[December-2022Free Braindump2go DBS-C01 PDF Dumps Download DBS-C01 155Q[Q245-Q292

December/2022 Latest Braindump2go DBS-C01 Exam Dumps with PDF and VCE Free Updated Today! Following are some new Braindump2go DBS-C01 Real Exam Questions!QUESTION 245The website of a manufacturing firm makes use of an Amazon Aurora PostgreSQL database cluster. Which settings will result in the LEAST amount of downtime for the application during failover? (Choose three.)A. Use the provided read and write Aurora endpoints to establish a connection to the Aurora DB cluster. B. Create an Amazon CloudWatch alert triggering a restore in another Availability Zone when the primary Aurora DB cluster is unreachable.C. Edit and enable Aurora DB cluster cache management in parameter groups.D. Set TCP keepalive parameters to a high value.E. Set JDBC connection string timeout variables to a low value.F. Set Java DNS caching timeouts to a high value.Answer: ACEExplanation:

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.BestPractices.html
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.cluster-cache-mgmt.html
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.BestPractices.html#AuroraPostgreSQL.Best
Practices.FastFailover.TCPKeepalivesQUESTION 246A company recently migrated its line-of-business (LOB) application to AWS.
The application uses an Amazon RDS for SQL Server DB instance as its database engine. The company must set up cross-Region disaster recovery for the application. The company needs a solution with the lowest possible RPO and RTO.Which solution will meet these requirements?A. Create a cross-Region read replica of the DB instance. Promote the read replica at the time of failover.
B. Set up SQL replication from the DB instance to an Amazon EC2 instance in the disaster recovery Region. Promote the EC2 instance as the primary server.C. Use AWS Database Migration Service (AWS KMS) for ongoing replication of the DB instance in the disaster recovery Region. Copy the snapshots to the disaster recovery Region.Answer: CExplanation:

https://aws.amazon.com/blogs/database/cross-region-disaster-recovery-of-amazon-rds-for-sql-server/QUESTION 247A company runs hundreds of Microsoft SQL Server databases on Windows servers in its on-premises data center. A database specialist needs to migrate these databases to Linux on AWS. Which combination of steps should the database specialist take to meet this requirement? (Choose three.)A. Install AWS Systems Manager Agent on the on-premises servers. Use Systems Manager Run Command to install the Windows to Linux replatforming assistant for Microsoft SQL Server Databases.B. Use AWS Systems Manager Run Command to install and configure the AWS Schema Conversion Tool on the on-premises servers.C. On the Amazon EC2 console, launch EC2 instances and select a Linux AMI that includes SQL Server.Install and configure AWS Systems Manager Agent on the EC2 instances.D. On the AWS Management Console, set up Amazon RDS for SQL Server DB instances with Linux as the operating system. Install AWS Systems Manager Agent on the DB instances by using an options group.E. Open the Windows to Linux replatforming assistant tool. Enter configuration details of the source and destination databases. Start migration.F. On the AWS Management Console, set up AWS Database Migration Service (AWS DMS) by entering details of the source SQL Server database and the destination SQL Server database on AWS.Start migration.Answer: ACEExplanation:

 $\underline{https://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/replatform-sql-server.html}$

https://dl.awsstatic.com/events/reinvent/2019/REPEAT 1 Leverage automation to re-platform SQL Server to Linux WIN322-R1.pdfQUESTION 248A company is running a blogging platform. A security audit determines that the Amazon RDS DB instance that is used by the platform is not configured to encrypt the data at rest. The company must encrypt the DB instance within 30 days. What should a database specialist do to meet this requirement with the LEAST amount of downtime? A. Create a read replica of the DB instance, and enable encryption. When the read replica is available, promote the read replica and update the endpoint that is used by the application. Delete the unencrypted DB instance. B. Take a snapshot of the DB instance. Make an encrypted copy of the snapshot. Restore the encrypted snapshot. When the new DB instance is available, update the endpoint that is used by the application. Delete the unencrypted DB instance. C. Create a new encrypted DB instance. Perform an initial data load, and set up logical replication between the two DB instances When the new DB instance is in sync with the source DB instance, update the endpoint that is used by the application. Delete the unencrypted DB cluster is available, update the endpoint that is used by the application to the cluster endpoint. Delete the unencrypted DB instance. Answer: CExplanation:

https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/encrypt-an-existing-amazon-rds-for-postgresql-db-instance.html When the new, encrypted copy of the DB instance becomes available, you can point your applications to the new database. However, if your project doesn't allow for significant downtime for this activity, you need an alternate approach that helps minimize

so that the cutover to the new, encrypted database can be done with minimal downtime. QUESTION 249An ecommerce company uses a backend application that stores data in an Amazon DynamoDB table. The backend application runs in a private subnet in a VPC and must connect to this table. The company must minimize any network latency that results from network connectivity issues, even during periods of heavy application usage. A database administrator also needs the ability to use a private connection to connect to the DynamoDB table from the application. Which solution will meet these requirements? A. Use network ACLs to ensure that any outgoing or incoming connections to any port except DynamoDB are deactivated. Encrypt API calls by using TLS.B. Create a VPC endpoint for DynamoDB in the application's VPC. Use the VPC endpoint to access the table.C. Create an AWS Lambda function that has access to DynamoDB. Restrict outgoing access only to this Lambda function from the application.D. Use a VPN to route all communication to DynamoDB through the company's own corporate network infrastructure. Answer: B Explanation:https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/vpc-endpoints-dynamodb.htmlQUESTION 250 A company's database specialist is building an Amazon RDS for Microsoft SQL Server DB instance to store hundreds of records in CSV format. A customer service tool uploads the records to an Amazon S3 bucket. An employee who previously worked at the company already created a custom stored procedure to map the necessary CSV fields to the database tables. The database specialist needs to implement a solution that reuses this previous work and minimizes operational overhead. Which solution will meet these requirements?A. Create an Amazon S3 event to invoke an AWS Lambda function. Configure the Lambda function to parse the .csv file and use a SQL client library to run INSERT statements to load the data into the tables.B. Write a custom .NET app that is hosted on Amazon EC2. Configure the .NET app to load the .csv file and call the custom stored procedure to insert the data into the tables.C. Download the .csv file from Amazon S3 to the RDS D drive by using an AWS msdb stored procedure. Call the custom stored procedure to insert the data from the RDS D drive into the tables.D. Create an Amazon S3 event to invoke AWS Step Functions to parse the .csv file and call the custom stored procedure to insert the data into the tables.Answer: CExplanation: https://www.mssqltips.com/sqlservertip/6619/rds-sql-server-data-import-from-amazon-s3/Amazon Web Service (AWS) recently announced a new feature of its Relational Database Service (RDS) for SQL Server. This feature allows a native integration between Amazon RDS SQL Server and Amazon S3. With this integration, it's now possible to import files from an Amazon S3 bucket into a local folder of the RDS instance. Similarly, files from that folder can be exported to S3. The RDS local folder path is D:S3.QUESTION 251A company hosts a 2 TB Oracle database in its on-premises data center. A database specialist is migrating the database from on premises to an Amazon Aurora PostgreSOL database on AWS. The database specialist identifies a problem that relates to compatibility Oracle stores metadata in its data dictionary in uppercase, but PostgreSQL stores the metadata in lowercase. The database specialist must resolve this problem to complete the migration. What is the MOST operationally efficient solution that meets these requirements?A. Override the default uppercase format of Oracle schema by encasing object names in quotation marks during creation.B. Use AWS Database Migration Service (AWS DMS) mapping rules with rule-action as convert-lowercase.C. Use the AWS Schema Conversion Tool conversion agent to convert the metadata from uppercase to lowercase.D. Use an AWS Glue job that is attached to an AWS Database Migration Service (AWS DMS) replication task to convert the metadata from uppercase to lowercase. Answer: BExplanation:

the downtime. This pattern uses the AWS Database Migration Service (AWS DMS) to migrate and continuously replicate the data

https://aws.amazon.com/premiumsupport/knowledge-center/dms-mapping-oracle-postgresql/QUESTION 252A financial company is running an Amazon Redshift cluster for one of its data warehouse solutions. The company needs to generate connection logs, user logs, and user activity logs. The company also must make these logs available for future analysis. Which combination of steps should a database specialist take to meet these requirements? (Choose two.)A. Edit the database configuration of the cluster by enabling audit logging. Direct the logging to a specified log group in Amazon CloudWatch Logs.B. Edit the database configuration of the cluster by enabling audit logging. Direct the logging to a specified Amazon S3 bucketC. Modify the cluster by enabling continuous delivery of AWS CloudTrail logs to Amazon S3.D. Create a new parameter group with the enable_user_activity_logging parameter set to true. Configure the cluster to use the new parameter group.E. Modify the system table to enable logging for each user. Answer: ADExplanation: AWS CloudWatch Logs are stored indefinitely and CloudWatch Log Insights is used to analyze the logs and query upon them.https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/AnalyzingLogData.html https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/WhatIsCloudWatchLogs.html "Log retention - By default, logs are kept indefinitely and never expire. You can adjust the retention policy for each log group, keeping the indefinite retention, or choosing a retention period between 10 years and one day."https://docs.aws.amazon.com/redshift/latest/mgmt/db-auditing.html QUESTION 253A company plans to migrate a MySQL-based application from an on-premises environment to AWS. The application performs database joins across several tables and uses indexes for faster query response times. The company needs the database to be highly available with automatic failover. Which solution on AWS will meet these requirements with the LEAST

operational overhead?A. Deploy an Amazon RDS DB instance with a read replica.B. Deploy an Amazon RDS Multi-AZ DB instance.C. Deploy Amazon DynamoDB global tables.D. Deploy multiple Amazon RDS DB instances. Use Amazon Route 53 DNS with failover health checks configured.Answer: BQUESTION 254A social media company is using Amazon DynamoDB to store user profile data and user activity data. Developers are reading and writing the data, causing the size of the tables to grow significantly. Developers have started to face performance bottlenecks with the tables.Which solution should a database specialist recommend to read items the FASTEST without consuming all the provisioned throughput for the tables?A. Use the Scan API operation in parallel with many workers to read all the items. Use the Query API operation to read multiple items that have a specific partition key and sort key. Use the GetItem API operation to read a single item.B. Use the Scan API operation with a filter expression that allows multiple items to be read. Use the Query API operation to read multiple items that have a specific partition key and sort key. Use the GetItem API operation to read a single item.C. Use the Scan API operation with a filter expression that allows multiple items to be read. Use the Query API operation to read a single item that has a specific primary key. Use the BatchGetItem API operation to read multiple items.D. Use the Scan API operation in parallel with many workers to read all the items. Use the Query API operation to read a single item that has a specific primary key Use the BatchGetItem API operation to read a multiple items. API operation to read and multiple items. API operation to read a single item that has a specific primary key Use the BatchGetItem API operation to read multiple items. API operation to read a single item that has a specific primary key Use the BatchGetItem API operation to read multiple items. API operation to read a single item that has a specific primary key Use the BatchGetItem API operation t

https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/SQLtoNoSQL.ReadData.htmlQUESTION 255A pharmaceutical company's drug search API is using an Amazon Neptune DB cluster. A bulk uploader process automatically updates the information in the database a few times each week. A few weeks ago during a bulk upload, a database specialist noticed that the database started torespond frequently with a ThrottlingException error. The problem also occurred with subsequent uploads. The database specialist must create a solution to prevent ThrottlingException errors for the database. The solution must minimize the downtime of the cluster. Which solution meets these requirements? A. Create a read replica that uses a larger instance size than the primary DB instance. Fail over the primary DB instance to the read replica. B. Add a read replica to each Availability Zone. Use an instance for the read replica that is the same size as the primary DB instance. Keep the traffic between the API and the database within the Availability Zone. C. Create a read replica that uses a larger instance size than the primary DB instance. Offload the reads from the primary DB instance.D. Take the latest backup, and restore it in a DB cluster of a larger size. Point the application to the newly created DB cluster. Answer: CExplanation:

https://docs.aws.amazon.com/neptune/latest/userguide/manage-console-add-replicas.htmlNeptune replicas connect to the same storage volume as the primary DB instance and support only read operations. Neptune replicas can offload read workloads from the primary DB instance.QUESTION 256A global company is developing an application across multiple AWS Regions. The company needs a database solution with low latency in each Region and automatic disaster recovery. The database must be deployed in an active-active configuration with automatic data synchronization between Regions. Which solution will meet these requirements with the LOWEST latency?A. Amazon RDS with cross-Region read replicasB. Amazon DynamoDB global tablesC. Amazon Aurora global databaseD. Amazon Athena and Amazon S3 with S3 Cross Region ReplicationAnswer: BQUESTION 257A pharmaceutical company uses Amazon Quantum Ledger Database (Amazon QLDB) to store its clinical trial data records. The company has an application that runs as AWS Lambda functions. The application is hosted in the private subnet in a VPC. The application does not have internet access and needs to read some of the clinical data records. The company is concerned that traffic between the QLDB ledger and the VPC could leave the AWS network. The company needs to secure access to the QLDB ledger and allow the VPC traffic to have read-only access. Which security strategy should a database specialist implement to meet these requirements?A. Move the QLDB ledger into a private database subnet inside the VPC. Run the Lambda functions inside the same VPC in an application private subnet. Ensure that the VPC route table allows read-only flow from the application subnet to the database subnet.B. Create an AWS PrivateLink VPC endpoint for the QLDB ledger. Attach a VPC policy to the VPC endpoint to allow read-only traffic for the Lambda functions that run inside the VPC.C. Add a security group to the QLDB ledger to allow access from the private subnets inside the VPC where the Lambda functions that access the QLDB ledger are running.D. Create a VPN connection to ensure pairing of the private subnet where the Lambda functions are running with the private subnet where the QLDB ledger is deployed. Answer: BExplanation: https://docs.aws.amazon.com/qldb/latest/developerguide/vpc-endpoints.html QUESTION 258A company's application development team wants to share an automated snapshot of its Amazon RDS database with another team. The database is encrypted with a custom AWS Key Management Service (AWS KMS) key under the "WeShare" AWS account. The application development team needs to share the DB snapshot under the "WeReceive" AWS account. Which combination of actions must the application development team take to meet these requirements? (Choose two.)A. Add access from the "WeReceive" account to the custom AWS KMS key policy of the sharing team.B. Make a copy of the DB snapshot, and set the encryption option to disable.C. Share the DB snapshot by setting the DB snapshot visibility option to public.D. Make a copy of

the DB snapshot, and set the encryption option to enable.E. Share the DB snapshot by using the default AWS KMS encryption key. Answer: ADExplanation: https://aws.amazon.com/premiumsupport/knowledge-center/rds-snapshots-share-account/QUESTION 259A company is using Amazon Redshift as its data warehouse solution. The Redshift cluster handles the following types of workloads:- Real-time inserts through Amazon Kinesis Data Firehose-Bulk inserts through COPY commands from Amazon S3-Analytics through SQL queriesRecently, the cluster has started to experience performance issues. Which combination of actions should a database specialist take to improve the cluster's performance? (Choose three.)A. Modify the Kinesis Data Firehose delivery stream to stream the data to Amazon S3 with a high buffer size and to load the data into Amazon Redshift by using the COPY command.B. Stream real-time data into Redshift temporary tables before loading the data into permanent tables.C. For bulk inserts, split input files on Amazon S3 into multiple files to match the number of slices on Amazon Redshift. Then use the COPY command to load data into Amazon Redshift.D. For bulk inserts, use the parallel parameter in the COPY command to enable multi-threading.E. Optimize analytics SQL queries to use sort keys.F. Avoid using temporary tables in analytics SQL queries.Answer: BCEExplanation:

https://aws.amazon.com/blogs/big-data/top-10-performance-tuning-techniques-for-amazon-redshift/Tip #6: Improving the efficiency of temporary tablesTip #9: Maintaining efficient data loadsAmazon Redshift best practices suggest using the COPY command to perform data loads of file-based data. Tip #3: Sort key recommendationSorting a table on an appropriate sort key can accelerate query performance, especially queries with range-restricted predicates, by requiring fewer table blocks to be read from disk. QUESTION 260An information management services company is storing JSON documents on premises. The company is using a MongoDB 3.6 database but wants to migrate to AWS. The solution must be compatible, scalable, and fully managed. The solution also must result in as little downtime as possible during the migration. Which solution meets these requirements? A. Create an AWS Database Migration Service (AWS DMS) replication instance, a source endpoint for MongoDB, and a target endpoint of Amazon DocumentDB (with MongoDB compatibility). B. Create an AWS Database Migration Service (AWS DMS) replication instance, a source endpoint for MongoDB, and a target endpoint of a MongoDB image that is hosted on Amazon EC2C. Use the mongodump and mongorestore tools to migrate the data from the source MongoDB deployment to Amazon DocumentDB (with MongoDB compatibility). D. Use the mongodump and mongorestore tools to migrate the data from the source MongoDB deployment to a MongoDB image that is hosted on Amazon EC2. Answer: AExplanation:

https://docs.aws.amazon.com/documentdb/latest/developerguide/docdb-migration.html#docdb-migration-approachesQUESTION 261A company stores critical data for a department in Amazon RDS for MySQL DB instances. The department was closed for 3 weeks and notified a database specialist that access to the RDS DB instances should not be granted to anyone during this time. To meet this requirement, the database specialist stopped all the DB instances used by the department but did not select the option to create a snapshot. Before the 3 weeks expired, the database specialist discovered that users could connect to the database successfully.What could be the reason for this?A. When stopping the DB instance, the option to create a snapshot should have been selected.B. When stopping the DB instance, the duration for stopping the DB instance should have been selected.C. Stopped DB instances will automatically restart if the number of attempted connections exceeds the threshold set.D. Stopped DB instances will automatically restart if the instance is not manually started after 7 days.Answer: DExplanation: https://aws.amazon.com/premiumsupport/knowledge-center/rds-stop-seven-days/QUESTION 262A company uses an on-premises Microsoft SQL Server database to host relational and JSON data and to run daily ETL and advanced analytics. The company wants to migrate the database to the AWS Cloud. Database specialist must choose one or more AWS services to run the company's

workloads. Which solution will meet these requirements in the MOST operationally efficient manner? A. Use Amazon Redshift for relational data. Use Amazon DynamoDB for JSON dataB. Use Amazon Redshift for relational data and JSON data. C. Use Amazon RDS for relational data. Use Amazon Neptune for JSON dataD. Use Amazon Redshift for relational data. Use Amazon S3 for JSON data. Answer: BExplanation: https://docs.aws.amazon.com/redshift/latest/dg/super-overview.htm QUESTION 263An online gaming company is using an Amazon DynamoDB table in on-demand mode to store game scores. After an intensive advertisement campaign in SouthAmerica, the average number of concurrent users rapidly increases from 100,000 to 500,000 in less than 10 minutes every day around 5 PM. The on-call software reliability engineer has observed that the application logs contain a high number of DynamoDB throttling exceptions caused by game score insertions around 5 PM. Customer service has also reported that several users are complaining about their scores not being registered. How should the database administrator remediate this issue at the lowest cost? A. Enable auto scaling and set the target usage rate to 90%. B. Switch the table to provisioned mode and enable auto scaling. C. Switch the table to provisioned mode and set the throughput to the peak value. D. Create a DynamoDB Accelerator cluster and use it to access the DynamoDB table. Answer: BQUESTION 264A 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.Answer: CExplanation: https://aws.amazon.com/about-aws/whats-new/2020/02/amazon-rds-data-api-now-supports-aws-privatelink/QUESTION 265A startup company in the travel industry wants to create an application that includes a personal travel assistant to display information for nearby airports based on user location. The application will use Amazon DynamoDB and must be able to access and display attributes such as airline names, arrival times, and flight numbers. However, the application must not be able to access or display pilot names or passenger counts. Which solution will meet these requirements MOST cost-effectively? A. Use a proxy tier between the application and DynamoDB to regulate access to specific tables, items, and attributes.B. Use IAM policies with a combination of IAM conditions and actions to implement fine-grained access control.C. Use DynamoDB resource policies to regulate access to specific tables, items, and attributes.D. Configure an AWS Lambda function to extract only allowed attributes from tables based on user profiles. Answer: BExplanation: https://aws.amazon.com/blogs/aws/fine-grained-access-control-for-amazon-dynamodb/ QUESTION 266A large IT hardware manufacturing company wants to deploy a MySQL database solution in the AWS Cloud. The solution should quickly create copies of the company's production databases for test purposes. The solution must deploy the test databases in minutes, and the test data should match the latest production data as closely as possible. Developers must also be able to make changes in the test database and delete the instances afterward. Which solution meets these requirements? A. Leverage Amazon RDS for MySQL with write-enabled replicas running on Amazon EC2. Create the test copies using a mysqidump backup from the RDS for MySQL DB instances and importing them into the new EC2 instances.B. Leverage Amazon Aurora MySQL. Use database cloning to create multiple test copies of the production DB clusters.C. Leverage Amazon Aurora MySQL. Restore previous production DB instance snapshots into new test copies of Aurora MySQL DB clusters to allow them to make changes.D. Leverage Amazon RDS for MySQL. Use database cloning to create multiple developer copies of the production DB instance. Answer: BQUESTION 267An 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.Answer: CExplanation: https://aws.amazon.com/premiumsupport/knowledge-center/dms-disable-public-access/QUESTION 268A 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. Answer: AExplanation: https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER PerfInsights.UsingDashboard.AnalyzeDBLoad.AdditionalMetri cs.MySQL.htmlQUESTION 269A 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). Answer: CExplanation: https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/AutoScaling.htmlQUESTION 270A 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. Answer: DQUESTION 271A 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. Answer: BExplanation: https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/AutoFailover.htmlYou can enable Multi-AZ only on Redis (cluster mode disabled) clusters that have at least one available read replica. Clusters without read replicas do not provide high availability or fault tolerance.QUESTION 272A 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. Answer: BCExplanation: https://aws.amazon.com/blogs/database/migrate-tde-enabled-sql-server-databases-to-amazon-rds-for-sql-server/QUESTION 273A 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

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_UpgradeDBInstance.Upgrading.htmlCall the modify-db-instance Amazon CLI command. Specify the name of your DB instance for the --db-instance-identifier option and true for the --auto-minor-version-upgrade option. Optionally, specify the --apply-immediately option to immediately enable this setting for your DB instance. Run a separate modify-db-instance command for each DB instance in the cluster.

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.Answer: ADExplanation:When Amazon RDS designates a minor engine version as the preferred minor engine version, each database that meets both of the following conditions is

https://docs.amazonaws.cn/en_us/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Updates.Patching.html#AuroraMySQL.Updates.AMVUQUESTION 274A 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

upgraded to the minor engine version automatically.

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. Answer: C Explanation:only creds in system manager secure parameter.QUESTION 275A 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. Answer: AEExplanation: https://aws.amazon.com/blogs/database/configuring-amazon-elasticache-for-redis-for-higher-availability/QUESTION 276A 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 instances.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. Answer: BEQUESTION 277A 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. Answer: DExplanation: Rotate your IAM credentials regularly.OUESTION 278A bank is using an Amazon RDS for MySOL DB instance in a proof of concept. A database specialist is evaluating automated database snapshots and cross-Region snapshot copies as -part of this proof of concept. After validating three automated snapshots successfully, the database specialist realizes that the fourth snapshot was not created. Which of the following are possible reasons why the snapshot was not created? (Choose two.)A. A copy of the automated snapshot for this DB instance is in progress within the same AWS Region.B. A copy of a manual snapshot for this DB instance is in progress for only certain databases within the DB instance.C. The RDS maintenance window is not specified.D. The DB instance is in the STORAGE_FULL state.E. RDS event notifications have not been enabled. Answer: ADQUESTION 279A 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. Answer: BQUESTION 280A 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: 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.Answer: ABCExplanation:

https://aws.amazon.com/premiumsupport/knowledge-center/oom-command-not-allowed-redis/QUESTION 281A 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. Answer: BDExplanation:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Appendix.SQLServer.CommonDBATasks.RenamingDB.html
QUESTION 282A 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.Answer: CEFQUESTION 283A 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 MemcachedB. Amazon NeptuneC. Amazon ElastiCache for RedisD. Amazon DynamoDB Accelerator (DAX)Answer: CExplanation:

https://aws.amazon.com/elasticache/redis-vs-memcached/

https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/elasticache-use-cases.html#elasticache-for-redis-use-cases-gaming QUESTION 284A 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. Answer: AExplanation:

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Auditing.htmlQUESTION 285A 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. Answer: A Explanation:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER Events.Messages.html#USER Events.Messages.security-group QUESTION 286A 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.Answer: ABCQUESTION 287A 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 periodD. Set up a scheduled job to create manual DB instance snapshots.Answer: CExplanation:You can restore a DB instance to a specific point in time (PITR), creating a new DB instance. To support PITR, your DB instances must have backup retention set to a nonzero value.

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/custom-backup-sqlserver.html

https://aws.amazon.com/blogs/database/setting-up-a-binlog-server-for-amazon-rds-mysql-and-mariadb-using-mariadb-maxscale/
"After you run the command, it's okay to enable backup retention on the RDS instance by using the AWS CLI or the console.

Enabling backup retention also enables binary logging."

https://aws.amazon.com/blogs/storage/point-in-time-recovery-and-continuous-backup-for-amazon-rds-with-aws-backup/
QUESTION 288A 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.Answer: ACEExplanation:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.Encryption.html#Overview.Encryption.Enabling QUESTION 289A company has an ecommerce website that runs on AWS. The website uses an Amazon RDS for MySQL database. A database specialist wants to enforce the use of temporary credentials to access the database. Which solution will meet this requirement?A. Use MySQL native database authentication.B. Use AWS Secrets Manager to rotate the credentials.C. Use AWS Identity and Access Management (IAM) database authentication.D. Use AWS Systems Manager Parameter Store for authentication. Answer: CQUESTION 290A 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. Answer: BExplanation: https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.BestPractices.htmlQUESTION 291A 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 SOL 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. Answer: BExplanation:

https://docs.aws.amazon.com/prescriptive-guidance/latest/migration-sql-server/ec2-fci.html}An FCI is generally preferable over an Always on availability group when: You're using SQL Server Standard edition instead of Enterprise edition.QUESTION 292A 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.Answer: BExplanation:

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_BestPractices.html#CHAP_BestPractices.Performance
For a full load task, we recommend that you drop primary key indexes, secondary indexes, referential integrity constraints, and data manipulation language (DML) triggers. Or you can delay their creation until after the full load tasks are complete. You don't need indexes during a full load task, and indexes incur maintenance overhead if they are present. Because the full load task loads groups of tables at a time, referential integrity constraints are violated. Similarly, insert, update, and delete triggers can cause errors, for example if a row insert is triggered for a previously bulk loaded table. Other types of triggers also affect performance due to added processing. https://docs.aws.amazon.com/dms/latest/userguide/CHAP_BestPractices.htmlResources From:1.2022 Latest Braindump2go
DBS-C01 Exam Dumps (PDF & VCE) Free Share:https://www.braindump2go.com/dbs-c01.html2.2022 Latest Braindump2go
DBS-C01 PDF and DBS-C01 VCE Dumps Free Share:

https://drive.google.com/drive/folders/11Uhzdg235eGRwUigG6XMx64UAN26dflw?usp=sharing3.2022 Free Braindump2go DBS-C01 Exam Questions Download:https://www.braindump2go.com/free-online-pdf/DBS-C01-PDF-Dumps(245-292).pdfFree Resources from Braindump2go,We Devoted to Helping You 100% Pass All Exams!