

Amazon

Exam Questions AWS-Certified-Database-Specialty

AWS Certified Database - Specialty



NEW QUESTION 1

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

Which approach meets these requirements?

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

Answer: C

Explanation:

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

NEW QUESTION 2

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 3

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 4

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 5

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.ht](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-global-database/>
<https://aws.amazon.com/about-aws/whats-new/2019/11/aurora-supports-in-place-conversion-to-global-database/>

NEW QUESTION 6

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 7

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 8

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 9

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 objec
- B. Encrypt the S3 object with an AWS KMS ke
- C. Set the KMS key to be rotated every 30 days by setting the EnableKeyRotation property to tru
- D. Use a CloudFormation custom resource to read the S3 object to extract the password.
- E. Create an AWS Lambda function to rotate the secre
- F. Modify the CloudFormation template to add an AWS::SecretsManager::RotationSchedule resourc
- 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 passwor
- 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-rotationssche](https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-secretsmanager-rotationsschedule.html)

NEW QUESTION 10

A database specialist is responsible for an Amazon RDS for MySQL DB instance with one read replica. The DB instance and the read replica are assigned to the default parameter group. The database team currently runs test queries against a read replica. The database team wants to create additional tables in the read replica that will only be accessible from the read replica to benefit the tests.

Which should the database specialist do to allow the database team to create the test tables?

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

Answer: D

Explanation:

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

NEW QUESTION 10

A business need 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 storag
- 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 dat
- D. Keep historical data on Amazon S3 and access it using the Amazon Redshift Spectrum laye
- 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 dat
- G. Keep historical data on Amazon S3 and access it using the Amazon Redshift Spectrum laye
- H. Enable Amazon Redshift Concurrency Scaling.
- I. Leverage an Amazon Redshift data warehouse solution using a dense storage instance to store the most recent dat
- J. Keep historical data on Amazon S3 and access it using the Amazon Redshift Spectrum laye
- 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 12

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

How should the company identify the source of the problem?

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

Answer: B

Explanation:

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

NEW QUESTION 14

A company wants to automate the creation of secure test databases with random credentials to be stored safely for later use. The credentials should have sufficient information about each test database to initiate a connection and perform automated credential rotations. The credentials should not be logged or stored anywhere in an unencrypted form.

Which steps should a Database Specialist take to meet these requirements using an AWS CloudFormation template?

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

Answer: C

NEW QUESTION 19

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 21

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 26

A company wants to migrate its on-premises MySQL databases to Amazon RDS for MySQL. To comply with the company's security policy, all databases must be encrypted at rest. RDS DB instance snapshots must also be shared across various accounts to provision testing and staging environments.

Which solution meets these requirements?

- A. Create an RDS for MySQL DB instance with an AWS Key Management Service (AWS KMS) customer managed CM
- B. Update the key policy to include the Amazon Resource Name (ARN) of the other AWS accounts as a principal, and then allow the kms:CreateGrant action.
- C. Create an RDS for MySQL DB instance with an AWS managed CM
- D. Create a new key policy to include the Amazon Resource Name (ARN) of the other AWS accounts as a principal, and then allow the kms:CreateGrant action.
- E. Create an RDS for MySQL DB instance with an AWS owned CM
- F. Create a new key policy to include the administrator user name of the other AWS accounts as a principal, and then allow the kms:CreateGrant action.
- G. Create an RDS for MySQL DB instance with an AWS CloudHSM key

H. Update the key policy to include the Amazon Resource Name (ARN) of the other AWS accounts as a principal, and then allow the kms:CreateGrant action.

Answer: A

Explanation:

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

NEW QUESTION 27

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 31

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 32

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 37

A company has an on-premises system that tracks various database operations that occur over the lifetime of a database, including database shutdown, deletion, creation, and backup.

The company recently moved two databases to Amazon RDS and is looking at a solution that would satisfy these requirements. The data could be used by other systems within the company.

Which solution will meet these requirements with minimal effort?

- A. Create an Amazon Cloudwatch Events rule with the operations that need to be tracked on Amazon RD
- B. Create an AWS Lambda function to act on these rules and write the output to the tracking systems.
- C. Create an AWS Lambda function to trigger on AWS CloudTrail API call
- D. Filter on specific RDS API calls and write the output to the tracking systems.
- E. Create RDS event subscription
- F. Have the tracking systems subscribe to specific RDS event system notifications.
- G. Write RDS logs to Amazon Kinesis Data Firehose
- H. Create an AWS Lambda function to act on these rules and write the output to the tracking systems.

Answer: C

NEW QUESTION 40

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 42

A large financial services company requires that all data be encrypted in transit. A Developer is attempting to connect to an Amazon RDS DB instance using the company VPC for the first time with credentials provided by a Database Specialist. Other members of the Development team can connect, but this user is consistently receiving an error indicating a communications link failure. The Developer asked the Database Specialist to reset the password a number of times, but the error persists.

Which step should be taken to troubleshoot this issue?

- A. Ensure that the database option group for the RDS DB instance allows ingress from the Developer machine's IP address
- B. Ensure that the RDS DB instance's subnet group includes a public subnet to allow the Developer to connect
- C. Ensure that the RDS DB instance has not reached its maximum connections limit
- D. Ensure that the connection is using SSL and is addressing the port where the RDS DB instance is listening for encrypted connections

Answer: D

Explanation:

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

NEW QUESTION 44

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 46

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

Which process should the database specialist recommend?

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

Answer: C

Explanation:

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

NEW QUESTION 51

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

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

Answer: BE

Explanation:

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

NEW QUESTION 55

A company needs a data warehouse solution that keeps data in a consistent, highly structured format. The company requires fast responses for end-user queries when looking at data from the current year, and users must have access to the full 15-year dataset, when needed. This solution also needs to handle a fluctuating number incoming queries. Storage costs for the 100 TB of data must be kept low.

Which solution meets these requirements?

- 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 and keep historical data on Amazon S3 and access it using the Amazon Redshift Spectrum layer.
- 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 and keep historical data on Amazon S3 and access it using the Amazon Redshift Spectrum layer.
- 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 and keep historical data on Amazon S3 and access it using the Amazon Redshift Spectrum layer.
- 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 60

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 65

A database professional is developing an application that will respond to single-instance requests. The program will query large amounts of client data and offer end users with results.

These reports may include a variety of fields. The database specialist wants to enable users to query the database using any of the fields offered.

During peak periods, the database's traffic volume will be significant yet changeable. However, the database will see little activity over the rest of the day.

Which approach will be the most cost-effective in meeting these requirements?

- A. Amazon DynamoDB with provisioned capacity mode and auto scaling
- B. Amazon DynamoDB with on-demand capacity mode
- C. Amazon Aurora with auto scaling enabled
- D. Amazon Aurora in a serverless mode

Answer: D

Explanation:

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

NEW QUESTION 69

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

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 73

In North America, a business launched a mobile game that swiftly expanded to 10 million daily active players. The game's backend is hosted on AWS and makes considerable use of a TTL-configured Amazon DynamoDB table.

When an item is added or changed, its TTL is set to 600 seconds plus the current epoch time. The game logic is reliant on the purging of outdated data in order to compute rewards points properly. At times, items from the table are read that are many hours beyond their TTL expiration.

How should a database administrator resolve this issue?

- A. Use a client library that supports the TTL functionality for DynamoDB.
- B. Include a query filter expression to ignore items with an expired TTL.
- C. Set the ConsistentRead parameter to true when querying the table.
- D. Create a local secondary index on the TTL attribute.

Answer: B

Explanation:

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

NEW QUESTION 77

A company is running an Amazon RDS for PostgreSQL DB instance and wants to migrate it to an Amazon Aurora PostgreSQL DB cluster. The current database is 1 TB in size. The migration needs to have minimal downtime.

What is the FASTEST way to accomplish this?

- A. Create an Aurora PostgreSQL DB cluster.
- B. Set up replication from the source RDS for PostgreSQL DB instance using AWS DMS to the target DB cluster.
- C. Use the pg_dump and pg_restore utilities to extract and restore the RDS for PostgreSQL DB instance to the Aurora PostgreSQL DB cluster.
- D. Create a database snapshot of the RDS for PostgreSQL DB instance and use this snapshot to create the Aurora PostgreSQL DB cluster.
- E. Migrate data from the RDS for PostgreSQL DB instance to an Aurora PostgreSQL DB cluster using an Aurora Replic.
- F. Promote the replica during the cutover.

Answer: D

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.Migrating.html> Migrating data from an RDS PostgreSQL DB instance to an Aurora PostgreSQL DB cluster by using an Aurora read replica. <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.Migrating.html#Aurora>
<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.Migrating.html#Aurora>

NEW QUESTION 82

A database specialist at a large multi-national financial company is in charge of designing the disaster recovery strategy for a highly available application that is in development. The application uses an Amazon DynamoDB table as its data store. The application requires a recovery time objective (RTO) of 1 minute and a recovery point objective (RPO) of 2 minutes.

Which operationally efficient disaster recovery strategy should the database specialist recommend for the DynamoDB table?

- A. Create a DynamoDB stream that is processed by an AWS Lambda function that copies the data to a DynamoDB table in another Region.
- B. Use a DynamoDB global table replica in another Region.
- C. Enable point-in-time recovery for both tables.
- D. Use a DynamoDB Accelerator table in another Region.
- E. Enable point-in-time recovery for the table.
- F. Create an AWS Backup plan and assign the DynamoDB table as a resource.

Answer: C

NEW QUESTION 83

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 85

An Amazon RDS EBS-optimized instance with Provisioned IOPS (PIOPS) storage is using less than half of its allocated IOPS over the course of several hours under constant load. The RDS instance exhibits multi-second read and write latency, and uses all of its maximum bandwidth for read throughput, yet the instance uses less than half of its CPU and RAM resources.

What should a Database Specialist do in this situation to increase performance and return latency to sub- second levels?

- A. Increase the size of the DB instance storage
- B. Change the underlying EBS storage type to General Purpose SSD (gp2)
- C. Disable EBS optimization on the DB instance
- D. Change the DB instance to an instance class with a higher maximum bandwidth

Answer: D

Explanation:

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

NEW QUESTION 88

A company has migrated a single MySQL database to Amazon Aurora. The production data is hosted in a DB cluster in VPC_PROD, and 12 testing environments are hosted in VPC_TEST using the same AWS account. Testing results in minimal changes to the test data. The Development team wants each environment refreshed nightly so each test database contains fresh production data every day.

Which migration approach will be the fastest and most cost-effective to implement?

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

Answer: A

NEW QUESTION 92

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 97

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 98

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

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

Answer: C

Explanation:

For details about using Aurora and Amazon Comprehend together, see [Using Amazon Comprehend for sentiment detection](#). Aurora machine learning uses a highly optimized integration between the Aurora database and the AWS machine learning (ML) services SageMaker and Amazon Comprehend.

<https://www.stackovercloud.com/2019/11/27/new-for-amazon-aurora-use-machine-learning-directly-from-your>

NEW QUESTION 101

A company is going through a security audit. The audit team has identified cleartext master user password in the AWS CloudFormation templates for Amazon RDS for MySQL DB instances. The audit team has flagged this as a security risk to the database team. What should a database specialist do to mitigate this risk?

- A. Change all the databases to use AWS IAM for authentication and remove all the cleartext passwords in CloudFormation templates.
- B. Use an AWS Secrets Manager resource to generate a random password and reference the secret in the CloudFormation template.
- C. Remove the passwords from the CloudFormation templates so Amazon RDS prompts for the password when the database is being created.
- D. Remove the passwords from the CloudFormation template and store them in a separate file.
- E. Replace the passwords by running CloudFormation using a sed command.

Answer: B

Explanation:

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

NEW QUESTION 102

A Database Specialist is migrating a 2 TB Amazon RDS for Oracle DB instance to an RDS for PostgreSQL DB instance using AWS DMS. The source RDS Oracle DB instance is in a VPC in the us-east-1 Region. The target RDS for PostgreSQL DB instance is in a VPC in the us-west-2 Region. Where should the AWS DMS replication instance be placed for the MOST optimal performance?

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

Answer: C

Explanation:

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

https://docs.aws.amazon.com/dms/latest/sbs/CHAP_SQLServer2Aurora.Steps.CreateReplicationInstance.html

NEW QUESTION 107

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

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

Answer: D

Explanation:

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

NEW QUESTION 108

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 111

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

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

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

Answer: B

Explanation:

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

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

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

NEW QUESTION 115

A large ecommerce company uses Amazon DynamoDB to handle the transactions on its web portal. Traffic patterns throughout the year are usually stable; however, a large event is planned. The company knows that traffic will increase by up to 10 times the normal load over the 3-day event. When sale prices are published during the event, traffic will spike rapidly.

How should a Database Specialist ensure DynamoDB can handle the increased traffic?

- A. Ensure the table is always provisioned to meet peak needs
- B. Allow burst capacity to handle the additional load
- C. Set an AWS Application Auto Scaling policy for the table to handle the increase in traffic
- D. Preprovision additional capacity for the known peaks and then reduce the capacity after the event

Answer: D

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/bp-partition-key-design.html#bp-partition> "DynamoDB provides some flexibility in your per-partition throughput provisioning by providing burst capacity. Whenever you're not fully using a partition's throughput, DynamoDB reserves a portion of that unused capacity for later bursts of throughput to handle usage spikes. DynamoDB currently retains up to 5 minutes (300 seconds) of unused read and write capacity. During an occasional burst of read or write activity, these extra capacity units can be consumed quickly—even faster than the per-second provisioned throughput capacity that you've defined for your table. DynamoDB can also consume burst capacity for background maintenance and other tasks without prior notice. Note that these burst capacity details might change in the future."

NEW QUESTION 119

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 120

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

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

Answer: A

Explanation:

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

NEW QUESTION 122

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 ke
- 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 ke
- D. Use the measurement time as the sort ke
- 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 ke
- G. Use the measurement time as the sort ke
- 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 ke
- 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 123

A company uses Amazon DynamoDB as the data store for its ecommerce website. The website receives little to no traffic at night, and the majority of the traffic occurs during the day. The traffic growth during peak hours is gradual and predictable on a daily basis, but it can be orders of magnitude higher than during off-peak hours. The company initially provisioned capacity based on its average volume during the day without accounting for the variability in traffic patterns. However, the website is experiencing a significant amount of throttling during peak hours. The company wants to reduce the amount of throttling while minimizing costs. What should a database specialist do to meet these requirements?

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

Answer: C

Explanation:

On-demand mode is a good option if any of the following are true: You create new tables with unknown workloads. You have unpredictable application traffic. You prefer the ease of paying for only what you use. <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/HowItWorks.ReadWriteCapacityMode.h> Amazon DynamoDB auto scaling uses the AWS Application Auto Scaling service to dynamically adjust provisioned throughput capacity on your behalf <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/AutoScaling.html>

NEW QUESTION 127

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 129

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 130

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": "Deny" and "Resource": "ProductionDatabase" in the policy.
- C. Create an AWS CloudFormation stack in XML format
- D. Set "DeletionPolicy" 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 134

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 136

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

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

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

Answer: C

Explanation:

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

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

"You can customize the order in which your Aurora Replicas are promoted to the primary instance after a failure by assigning each replica a priority. Priorities range from 0 for the first priority to 15 for the last priority. If the primary instance fails, Amazon RDS promotes the Aurora Replica with the better priority to the new primary instance. You can modify the priority of an Aurora Replica at any time. Modifying the priority doesn't trigger a failover. More than one Aurora Replica can share the same priority, resulting in promotion tiers. If two or more Aurora Replicas share the same priority, then Amazon RDS promotes the replica that is largest in size. If two or more Aurora Replicas share the same priority and size, then Amazon RDS promotes an arbitrary replica in the same promotion tier. "

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

NEW QUESTION 139

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 140

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 144

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 148

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 153

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 156

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

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

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

Answer: ABC

NEW QUESTION 160

To meet new data compliance requirements, a company needs to keep critical data durably stored and readily accessible for 7 years. Data that is more than 1 year old is considered archival data and must automatically be moved out of the Amazon Aurora MySQL DB cluster every week. On average, around 10 GB of new data is added to the database every month. A database specialist must choose the most operationally efficient solution to migrate the archival data to Amazon S3. Which solution meets these requirements?

- A. Create a custom script that exports archival data from the DB cluster to Amazon S3 using a SQL view, then deletes the archival data from the DB cluster
- B. Launch an Amazon EC2 instance with a weekly cron job to execute the custom script.
- C. Configure an AWS Lambda function that exports archival data from the DB cluster to Amazon S3 using a SELECT INTO OUTFILE S3 statement, then deletes the archival data from the DB cluster
- D. Schedule the Lambda function to run weekly using Amazon EventBridge (Amazon CloudWatch Events).
- E. Configure two AWS Lambda functions: one that exports archival data from the DB cluster to Amazon S3 using the mysqldump utility, and another that deletes the archival data from the DB cluster
- F. Schedule both Lambda functions to run weekly using Amazon EventBridge (Amazon CloudWatch Events).
- G. Use AWS Database Migration Service (AWS DMS) to continually export the archival data from the DB cluster to Amazon S3. Configure an AWS Data Pipeline process to run weekly that executes a custom SQL script to delete the archival data from the DB cluster.

Answer: B

Explanation:

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

NEW QUESTION 165

A retail company manages a web application that stores data in an Amazon DynamoDB table. The company is undergoing account consolidation efforts. A database engineer needs to migrate the DynamoDB table from the current AWS account to a new AWS account. Which strategy meets these requirements with the LEAST amount of administrative work?

- A. Use AWS Glue to crawl the data in the DynamoDB tabl
- B. Create a job using an available blueprint to export the data to Amazon S3. Import the data from the S3 file to a DynamoDB table in the new account.
- C. Create an AWS Lambda function to scan the items of the DynamoDB table in the current account and write to a file in Amazon S3. Create another Lambda function to read the S3 file and restore the items of a DynamoDB table in the new account.
- D. Use AWS Data Pipeline in the current account to export the data from the DynamoDB table to a file in Amazon S3. Use Data Pipeline to import the data from the S3 file to a DynamoDB table in the new account.
- E. Configure Amazon DynamoDB Streams for the DynamoDB table in the current account.
- F. Create an AWS Lambda function to read from the stream and write to a file in Amazon S3. Create another Lambda function to read the S3 file and restore the items to a DynamoDB table in the new account.

Answer: C

Explanation:

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

NEW QUESTION 167

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

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

Answer: A

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.Reference.html> "This wait most often arises when there is a very high rate of commit activity on the system. You can sometimes alleviate this wait by modifying applications to commit transactions in batches. "
<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/apg-waits.xactsync.html>

NEW QUESTION 168

Developers have requested a new Amazon Redshift cluster so they can load new third-party marketing data. The new cluster is ready and the user credentials are given to the developers. The developers indicate that their copy jobs fail with the following error message: "Amazon Invalid operation: S3ServiceException:Access Denied,Status 403,Error AccessDenied." The developers need to load this data soon, so a database specialist must act quickly to solve this issue. What is the MOST secure solution?

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

Answer: B

Explanation:

<https://docs.aws.amazon.com/redshift/latest/gsg/rs-gsg-create-an-iam-role.html>
"Now that you have created the new role, your next step is to attach it to your cluster. You can attach the role when you launch a new cluster or you can attach it to an existing cluster. In the next step, you attach the role to a new cluster."
https://docs.aws.amazon.com/redshift/latest/dg/copy-usage_notes-access-permissions.html

NEW QUESTION 169

A company is building a software as a service application. As part of the new user sign-on workflow, a Python script invokes the CreateTable operation using the Amazon DynamoDB API. After the call returns, the script attempts to call PutItem. Occasionally, the PutItem request fails with a ResourceNotFoundException error, which causes the workflow to fail. The development team has confirmed that the same table name is used in the two API calls. How should a database specialist fix this issue?

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

Answer: C

Explanation:

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

NEW QUESTION 172

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

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

Answer: C

Explanation:

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_BestPractices.html#CHAP_BestPractices.LOBS, "AWS DMS migrates LOB data in two phases: 1. AWS DMS creates a new row in the target table and populates the row with all data except the associated LOB value. 2. AWS DMS updates the row in the target table with the LOB data." This means that we would need two tasks, one per phase and use limited LOB mode for best performance.

NEW QUESTION 176

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

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

Answer: C

Explanation:

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

NEW QUESTION 180

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

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

Answer: A

Explanation:

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

NEW QUESTION 181

Amazon DynamoDB global tables are being used by a business to power an online gaming game. The game is played by gamers from all around the globe. As the game became popularity, the amount of queries to DynamoDB substantially rose. Recently, gamers have complained about the game's condition being inconsistent between nations. A database professional notices that the ReplicationLatency metric for many replica tables is set to an abnormally high value. Which strategy will resolve the issue?

- A. Configure all replica tables to use DynamoDB auto scaling.
- B. Configure a DynamoDB Accelerator (DAX) cluster on each of the replicas.
- C. Configure the primary table to use DynamoDB auto scaling and the replica tables to use manually provisioned capacity.
- D. Configure the table-level write throughput limit service quota to a higher value.

Answer: A

Explanation:

https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/V2globaltables_reqs_bestpractices.html

NEW QUESTION 183

A database specialist must create nightly backups of an Amazon DynamoDB table in a mission-critical workload as part of a disaster recovery strategy. Which backup methodology should the database specialist use to MINIMIZE management overhead?

- A. Install the AWS CLI on an Amazon EC2 instanc
- B. Write a CLI command that creates a backup of the DynamoDB tabl
- C. Create a scheduled job or task that executes the command on a nightly basis.
- D. Create an AWS Lambda function that creates a backup of the DynamoDB tabl
- E. Create an Amazon CloudWatch Events rule that executes the Lambda function on a nightly basis.
- F. Create a backup plan using AWS Backup, specify a backup frequency of every 24 hours, and give the plan a nightly backup window.
- G. Configure DynamoDB backup and restore for an on-demand backup frequency of every 24 hours.

Answer: C

Explanation:

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

NEW QUESTION 184

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

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

Which of the following activities should the Database Specialist take to resolve the issue? (Select two.)

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

Answer: BE

Explanation:

<https://docs.aws.amazon.com/neptune/latest/userguide/bulk-load-tutorial-IAM.html> <https://docs.aws.amazon.com/neptune/latest/userguide/bulk-load-data.html>
“An IAM role for the Neptune DB instance to assume that has an IAM policy that allows access to the data files in the S3 bucket. The policy must grant Read and List permissions.” “An Amazon S3 VPC endpoint. For more information, see the Creating an Amazon S3 VPC Endpoint section.”

NEW QUESTION 189

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

Which solution satisfies these criteria?

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

Answer: D

NEW QUESTION 193

A company is running an on-premises application comprised of a web tier, an application tier, and a MySQL database tier. The database is used primarily during business hours with random activity peaks throughout the day. A database specialist needs to improve the availability and reduce the cost of the MySQL database tier as part of the company's migration to AWS.

Which MySQL database option would meet these requirements?

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

Answer: C

NEW QUESTION 194

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

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 196

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 environmen
- B. Reference the environment-specific parameter file in the CloudFormation template.
- C. Encrypt the master password using an AWS KMS ke
- 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 200

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

- A. Create an Amazon CloudWatch dashboard to show the number of connections, CPU usage, and disk space consumptio
- 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 databas
- F. Use the metrics to identify any queries that are related to spikes in the graph during the next slow period.

Answer: D

Explanation:

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

NEW QUESTION 201

A company uses an Amazon RDS for PostgreSQL DB instance for its customer relationship management (CRM) system. New compliance requirements specify that the database must be encrypted at rest. Which action will meet these requirements?

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

Answer: A

Explanation:

<https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/encrypt-an-existing-amazon-rds-for-postgresq> You can enable encryption for an Amazon RDS DB instance when you create it, but not after it's created. However, you can add encryption to an unencrypted DB instance by creating a snapshot of your DB instance, and then creating an encrypted copy of that snapshot. You can then restore a DB instance from the encrypted snapshot to get an encrypted copy of your original DB instance. The pattern uses AWS Database Migration Service (AWS DMS) to migrate data and AWS Key Management Service (AWS KMS) for encryption.

NEW QUESTION 204

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 perio
- B. Assign the DB instances to the backup plan.
- C. Modify the DB instances to enable the automated backup optio
- 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 209

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

Which approach should the Database Specialist take to securely manage the database credentials?

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

Answer: C

NEW QUESTION 211

A company is running a website on Amazon EC2 instances deployed in multiple Availability Zones (AZs). The site performs a high number of repetitive reads and writes each second on an Amazon RDS for MySQL Multi- AZ DB instance with General Purpose SSD (gp2) storage. After comprehensive testing and analysis, a database specialist discovers that there is high read latency and high CPU utilization on the DB instance.

Which approach should the database specialist to take to resolve this issue without changing the application?

- A. Implementing sharding to distribute the load to multiple RDS for MySQL databases.
- B. Use the same RDS for MySQL instance class with Provisioned IOPS (PIOPS) storage.
- C. Add an RDS for MySQL read replica.
- D. Modify the RDS for MySQL database class to a bigger size and implement Provisioned IOPS (PIOPS).

Answer: D

NEW QUESTION 213

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

AWS-Certified-Database-Specialty Practice Exam Features:

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

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