

AWS Certified Database - Specialty

1. A company has deployed an e-commerce web application in a new AWS account. An Amazon RDS for MySQL Multi-AZ DB instance is part of this deployment with a database-1.xxxxxxxxxxxx.us-east-1.rds.amazonaws.com endpoint listening on port 3306. The company's Database Specialist is able to log in to MySQL and run queries from the bastion host using these details.

When users try to utilize the application hosted in the AWS account, they are presented with a generic error message. The application servers are logging a "could not connect to server: Connection times out" error message to Amazon CloudWatch Logs.

What is the cause of this error?

A. The user name and password the application is using are incorrect.

B. The security group assigned to the application servers does not have the necessary rules to allow inbound connections from the DB instance.

C. The security group assigned to the DB instance does not have the necessary rules to allow inbound connections from the application servers.

D. The user name and password are correct, but the user is not authorized to use the DB instance.

Answer(s): C

2. An AWS CloudFormation stack that included an Amazon RDS DB instance was accidentally deleted and recent data was lost. A Database Specialist needs to add RDS settings to the CloudFormation template to reduce the chance of accidental instance data loss in the future.

Which settings will meet this requirement? (Choose three.)

A. Set DeletionProtection to True

B. Set MultiAZ to True

C. Set TerminationProtection to True

D. Set DeleteAutomatedBackups to False

E. Set DeletionPolicy to Delete

F. Set DeletionPolicy to Retain

Answer(s): A D F

3. 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(s): B

4. A company is deploying a solution in Amazon Aurora by migrating from an on-premises system. The IT department has established an AWS Direct Connect link from the company's data center. The company's Database Specialist has selected the option to require SSL/TLS for connectivity to prevent plaintext data from being set over the network. The migration appears to be working successfully, and the data can be queried from a desktop machine.

Two Data Analysts have been asked to query and validate the data in the new Aurora DB cluster.

Both Analysts are unable to connect to Aurora. Their user names and passwords have been verified as valid and the Database Specialist can connect to the DB cluster using their accounts. The Database Specialist also verified that the security group configuration allows network from all corporate IP addresses.

What should the Database Specialist do to correct the Data Analysts' inability to connect?

A. Restart the DB cluster to apply the SSL change.

B. Instruct the Data Analysts to download the root certificate and use the SSL certificate on the connection string to connect.

C. Add explicit mappings between the Data Analysts' IP addresses and the instance in the security group assigned to the DB cluster.

D. Modify the Data Analysts' local client firewall to allow network traffic to AWS.

Answer(s): B

5. 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(s): C

6. 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 RDS. Create an AWS Lambda function to act on these rules and write the output to the tracking systems.

B. Create an AWS Lambda function to trigger on AWS CloudTrail API calls. Filter on specific RDS API calls and write the output to the tracking systems.

C. Create RDS event subscriptions. Have the tracking systems subscribe to specific RDS event system notifications.

D. Write RDS logs to Amazon Kinesis Data Firehose. Create an AWS Lambda function to act on these rules and write the output to the tracking systems.

Answer(s): C

7. 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. Restrict permissions on the bucket to the IAM role associated with the instance profile only. Modify the application to download the text file and retrieve the credentials on start up. Update the text file every 60 days.

B. Configure IAM database authentication for the application to connect to the database. Create an IAM user and map it to a separate database user for each ecommerce user. Require users to update their passwords every 60 days.

C. Store the credentials in AWS Secrets Manager. Restrict permissions on the secret to only the IAM role associated with the instance profile. Modify the application to retrieve the credentials from Secrets Manager on start up. Configure the rotation interval to 60 days.

D. Store the credentials in an encrypted text file in the application AMI. Use AWS KMS to store the key for decrypting the text file. Modify the application to decrypt the text file and retrieve the credentials on start up. Update the text file and publish a new AMI every 60 days.

Answer(s): C

8. 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(s): A C F

9. 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. Set up replication from the source RDS for PostgreSQL DB instance using AWS DMS to the target DB cluster.

B. Use the `pg_dump` and `pg_restore` utilities to extract and restore the RDS for PostgreSQL DB instance to the Aurora PostgreSQL DB cluster.

C. Create a database snapshot of the RDS for PostgreSQL DB instance and use this snapshot to create the Aurora PostgreSQL DB cluster.

D. Migrate data from the RDS for PostgreSQL DB instance to an Aurora PostgreSQL DB cluster using an Aurora Replica. Promote the replica during the cutover.

Answer(s): D

10. 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 use-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(s): C

11. The Development team recently executed a database script containing several data definition language (DDL) and data manipulation language (DML) statements on an Amazon Aurora MySQL DB cluster. The release accidentally deleted thousands of rows from an important table and broke some application functionality. This was discovered 4 hours after the release. Upon investigation, a Database Specialist tracked the issue to a DELETE command in the script with an incorrect WHERE clause filtering the wrong set of rows.

The Aurora DB cluster has Backtrack enabled with an 8-hour backtrack window. The Database Administrator also took a manual snapshot of the DB cluster before the release started. The database needs to be returned to the correct state as quickly as possible to resume full application functionality. Data loss must be minimal. How can the Database Specialist accomplish this?

A. Quickly rewind the DB cluster to a point in time before the release using Backtrack.

B. Perform a point-in-time recovery (PITR) of the DB cluster to a time before the release and copy the deleted rows from the restored database to the original database.

C. Restore the DB cluster using the manual backup snapshot created before the release and change the application configuration settings to point to the new DB cluster.

D. Create a clone of the DB cluster with Backtrack enabled. Rewind the cloned cluster to a point in time before the release. Copy deleted rows from the clone to the original database.

Answer(s): B

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

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

A. Review the stack drift before modifying the template

B. Create and review a change set before applying it

C. Export the database resources as stack outputs

D. Define the database resources in a nested stack

E. Set a stack policy for the database resources

Answer(s): B E

13. 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(s): A C E

14. A company is hosting critical business data in an Amazon Redshift cluster. Due to the sensitive nature of the data, the cluster is encrypted at rest using AWS KMS. As a part of disaster recovery requirements, the company needs to copy the Amazon Redshift snapshots to another Region.

Which steps should be taken in the AWS Management Console to meet the disaster recovery requirements?

A. Create a new KMS customer master key in the source Region. Switch to the destination Region, enable Amazon Redshift cross-Region snapshots, and use the KMS key of the source Region.

B. Create a new IAM role with access to the KMS key. Enable Amazon Redshift cross-Region replication using the new IAM role, and use the KMS key of the source Region.

C. Enable Amazon Redshift cross-Region snapshots in the source Region, and create a snapshot copy grant and use a KMS key in the destination Region.

D. Create a new KMS customer master key in the destination Region and create a new IAM role with access to the new KMS key. Enable Amazon Redshift cross-Region replication in the source Region and use the KMS key of the destination Region.

Answer(s): C

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

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

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

A. Split up the DB cluster into two different clusters: one for OLTP and the other for reporting. Monitor and set up replication between the two clusters to keep data consistent.

B. Review all evaluate the peak combined workload. Ensure that utilization of the DB cluster node is at an acceptable level. Adjust the number of instances, if necessary.

C. Use the stop cluster functionality to stop all the nodes of the DB cluster during times of minimal workload. The cluster can be restarted again depending on the workload at the time.

D. Set up automatic scaling on the DB cluster. This will allow the number of reader nodes to adjust automatically to the reporting workload, when needed.

Answer(s): D

16. A company is running a finance application on an Amazon RDS for MySQL DB instance. The application is governed by multiple financial regulatory agencies. The RDS DB instance is set up with security groups to allow access to certain Amazon EC2 servers only. AWS KMS is used for

encryption at rest.

Which step will provide additional security?

A. Set up NACLs that allow the entire EC2 subnet to access the DB instance

B. Disable the master user account

C. Set up a security group that blocks SSH to the DB instance

D. Set up RDS to use SSL for data in transit

Answer(s): D

17. 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. Provision enough instances to support high demand.

B. Leverage an Amazon Redshift data warehouse solution using a dense storage instance to store the most recent data. Keep historical data on Amazon S3 and access it using the Amazon Redshift Spectrum layer. 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. Keep historical data on Amazon S3 and access it using the Amazon Redshift Spectrum layer. Enable Amazon Redshift Concurrency Scaling.

D. Leverage an Amazon Redshift data warehouse solution using a dense storage instance to store the most recent data. Keep historical data on Amazon S3 and access it using the Amazon Redshift Spectrum layer. Leverage Amazon Redshift elastic resize.

Answer(s): C

18. A gaming company wants to deploy a game in multiple Regions. The company plans to save local high scores in Amazon DynamoDB tables in each Region. A Database Specialist needs to design a solution to automate the deployment of the database with identical configurations in additional Regions, as needed. The solution should also automate configuration changes across all Regions.

Which solution would meet these requirements and deploy the DynamoDB tables?

A. Create an AWS CLI command to deploy the DynamoDB table to all the Regions and save it for future deployments.

B. Create an AWS CloudFormation template and deploy the template to all the Regions.

C. Create an AWS CloudFormation template and use a stack set to deploy the template to all the Regions.

D. Create DynamoDB tables using the AWS Management Console in all the Regions and create a step-by-step guide for future deployments.

Answer(s): C

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

How can the Database Specialists accomplish this?

A. Enable the option to push all database logs to Amazon CloudWatch for advanced analysis

B. Create appropriate Amazon CloudWatch dashboards to contain specific periods of time

C. Enable Amazon RDS Performance Insights and review the appropriate dashboard

D. Enable Enhanced Monitoring with the appropriate settings

Answer(s): C

20. 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(s): B
