

Exam question from Amazon's AWS Certified Database - Specialty

Question #: 1

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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.xxxxxxxxxx.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.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 2

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 3

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 4

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 5

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 6

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 7

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 8

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 9

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 10

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 11

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 12

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 13

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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.

Show Suggested Answer

Exam question from Amazon's AWS Certified Database - Specialty

Question #: 14

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 15

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 16

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 17

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 18

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 19

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 20

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 21

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 22

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A Database Specialist is setting up a new Amazon Aurora DB cluster with one primary instance and three Aurora Replicas for a highly intensive, business-critical application. The Aurora DB cluster has one medium-sized primary instance, one large-sized replica, and two medium sized replicas. The Database Specialist did not assign a promotion tier to the replicas.

In the event of a primary failure, what will occur?

- A. Aurora will promote an Aurora Replica that is of the same size as the primary instance
- B. Aurora will promote an arbitrary Aurora Replica
- C. Aurora will promote the largest-sized Aurora Replica
- D. Aurora will not promote an Aurora Replica

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 23

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company is running its line of business application on AWS, which uses Amazon RDS for MySQL at the persistent data store. The company wants to minimize downtime when it migrates the database to Amazon Aurora.

Which migration method should a Database Specialist use?

- A. Take a snapshot of the RDS for MySQL DB instance and create a new Aurora DB cluster with the option to migrate snapshots.
- B. Make a backup of the RDS for MySQL DB instance using the mysqldump utility, create a new Aurora DB cluster, and restore the backup.
- C. Create an Aurora Replica from the RDS for MySQL DB instance and promote the Aurora DB cluster.
- D. Create a clone of the RDS for MySQL DB instance and promote the Aurora DB cluster.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 24

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 25

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company is using Amazon RDS for MySQL to redesign its business application. A Database Specialist has noticed that the Development team is restoring their MySQL database multiple times a day when Developers make mistakes in their schema updates. The Developers sometimes need to wait hours for the restores to complete.

Multiple team members are working on the project, making it difficult to find the correct restore point for each mistake.

Which approach should the Database Specialist take to reduce downtime?

- A. Deploy multiple read replicas and have the team members make changes to separate replica instances
- B. Migrate to Amazon RDS for SQL Server, take a snapshot, and restore from the snapshot
- C. Migrate to Amazon Aurora MySQL and enable the Aurora Backtrack feature
- D. Enable the Amazon RDS for MySQL Backtrack feature

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 26

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A media company is using Amazon RDS for PostgreSQL to store user data. The RDS DB instance currently has a publicly accessible setting enabled and is hosted in a public subnet. Following a recent AWS Well-Architected Framework review, a Database Specialist was given new security requirements.

- ⇒ Only certain on-premises corporate network IPs should connect to the DB instance.
- ⇒ Connectivity is allowed from the corporate network only.

Which combination of steps does the Database Specialist need to take to meet these new requirements? (Choose three.)

- A. Modify the `pg_hba.conf` file. Add the required corporate network IPs and remove the unwanted IPs.
- B. Modify the associated security group. Add the required corporate network IPs and remove the unwanted IPs.
- C. Move the DB instance to a private subnet using AWS DMS.
- D. Enable VPC peering between the application host running on the corporate network and the VPC associated with the DB instance.
- E. Disable the publicly accessible setting.
- F. Connect to the DB instance using private IPs and a VPN.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 27

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 28

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company with branch offices in Portland, New York, and Singapore has a three-tier web application that leverages a shared database. The database runs on Amazon RDS for MySQL and is hosted in the us-west-2 Region. The application has a distributed front end deployed in the us-west-2, ap-southeast-1, and us-east-2 Regions.

This front end is used as a dashboard for Sales Managers in each branch office to see current sales statistics. There are complaints that the dashboard performs more slowly in the Singapore location than it does in Portland or New York. A solution is needed to provide consistent performance for all users in each location.

Which set of actions will meet these requirements?

- A. Take a snapshot of the instance in the us-west-2 Region. Create a new instance from the snapshot in the ap-southeast-1 Region. Reconfigure the ap-southeast-1 front-end dashboard to access this instance.
- B. Create an RDS read replica in the ap-southeast-1 Region from the primary RDS DB instance in the us-west-2 Region. Reconfigure the ap-southeast-1 front-end dashboard to access this instance.
- C. Create a new RDS instance in the ap-southeast-1 Region. Use AWS DMS and change data capture (CDC) to update the new instance in the ap-southeast-1 Region. Reconfigure the ap-southeast-1 front-end dashboard to access this instance.
- D. Create an RDS read replica in the us-west-2 Region where the primary instance resides. Create a read replica in the ap-southeast-1 Region from the read replica located on the us-west-2 Region. Reconfigure the ap-southeast-1 front-end dashboard to access this instance.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 29

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

Which approach will MOST effectively meet these requirements?

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

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 30

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company is planning to close for several days. A Database Specialist needs to stop all applications along with the DB instances to ensure employees do not have access to the systems during this time. All databases are running on Amazon RDS for MySQL.

The Database Specialist wrote and ran a script to stop all the DB instances. When reviewing the logs, the Database Specialist found that Amazon RDS DB instances with read replicas did not stop.

How should the Database Specialist edit the script to fix this issue?

- A. Stop the source instances before stopping their read replicas
- B. Delete each read replica before stopping its corresponding source instance
- C. Stop the read replicas before stopping their source instances
- D. Use the AWS CLI to stop each read replica and source instance at the same time

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 31

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

Which database solution meets these requirements?

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 32

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

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

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 33

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

Which solution meets these requirements in the MOST efficient way?

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

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 34

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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. 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.
- B. 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. Configure an Amazon RDS event subscription to send a notification to an Amazon SNS topic to which the Lambda function is subscribed.
- C. Configure one Aurora Replica to have the same instance class as the primary DB instance. Implement Aurora PostgreSQL DB cluster cache management. Set the failover priority to tier-0 for the primary DB instance and one replica with the same instance class. Set the failover priority to tier-1 for the other replicas.
- D. Configure both Aurora Replicas to have the same instance class as the primary DB instance. Implement Aurora PostgreSQL DB cluster cache management. Set the failover priority to tier-0 for the primary DB instance and to tier-1 for the replicas.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 35

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 36

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company is using Amazon with Aurora Replicas for read-only workload scaling. A Database Specialist needs to split up two read-only applications so each application always connects to a dedicated replica. The Database Specialist wants to implement load balancing and high availability for the read-only applications. Which solution meets these requirements?

- A. Use a specific instance endpoint for each replica and add the instance endpoint to each read-only application connection string.
- B. Use reader endpoints for both the read-only workload applications.
- C. Use a reader endpoint for one read-only application and use an instance endpoint for the other read-only application.
- D. Use custom endpoints for the two read-only applications.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 37

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

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

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

Which choice of keys is recommended for the DynamoDB table?

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 38

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A Database Specialist migrated an existing production MySQL database from on-premises to an Amazon RDS for MySQL DB instance. However, after the migration, the database needed to be encrypted at rest using AWS KMS. Due to the size of the database, reloading, the data into an encrypted database would be too time-consuming, so it is not an option.

How should the Database Specialist satisfy this new requirement?

- A. Create a snapshot of the unencrypted RDS DB instance. Create an encrypted copy of the unencrypted snapshot. Restore the encrypted snapshot copy.
- B. Modify the RDS DB instance. Enable the AWS KMS encryption option that leverages the AWS CLI.
- C. Restore an unencrypted snapshot into a MySQL RDS DB instance that is encrypted.
- D. Create an encrypted read replica of the RDS DB instance. Promote it the master.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 39

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

What is the most likely reason for this?

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

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 40

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

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

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

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 41

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company has an Amazon RDS Multi-AZ DB instances that is 200 GB in size with an RPO of 6 hours. To meet the company's disaster recovery policies, the database backup needs to be copied into another Region. The company requires the solution to be cost-effective and operationally efficient.

What should a Database Specialist do to copy the database backup into a different Region?

- A. Use Amazon RDS automated snapshots and use AWS Lambda to copy the snapshot into another Region
- B. Use Amazon RDS automated snapshots every 6 hours and use Amazon S3 cross-Region replication to copy the snapshot into another Region
- C. Create an AWS Lambda function to take an Amazon RDS snapshot every 6 hours and use a second Lambda function to copy the snapshot into another Region
- D. Create a cross-Region read replica for Amazon RDS in another Region and take an automated snapshot of the read replica

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 42

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 43

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 44

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 45

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

An ecommerce company has tasked a Database Specialist with creating a reporting dashboard that visualizes critical business metrics that will be pulled from the core production database running on Amazon Aurora. Data that is read by the dashboard should be available within 100 milliseconds of an update.

The Database Specialist needs to review the current configuration of the Aurora DB cluster and develop a cost-effective solution. The solution needs to accommodate the unpredictable read workload from the reporting dashboard without any impact on the write availability and performance of the DB cluster.

Which solution meets these requirements?

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 46

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 47

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

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

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 48

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company is looking to migrate a 1 TB Oracle database from on-premises to an Amazon Aurora PostgreSQL DB cluster. The company's Database Specialist discovered that the Oracle database is storing 100 GB of large binary objects (LOBs) across multiple tables. The Oracle database has a maximum LOB size of 500 MB with an average LOB size of 350 MB. The Database Specialist has chosen AWS DMS to migrate the data with the largest replication instances.

How should the Database Specialist optimize the database migration using AWS DMS?

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 49

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A Database Specialist is designing a disaster recovery strategy for a production Amazon DynamoDB table. The table uses provisioned read/write capacity mode, global secondary indexes, and time to live (TTL). The Database Specialist has restored the latest backup to a new table.

To prepare the new table with identical settings, which steps should be performed? (Choose two.)

- A. Re-create global secondary indexes in the new table
- B. Define IAM policies for access to the new table
- C. Define the TTL settings
- D. Encrypt the table from the AWS Management Console or use the update-table command
- E. Set the provisioned read and write capacity

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 50

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

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

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 51

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company runs online transaction processing (OLTP) workloads on an Amazon RDS for PostgreSQL Multi-AZ DB instance. Tests were run on the database after work hours, which generated additional database logs. The free storage of the RDS DB instance is low due to these additional logs.

What should the company do to address this space constraint issue?

- A. Log in to the host and run the `rm $PGDATA/pg_logs/*` command
- B. Modify the `rds.log_retention_period` parameter to 1440 and wait up to 24 hours for database logs to be deleted
- C. Create a ticket with AWS Support to have the logs deleted
- D. Run the `SELECT rds_rotate_error_log()` stored procedure to rotate the logs

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 52

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 53

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A gaming company is designing a mobile gaming app that will be accessed by many users across the globe. The company wants to have replication and full support for multi-master writes. The company also wants to ensure low latency and consistent performance for app users.

Which solution meets these requirements?

- A. Use Amazon DynamoDB global tables for storage and enable DynamoDB automatic scaling
- B. Use Amazon Aurora for storage and enable cross-Region Aurora Replicas
- C. Use Amazon Aurora for storage and cache the user content with Amazon ElastiCache
- D. Use Amazon Neptune for storage

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 54

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 55

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

Which approach should the Database Specialist take?

- A. Dump all the tables from the Oracle database into an Amazon S3 bucket using datapump (expdp). Run data transformations in AWS Glue. Load the data from the S3 bucket to the Aurora DB cluster.
- B. Order an AWS Snowball appliance and copy the Oracle backup to the Snowball appliance. Once the Snowball data is delivered to Amazon S3, create a new Aurora DB cluster. Enable the S3 integration to migrate the data directly from Amazon S3 to Amazon RDS.
- C. Use the AWS Schema Conversion Tool (AWS SCT) to help rewrite database objects to MySQL during the schema migration. Use AWS DMS to perform the full load and change data capture (CDC) tasks.
- D. Use AWS Server Migration Service (AWS SMS) to import the Oracle virtual machine image as an Amazon EC2 instance. Use the Oracle Logical Dump utility to migrate the Oracle data from Amazon EC2 to an Aurora DB cluster.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 56

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A marketing company is using Amazon DocumentDB and requires that database audit logs be enabled. A Database Specialist needs to configure monitoring so that all data definition language (DDL) statements performed are visible to the Administrator. The Database Specialist has set the audit_logs parameter to enabled in the cluster parameter group.

What should the Database Specialist do to automatically collect the database logs for the Administrator?

- A. Enable DocumentDB to export the logs to Amazon CloudWatch Logs
- B. Enable DocumentDB to export the logs to AWS CloudTrail
- C. Enable DocumentDB Events to export the logs to Amazon CloudWatch Logs
- D. Configure an AWS Lambda function to download the logs using the download-db-log-file-portion operation and store the logs in Amazon S3

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 57

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

What is the quickest way for the company to gather data on the migration compatibility?

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 58

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

An ecommerce company is using Amazon DynamoDB as the backend for its order-processing application. The steady increase in the number of orders is resulting in increased DynamoDB costs. Order verification and reporting perform many repeated GetItem functions that pull similar datasets, and this read activity is contributing to the increased costs. The company wants to control these costs without significant development efforts.

How should a Database Specialist address these requirements?

- A. Use AWS DMS to migrate data from DynamoDB to Amazon DocumentDB
- B. Use Amazon DynamoDB Streams and Amazon Kinesis Data Firehose to push the data into Amazon Redshift
- C. Use an Amazon ElastiCache for Redis in front of DynamoDB to boost read performance
- D. Use DynamoDB Accelerator to offload the reads

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 59

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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 templates. Deploy a stack with the DB cluster for each development group. Delete the stack at the end of the development cycle.
- B. Use the Aurora DB cloning feature. Deploy a single development and test Aurora DB instance, and create clone instances for the development groups. Delete the clones at the end of the development cycle.
- C. Use Aurora Replicas. From the master automatic pause compute capacity option, create replicas for each development group, and promote each replica to master. Delete the replicas at the end of the development cycle.
- D. Use Aurora Serverless. Restore current Aurora snapshot and deploy to a serverless cluster for each development group. Enable the option to pause the compute capacity on the cluster and set an appropriate timeout.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 60

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company has multiple applications serving data from a secure on-premises database. The company is migrating all applications and databases to the AWS Cloud. The IT Risk and Compliance department requires that auditing be enabled on all secure databases to capture all log ins, log outs, failed logins, permission changes, and database schema changes. A Database Specialist has recommended Amazon Aurora MySQL as the migration target, and leveraging the Advanced Auditing feature in Aurora.

Which events need to be specified in the Advanced Auditing configuration to satisfy the minimum auditing requirements? (Choose three.)

- A. CONNECT
- B. QUERY_DCL
- C. QUERY_DDL
- D. QUERY_DML
- E. TABLE
- F. QUERY

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 61

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

Which solution will meet these requirements at the lowest cost?

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 62

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

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

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 63

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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 loading. The Database Specialist needs to modify the workload to load the data slowly.
- B. The scaling of Aurora storage cannot catch up with the data loading. The Database Specialist needs to enable Aurora storage scaling.
- C. The local storage used to store temporary tables is full. The Database Specialist needs to scale up the instance.
- D. The local storage used to store temporary tables is full. The Database Specialist needs to enable local storage scaling.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 64

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A financial company wants to store sensitive user data in an Amazon Aurora PostgreSQL DB cluster. The database will be accessed by multiple applications across the company. The company has mandated that all communications to the database be encrypted and the server identity must be validated. Any non-SSL- based connections should be disallowed access to the database.

Which solution addresses these requirements?

- A. Set the `rds.force_ssl=0` parameter in DB parameter groups. Download and use the Amazon RDS certificate bundle and configure the PostgreSQL connection string with `sslmode=allow`.
- B. Set the `rds.force_ssl=1` parameter in DB parameter groups. Download and use the Amazon RDS certificate bundle and configure the PostgreSQL connection string with `sslmode=disable`.
- C. Set the `rds.force_ssl=0` parameter in DB parameter groups. Download and use the Amazon RDS certificate bundle and configure the PostgreSQL connection string with `sslmode=verify-ca`.
- D. Set the `rds.force_ssl=1` parameter in DB parameter groups. Download and use the Amazon RDS certificate bundle and configure the PostgreSQL connection string with `sslmode=verify-full`.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 65

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company is using 5 TB Amazon RDS DB instances and needs to maintain 5 years of monthly database backups for compliance purposes. A Database Administrator must provide Auditors with data within 24 hours.

Which solution will meet these requirements and is the MOST operationally efficient?

- A. Create an AWS Lambda function to run on the first day of every month to take a manual RDS snapshot. Move the snapshot to the company's Amazon S3 bucket.
- B. Create an AWS Lambda function to run on the first day of every month to take a manual RDS snapshot.
- C. Create an RDS snapshot schedule from the AWS Management Console to take a snapshot every 30 days.
- D. Create an AWS Lambda function to run on the first day of every month to create an automated RDS snapshot.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 66

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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 values. Then, create the secret with the user name and password set to the same default values. Add a Secret Target Attachment resource with the SecretId and TargetId properties set to the Amazon Resource Names (ARNs) of the secret and the database. Finally, update the secret's password value with a randomly generated string set by the GenerateSecretString property.
- B. Add a Mapping property from the database Amazon Resource Name (ARN) to the secret ARN. Then, create the secret with a chosen user name and a randomly generated password set by the GenerateSecretString property. Add the database with the MasterUserName and MasterUserPassword properties set to the user name of the secret.
- C. Add a resource of type `AWS::SecretsManager::Secret` and specify the `GenerateSecretString` property. Then, define the database user name in the `SecureStringTemplate` template. Create a resource for the database and reference the secret string for the `MasterUserName` and `MasterUserPassword` properties. 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.
- D. Create the secret with a chosen user name and a randomly generated password set by the `GenerateSecretString` property. 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 ARN. Then, create a database with the `MasterUserName` and `MasterUserPassword` properties set to the previously created values in the secret.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 67

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 68

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A Database Specialist is working with a company to launch a new website built on Amazon Aurora with several Aurora Replicas. This new website will replace an on-premises website connected to a legacy relational database. Due to stability issues in the legacy database, the company would like to test the resiliency of Aurora.

Which action can the Database Specialist take to test the resiliency of the Aurora DB cluster?

- A. Stop the DB cluster and analyze how the website responds
- B. Use Aurora fault injection to crash the master DB instance
- C. Remove the DB cluster endpoint to simulate a master DB instance failure
- D. Use Aurora Backtrack to crash the DB cluster

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 69

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

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

- A. Create an Amazon CloudWatch dashboard to show the number of connections, CPU usage, and disk space consumption. Watch these dashboards during the next slow period.
- B. 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.
- C. 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.
- D. Enable Amazon RDS Performance Insights on the PostgreSQL database. Use the metrics to identify any queries that are related to spikes in the graph during the next slow period.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 70

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company has a web-based survey application that uses Amazon DynamoDB. During peak usage, when survey responses are being collected, a Database Specialist sees the ProvisionedThroughputExceededException error.

What can the Database Specialist do to resolve this error? (Choose two.)

- A. Change the table to use Amazon DynamoDB Streams
- B. Purchase DynamoDB reserved capacity in the affected Region
- C. Increase the write capacity units for the specific table
- D. Change the table capacity mode to on-demand
- E. Change the table type to throughput optimized

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 71

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company is running a two-tier ecommerce application in one AWS account. The web server is deployed using an Amazon RDS for MySQL Multi-AZ DB instance. A Developer mistakenly deleted the database in the production environment. The database has been restored, but this resulted in hours of downtime and lost revenue. Which combination of changes in existing IAM policies should a Database Specialist make to prevent an error like this from happening in the future? (Choose three.)

- A. Grant least privilege to groups, users, and roles
- B. Allow all users to restore a database from a backup that will reduce the overall downtime to restore the database
- C. Enable multi-factor authentication for sensitive operations to access sensitive resources and API operations
- D. Use policy conditions to restrict access to selective IP addresses
- E. Use AccessList Controls policy type to restrict users for database instance deletion
- F. Enable AWS CloudTrail logging and Enhanced Monitoring

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 72

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 73

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

What is most secure solution to store the master password?

- A. Store the master password in a parameter file in each environment. Reference the environment-specific parameter file in the CloudFormation template.
- B. Encrypt the master password using an AWS KMS key. Store the encrypted master password in the CloudFormation template.
- C. Use the secretsmanager dynamic reference to retrieve the master password stored in AWS Secrets Manager and enable automatic rotation.
- D. Use the ssm dynamic reference to retrieve the master password stored in the AWS Systems Manager Parameter Store and enable automatic rotation.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 74

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

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

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

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 75

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 76

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 77

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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)

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 78

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 79

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 80

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

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

- A. Migrate the workloads from Amazon RDS for MySQL to Amazon DynamoDB
- B. Enable cross-Region read replicas and direct read traffic to them when Amazon RDS is down
- C. Enable a read replica and direct read traffic to it when Amazon RDS is down
- D. Enable an Amazon RDS for MySQL Multi-AZ configuration

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 81

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A Database Specialist must create a read replica to isolate read-only queries for an Amazon RDS for MySQL DB instance. Immediately after creating the read replica, users that query it report slow response times.

What could be causing these slow response times?

- A. New volumes created from snapshots load lazily in the background
- B. Long-running statements on the master
- C. Insufficient resources on the master
- D. Overload of a single replication thread by excessive writes on the master

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 82

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company developed an AWS CloudFormation template used to create all new Amazon DynamoDB tables in its AWS account. The template configures provisioned throughput capacity using hard-coded values. The company wants to change the template so that the tables it creates in the future have independently configurable read and write capacity units assigned.

Which solution will enable this change?

- A. Add values for the rcuCount and wcuCount parameters to the Mappings section of the template. Configure DynamoDB to provision throughput capacity using the stack's mappings.
- B. Add values for two Number parameters, rcuCount and wcuCount, to the template. Replace the hard-coded values with calls to the Ref intrinsic function, referencing the new parameters.
- C. Add values for the rcuCount and wcuCount parameters as outputs of the template. Configure DynamoDB to provision throughput capacity using the stack outputs.
- D. Add values for the rcuCount and wcuCount parameters to the Mappings section of the template. Replace the hard-coded values with calls to the Ref intrinsic function, referencing the new parameters.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 83

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A retail company with its main office in New York and another office in Tokyo plans to build a database solution on AWS. The company's main workload consists of a mission-critical application that updates its application data in a data store. The team at the Tokyo office is building dashboards with complex analytical queries using the application data. The dashboards will be used to make buying decisions, so they need to have access to the application data in less than 1 second.

Which solution meets these requirements?

- A. Use an Amazon RDS DB instance deployed in the us-east-1 Region with a read replica instance in the ap-northeast-1 Region. Create an Amazon ElastiCache cluster in the ap-northeast-1 Region to cache application data from the replica to generate the dashboards.
- B. Use an Amazon DynamoDB global table in the us-east-1 Region with replication into the ap-northeast-1 Region. Use Amazon QuickSight for displaying dashboard results.
- C. Use an Amazon RDS for MySQL DB instance deployed in the us-east-1 Region with a read replica instance in the ap-northeast-1 Region. Have the dashboard application read from the read replica.
- D. Use an Amazon Aurora global database. Deploy the writer instance in the us-east-1 Region and the replica in the ap-northeast-1 Region. Have the dashboard application read from the replica ap-northeast-1 Region.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 84

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company is using Amazon RDS for PostgreSQL. The Security team wants all database connection requests to be logged and retained for 180 days. The RDS for PostgreSQL DB instance is currently using the default parameter group. A Database Specialist has identified that setting the `log_connections` parameter to 1 will enable connections logging.

Which combination of steps should the Database Specialist take to meet the logging and retention requirements? (Choose two.)

- A. Update the `log_connections` parameter in the default parameter group
- B. Create a custom parameter group, update the `log_connections` parameter, and associate the parameter with the DB instance
- C. Enable publishing of database engine logs to Amazon CloudWatch Logs and set the event expiration to 180 days
- D. Enable publishing of database engine logs to an Amazon S3 bucket and set the lifecycle policy to 180 days
- E. Connect to the RDS PostgreSQL host and update the `log_connections` parameter in the `postgresql.conf` file

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 85

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A Database Specialist is creating a new Amazon Neptune DB cluster, and is attempting to load data from Amazon S3 into the Neptune DB cluster using the Neptune bulk loader API. The Database Specialist receives the following error:

```
`Unable to connect to s3 endpoint. Provided source = s3://mybucket/graphdata/ and region = us-east-1. Please verify your S3 configuration.`
```

Which combination of actions should the Database Specialist take to troubleshoot the problem? (Choose 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

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 86

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A database specialist manages a critical Amazon RDS for MySQL DB instance for a company. The data stored daily could vary from .01% to 10% of the current database size. The database specialist needs to ensure that the DB instance storage grows as needed.

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

- A. Configure RDS Storage Auto Scaling.
- B. Configure RDS instance Auto Scaling.
- C. Modify the DB instance allocated storage to meet the forecasted requirements.
- D. Monitor the Amazon CloudWatch FreeStorageSpace metric daily and add storage as required.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 87

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company is due for renewing its database license. The company wants to migrate its 80 TB transactional database system from on-premises to the AWS Cloud. The migration should incur the least possible downtime on the downstream database applications. The company's network infrastructure has limited network bandwidth that is shared with other applications.

Which solution should a database specialist use for a timely migration?

- A. Perform a full backup of the source database to AWS Snowball Edge appliances and ship them to be loaded to Amazon S3. Use AWS DMS to migrate change data capture (CDC) data from the source database to Amazon S3. Use a second AWS DMS task to migrate all the S3 data to the target database.
- B. Perform a full backup of the source database to AWS Snowball Edge appliances and ship them to be loaded to Amazon S3. Periodically perform incremental backups of the source database to be shipped in another Snowball Edge appliance to handle syncing change data capture (CDC) data from the source to the target database.
- C. Use AWS DMS to migrate the full load of the source database over a VPN tunnel using the internet for its primary connection. Allow AWS DMS to handle syncing change data capture (CDC) data from the source to the target database.
- D. Use the AWS Schema Conversion Tool (AWS SCT) to migrate the full load of the source database over a VPN tunnel using the internet for its primary connection. Allow AWS SCT to handle syncing change data capture (CDC) data from the source to the target database.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 88

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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 replica. Reboot the read replica. Connect to the read replica and create the tables.
- B. Change the `read_only` parameter to false (`read_only=0`) in the default parameter group of the read replica. Perform a reboot without failover. Connect to the read replica and create the tables using the `local_only` MySQL option.
- C. Change the `read_only` parameter to false (`read_only=0`) in the default parameter group. Reboot the read replica. Connect to the read replica and create the tables.
- D. Create a new DB parameter group. Change the `read_only` parameter to false (`read_only=0`). Associate the read replica with the new group. Reboot the read replica. Connect to the read replica and create the tables.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 89

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 90

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 91

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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. Enable point-in-time recovery for both tables.
- C. Use a DynamoDB Accelerator table in another Region. Enable point-in-time recovery for the table.
- D. Create an AWS Backup plan and assign the DynamoDB table as a resource.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 92

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

- A. Deploy a new RDS for MySQL DB instance and configure it for access from the on-premises data center. Use the mysqldump utility to create an initial snapshot from the on-premises MySQL server, and copy it to an Amazon S3 bucket. Import the snapshot into the DB instance utilizing the MySQL utilities running on an Amazon EC2 instance. Immediately point the application to the DB instance.
- B. Deploy a new Amazon EC2 instance, install the MySQL software on the EC2 instance, and configure networking for access from the on-premises data center. Use the mysqldump utility to create a snapshot of the on-premises MySQL server. Copy the snapshot into the EC2 instance and restore it into the EC2 MySQL instance. Use AWS DMS to migrate data into a new RDS for MySQL DB instance. Point the application to the DB instance.
- C. Deploy a new Amazon EC2 instance, install the MySQL software on the EC2 instance, and configure networking for access from the on-premises data center. Use the mysqldump utility to create a snapshot of the on-premises MySQL server. Copy the snapshot into an Amazon S3 bucket and import the snapshot into a new RDS for MySQL DB instance using the MySQL utilities running on an EC2 instance. Point the application to the DB instance.
- D. Deploy a new RDS for MySQL DB instance and configure it for access from the on-premises data center. Use the mysqldump utility to create an initial snapshot from the on-premises MySQL server, and copy it to an Amazon S3 bucket. Import the snapshot into the DB instance using the MySQL utilities running on an Amazon EC2 instance. Establish replication into the new DB instance using MySQL replication. Stop application access to the on-premises MySQL server and let the remaining transactions replicate over. Point the application to the DB instance.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 93

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

Which solution meets these requirements?

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

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 94

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A database specialist is building a system that uses a static vendor dataset of postal codes and related territory information that is less than 1 GB in size. The dataset is loaded into the application's cache at start up. The company needs to store this data in a way that provides the lowest cost with a low application startup time. Which approach will meet these requirements?

- A. Use an Amazon RDS DB instance. Shut down the instance once the data has been read.
- B. Use Amazon Aurora Serverless. Allow the service to spin resources up and down, as needed.
- C. Use Amazon DynamoDB in on-demand capacity mode.
- D. Use Amazon S3 and load the data from flat files.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 95

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

How can this solution be implemented?

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 96

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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. Replace the passwords by running CloudFormation using a sed command.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 97

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

What should the database specialist do to connect to the new, restored Amazon DocumentDB cluster?

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

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 98

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

Which AWS solution meets these requirements?

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

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 99

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company needs to migrate Oracle Database Standard Edition running on an Amazon EC2 instance to an Amazon RDS for Oracle DB instance with Multi-AZ. The database supports an ecommerce website that runs continuously. The company can only provide a maintenance window of up to 5 minutes. Which solution will meet these requirements?

- A. Configure Oracle Real Application Clusters (RAC) on the EC2 instance and the RDS DB instance. Update the connection string to point to the RAC cluster. Once the EC2 instance and RDS DB instance are in sync, fail over from Amazon EC2 to Amazon RDS.
- B. Export the Oracle database from the EC2 instance using Oracle Data Pump and perform an import into Amazon RDS. Stop the application for the entire process. When the import is complete, change the database connection string and then restart the application.
- C. Configure AWS DMS with the EC2 instance as the source and the RDS DB instance as the destination. Stop the application when the replication is in sync, change the database connection string, and then restart the application.
- D. Configure AWS DataSync with the EC2 instance as the source and the RDS DB instance as the destination. Stop the application when the replication is in sync, change the database connection string, and then restart the application.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 100

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 101

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 102

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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 Server. Use a CloudWatch dashboard to identify the root cause.
- B. Enable RDS Performance Insights and determine which query is creating the problem. Request changes to the query to address the problem.
- C. Use AWS X-Ray deployed with Amazon RDS to track query system traces.
- D. Create a support request and work with AWS Support to identify the source of the issue.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 103

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

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

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 104

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 105

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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. Copy the encrypted snapshot to Amazon S3. Restore the DB instance from the encrypted snapshot using Amazon S3.
- B. Create a new RDS for MySQL DB instance with encryption enabled. Restore the unencrypted snapshot to this DB instance.
- C. Create a snapshot of the unencrypted DB instance. Create an encrypted copy of the snapshot. Restore the DB instance from the encrypted snapshot.
- D. Temporarily shut down the unencrypted DB instance. Enable AWS KMS encryption in the AWS Management Console using an AWS managed CMK. Restart the DB instance in an encrypted state.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 106

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company is migrating its on-premises database workloads to the AWS Cloud. A database specialist performing the move has chosen AWS DMS to migrate an Oracle database with a large table to Amazon RDS. The database specialist notices that AWS DMS is taking significant time to migrate the data.

Which actions would improve the data migration speed? (Choose three.)

- A. Create multiple AWS DMS tasks to migrate the large table.
- B. Configure the AWS DMS replication instance with Multi-AZ.
- C. Increase the capacity of the AWS DMS replication server.
- D. Establish an AWS Direct Connect connection between the on-premises data center and AWS.
- E. Enable an Amazon RDS Multi-AZ configuration.
- F. Enable full large binary object (LOB) mode to migrate all LOB data for all large tables.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 107

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

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

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

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 108

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company has a 20 TB production Amazon Aurora DB cluster. The company runs a large batch job overnight to load data into the Aurora DB cluster. To ensure the company's development team has the most up-to-date data for testing, a copy of the DB cluster must be available in the shortest possible time after the batch job completes.

How should this be accomplished?

- A. Use the AWS CLI to schedule a manual snapshot of the DB cluster. Restore the snapshot to a new DB cluster using the AWS CLI.
- B. Create a dump file from the DB cluster. Load the dump file into a new DB cluster.
- C. Schedule a job to create a clone of the DB cluster at the end of the overnight batch process.
- D. Set up a new daily AWS DMS task that will use cloning and change data capture (CDC) on the DB cluster to copy the data to a new DB cluster. Set up a time for the AWS DMS stream to stop when the new cluster is current.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 109

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company has two separate AWS accounts: one for the business unit and another for corporate analytics. The company wants to replicate the business unit data stored in Amazon RDS for MySQL in us-east-1 to its corporate analytics Amazon Redshift environment in us-west-1. The company wants to use AWS DMS with Amazon RDS as the source endpoint and Amazon Redshift as the target endpoint.

Which action will allow AVS DMS to perform the replication?

- A. Configure the AWS DMS replication instance in the same account and Region as Amazon Redshift.
- B. Configure the AWS DMS replication instance in the same account as Amazon Redshift and in the same Region as Amazon RDS.
- C. Configure the AWS DMS replication instance in its own account and in the same Region as Amazon Redshift.
- D. Configure the AWS DMS replication instance in the same account and Region as Amazon RDS.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 110

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A database specialist is managing an application in the us-west-1 Region and wants to set up disaster recovery in the us-east-1 Region. The Amazon Aurora MySQL DB cluster needs an RPO of 1 minute and an RTO of 2 minutes.

Which approach meets these requirements with no negative performance impact?

- A. Enable synchronous replication.
- B. Enable asynchronous binlog replication.
- C. Create an Aurora Global Database.
- D. Copy Aurora incremental snapshots to the us-east-1 Region.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 111

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

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

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 112

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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).

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 113

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

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

What should a database specialist do to meet these requirements?

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 114

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 115

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 116

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company uses Amazon Aurora for secure financial transactions. The data must always be encrypted at rest and in transit to meet compliance requirements. Which combination of actions should a database specialist take to meet these requirements? (Choose two.)

- A. Create an Aurora Replica with encryption enabled using AWS Key Management Service (AWS KMS). Then promote the replica to master.
- B. Use SSL/TLS to secure the in-transit connection between the financial application and the Aurora DB cluster.
- C. Modify the existing Aurora DB cluster and enable encryption using an AWS Key Management Service (AWS KMS) encryption key. Apply the changes immediately.
- D. Take a snapshot of the Aurora DB cluster and encrypt the snapshot using an AWS Key Management Service (AWS KMS) encryption key. Restore the snapshot to a new DB cluster and update the financial application database endpoints.
- E. Use AWS Key Management Service (AWS KMS) to secure the in-transit connection between the financial application and the Aurora DB cluster.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 117

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

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

- A. Implement 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).

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 118

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A banking company recently launched an Amazon RDS for MySQL DB instance as part of a proof-of-concept project. A database specialist has configured automated database snapshots. As a part of routine testing, the database specialist noticed one day that the automated database snapshot was not created.

Which of the following are possible reasons why the snapshot was not created? (Choose two.)

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 119

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

What should the database specialist do to achieve this? (Choose two.)

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 120

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A large company has a variety of Amazon DB clusters. Each of these clusters has various configurations that adhere to various requirements. Depending on the team and use case, these configurations can be organized into broader categories.

A database administrator wants to make the process of storing and modifying these parameters more systematic. The database administrator also wants to ensure that changes to individual categories of configurations are automatically applied to all instances when required.

Which AWS service or feature will help automate and achieve this objective?

- A. AWS Systems Manager Parameter Store
- B. DB parameter group
- C. AWS Config
- D. AWS Secrets Manager

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 121

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 122

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

Which solution meets these requirements?

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

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 123

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company has applications running on Amazon EC2 instances in a private subnet with no internet connectivity. The company deployed a new application that uses Amazon DynamoDB, but the application cannot connect to the DynamoDB tables. A developer already checked that all permissions are set correctly.

What should a database specialist do to resolve this issue while minimizing access to external resources?

- A. Add a route to an internet gateway in the subnet's route table.
- B. Add a route to a NAT gateway in the subnet's route table.
- C. Assign a new security group to the EC2 instances with an outbound rule to ports 80 and 443.
- D. Create a VPC endpoint for DynamoDB and add a route to the endpoint in the subnet's route table.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 124

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

How should the database engineer meet this requirement?

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 125

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 126

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A large gaming company is creating a centralized solution to store player session state for multiple online games. The workload required key-value storage with low latency and will be an equal mix of reads and writes. Data should be written into the AWS Region closest to the user across the games' geographically distributed user base. The architecture should minimize the amount of overhead required to manage the replication of data between Regions.

Which solution meets these requirements?

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 127

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 128

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

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

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

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 129

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company's ecommerce website uses Amazon DynamoDB for purchase orders. Each order is made up of a Customer ID and an Order ID. The DynamoDB table uses the Customer ID as the partition key and the Order ID as the sort key.

To meet a new requirement, the company also wants the ability to query the table by using a third attribute named Invoice ID. Queries using the Invoice ID must be strongly consistent. A database specialist must provide this capability with optimal performance and minimal overhead.

What should the database administrator do to meet these requirements?

- A. Add a global secondary index on Invoice ID to the existing table.
- B. Add a local secondary index on Invoice ID to the existing table.
- C. Recreate the table by using the latest snapshot while adding a local secondary index on Invoice ID.
- D. Use the partition key and a FilterExpression parameter with a filter on Invoice ID for all queries.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 130

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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 CMK. 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.
- B. Create an RDS for MySQL DB instance with an AWS managed CMK. 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.
- C. Create an RDS for MySQL DB instance with an AWS owned CMK. 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.
- D. Create an RDS for MySQL DB instance with an AWS CloudHSM key. 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.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 131

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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 table. 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.
- B. 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.
- C. 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.
- D. Configure Amazon DynamoDB Streams for the DynamoDB table in the current account. 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.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 132

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

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

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

A.

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

B.

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

C.

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

D.

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

Exam question from Amazon's AWS Certified Database - Specialty

Question #: 133

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company has an application that uses an Amazon DynamoDB table to store user data. Every morning, a single-threaded process calls the DynamoDB API Scan operation to scan the entire table and generate a critical start-of-day report for management. A successful marketing campaign recently doubled the number of items in the table, and now the process takes too long to run and the report is not generated in time.

A database specialist needs to improve the performance of the process. The database specialist notes that, when the process is running, 15% of the table's provisioned read capacity units (RCUs) are being used.

What should the database specialist do?

- A. Enable auto scaling for the DynamoDB table.
- B. Use four threads and parallel DynamoDB API Scan operations.
- C. Double the table's provisioned RCUs.
- D. Set the Limit and Offset parameters before every call to the API.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 134

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 135

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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. Launch an Amazon EC2 instance with a weekly cron job to execute the custom script.
- B. 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. Schedule the Lambda function to run weekly using Amazon EventBridge (Amazon CloudWatch Events).
- C. 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. Schedule both Lambda functions to run weekly using Amazon EventBridge (Amazon CloudWatch Events).
- D. 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.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 136

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

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

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 137

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

How should the company perform this data load?

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 138

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 139

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

What should the team do to meet this requirement?

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 140

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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 instance. Write a CLI command that creates a backup of the DynamoDB table. Create a scheduled job or task that runs the command on a nightly basis.
- B. Create an AWS Lambda function that creates a backup of the DynamoDB table. Create an Amazon CloudWatch Events rule that runs the Lambda function on a nightly basis.
- C. Create a backup plan using AWS Backup, specify a backup frequency of every 24 hours, and give the plan a nightly backup window.
- D. Configure DynamoDB backup and restore for an on-demand backup frequency of every 24 hours.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 141

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company is using a Single-AZ Amazon RDS for MySQL DB instance for development. The DB instance is experiencing slow performance when queries run. 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 run 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).

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 142

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company has an AWS CloudFormation template written in JSON that is used to launch new Amazon RDS for MySQL DB instances. The security team has asked a database specialist to ensure that the master password is automatically rotated every 30 days for all new DB instances that are launched using the template. What is the MOST operationally efficient solution to meet these requirements?

- A. Save the password in an Amazon S3 object. Encrypt the S3 object with an AWS KMS key. Set the KMS key to be rotated every 30 days by setting the `EnableKeyRotation` property to true. Use a CloudFormation custom resource to read the S3 object to extract the password.
- B. Create an AWS Lambda function to rotate the secret. Modify the CloudFormation template to add an `AWS::SecretsManager::RotationSchedule` resource. Configure the `RotationLambdaARN` value and, for the `RotationRules` property, set the `AutomaticallyAfterDays` parameter to 30.
- C. Modify the CloudFormation template to use the AWS KMS key as the database password. Configure an Amazon EventBridge rule to invoke the KMS API to rotate the key every 30 days by setting the `ScheduleExpression` parameter to `*** / 30 ***`.
- D. Integrate the Amazon RDS for MySQL DB instances with AWS IAM and centrally manage the master database user password.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 143

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

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

Which database engine meets these requirements?

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 144

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

An online retail company is planning a multi-day flash sale that must support processing of up to 5,000 orders per second. The number of orders and exact schedule for the sale will vary each day. During the sale, approximately 10,000 concurrent users will look at the deals before buying items. Outside of the sale, the traffic volume is very low. The acceptable performance for read/write queries should be under 25 ms. Order items are about 2 KB in size and have a unique identifier. The company requires the most cost-effective solution that will automatically scale and is highly available.

Which solution meets these requirements?

- A. Amazon DynamoDB with on-demand capacity mode
- B. Amazon Aurora with one writer node and an Aurora Replica with the parallel query feature enabled
- C. Amazon DynamoDB with provisioned capacity mode with 5,000 write capacity units (WCUs) and 10,000 read capacity units (RCUs)
- D. Amazon Aurora with one writer node and two cross-Region Aurora Replicas

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 145

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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 bookings. Associate an AWS Lambda function to capture changes and push the booking data to the RDS for MySQL DB instance used by the CRM.
- B. Use Amazon DynamoDB to accept the bookings. Enable DynamoDB Streams and associate an AWS Lambda function to capture changes and push the booking data to an Amazon SQS queue. 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.
- C. Use Amazon ElastiCache for Redis to accept the bookings. Associate an AWS Lambda function to capture changes and push the booking data to an Amazon Redshift database used by the CRM.
- D. Use Amazon DynamoDB to accept the bookings. 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.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 146

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

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

What is the MOST likely reason for this occurrence?

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 147

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A financial company recently launched a portfolio management solution. The backend of the application is powered by Amazon Aurora with MySQL compatibility. The company requires an RTO of 5 minutes and an RPO of 5 minutes. A database specialist must configure an efficient disaster recovery solution with minimal replication lag.

Which approach should the database specialist take to meet 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.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 148

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company hosts an internal file-sharing application running on Amazon EC2 instances in VPC_A. This application is backed by an Amazon ElastiCache cluster, which is in VPC_B and peered with VPC_A. The company migrates its application instances from VPC_A to VPC_B. Logs indicate that the file-sharing application no longer can connect to the ElastiCache cluster.

What should a database specialist do to resolve this issue?

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

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 149

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A database specialist must load 25 GB of data files from a company's on-premises storage to an Amazon Neptune database.

Which approach to load the data is FASTEST?

- A. Upload the data to Amazon S3 and use the Loader command to load the data from Amazon S3 into the Neptune database.
- B. Write a utility to read the data from the on-premises storage and run INSERT statements in a loop to load the data into the Neptune database.
- C. Use the AWS CLI to load the data directly from the on-premises storage into the Neptune database.
- D. Use AWS DataSync to load the data directly from the on-premises storage into the Neptune database.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 150

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A finance company needs to make sure that its MySQL database backups are available for the most recent 90 days. All of the MySQL databases are hosted on Amazon RDS for MySQL DB instances. A database specialist must implement a solution that meets the backup retention requirement with the least possible development effort.

Which approach should the database specialist take?

- A. Use AWS Backup to build a backup plan for the required retention period. Assign the DB instances to the backup plan.
- B. Modify the DB instances to enable the automated backup option. Select the required backup retention period.
- C. 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.
- D. Use AWS Lambda to schedule a daily manual snapshot of the DB instances. Delete snapshots that exceed the retention requirement.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 151

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

An online advertising company uses an Amazon DynamoDb table as its data store. The table has Amazon DynamoDB Streams enabled and has a global secondary index on one of the keys. The table is encrypted using an AWS Key Management Service (AWS KMS) customer managed key.

The company has decided to expand its operations globally and wants to replicate the database in a different AWS Region by using DynamoDB global tables.

Upon review, an administrator notices the following:

- ⇒ 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 configurations will block the creation of a global table or the creation of a replica in the new Region? (Choose 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.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 152

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A large automobile company is migrating the database of a critical financial application to Amazon DynamoDB. The company's risk and compliance policy requires that every change in the database be recorded as a log entry for audits. The system is anticipating more than 500,000 log entries each minute. Log entries should be stored in batches of at least 100,000 records in each file in Apache Parquet format.

How should a database specialist implement these requirements with DynamoDB?

- A. Enable Amazon DynamoDB Streams on the table. Create an AWS Lambda function triggered by the stream. Write the log entries to an Amazon S3 object.
- B. Create a backup plan in AWS Backup to back up the DynamoDB table once a day. Create an AWS Lambda function that restores the backup in another table and compares both tables for changes. Generate the log entries and write them to an Amazon S3 object.
- C. Enable AWS CloudTrail logs on the table. Create an AWS Lambda function that reads the log files once an hour and filters DynamoDB API actions. Write the filtered log files to Amazon S3.
- D. Enable Amazon DynamoDB Streams on the table. Create an AWS Lambda function triggered by the stream. Write the log entries to an Amazon Kinesis Data Firehose delivery stream with buffering and Amazon S3 as the destination.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 153

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company released a mobile game that quickly grew to 10 million daily active users in North America. The game's backend is hosted on AWS and makes extensive use of an Amazon DynamoDB table that is configured with a TTL attribute.

When an item is added or updated, its TTL is set to the current epoch time plus 600 seconds. The game logic relies on old data being purged so that it can calculate rewards points accurately. Occasionally, items are read from the table that are several hours past their TTL expiry.

How should a database specialist fix 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.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 154

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A development team at an international gaming company is experimenting with Amazon DynamoDB to store in-game events for three mobile games. The most popular game hosts a maximum of 500,000 concurrent users, and the least popular game hosts a maximum of 10,000 concurrent users. The average size of an event is 20 KB, and the average user session produces one event each second. Each event is tagged with a time in milliseconds and a globally unique identifier.

The lead developer created a single DynamoDB table for the events with the following schema:

- ⇒ Partition key: game name
- ⇒ Sort key: event identifier
- ⇒ Local secondary index: player identifier
- ⇒ Event time

The tests were successful in a small-scale development environment. However, when deployed to production, new events stopped being added to the table and the logs show DynamoDB failures with the `ItemCollectionSizeLimitExceededException` error code.

Which design change should a database specialist recommend to the development team?

- A. Use the player identifier as the partition key. Use the event time as the sort key. Add a global secondary index with the game name as the partition key and the event time as the sort key.
- B. Create two tables. Use the game name as the partition key in both tables. Use the event time as the sort key for the first table. Use the player identifier as the sort key for the second table.
- C. Replace the sort key with a compound value consisting of the player identifier collated with the event time, separated by a dash. Add a local secondary index with the player identifier as the sort key.
- D. Create one table for each game. Use the player identifier as the partition key. Use the event time as the sort key.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 155

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

An ecommerce company recently migrated one of its SQL Server databases to an Amazon RDS for SQL Server Enterprise Edition DB instance. The company expects a spike in read traffic due to an upcoming sale. A database specialist must create a read replica of the DB instance to serve the anticipated read traffic.

Which actions should the database specialist take before creating the read replica? (Choose two.)

- A. Identify a potential downtime window and stop the application calls to the source DB instance.
- B. Ensure that automatic backups are enabled for the source DB instance.
- C. Ensure that the source DB instance is a Multi-AZ deployment with Always ON Availability Groups.
- D. Ensure that the source DB instance is a Multi-AZ deployment with SQL Server Database Mirroring (DBM).
- E. Modify the read replica parameter group setting and set the value to 1.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 156

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company is running a two-tier ecommerce application in one AWS account. The application is backed by an Amazon RDS for MySQL Multi-AZ DB instance. A developer mistakenly deleted the DB instance in the production environment. The company restores the database, but this event results in hours of downtime and lost revenue.

Which combination of changes would minimize the risk of this mistake occurring in the future? (Choose three.)

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 157

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A financial services company uses Amazon RDS for Oracle with Transparent Data Encryption (TDE). The company is required to encrypt its data at rest at all times. The key required to decrypt the data has to be highly available, and access to the key must be limited. As a regulatory requirement, the company must have the ability to rotate the encryption key on demand. The company must be able to make the key unusable if any potential security breaches are spotted. The company also needs to accomplish these tasks with minimum overhead.

What should the database administrator use to set up the encryption to meet these requirements?

- A. AWS CloudHSM
- B. AWS Key Management Service (AWS KMS) with an AWS managed key
- C. AWS Key Management Service (AWS KMS) with server-side encryption
- D. AWS Key Management Service (AWS KMS) CMK with customer-provided material

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 158

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company is setting up a new Amazon RDS for SQL Server DB instance. The company wants to enable SQL Server auditing on the database. Which combination of steps should a database specialist take to meet this requirement? (Choose two.)

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

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 159

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A database specialist is creating an AWS CloudFormation stack. The database specialist wants to prevent accidental deletion of an Amazon RDS ProductionDatabase resource in the stack.

Which solution will meet this requirement?

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

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 160

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

An ecommerce company migrates an on-premises MongoDB database to Amazon DocumentDB (with MongoDB compatibility). After the migration, a database specialist realizes that encryption at rest has not been turned on for the Amazon DocumentDB cluster.

What should the database specialist do to enable encryption at rest for the Amazon DocumentDB cluster?

- A. Take a snapshot of the Amazon DocumentDB cluster. Restore the unencrypted snapshot as a new cluster while specifying the encryption option, and provide an AWS Key Management Service (AWS KMS) key.
- B. Enable encryption for the Amazon DocumentDB cluster on the AWS Management Console. Reboot the cluster.
- C. Modify the Amazon DocumentDB cluster by using the `modify-db-cluster` command with the `--storage-encrypted` parameter set to `true`.
- D. Add a new encrypted instance to the Amazon DocumentDB cluster, and then delete an unencrypted instance from the cluster. Repeat until all instances are encrypted.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 161

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company that analyzes the stock market has two offices: one in the us-east-1 Region and another in the eu-west-2 Region. The company wants to implement an AWS database solution that can provide fast and accurate updates.

The office in eu-west-2 has dashboards with complex analytical queries to display the data. The company will use these dashboards to make buying decisions, so the dashboards must have access to the application data in less than 1 second.

Which solution meets these requirements and provides the MOST up-to-date dashboard?

- A. Deploy an Amazon RDS DB instance in us-east-1 with a read replica instance in eu-west-2. Create an Amazon ElastiCache cluster in eu-west-2 to cache data from the read replica to generate the dashboards.
- B. Use an Amazon DynamoDB global table in us-east-1 with replication into eu-west-2. Use multi-active replication to ensure that updates are quickly propagated to eu-west-2.
- C. Use an Amazon Aurora global database. Deploy the primary DB cluster in us-east-1. Deploy the secondary DB cluster in eu-west-2. Configure the dashboard application to read from the secondary cluster.
- D. Deploy an Amazon RDS for MySQL DB instance in us-east-1 with a read replica instance in eu-west-2. Configure the dashboard application to read from the read replica.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 162

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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. Use SQL functions to extract sentiment analysis.
- D. Set up Aurora native integration with Amazon SageMaker. Use SQL functions to extract sentiment analysis.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 163

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A bank plans to use an Amazon RDS for MySQL DB instance. The database should support read-intensive traffic with very few repeated queries.

Which solution meets these requirements?

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

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 164

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A database specialist has a fleet of Amazon RDS DB instances that use the default DB parameter group. The database specialist needs to associate a custom parameter group with some of the DB instances.

After the database specialist makes this change, when will the instances be assigned to this new parameter group?

- A. Instantaneously after the change is made to the parameter group
- B. In the next scheduled maintenance window of the DB instances
- C. After the DB instances are manually rebooted
- D. Within 24 hours after the change is made to the parameter group

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 165

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company is planning on migrating a 500-GB database from Oracle to Amazon Aurora PostgreSQL using the AWS Schema Conversion Tool (AWS SCT) and AWS DMS. The database does not have any stored procedures to migrate but has some tables that are large or partitioned. The application is critical for business so a migration with minimal downtime is preferred.

Which combination of steps should a database specialist take to accelerate the migration process? (Choose three.)

- A. Use the AWS SCT data extraction agent to migrate the schema from Oracle to Aurora PostgreSQL.
- B. For the large tables, change the setting for the maximum number of tables to load in parallel and perform a full load using AWS DMS.
- C. For the large tables, create a table settings rule with a parallel load option in AWS DMS, then perform a full load using DMS.
- D. Use AWS DMS to set up change data capture (CDC) for continuous replication until the cutover date.
- E. Use AWS SCT to convert the schema from Oracle to Aurora PostgreSQL.
- F. Use AWS DMS to convert the schema from Oracle to Aurora PostgreSQL and for continuous replication.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 166

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company is migrating an IBM Informix database to a Multi-AZ deployment of Amazon RDS for SQL Server with Always On Availability Groups (AGs). SQL Server Agent jobs on the Always On AG listener run at 5-minute intervals to synchronize data between the Informix database and the SQL Server database. Users experience hours of stale data after a successful failover to the secondary node with minimal latency.

What should a database specialist do to ensure that users see recent data after a failover?

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 167

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

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

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 168

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company is running its production databases in a 3 TB Amazon Aurora MySQL DB cluster. The DB cluster is deployed to the us-east-1 Region. For disaster recovery (DR) purposes, the company's database specialist needs to make the DB cluster rapidly available in another AWS Region to cover the production load with an RTO of less than 2 hours.

What is the MOST operationally efficient solution to meet 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. Restore the snapshot to an appropriately sized DB cluster in the DR Region.
- B. Add a cross-Region read replica in the DR Region with the same instance type as the current primary instance. If the read replica in the DR Region needs to be used for production, promote the read replica to become a standalone DB cluster.
- C. Create a smaller DB cluster in the DR Region. 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.
- D. Create an Aurora global database that spans two Regions. Use AWS Database Migration Service (AWS DMS) to migrate the existing database to the new global database.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 169

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company has an on-premises SQL Server database. The users access the database using Active Directory authentication. The company successfully migrated its database to Amazon RDS for SQL Server. However, the company is concerned about user authentication in the AWS Cloud environment.

Which solution should a database specialist provide for the user to authenticate?

- A. Deploy Active Directory Federation Services (AD FS) on premises and configure it with an on-premises Active Directory. Set up delegation between the on-premises AD FS and AWS Security Token Service (AWS STS) to map user identities to a role using the `AmazonRDSDirectoryServiceAccess` managed IAM policy.
- B. Establish a forest trust between the on-premises Active Directory and AWS Directory Service for Microsoft Active Directory. Use AWS SSO to configure an Active Directory user delegated to access the databases in RDS for SQL Server.
- C. Use Active Directory Connector to redirect directory requests to the company's on-premises Active Directory without caching any information in the cloud. Use the RDS master user credentials to connect to the DB instance and configure SQL Server logins and users from the Active Directory users and groups.
- D. Establish a forest trust between the on-premises Active Directory and AWS Directory Service for Microsoft Active Directory. Ensure RDS for SQL Server is using mixed mode authentication. Use the RDS master user credentials to connect to the DB instance and configure SQL Server logins and users from the Active Directory users and groups.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 170

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company uses an Amazon Redshift cluster to run its analytical workloads. Corporate policy requires that the company's data be encrypted at rest with customer managed keys. The company's disaster recovery plan requires that backups of the cluster be copied into another AWS Region on a regular basis.

How should a database specialist automate the process of backing up the cluster data in compliance with these policies?

- A. Copy the AWS Key Management Service (AWS KMS) customer managed key from the source Region to the destination Region. Set up an AWS Glue job in the source Region to copy the latest snapshot of the Amazon Redshift cluster from the source Region to the destination Region. Use a time-based schedule in AWS Glue to run the job on a daily basis.
- B. Create a new AWS Key Management Service (AWS KMS) customer managed key in the destination Region. Create a snapshot copy grant in the destination Region specifying the new key. In the source Region, configure cross-Region snapshots for the Amazon Redshift cluster specifying the destination Region, the snapshot copy grant, and retention periods for the snapshot.
- C. Copy the AWS Key Management Service (AWS KMS) customer-managed key from the source Region to the destination Region. Create Amazon S3 buckets in each Region using the keys from their respective Regions. Use Amazon EventBridge (Amazon CloudWatch Events) to schedule an AWS Lambda function in the source Region to copy the latest snapshot to the S3 bucket in that Region. Configure S3 Cross-Region Replication to copy the snapshots to the destination Region, specifying the source and destination KMS key IDs in the replication configuration.
- D. Use the same customer-supplied key materials to create a CMK with the same private key in the destination Region. Configure cross-Region snapshots in the source Region targeting the destination Region. Specify the corresponding CMK in the destination Region to encrypt the snapshot.

Show Suggested Answer

Exam question from Amazon's AWS Certified Database - Specialty

Question #: 171

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A database specialist is launching a test graph database using Amazon Neptune for the first time. The database specialist needs to insert millions of rows of test observations from a .csv file that is stored in Amazon S3. The database specialist has been using a series of API calls to upload the data to the Neptune DB instance. Which combination of steps would allow the database specialist to upload the data faster? (Choose 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.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 172

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company is using Amazon DynamoDB global tables for an online gaming application. The game has players around the world. As the game has become more popular, the volume of requests to DynamoDB has increased significantly. Recently, players have reported that the game state is inconsistent between players in different countries. A database specialist observes that the ReplicationLatency metric for some of the replica tables is too high.

Which approach will alleviate the problem?

- 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.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 173

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company runs a MySQL database for its ecommerce application on a single Amazon RDS DB instance. Application purchases are automatically saved to the database, which causes intensive writes. Company employees frequently generate purchase reports. The company needs to improve database performance and reduce downtime due to patching for upgrades.

Which approach will meet these requirements with the LEAST amount of operational overhead?

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

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 174

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

An ecommerce company is migrating its core application database to Amazon Aurora MySQL. The company is currently performing online transaction processing (OLTP) stress testing with concurrent database sessions. During the first round of tests, a database specialist noticed slow performance for some specific write operations.

Reviewing Amazon CloudWatch metrics for the Aurora DB cluster showed 90% CPU utilization.

Which steps should the database specialist take to MOST effectively identify the root cause of high CPU utilization and slow performance? (Choose two.)

- A. Enable Enhanced Monitoring at less than 30 seconds of granularity to review the operating system metrics before the next round of tests.
- B. Review the VolumeBytesUsed metric in CloudWatch to see if there is a spike in write I/O.
- C. Review Amazon RDS Performance Insights to identify the top SQL statements and wait events.
- D. Review Amazon RDS API calls in AWS CloudTrail to identify long-running queries.
- E. Enable Advance Auditing to log QUERY events in Amazon CloudWatch before the next round of tests.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 175

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

An online advertising company is implementing an application that displays advertisements to its users. The application uses an Amazon DynamoDB table as a data store.

The application also uses a DynamoDB Accelerator (DAX) cluster to cache its reads. Most of the reads are from the GetItem query and the BatchGetItem query. Consistency of reads is not a requirement for this application.

Upon deployment, the application cache is not performing as expected. Specific strongly consistent queries that run against the DAX cluster are taking many milliseconds to respond instead of microseconds.

How can the company improve the cache behavior to increase application performance?

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 176

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company is running its critical production workload on a 500 GB Amazon Aurora MySQL DB cluster. A database engineer must move the workload to a new Amazon Aurora Serverless MySQL DB cluster without data loss.

Which solution will accomplish the move with the LEAST downtime and the LEAST application impact?

- A. Modify the existing DB cluster and update the Aurora configuration to `Serverless`.
- B. Create a snapshot of the existing DB cluster and restore it to a new Aurora Serverless DB cluster.
- C. Create an Aurora Serverless replica from the existing DB cluster and promote it to primary when the replica lag is minimal.
- D. Replicate the data between the existing DB cluster and a new Aurora Serverless DB cluster by using AWS Database Migration Service (AWS DMS) with change data capture (CDC) enabled.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 177

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company is building a web application on AWS. The application requires the database to support read and write operations in multiple AWS Regions simultaneously. The database also needs to propagate data changes between Regions as the changes occur. The application must be highly available and must provide latency of single-digit milliseconds.

Which solution meets these requirements?

- 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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 178

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company is using Amazon Neptune as the graph database for one of its products. The company's data science team accidentally created large amounts of temporary information during an ETL process. The Neptune DB cluster automatically increased the storage space to accommodate the new data, but the data science team deleted the unused information.

What should a database specialist do to avoid unnecessary charges for the unused cluster volume space?

- A. Take a snapshot of the cluster volume. Restore the snapshot in another cluster with a smaller volume size.
- B. Use the AWS CLI to turn on automatic resizing of the cluster volume.
- C. Export the cluster data into a new Neptune DB cluster.
- D. Add a Neptune read replica to the cluster. Promote this replica as a new primary DB instance. Reset the storage space of the cluster.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 179

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A database specialist is responsible for designing a highly available solution for online transaction processing (OLTP) using Amazon RDS for MySQL production databases. Disaster recovery requirements include a cross-Region deployment along with an RPO of 5 minutes and RTO of 30 minutes.

What should the database specialist do to align to the high availability and disaster recovery requirements?

- A. Use a Multi-AZ deployment in each Region.
- B. Use read replica deployments in all Availability Zones of the secondary Region.
- C. Use Multi-AZ and read replica deployments within a Region.
- D. Use Multi-AZ and deploy a read replica in a secondary Region.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 180

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A media company wants to use zero-downtime patching (ZDP) for its Amazon Aurora MySQL database. Multiple processing applications are using SSL certificates to connect to database endpoints and the read replicas.

Which factor will have the LEAST impact on the success of ZDP?

- A. Binary logging is enabled, or binary log replication is in progress.
- B. Current SSL connections are open to the database.
- C. Temporary tables or table locks are in use.
- D. The value of the `lower_case_table_names` server parameter was set to 0 when the tables were created.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 181

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

Which solution should the database specialist recommend?

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 182

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company uses a single-node Amazon RDS for MySQL DB instance for its production database. The DB instance runs in an AWS Region in the United States. A week before a big sales event, a new maintenance update is available for the DB instance. The maintenance update is marked as required. The company wants to minimize downtime for the DB instance and asks a database specialist to make the DB instance highly available until the sales event ends.

Which solution will meet these requirements?

- A. Defer the maintenance update until the sales event is over.
- B. Create a read replica with the latest update. Initiate a failover before the sales event.
- C. Create a read replica with the latest update. Transfer all read-only traffic to the read replica during the sales event.
- D. Convert the DB instance into a Multi-AZ deployment. Apply the maintenance update.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 183

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company is migrating a database in an Amazon RDS for SQL Server DB instance from one AWS Region to another. The company wants to minimize database downtime during the migration.

Which strategy should the company choose for this cross-Region migration?

- A. Back up the source database using native backup to an Amazon S3 bucket in the same Region. Then restore the backup in the target Region.
- B. Back up the source database using native backup to an Amazon S3 bucket in the same Region. Use Amazon S3 Cross-Region Replication to copy the backup to an S3 bucket in the target Region. Then restore the backup in the target Region.
- C. Configure AWS Database Migration Service (AWS DMS) to replicate data between the source and the target databases. Once the replication is in sync, terminate the DMS task.
- D. Add an RDS for SQL Server cross-Region read replica in the target Region. Once the replication is in sync, promote the read replica to master.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 184

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A financial company is hosting its web application on AWS. The application's database is hosted on Amazon RDS for MySQL with automated backups enabled. The application has caused a logical corruption of the database, which is causing the application to become unresponsive. The specific time of the corruption has been identified, and it was within the backup retention period.

How should a database specialist recover the database to the most recent point before corruption?

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 185

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A database specialist is designing an application to answer one-time queries. The application will query complex customer data and provide reports to end users. These reports can include many fields. The database specialist wants to give users the ability to query the database by using any of the provided fields. The database's traffic volume will be high but variable during peak times. However, the database will not have much traffic at other times during the day. Which solution will meet these requirements MOST cost-effectively?

- 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

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 186

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A financial services company runs an on-premises MySQL database for a critical application. The company is dissatisfied with its current database disaster recovery (DR) solution. The application experiences a significant amount of downtime whenever the database fails over to its DR facility. The application also experiences slower response times when reports are processed on the same database. To minimize the downtime in DR situations, the company has decided to migrate the database to AWS. The company requires a solution that is highly available and the most cost-effective.

Which solution meets these requirements?

- A. Create an Amazon RDS for MySQL Multi-AZ DB instance and configure a read replica in a different Availability Zone. Configure the application to reference the replica instance endpoint and report queries to reference the primary DB instance endpoint.
- B. Create an Amazon RDS for MySQL Multi-AZ DB instance and configure a read replica in a different Availability Zone. Configure the application to reference the primary DB instance endpoint and report queries to reference the replica instance endpoint.
- C. Create an Amazon Aurora DB cluster and configure an Aurora Replica in a different Availability Zone. Configure the application to reference the cluster endpoint and report queries to reference the reader endpoint.
- D. Create an Amazon Aurora DB cluster and configure an Aurora Replica in a different Availability Zone. Configure the application to reference the primary DB instance endpoint and report queries to reference the replica instance endpoint.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 187

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company with 500,000 employees needs to supply its employee list to an application used by human resources. Every 30 minutes, the data is exported using the LDAP service to load into a new Amazon DynamoDB table. The data model has a base table with Employee ID for the partition key and a global secondary index with Organization ID as the partition key.

While importing the data, a database specialist receives ProvisionedThroughputExceededException errors. After increasing the provisioned write capacity units (WCUs) to 50,000, the specialist receives the same errors. Amazon CloudWatch metrics show a consumption of 1,500 WCUs.

What should the database specialist do to address the issue?

- A. Change the data model to avoid hot partitions in the global secondary index.
- B. Enable auto scaling for the table to automatically increase write capacity during bulk imports.
- C. Modify the table to use on-demand capacity instead of provisioned capacity.
- D. Increase the number of retries on the bulk loading application.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 188

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company has an application that uses an Amazon DynamoDB table as its data store. During normal business days, the throughput requirements from the application are uniform and consist of 5 standard write calls per second to the DynamoDB table. Each write call has 2 KB of data.

For 1 hour each day, the company runs an additional automated job on the DynamoDB table that makes 20 write requests per second. No other application writes to the DynamoDB table. The DynamoDB table does not have to meet any additional capacity requirements.

How should a database specialist configure the DynamoDB table's capacity to meet these requirements MOST cost-effectively?

- A. Use DynamoDB provisioned capacity with 5 WCUs and auto scaling.
- B. Use DynamoDB provisioned capacity with 5 WCUs and a write-through cache that DynamoDB Accelerator (DAX) provides.
- C. Use DynamoDB provisioned capacity with 10 WCUs and auto scaling.
- D. Use DynamoDB provisioned capacity with 10 WCUs and no auto scaling.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 189

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company wants to build a new invoicing service for its cloud-native application on AWS. The company has a small development team and wants to focus on service feature development and minimize operations and maintenance as much as possible. The company expects the service to handle billions of requests and millions of new records every day. The service feature requirements, including data access patterns are well-defined. The service has an availability target of 99.99% with a milliseconds latency requirement. The database for the service will be the system of record for invoicing data.

Which database solution meets these requirements at the LOWEST cost?

- A. Amazon Neptune
- B. Amazon Aurora PostgreSQL Serverless
- C. Amazon RDS for PostgreSQL
- D. Amazon DynamoDB

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 190

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

Application developers have reported that an application is running slower as more users are added. The application database is running on an Amazon Aurora DB cluster with an Aurora Replica. The application is written to take advantage of read scaling through reader endpoints. A database specialist looks at the performance metrics of the database and determines that, as new users were added to the database, the primary instance CPU utilization steadily increased while the Aurora Replica CPU utilization remained steady.

How can the database specialist improve database performance while ensuring minimal downtime?

- A. Modify the Aurora DB cluster to add more replicas until the overall load stabilizes. Then, reduce the number of replicas once the application meets service level objectives.
- B. Modify the primary instance to a larger instance size that offers more CPU capacity.
- C. Modify a replica to a larger instance size that has more CPU capacity. Then, promote the modified replica.
- D. Restore the Aurora DB cluster to one that has an instance size with more CPU capacity. Then, swap the names of the old and new DB clusters.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 191

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company's development team needs to have production data restored in a staging AWS account. The production database is running on an Amazon RDS for PostgreSQL Multi-AZ DB instance, which has AWS KMS encryption enabled using the default KMS key. A database specialist planned to share the most recent automated snapshot with the staging account, but discovered that the option to share snapshots is disabled in the AWS Management Console.

What should the database specialist do to resolve this?

- A. Disable automated backups in the DB instance. Share both the automated snapshot and the default KMS key with the staging account. Restore the snapshot in the staging account and enable automated backups.
- B. Copy the automated snapshot specifying a custom KMS encryption key. Share both the copied snapshot and the custom KMS encryption key with the staging account. Restore the snapshot to the staging account within the same Region.
- C. Modify the DB instance to use a custom KMS encryption key. Share both the automated snapshot and the custom KMS encryption key with the staging account. Restore the snapshot in the staging account.
- D. Copy the automated snapshot while keeping the default KMS key. Share both the snapshot and the default KMS key with the staging account. Restore the snapshot in the staging account.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 192

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A software-as-a-service (SaaS) company is using an Amazon Aurora Serverless DB cluster for its production MySQL database. The DB cluster has general logs and slow query logs enabled. A database engineer must use the most operationally efficient solution with minimal resource utilization to retain the logs and facilitate interactive search and analysis.

Which solution meets these requirements?

- A. Use an AWS Lambda function to ship database logs to an Amazon S3 bucket. Use Amazon Athena and Amazon QuickSight to search and analyze the logs.
- B. Download the logs from the DB cluster and store them in Amazon S3 by using manual scripts. Use Amazon Athena and Amazon QuickSight to search and analyze the logs.
- C. Use an AWS Lambda function to ship database logs to an Amazon S3 bucket. Use Amazon Elasticsearch Service (Amazon ES) and Kibana to search and analyze the logs.
- D. Use Amazon CloudWatch Logs Insights to search and analyze the logs when the logs are automatically uploaded by the DB cluster.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 193

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A retail company uses Amazon Redshift Spectrum to run complex analytical queries on objects that are stored in an Amazon S3 bucket. The objects are joined with multiple dimension tables that are stored in an Amazon Redshift database. The company uses the database to create monthly and quarterly aggregated reports. Users who attempt to run queries are reporting the following error message: error: Spectrum Scan Error: Access throttled

Which solution will resolve this error?

- A. Check file sizes of fact tables in Amazon S3, and look for large files. Break up large files into smaller files of equal size between 100 MB and 1 GB
- B. Reduce the number of queries that users can run in parallel.
- C. Check file sizes of fact tables in Amazon S3, and look for small files. Merge the small files into larger files of at least 64 MB in size.
- D. Review and optimize queries that submit a large aggregation step to Redshift Spectrum.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 194

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

Which solution meets this requirement with the MOST operational efficiency?

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

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 195

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

An application reads and writes data to an Amazon RDS for MySQL DB instance. A new reporting dashboard needs read-only access to the database. When the application and reports are both under heavy load, the database experiences performance degradation. A database specialist needs to improve the database performance. What should the database specialist do to meet these requirements?

- A. Create a read replica of the DB instance. Configure the reports to connect to the replication instance endpoint.
- B. Create a read replica of the DB instance. Configure the application and reports to connect to the cluster endpoint.
- C. Enable Multi-AZ deployment. Configure the reports to connect to the standby replica.
- D. Enable Multi-AZ deployment. Configure the application and reports to connect to the cluster endpoint.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 196

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company is loading sensitive data into an Amazon Aurora MySQL database. To meet compliance requirements, the company needs to enable audit logging on the Aurora MySQL DB cluster to audit database activity. This logging will include events such as connections, disconnections, queries, and tables queried. The company also needs to publish the DB logs to Amazon CloudWatch to perform real-time data analysis.

Which solution meets these requirements?

- A. Modify the default option group parameters to enable Advanced Auditing. Restart the database for the changes to take effect.
- B. Create a custom DB cluster parameter group. Modify the parameters for Advanced Auditing. Modify the cluster to associate the new custom DB parameter group with the Aurora MySQL DB cluster.
- C. Take a snapshot of the database. Create a new DB instance, and enable custom auditing and logging to CloudWatch. Deactivate the DB instance that has no logging.
- D. Enable AWS CloudTrail for the DB instance. Create a filter that provides only connections, disconnections, queries, and tables queried.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 197

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company has an on-premises production Microsoft SQL Server with 250 GB of data in one database. A database specialist needs to migrate this on-premises SQL Server to Amazon RDS for SQL Server. The nightly native SQL Server backup file is approximately 120 GB in size. The application can be down for an extended period of time to complete the migration. Connectivity between the on-premises environment and AWS can be initiated from on-premises only.

How can the database be migrated from on-premises to Amazon RDS with the LEAST amount of effort?

- A. Back up the SQL Server database using a native SQL Server backup. Upload the backup files to Amazon S3. Download the backup files on an Amazon EC2 instance and restore them from the EC2 instance into the new production RDS instance.
- B. Back up the SQL Server database using a native SQL Server backup. Upload the backup files to Amazon S3. Restore the backup files from the S3 bucket into the new production RDS instance.
- C. Provision and configure AWS DMS. Set up replication between the on-premises SQL Server environment to replicate the database to the new production RDS instance.
- D. Back up the SQL Server database using AWS Backup. Once the backup is complete, restore the completed backup to an Amazon EC2 instance and move it to the new production RDS instance.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 198

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A database specialist needs to delete user data and sensor data 1 year after it was loaded in an Amazon DynamoDB table. TTL is enabled on one of the attributes. The database specialist monitors TTL rates on the Amazon CloudWatch metrics for the table and observes that items are not being deleted as expected.

What is the MOST likely reason that the items are not being deleted?

- A. The TTL attribute's value is set as a Number data type.
- B. The TTL attribute's value is set as a Binary data type.
- C. The TTL attribute's value is a timestamp in the Unix epoch time format in seconds.
- D. The TTL attribute's value is set with an expiration of 1 year.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 199

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company has deployed an application that uses an Amazon RDS for MySQL DB cluster. The DB cluster uses three read replicas. The primary DB instance is an 8XL-sized instance, and the read replicas are each XL-sized instances.

Users report that database queries are returning stale data. The replication lag indicates that the replicas are 5 minutes behind the primary DB instance. Status queries on the replicas show that the SQL_THREAD is 10 binlogs behind the IO_THREAD and that the IO_THREAD is 1 binlog behind the primary.

Which changes will reduce the lag? (Choose two.)

- A. Deploy two additional read replicas matching the existing replica DB instance size.
- B. Migrate the primary DB instance to an Amazon Aurora MySQL DB cluster and add three Aurora Replicas.
- C. Move the read replicas to the same Availability Zone as the primary DB instance.
- D. Increase the instance size of the primary DB instance within the same instance class.
- E. Increase the instance size of the read replicas to the same size and class as the primary DB instance.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 200

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company is using Amazon Aurora MySQL as the database for its retail application on AWS. The company receives a notification of a pending database upgrade and wants to ensure upgrades do not occur before or during the most critical time of year. Company leadership is concerned that an Amazon RDS maintenance window will cause an outage during data ingestion.

Which step can be taken to ensure that the application is not interrupted?

- A. Disable weekly maintenance on the DB cluster.
- B. Clone the DB cluster and migrate it to a new copy of the database.
- C. Choose to defer the upgrade and then find an appropriate down time for patching.
- D. Set up an Aurora Replica and promote it to primary at the time of patching.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 201

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

An ecommerce company uses Amazon DynamoDB as the backend for its payments system. A new regulation requires the company to log all data access requests for financial audits. For this purpose, the company plans to use AWS logging and save logs to Amazon S3

How can a database specialist activate logging on the database?

- A. Use AWS CloudTrail to monitor DynamoDB control-plane operations. Create a DynamoDB stream to monitor data-plane operations. Pass the stream to Amazon Kinesis Data Streams. Use that stream as a source for Amazon Kinesis Data Firehose to store the data in an Amazon S3 bucket.
- B. Use AWS CloudTrail to monitor DynamoDB data-plane operations. Create a DynamoDB stream to monitor control-plane operations. Pass the stream to Amazon Kinesis Data Streams. Use that stream as a source for Amazon Kinesis Data Firehose to store the data in an Amazon S3 bucket.
- C. Create two trails in AWS CloudTrail. Use Trail1 to monitor DynamoDB control-plane operations. Use Trail2 to monitor DynamoDB data-plane operations.
- D. Use AWS CloudTrail to monitor DynamoDB data-plane and control-plane operations.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 202

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A vehicle insurance company needs to choose a highly available database to track vehicle owners and their insurance details. The persisted data should be immutable in the database, including the complete and sequenced history of changes over time with all the owners and insurance transfer details for a vehicle.

The data should be easily verifiable for the data lineage of an insurance claim.

Which approach meets these requirements with MINIMAL effort?

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

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 203

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company stores session history for its users in an Amazon DynamoDB table. The company has a large user base and generates large amounts of session data. Teams analyze the session data for 1 week, and then the data is no longer needed. A database specialist needs to design an automated solution to purge session data that is more than 1 week old.

Which strategy meets these requirements with the MOST operational efficiency?

- A. Create an AWS Step Functions state machine with a DynamoDB DeleteItem operation that uses the ConditionExpression parameter to delete items older than a week. Create an Amazon EventBridge (Amazon CloudWatch Events) scheduled rule that runs the Step Functions state machine on a weekly basis.
- B. Create an AWS Lambda function to delete items older than a week from the DynamoDB table. Create an Amazon EventBridge (Amazon CloudWatch Events) scheduled rule that triggers the Lambda function on a weekly basis.
- C. Enable Amazon DynamoDB Streams on the table. Use a stream to invoke an AWS Lambda function to delete items older than a week from the DynamoDB table
- D. Enable TTL on the DynamoDB table and set a Number data type as the TTL attribute. DynamoDB will automatically delete items that have a TTL that is less than the current time.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 204

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company conducted a security audit of its AWS infrastructure. The audit identified that data was not encrypted in transit between application servers and a MySQL database that is hosted in Amazon RDS.

After the audit, the company updated the application to use an encrypted connection. To prevent this problem from occurring again, the company's database team needs to configure the database to require in-transit encryption for all connections.

Which solution will meet this requirement?

- A. Update the parameter group in use by the DB instance, and set the `require_secure_transport` parameter to ON.
- B. Connect to the database, and use `ALTER USER` to enable the `REQUIRE SSL` option on the database user.
- C. Update the security group in use by the DB instance, and remove port 80 to prevent unencrypted connections from being established.
- D. Update the DB instance, and enable the `Require Transport Layer Security` option.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 205

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A database specialist is designing an enterprise application for a large company. The application uses Amazon DynamoDB with DynamoDB Accelerator (DAX).

The database specialist observes that most of the queries are not found in the DAX cache and that they still require DynamoDB table reads.

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

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 206

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company plans to use AWS Database Migration Service (AWS DMS) to migrate its database from one Amazon EC2 instance to another EC2 instance as a full load task. The company wants the database to be inactive during the migration. The company will use a `dms.t3.medium` instance to perform the migration and will use the default settings for the migration.

Which solution will MOST improve the performance of the data migration?

- A. Increase the number of tables that are loaded in parallel.
- B. Drop all indexes on the source tables.
- C. Change the processing mode from the batch optimized apply option to transactional mode.
- D. Enable Multi-AZ on the target database while the full load task is in progress.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 207

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

Which solution will meet these requirements?

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 208

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company has a 4 |Ç|' on-premises Oracle Real Application Clusters (RAC) database. The company wants to migrate the database to AWS and reduce licensing costs. The company's application team wants to store JSON payloads that expire after 28 hours. The company has development capacity if code changes are required.

Which solution meets these requirements?

- A. Use Amazon DynamoDB and leverage the Time to Live (TTL) feature to automatically expire the data.
- B. Use Amazon RDS for Oracle with Multi-AZ. Create an AWS Lambda function to purge the expired data. Schedule the Lambda function to run daily using Amazon EventBridge.
- C. Use Amazon DocumentDB with a read replica in a different Availability Zone. Use DocumentDB change streams to expire the data.
- D. Use Amazon Aurora PostgreSQL with Multi-AZ and leverage the Time to Live (TTL) feature to automatically expire the data.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 209

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

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

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 210

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A software company uses an Amazon RDS for MySQL Multi-AZ DB instance as a data store for its critical applications. During an application upgrade process, a database specialist runs a custom SQL script that accidentally removes some of the default permissions of the master user.

What is the MOST operationally efficient way to restore the default permissions of the master user?

- A. Modify the DB instance and set a new master user password.
- B. Use AWS Secrets Manager to modify the master user password and restart the DB instance.
- C. Create a new master user for the DB instance.
- D. Review the IAM user that owns the DB instance, and add missing permissions.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 211

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

What should a database specialist do so that point-in-time recovery can be successful?

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

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 212

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company has a database fleet that includes an Amazon RDS for MySQL DB instance. During an audit, the company discovered that the data that is stored on the DB instance is unencrypted.

A database specialist must enable encryption for the DB instance. The database specialist also must encrypt all connections to the DB instance.

Which combination of actions should the database specialist take to meet these requirements? (Choose three.)

- A. In the RDS console, choose "Enable encryption" to encrypt the DB instance by using an AWS Key Management Service (AWS KMS) key.
- B. Encrypt the read replica of the unencrypted DB instance by using an AWS Key Management Service (AWS KMS) key. Fail over the read replica to the primary DB instance.
- C. Create a snapshot of the unencrypted DB instance. Encrypt the snapshot by using an AWS Key Management Service (AWS KMS) key. Restore the DB instance from the encrypted snapshot. Delete the original DB instance.
- D. Require SSL connections for applicable database user accounts.
- E. Use SSL/TLS from the application to encrypt a connection to the DB instance.
- F. Enable SSH encryption on the DB instance.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 213

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

Which solution will meet this requirement?

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

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 214

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A manufacturing company has an inventory system that stores information in an Amazon Aurora MySQL DB cluster. The database tables are partitioned. The database size has grown to 3 TB. Users run one-time queries by using a SQL client. Queries that use an equijoin to join large tables are taking a long time to run. Which action will improve query performance with the LEAST operational effort?

- A. Migrate the database to a new Amazon Redshift data warehouse.
- B. Enable hash joins on the database by setting the variable `optimizer_switch` to `hash_join=on`.
- C. Take a snapshot of the DB cluster. Create a new DB instance by using the snapshot, and enable parallel query mode.
- D. Add an Aurora read replica.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 215

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company is running a business-critical application on premises by using Microsoft SQL Server. A database specialist is planning to migrate the instance with several databases to the AWS Cloud. The database specialist will use SQL Server Standard edition hosted on Amazon EC2 Windows instances. The solution must provide high availability and must avoid a single point of failure in the SQL Server deployment architecture.

Which solution will meet these requirements?

- A. Create Amazon RDS for SQL Server Multi-AZ DB instances. Use Amazon S3 as a shared storage option to host the databases.
- B. Set up Always On Failover Cluster Instances as a single SQL Server instance. Use Multi-AZ Amazon FSx for Windows File Server as a shared storage option to host the databases.
- C. Set up Always On availability groups to group one or more user databases that fail over together across multiple SQL Server instances. Use Multi-AZ Amazon FSx for Windows File Server as a shared storage option to host the databases.
- D. Create an Application Load Balancer to distribute database traffic across multiple EC2 instances in multiple Availability Zones. Use Amazon S3 as a shared storage option to host the databases.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 216

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company is planning to use Amazon RDS for SQL Server for one of its critical applications. The company's security team requires that the users of the RDS for SQL Server DB instance are authenticated with on-premises Microsoft Active Directory credentials.

Which combination of steps should a database specialist take to meet this requirement? (Choose three.)

- A. Extend the on-premises Active Directory to AWS by using AD Connector.
- B. Create an IAM user that uses the AmazonRDSDirectoryServiceAccess managed IAM policy.
- C. Create a directory by using AWS Directory Service for Microsoft Active Directory.
- D. Create an Active Directory domain controller on Amazon EC2.
- E. Create an IAM role that uses the AmazonRDSDirectoryServiceAccess managed IAM policy.
- F. Create a one-way forest trust from the AWS Directory Service for Microsoft Active Directory directory to the on-premises Active Directory.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 217

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company is developing an application that performs intensive in-memory operations on advanced data structures such as sorted sets. The application requires sub-millisecond latency for reads and writes. The application occasionally must run a group of commands as an ACID-compliant operation. A database specialist is setting up the database for this application. The database specialist needs the ability to create a new database cluster from the latest backup of the production cluster.

Which type of cluster should the database specialist create to meet these requirements?

- A. Amazon ElastiCache for Memcached
- B. Amazon Neptune
- C. Amazon ElastiCache for Redis
- D. Amazon DynamoDB Accelerator (DAX)

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 218

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

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

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 219

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company requires near-real-time notifications when changes are made to Amazon RDS DB security groups.

Which solution will meet this requirement with the LEAST operational overhead?

- A. Configure an RDS event notification subscription for DB security group events.
- B. Create an AWS Lambda function that monitors DB security group changes. Create an Amazon Simple Notification Service (Amazon SNS) topic for notification.
- C. Turn on AWS CloudTrail. Configure notifications for the detection of changes to DB security groups.
- D. Configure an Amazon CloudWatch alarm for RDS metrics about changes to DB security groups.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 220

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A development team asks a database specialist to create a copy of a production Amazon RDS for MySQL DB instance every morning. The development team will use the copied DB instance as a testing environment for development. The original DB instance and the copy will be hosted in different VPCs of the same AWS account. The development team wants the copy to be available by 6 AM each day and wants to use the same endpoint address each day. Which combination of steps should the database specialist take to meet these requirements MOST cost-effectively? (Choose three.)

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

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 221

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A software company is conducting a security audit of its three-node Amazon Aurora MySQL DB cluster.

Which finding is a security concern that needs to be addressed?

- A. The AWS account root user does not have the minimum privileges required for client applications.
- B. Encryption in transit is not configured for all Aurora native backup processes.
- C. Each Aurora DB cluster node is not in a separate private VPC with restricted access.
- D. The IAM credentials used by the application are not rotated regularly.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 222

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company has an AWS CloudFormation stack that defines an Amazon RDS DB instance. The company accidentally deletes the stack and loses recent data from the DB instance. A database specialist must change the CloudFormation template for the RDS resource to reduce the chance of accidental data loss from the DB instance in the future.

Which combination of actions should the database specialist take to meet this requirement? (Choose three.)

- A. Set the DeletionProtection property to True.
- B. Set the MultiAZ property to True.
- C. Set the TerminationProtection property to True.
- D. Set the DeleteAutomatedBackups property to False.
- E. Set the DeletionPolicy attribute to No.
- F. Set the DeletionPolicy attribute to Retain.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 223

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company has branch offices in the United States and Singapore. The company has a three-tier web application that uses a shared database. The database runs on an Amazon RDS for MySQL DB instance that is hosted in the us-west-2 Region. The application has a distributed front end that is deployed in us-west-2 and in the ap-southeast-1 Region. The company uses this front end as a dashboard that provides statistics to sales managers in each branch office.

The dashboard loads more slowly in the Singapore branch office than in the United States branch office. The company needs a solution so that the dashboard loads consistently for users in each location.

Which solution will meet these requirements in the MOST operationally efficient way?

- A. Take a snapshot of the DB instance in us-west-2. Create a new DB instance in ap-southeast-2 from the snapshot. Reconfigure the ap-southeast-1 front-end dashboard to access the new DB instance.
- B. Create an RDS read replica in ap-southeast-1 from the primary DB instance in us-west-2. Reconfigure the ap-southeast-1 front-end dashboard to access the read replica.
- C. Create a new DB instance in ap-southeast-1. Use AWS Database Migration Service (AWS DMS) and change data capture (CDC) to update the new DB instance in ap-southeast-1. Reconfigure the ap-southeast-1 front-end dashboard to access the new DB instance.
- D. Create an RDS read replica in us-west-2, where the primary DB instance resides. Create a read replica in ap-southeast-1 from the read replica in us-west-2. Reconfigure the ap-southeast-1 front-end dashboard to access the read replica in ap-southeast-1.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 224

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company is using an Amazon ElastiCache for Redis cluster to host its online shopping website. Shoppers receive the following error when the website's application queries the cluster:

```
OOM command not allowed when used memory > 'maxmemory'
```

Which solutions will resolve this memory issues with the LEAST amount of effort? (Choose three.)

- A. Reduce the TTL value for keys on the node.
- B. Choose a larger node type.
- C. Test different values in the parameter group for the maxmemory-policy parameter to find the ideal value to use.
- D. Increase the number of nodes.
- E. Monitor the EngineCPUUtilization Amazon CloudWatch metric. Create an AWS Lambda function to delete keys on nodes when a threshold is reached.
- F. Increase the TTL value for keys on the node.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 225

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company uses Microsoft SQL Server on Amazon RDS in a Multi-AZ deployment as the database engine for its application. The company was recently acquired by another company. A database specialist must rename the database to follow a new naming standard.

Which combination of steps should the database specialist take to rename the database? (Choose two.)

- A. Turn off automatic snapshots for the DB instance. Rename the database with the `rdsadmin.dbo.rds_modify_db_name` stored procedure. Turn on the automatic snapshots.
- B. Turn off Multi-AZ for the DB instance. Rename the database with the `rdsadmin.dbo.rds_modify_db_name` stored procedure. Turn on Multi-AZ Mirroring.
- C. Delete all existing snapshots for the DB instance. Use the `rdsadmin.dbo.rds_modify_db_name` stored procedure.
- D. Update the application with the new database connection string.
- E. Update the DNS record for the DB instance.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 226

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

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

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

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 227

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A database specialist wants to ensure that an Amazon Aurora DB cluster is always automatically upgraded to the most recent minor version available. Noticing that there is a new minor version available, the database specialist has issues an AWS CLI command to enable automatic minor version updates. The command runs successfully, but checking the Aurora DB cluster indicates that no update to the Aurora version has been made.

What might account for this? (Choose two.)

- A. The new minor version has not yet been designated as preferred and requires a manual upgrade.
- B. Configuring automatic upgrades using the AWS CLI is not supported. This must be enabled expressly using the AWS Management Console.
- C. Applying minor version upgrades requires sufficient free space.
- D. The AWS CLI command did not include an `apply-immediately` parameter.
- E. Aurora has detected a breaking change in the new minor version and has automatically rejected the upgrade.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 228

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A security team is conducting an audit for a financial company. The security team discovers that the database credentials of an Amazon RDS for MySQL DB instance are hardcoded in the source code. The source code is stored in a shared location for automatic deployment and is exposed to all users who can access the location.

A database specialist must use encryption to ensure that the credentials are not visible in the source code.

Which solution will meet these requirements?

- A. Use an AWS Key Management Service (AWS KMS) key to encrypt the most recent database backup. Restore the backup as a new database to activate encryption.
- B. Store the source code to access the credentials in an AWS Systems Manager Parameter Store secure string parameter that is encrypted by AWS Key Management Service (AWS KMS). Access the code with calls to Systems Manager.
- C. Store the credentials in an AWS Systems Manager Parameter Store secure string parameter that is encrypted by AWS Key Management Service (AWS KMS). Access the credentials with calls to Systems Manager.
- D. Use an AWS Key Management Service (AWS KMS) key to encrypt the DB instance at rest. Activate RDS encryption in transit by using SSL certificates.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 229

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

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

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

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 230

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company's database specialist implements an AWS Database Migration Service (AWS DMS) task for change data capture (CDC) to replicate data from an on-premises Oracle database to Amazon S3. When usage of the company's application increases, the database specialist notices multiple hours of latency with the CDC.

Which solutions will reduce this latency? (Choose two.)

- A. Configure the DMS task to run in full large binary object (LOB) mode.
- B. Configure the DMS task to run in limited large binary object (LOB) mode.
- C. Create a Multi-AZ replication instance.
- D. Load tables in parallel by creating multiple replication instances for sets of tables that participate in common transactions.
- E. Replicate tables in parallel by creating multiple DMS tasks for sets of tables that do not participate in common transactions.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 231

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

An ecommerce company is running AWS Database Migration Service (AWS DMS) to replicate an on-premises Microsoft SQL Server database to Amazon RDS for SQL Server. The company has set up an AWS Direct Connect connection from its on-premises data center to AWS. During the migration, the company's security team receives an alarm that is related to the migration. The security team mandates that the DMS replication instance must not be accessible from public IP addresses.

What should a database specialist do to meet this requirement?

- A. Set up a VPN connection to encrypt the traffic over the Direct Connect connection.
- B. Modify the DMS replication instance by disabling the publicly accessible option.
- C. Delete the DMS replication instance. Recreate the DMS replication instance with the publicly accessible option disabled.
- D. Create a new replication VPC subnet group with private subnets. Modify the DMS replication instance by selecting the newly created VPC subnet group.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 232

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company is using an Amazon Aurora MySQL database with Performance Insights enabled. A database specialist is checking Performance Insights and observes an alert message that starts with the following phrase: `Performance Insights is unable to collect SQL Digest statistics on new queries`}

Which action will resolve this alert message?

- A. Truncate the events_statements_summary_by_digest table.
- B. Change the AWS Key Management Service (AWS KMS) key that is used to enable Performance Insights.
- C. Set the value for the performance_schema parameter in the parameter group to 1.
- D. Disable and reenable Performance Insights to be effective in the next maintenance window.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 233

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A bike rental company operates an application to track its bikes. The application receives location and condition data from bike sensors. The application also receives rental transaction data from the associated mobile app.

The application uses Amazon DynamoDB as its database layer. The company has configured DynamoDB with provisioned capacity set to 20% above the expected peak load of the application. On an average day, DynamoDB used 22 billion read capacity units (RCUs) and 60 billion write capacity units (WCUs). The application is running well. Usage changes smoothly over the course of the day and is generally shaped like a bell curve. The timing and magnitude of peaks vary based on the weather and season, but the general shape is consistent.

Which solution will provide the MOST cost optimization of the DynamoDB database layer?

- A. Change the DynamoDB tables to use on-demand capacity.
- B. Use AWS Auto Scaling and configure time-based scaling.
- C. Enable DynamoDB capacity-based auto scaling.
- D. Enable DynamoDB Accelerator (DAX).

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 234

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company has a quarterly customer survey. The survey uses an Amazon EC2 instance that is hosted in a public subnet to host a customer survey website. The company uses an Amazon RDS DB instance that is hosted in a private subnet in the same VPC to store the survey results.

The company takes a snapshot of the DB instance after a survey is complete, deletes the DB instance, and then restores the DB instance from the snapshot when the survey needs to be conducted again. A database specialist discovers that the customer survey website times out when it attempts to establish a connection to the restored DB instance.

What is the root cause of this problem?

- A. The VPC peering connection has not been configured properly for the EC2 instance to communicate with the DB instance.
- B. The route table of the private subnet that hosts the DB instance does not have a NAT gateway configured for communication with the EC2 instance.
- C. The public subnet that hosts the EC2 instance does not have an internet gateway configured for communication with the DB instance.
- D. The wrong security group was associated with the new DB instance when it was restored from the snapshot.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 235

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A company wants to improve its ecommerce website on AWS. A database specialist decides to add Amazon ElastiCache for Redis in the implementation stack to ease the workload off the database and shorten the website response times. The database specialist must also ensure the ecommerce website is highly available within the company's AWS Region.

How should the database specialist deploy ElastiCache to meet this requirement?

- A. Launch an ElastiCache for Redis cluster using the AWS CLI with the `-cluster-enabled` switch.
- B. Launch an ElastiCache for Redis cluster and select read replicas in different Availability Zones.
- C. Launch two ElastiCache for Redis clusters in two different Availability Zones. Configure Redis streams to replicate the cache from the primary cluster to another.
- D. Launch an ElastiCache cluster in the primary Availability Zone and restore the cluster's snapshot to a different Availability Zone during disaster recovery.

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 236

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

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

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

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

Show Suggested Answer



Exam question from Amazon's AWS Certified Database - Specialty

Question #: 237

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

An IT company wants to reduce its database operation costs in its development environment. The company's workflow creates an Amazon Aurora MySQL DB cluster for each development group. The DB clusters are used for only 8 hours a day. The DB clusters can be deleted at the end of a development cycle, which lasts 2 weeks. Which solution will meet these requirements MOST cost-effectively?

- A. Use AWS CloudFormation templates. Deploy a stack with a DB cluster for each development group. Delete the stack at the end of each development cycle.
- B. Use the Aurora cloning feature. Deploy a single development and test Aurora DB instance. Create clone instances for the development groups. Delete the clones at the end of each development cycle.
- C. Use Aurora Replicas. From the primary writer instance, create read replicas for each development group. Promote each read replica to a standalone DB cluster. Delete the standalone DB cluster at the end of each development cycle.
- D. Use Aurora Serverless. Restore a current Aurora snapshot to an Aurora Serverless cluster for each development group. Select the option to pause the compute capacity on the cluster after a specified amount of time with no activity. Delete the Aurora Serverless cluster at the end of each development cycle.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 238

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

A gaming company uses Amazon Aurora Serverless for one of its internal applications. The company's developers use Amazon RDS Data API to work with the Aurora Serverless DB cluster. After a recent security review, the company is mandating security enhancements. A database specialist must ensure that access to RDS Data API is private and never passes through the public internet.

What should the database specialist do to meet this requirement?

- A. Modify the Aurora Serverless cluster by selecting a VPC with private subnets.
- B. Modify the Aurora Serverless cluster by unchecking the publicly accessible option.
- C. Create an interface VPC endpoint that uses AWS PrivateLink for RDS Data API.
- D. Create a gateway VPC endpoint for RDS Data API.

Show Suggested Answer





Exam question from Amazon's AWS Certified Database - Specialty

Question #: 239

Topic #: 1

[\[All AWS Certified Database - Specialty Questions\]](#)

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

Which solution will meet these requirements MOST cost-effectively?

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

Show Suggested Answer

