transformations in AWS Glue
- B. Load the data from the S3 bucket to the Aurora DB cluster.

- C. Order an AWS Snowball appliance and copy the Oracle backup to the Snowball appliance
- D. Once the Snowball data is delivered to Amazon S3, create a new Aurora DB cluster
- E. Enable the S3 integration to migrate the data directly from Amazon S3 to Amazon RDS.
- F. Use the AWS Schema Conversion Tool (AWS SCT) to help rewrite database objects to MySQL during the schema migration
- G. Use AWS DMS to perform the full load and change data capture (CDC) tasks.
- H. Use AWS Server Migration Service (AWS SMS) to import the Oracle virtual machine image as an Amazon EC2 instance
- I. Use the Oracle Logical Dump utility to migrate the Oracle data from Amazon EC2 to an Aurora DB cluster.

Answer: C

Explanation:

<https://aws.amazon.com/blogs/database/migrating-oracle-databases-with-near-zero-downtime-using-aws-dms/>

NEW QUESTION 167

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

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

Answer: AE

Explanation:

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

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

NEW QUESTION 171

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 175

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 179

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

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

- A. Review the stack drift before modifying the template
- B. Create and review a change set before applying it

- C. Export the database resources as stack outputs
- D. Define the database resources in a nested stack
- E. Set a stack policy for the database resources

Answer: BE

Explanation:

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

NEW QUESTION 180

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

Which AWS solution meets these requirements?

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

Answer: D

Explanation:

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

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/USER_Events.Messages.html We can use stored procedure to invoke Lambda function -

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

NEW QUESTION 182

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

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

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

Answer: AC

NEW QUESTION 187

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 188

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 191

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

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

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

Answer: AE

Explanation:

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

NEW QUESTION 193

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 197

A financial organization must ensure that the most current 90 days of MySQL database backups are accessible. Amazon RDS for MySQL DB instances are used to host all MySQL databases. A database expert must create

a solution that satisfies the criteria for backup retention with the least amount of development work feasible. Which strategy should the database administrator take?

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

Answer: A

Explanation:

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

NEW QUESTION 199

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual AWS-Certified-Database-Specialty Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the AWS-Certified-Database-Specialty Product From:

<https://www.2passeasy.com/dumps/AWS-Certified-Database-Specialty/>

Money Back Guarantee

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