AWS-Certified-Database-Specialty Dumps

AWS Certified Database - Specialty

https://www.certleader.com/AWS-Certified-Database-Specialty-dumps.html

NEW QUESTION 1
A company migrated one of its business-critical database workloads to an Amazon Aurora Multi-AZ DB cluster. The company requires a very low RTO and needs to improve the application recovery time after database failovers.
Which approach meets these requirements?

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

Answer: C

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

NEW QUESTION 2
A 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 tabl
E. Then use the restore-table-from-backup command and modify the partition key.

Answer: A

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

NEW QUESTION 3
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 AW
B. Leverage AWS SCT and apply the converted schema to Amazon Redshif
C. Once complete, start an AWS DMS task to move the data from the source to Amazon S3. Use AWS Glue to load the data from Amazon S3 to Amazon Redshift.
D. Leverage AWS SCT and apply the converted schema to Amazon Redshif
E. Start an AWS DMS task with two AWS Snowball Edge devices to copy data from on-premises to Amazon S3 with AWS KMS encryptio
F. Use AWS DMS to finish copying data to Amazon Redshift.
G. Leverage AWS SCT and apply the converted schema to Amazon Redshif
H. Once complete, use a fleet of 10 TB dedicated encrypted drives using the AWS Import/Export feature to copy data from on-premises to Amazon S3 with AWS KMS encryptio
I. Use AWS Glue to load the data to Amazon redshift.
J. Set up a VPN tunnel for encrypting data over the network from the data center to AW
K. Leverage a native database export feature to export the data and compress the file
L. Use the aws S3 cp multi-port upload command to upload these files to Amazon S3 with AWS KMS encryptio
M. Once complete, load the data to Amazon Redshift using AWS Glue.

Answer: B

Explanation:
https://aws.amazon.com/blogs/database/new-aws-dms-and-aws-snowball-integration-enables-mass-database-mi

NEW QUESTION 4
Recently, a financial institution created a portfolio management service. The application's backend is powered by Amazon Aurora, which supports MySQL.
The firm demands a response time of five minutes and a response time of five minutes. A database professional must create a disaster recovery system that is both efficient and has a low replication latency.
How should the database professional tackle these requirements?

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

Answer: B

Explanation:

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

NEW QUESTION 5
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 afterthe-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 bucke
B. Set a lifecycle policy to expire the objects after 90 days.
C. Modify the RDS databases to publish log to Amazon CloudWatch Log
D. Change the log retention policy for each log group to expire the events after 90 days.
E. Write a stored procedure in each RDS database to download the logs and consolidate the log files in an Amazon S3 bucke
F. Set a lifecycle policy to expire the objects after 90 days.
G. Create an AWS Lambda function to download the logs from the RDS databases and publish the logs to Amazon CloudWatch Log
H. Change the log retention policy for the log group to expire the events after 90 days.

Answer: B

Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_LogAccess.html
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_LogAccess.Procedural.UploadtoCloudWat
https://aws.amazon.com/premiumsupport/knowledge-center/rds-aurora-mysql-logs-cloudwatch/
https://docs.aws.amazon.com/AmazonCloudWatchLogs/latest/APIReference/API_PutRetentionPolicy.html

NEW QUESTION 6
A gaming company wants to deploy a game in multiple Regions. The company plans to save local high scores in Amazon DynamoDB tables in each Region. A Database Specialist needs to design a solution to automate the deployment of the database with identical configurations in additional Regions, as needed. The solution should also automate configuration changes across all Regions.
Which solution would meet these requirements and deploy the DynamoDB tables?

A. Create an AWS CLI command to deploy the DynamoDB table to all the Regions and save it for future deployments.
B. Create an AWS CloudFormation template and deploy the template to all the Regions.
C. Create an AWS CloudFormation template and use a stack set to deploy the template to all the Regions.
D. Create DynamoDB tables using the AWS Management Console in all the Regions and create a step-by- step guide for future deployments.

Answer: C

Explanation:
https://aws.amazon.com/blogs/aws/use-cloudformation-stacksets-to-provision-resources-across-multiple-aws-ac
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stacksets-concepts.html

NEW QUESTION 7
An AWS CloudFormation stack that included an Amazon RDS DB instance was accidentally deleted and recent data was lost. A Database Specialist needs to add RDS settings to the CloudFormation template to reduce the chance of accidental instance data loss in the future.
Which settings will meet this requirement? (Choose three.)

A. Set DeletionProtection to True
B. Set MultiAZ to True
C. Set TerminationProtection to True
D. Set DeleteAutomatedBackups to False
E. Set DeletionPolicy to Delete
F. Set DeletionPolicy to Retain

Answer: ACF

NEW QUESTION 8
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 30days.
D. Create an AWS Lambda function to run on the first day of every month to create an automated RDS snapshot.

Answer: A

Explanation:
Unlike automated backups, manual snapshots aren't subject to the backup retention period. Snapshots don't expire. For very long-term backups of MariaDB, MySQL, and PostgreSQL data, we recommend exporting snapshot data to Amazon S3. If the major version of your DB engine is no longer supported, you can't restore to that version from a snapshot. https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_CreateSnapshot.html

NEW QUESTION 9
A company wants to automate the creation of secure test databases with random credentials to be stored safely for later use. The credentials should have sufficient information about each test database to initiate a connection and perform automated credential rotations. The credentials should not be logged or stored

anywhere in an unencrypted form.
Which steps should a Database Specialist take to meet these requirements using an AWS CloudFormation template?

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

Answer: C

NEW QUESTION 10
A banking company recently launched an Amazon RDS for MySQL DB instance as part of a proof-of-concept project. A database specialist has configured automated database snapshots. As a part of routine testing, the database specialist noticed one day that the automated database snapshot was not created. Which of the following are possible reasons why the snapshot was not created? (Choose two.)

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

Answer: AD

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

NEW QUESTION 10
A business maintains a SQL Server database on-premises. Active Directory authentication is used to provide users access to the database. The organization transferred their database successfully to Amazon RDS for SQL Server. The organization, however, has reservations regarding user authentication in the AWS Cloud environment.
Which authentication solution should a database professional provide?

A. Deploy Active Directory Federation Services (AD FS) on premises and configure it with an on-premises Active Director
B. Set up delegation between the on- premises AD FS and AWS Security Token Service (AWS STS) to map user identities to a role using theAmazonRDSDirectoryServiceAccess managed IAM policy.
C. Establish a forest trust between the on-premises Active Directory and AWS Directory Service for Microsoft Active Director
D. Use AWS SSO to configure an Active Directory user delegated to access the databases in RDS for SQL Server.
E. Use Active Directory Connector to redirect directory requests to the company€™s on-premises Active Directory without caching any information in the clou
F. 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.
G. Establish a forest trust between the on-premises Active Directory and AWS Directory Service for Microsoft Active Director
H. Ensure RDS for SQL Server is using mixed mode authenticatio
I. 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.

Answer: D

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

NEW QUESTION 12
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.

Answer: A

NEW QUESTION 13
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 Regio
B. Switch to the destination Region, enable Amazon Redshift cross-Region snapshots, and use the KMS key of the source Region.

C. Create a new IAM role with access to the KMS ke
D. Enable Amazon Redshift cross-Region replication using the new IAM role, and use the KMS key of the source Region.
E. 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.
F. Create a new KMS customer master key in the destination Region and create a new IAM role with access to the new KMS ke
G. Enable Amazon Redshift cross-Region replication in the source Region anduse the KMS key of the destination Region.

Answer: C

Explanation:
If you want to enable cross-Region snapshot copy for an AWS KMS–encrypted cluster, you must configure a snapshot copy grant for a root key in the destination AWS Region Source-Region : configure a cross-Region snapshot for an AWS KMS–encrypted cluster In Destination AWS Region : choose the AWS Region to which to copy snapshots.
https://docs.aws.amazon.com/redshift/latest/mgmt/managing-snapshots-console.html#xregioncopy-kms-encrypt

NEW QUESTION 16
A company is building a new web platform where user requests trigger an AWS Lambda function that performs an insert into an Amazon Aurora MySQL DB cluster. Initial tests with less than 10 users on the new platform yielded successful execution and fast response times. However, upon more extensive tests with the actual target of 3,000 concurrent users, Lambda functions are unable to connect to the DB cluster and receive too many connections errors.
Which of the following will resolve this issue?

A. Edit the my.cnf file for the DB cluster to increase max_connections
B. Increase the instance size of the DB cluster
C. Change the DB cluster to Multi-AZ
D. Increase the number of Aurora Replicas

Answer: B

Explanation:
Max_connection is a formula in RDS parameter group:
GREATEST({log(DBInstanceClassMemory/805306368)*45},{log(DBInstanceClassMemory/8187281408)*100
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Managing.Performance.htm You can increase the maximum number of connections to your Aurora MySQL DB instance by scaling the instance up to a DB instance class with more memory, or by setting a larger value for the max_connections parameter in the DB parameter group for your instance, up to 16,000. You must change a larger value for the max_connections parameter in the DB parameter group, not edit my.cnf, it is not physical server hosting MySQL.

NEW QUESTION 17
A company has a database monitoring solution that uses Amazon CloudWatch for its Amazon RDS for SQL Server environment. The cause of a recent spike in CPU utilization was not determined using the standard metrics that were collected. The CPU spike caused the application to perform poorly, impacting users. A Database Specialist needs to determine what caused the CPU spike.
Which combination of steps should be taken to provide more visibility into the processes and queries running during an increase in CPU load? (Choose two.)

A. Enable Amazon CloudWatch Events and view the incoming T-SQL statements causing the CPU to spike.
B. Enable Enhanced Monitoring metrics to view CPU utilization at the RDS SQL Server DB instance level.
C. Implement a caching layer to help with repeated queries on the RDS SQL Server DB instance.
D. Use Amazon QuickSight to view the SQL statement being run.
E. Enable Amazon RDS Performance Insights to view the database load and filter the load by waits, SQL statements, hosts, or users.

Answer: BE

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/rds-instance-high-cpu/ "Several factors can cause an increase in CPU utilization. For example, user-initiated heavy workloads, analytic queries, prolonged deadlocks and lock waits, multiple concurrent transactions, long-running transactions, or other processes that utilize CPU resources. First, you can identify the source of the CPU usage by: Using Enhanced Monitoring Using Performance Insights"

NEW QUESTION 21
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 will the appropriate settings

Answer: C

Explanation:
 https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_PerfInsights.Enabling.html https://aws.amazon.com/rds/performance-insights/
https://aws.amazon.com/blogs/database/tuning-amazon-rds-for-mysql-with-performance-insights/

NEW QUESTION 24
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 fil
B. Add the required corporate network IPs and remove the unwanted IPs.
C. Modify the associated security grou

D. Add the required corporate network IPs and remove the unwanted IPs.
E. Move the DB instance to a private subnet using AWS DMS.
F. Enable VPC peering between the application host running on the corporate network and the VPC associated with the DB instance.
G. Disable the publicly accessible setting.
H. Connect to the DB instance using private IPs and a VPN.

Answer: BEF

Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_VPC.WorkingWithRDSInstanceinaVPC.ht

NEW QUESTION 25
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.

Answer: C

Explanation:
https://aws.amazon.com/blogs/database/best-practices-for-migrating-rds-for-mysql-databases-to-amazon-aurora/
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.Migrating.html#Aurora

NEW QUESTION 26
A company has a heterogeneous six-node production Amazon Aurora DB cluster that handles online transaction processing (OLTP) for the core business and OLAP reports for the human resources department. To match compute resources to the use case, the company has decided to have the reporting workload for the human resources department be directed to two small nodes in the Aurora DB cluster, while every other workload goes to four large nodes in the same DB cluster.
Which option would ensure that the correct nodes are always available for the appropriate workload while meeting these requirements?

A. Use the writer endpoint for OLTP and the reader endpoint for the OLAP reporting workload.
B. Use automatic scaling for the Aurora Replica to have the appropriate number of replicas for the desired workload.
C. Create additional readers to cater to the different scenarios.
D. Use custom endpoints to satisfy the different workloads.

Answer: D

Explanation:
https://aws.amazon.com/about-aws/whats-new/2018/11/amazon-aurora-simplifies-workload-management-with-c You can now create custom endpoints for Amazon Aurora databases. This allows you to distribute and load balance workloads across different sets of database instances in your Aurora cluster. For example, you may provision a set of Aurora Replicas to use an instance type with higher memory capacity in order to run an analytics workload. A custom endpoint can then help you route the analytics workload to these
appropriately-configured instances, while keeping other instances in your cluster isolated from this workload. As you add or remove instances from the custom endpoint to match your workload, the endpoint helps spread the load around.

NEW QUESTION 29
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 instanc
B. Copy the encrypted snapshot to Amazon S3. Restore the DB instance from the encrypted snapshot using Amazon S3.
C. Create a new RDS for MySQL DB instance with encryption enable
D. Restore the unencrypted snapshot to this DB instance.
E. Create a snapshot of the unencrypted DB instanc
F. Create an encrypted copy of the snapsho
G. Restore the DB instance from the encrypted snapshot.
H. Temporarily shut down the unencrypted DB instanc
I. Enable AWS KMS encryption in the AWS Management Console using an AWS managed CM
J. Restart the DB instance in an encrypted state.

Answer: C

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

NEW QUESTION 32
A company is using a Single-AZ Amazon RDS for MySQL DB instance for development. The DB instance is experiencing slow performance when queries are executed. Amazon CloudWatch metrics indicate that the instance requires more I/O capacity.
Which actions can a database specialist perform to resolve this issue? (Choose two.)

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

E. Convert from General Purpose to Provisioned IOPS (PIOPS).

Answer: BE

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

NEW QUESTION 35
A significant automotive manufacturer is switching a mission-critical finance application's database to Amazon DynamoDB. According to the company's risk and compliance policy, any update to the database must be documented as a log entry for auditing purposes. Each minute, the system anticipates about 500,000 log entries. Log entries should be kept in Apache Parquet files in batches of at least 100,000 records per file.
How could a database professional approach these needs while using DynamoDB?

A. Enable Amazon DynamoDB Streams on the tabl
B. Create an AWS Lambda function triggered by the strea
C. Write the log entries to an Amazon S3 object.
D. Create a backup plan in AWS Backup to back up the DynamoDB table once a da
E. Create an AWS Lambda function that restores the backup in another table and compares both tables for change
F. Generate the log entries and write them to an Amazon S3 object.
G. Enable AWS CloudTrail logs on the tabl
H. Create an AWS Lambda function that reads the log files once an hour and filters DynamoDB API action
I. Write the filtered log files to Amazon S3.
J. Enable Amazon DynamoDB Streams on the tabl
K. Create an AWS Lambda function triggered by the strea
L. Write the log entries to an Amazon Kinesis Data Firehose delivery stream with buffering and Amazon S3 as the destination.

Answer: D

NEW QUESTION 36
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

Answer: D

Explanation:
https://aws.amazon.com/blogs/database/amazon-dynamodb-gaming-use-cases-and-design-patterns/ "EA uses the user ID as the partition key and primary key (a 1:1 modeling pattern)."
https://aws.amazon.com/blogs/database/choosing-the-right-dynamodb-partition-key/
"Partition key and sort key: Referred to as a composite primary key, this type of key is composed of two attributes. The first attribute is the partition key, and the second attribute is the sort key."

NEW QUESTION 40
A retail company is about to migrate its online and mobile store to AWS. The company's CEO has strategic plans to grow the brand globally. A Database Specialist has been challenged to provide predictable read and write database performance with minimal operational overhead.
What should the Database Specialist do to meet these requirements?

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

Answer: A

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

NEW QUESTION 44
A company is running an Amazon RDS for PostgeSQL 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 cluste
B. Set up replication from the source RDS for PostgreSQL DB instance using AWS DMS to the target DB cluster.
C. Use the pg_dump and pg_restore utilities to extract and restore the RDS for PostgreSQL DB instance to the Aurora PostgreSQL DB cluster.
D. Create a database snapshot of the RDS for PostgreSQL DB instance and use this snapshot to create the Aurora PostgreSQL DB cluster.
E. Migrate data from the RDS for PostgreSQL DB instance to an Aurora PostgreSQL DB cluster using anAurora Replic
F. Promote the replica during the cutover.

Answer: D

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

NEW QUESTION 46
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 Regio
C. Enable point-in-time recovery for both tables.
D. Use a DynamoDB Accelerator table in another Regio
E. Enable point-in-time recovery for the table.
F. Create an AWS Backup plan and assign the DynamoDB table as a resource.

Answer: C

NEW QUESTION 47
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 cluste
B. Verify the datatype of the columns.
C. Use the table metrics of the AWS DMS task created for migrating the data to verify the statistics for the tables being migrated and to verify that the data definition language (DDL) statements are completed.
D. Enable the AWS Schema Conversion Tool (AWS SCT) premigration validation and review the premigration checklist to make sure there are no issues with the conversion.
E. Enable AWS DMS data validation on the task so the AWS DMS task compares the source and target records, and reports any mismatches.

Answer: D

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

NEW QUESTION 50
An worldwide gaming company's development team is experimenting with using Amazon DynamoDB to store in-game events for three mobile titles. Maximum concurrent users for the most popular game is 500,000, while the least popular game is 10,000. The typical event is 20 KB in size, while the average user session generates one event each second. Each event is assigned a millisecond time stamp and a globally unique identification.
The lead developer generated a single DynamoDB database with the following structure for the events:
 Partition key: game name
 Sort key: event identifier
 Local secondary index: player identifier
 Event time
In a small-scale development setting, the tests were successful. When the application was deployed to production, however, new events were not being added to the database, and the logs indicated DynamoDB failures with the ItemCollectionSizeLimitExceededException issue code.
Which design modification should a database professional offer to the development team?

A. Use the player identifier as the partition ke
B. Use the event time as the sort ke
C. Add a global secondary index with the game name as the partition key and the event time as the sort key.
D. Create two table
E. Use the game name as the partition key in both table
F. Use the event time as the sort key for the first tabl
G. Use the player identifier as the sort key for the second table.
H. Replace the sort key with a compound value consisting of the player identifier collated with the event time, separated by a das
I. Add a local secondary index with the player identifier as the sort key.
J. Create one table for each gam
K. Use the player identifier as the partition ke
L. Use the event time as the sort key.

Answer: D

NEW QUESTION 51
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.

Answer: CDE

NEW QUESTION 54
An Amazon RDS EBS-optimized instance with Provisioned IOPS (PIOPS) storage is using less than half of its allocated IOPS over the course of several hours under constant load. The RDS instance exhibits multi-second read and write latency, and uses all of its maximum bandwidth for read throughput, yet the instance uses less than half of its CPU and RAM resources.
What should a Database Specialist do in this situation to increase performance and return latency to sub- second levels?

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

Answer: D

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

NEW QUESTION 59
After restoring an Amazon RDS snapshot from 3 days ago, a company's Development team cannot connect to the restored RDS DB instance. What is the likely cause of this problem?

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

Answer: C

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

NEW QUESTION 60
A software development company is using Amazon Aurora MySQL DB clusters for several use cases, including development and reporting. These use cases place unpredictable and varying demands on the Aurora DB clusters, and can cause momentary spikes in latency. System users run ad-hoc queries sporadically throughout the week. Cost is a primary concern for the company, and a solution that does not require significant rework is needed.
Which solution meets these requirements?

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

Answer: A

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

NEW QUESTION 61
Amazon RDS for Oracle with Transparent Data Encryption is used by a financial services organization (TDE). At all times, the organization is obligated to encrypt its data at rest. The decryption key must be widely distributed, and access to the key must be restricted. The organization must be able to rotate the encryption key on demand to comply with regulatory requirements. If any possible security vulnerabilities are discovered, the organization must be able to disable the key. Additionally, the company's overhead must be kept to a minimal.
What method should the database administrator use to configure the encryption to fulfill these specifications?

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

Answer: D

Explanation:
https://docs.aws.amazon.com/whitepapers/latest/kms-best-practices/aws-managed-and-customer-managed-cmks

NEW QUESTION 66
An electric utility company wants to store power plant sensor data in an Amazon DynamoDB table. The utility company has over 100 power plants and each power plant has over 200 sensors that send data every 2 seconds. The sensor data includes time with milliseconds precision, a value, and a fault attribute if the sensor is malfunctioning. Power plants are identified by a globally unique identifier. Sensors are identified by a unique identifier within each power plant. A database specialist needs to design the table to support an efficient method of finding all faulty sensors within a given power plant.
Which schema should the database specialist use when creating the DynamoDB table to achieve the fastest query time when looking for faulty sensors?

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

Answer: D

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

NEW QUESTION 67
A company uses Amazon DynamoDB as the data store for its ecommerce website. The website receives little to no traffic at night, and the majority of the traffic occurs during the day. The traffic growth during peak hours is gradual and predictable on a daily basis, but it can be orders of magnitude higher than during off-peak hours.
The company initially provisioned capacity based on its average volume during the day without accounting for the variability in traffic patterns. However, the website is experiencing a significant amount of throttling during peak hours. The company wants to reduce the amount of throttling while minimizing costs.
What should a database specialist do to meet these requirements?

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

Answer: C

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

NEW QUESTION 70
A business uses Amazon EC2 instances in VPC A to serve an internal file-sharing application. This application is supported by an Amazon ElastiCache cluster in VPC B that is peering with VPC A. The corporation migrates the instances of its applications from VPC A to VPC B. The file-sharing application is no longer able to connect to the ElastiCache cluster, as shown by the logs.
What is the best course of action for a database professional to take in order to remedy this issue?

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

Answer: D

Explanation:
https://docs.aws.amazon.com/vpc/latest/peering/vpc-peering-security-groups.html

NEW QUESTION 74
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

Answer: C

Explanation:
System snapshot can't fulfill 6 hours requirement. You need to control it by script
https://aws.amazon.com/blogs/database/%C2%AD%C2%AD%C2%ADautomating-cross-region-cross-account

NEW QUESTION 77
A ride-hailing application stores bookings in a persistent Amazon RDS for MySQL DB instance. This program is very popular, and the corporation anticipates a tenfold rise in the application's user base over the next several months. The application receives a higher volume of traffic in the morning and evening.
This application is divided into two sections:

An internal booking component that takes online reservations in response to concurrent user queries. A component of a third-party customer relationship management (CRM) system that customer service
professionals utilize. Booking data is accessed using queries in the CRM.
To manage this workload effectively, a database professional must create a cost-effective database system. Which solution satisfies these criteria?

A. Use Amazon ElastiCache for Redis to accept the booking
B. Associate an AWS Lambda function to capture changes and push the booking data to the RDS for MySQL DB instance used by the CRM.
C. Use Amazon DynamoDB to accept the booking
D. Enable DynamoDB Streams and associate an AWS Lambda function to capture changes and push the booking data to an Amazon SQS queu
E. This triggers another Lambda function that pulls data from Amazon SQS and writes it to the RDS for MySQL DB instance used by the CRM.
F. Use Amazon ElastiCache for Redis to accept the booking
G. Associate an AWS Lambda function to capture changes and push the booking data to an Amazon Redshift database used by the CRM.
H. Use Amazon DynamoDB to accept the booking
I. Enable DynamoDB Streams and associate an AWS Lambda function to capture changes and push the booking data to Amazon Athena, which is used by the CRM.

Answer: B

Explanation:
"AWS Lambda function to capture changes" capture changes to what? ElastiCache? The main use of ElastiCache is to cache frequently read data. Also "the company expects a tenfold increase in the user base" and "correspond to simultaneous requests from users"

NEW QUESTION 80
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 reportin
B. Monitor and set up replication between the two clusters to keep data consistent.
C. Review all evaluate the peak combined workloa
D. Ensure that utilization of the DB cluster node is at an acceptable leve
E. Adjust the number of instances, if necessary.
F. Use the stop cluster functionality to stop all the nodes of the DB cluster during times of minimal workloa
G. The cluster can be restarted again depending on the workload at the time.
H. Set up automatic scaling on the DB cluste
I. This will allow the number of reader nodes to adjust automatically to the reporting workload, when needed.

Answer: D

NEW QUESTION 82
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 cluste
B. Launch an Amazon EC2 instance with a weekly cron job to execute the custom script.
C. Configure an AWS Lambda function that exports archival data from the DB cluster to Amazon S3 using a SELECT INTO OUTFILE S3 statement, then deletes the archival data from the DB cluste
D. Schedule the Lambda function to run weekly using Amazon EventBridge (Amazon CloudWatch Events).
E. Configure two AWS Lambda functions: one that exports archival data from the DB cluster to Amazon S3 using the mysqldump utility, and another that deletes the archival data from the DB cluste
F. Schedule both Lambda functions to run weekly using Amazon EventBridge (Amazon CloudWatch Events).
G. Use AWS Database Migration Service (AWS DMS) to continually export the archival data from the DB cluster to Amazon S3. Configure an AWS Data Pipeline process to run weekly that executes a custom SQL script to delete the archival data from the DB cluster.

Answer: B

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

NEW QUESTION 87
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 cluste
B. Identify the gaps and compatibility of the objects migrated by comparing row counts from source and target tables.
C. Run AWS DMS from the Db2 database to an Aurora DB cluste
D. Identify the gaps and compatibility of the objects migrated by comparing the row counts from source and target tables.
E. Run native PostgreSQL logical replication from the Db2 database to an Aurora DB cluster to evaluate the migration compatibility.
F. 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.

Answer: D

NEW QUESTION 88

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

Answer: C

Explanation:

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

NEW QUESTION 90

A Database Specialist is creating a new Amazon Neptune DB cluster, and is attempting to load fata 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

Answer: BD

NEW QUESTION 93

A company has an application that uses an Amazon DynamoDB table to store user data. Every morning, a single-threaded process calls the DynamoDB API Scan operation to scan the entire table and generate a critical start-of-day report for management. A successful marketing campaign recently doubled the number of items in the table, and now the process takes too long to run and the report is not generated in time.
A database specialist needs to improve the performance of the process. The database specialist notes that, when the process is running, 15% of the table's provisioned read capacity units (RCUs) are being used.
What should the database specialist do?

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

Answer: B

Explanation:

https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Scan.html#Scan.ParallelScan

NEW QUESTION 96

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

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

Answer: C

Explanation:

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

NEW QUESTION 101

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

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

Answer: C

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

NEW QUESTION 106
A database specialist deployed an Amazon RDS DB instance in Dev-VPC1 used by their development team. Dev-VPC1 has a peering connection with Dev-VPC2 that belongs to a different development team in the same department. The networking team confirmed that the routing between VPCs is correct; however, the database engineers in Dev-VPC2 are getting a timeout connections error when trying to connect to the database in Dev- VPC1.
What is likely causing the timeouts?

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

Answer: C

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

NEW QUESTION 108
A company runs a customer relationship management (CRM) system that is hosted on-premises with a MySQL database as the backend. A custom stored procedure is used to send email notifications to another system when data is inserted into a table. The company has noticed that the performance of the CRM system has decreased due to database reporting applications used by various teams. The company requires an AWS solution that would reduce maintenance, improve performance, and accommodate the email notification feature.
Which AWS solution meets these requirements?

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

Answer: D

Explanation:
RDS event subscriptions do not cover "data is inserted into a table" - see
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/USER_Events.Messages.html We can use stored procedure to invoke Lambda function -
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Integrating.Lambda.html

NEW QUESTION 110
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

Answer: C

NEW QUESTION 114
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.

Answer: D

Explanation:
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/vpc-endpoints-dynamodb.html

NEW QUESTION 116
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

Answer: BC

Explanation:
The following items need to be reconfigured after restoring the DynamoDB table.
--AutoScaling policy
--IAM policy
--CloudWatch settings
--Tags
--Stream settings
--TTL
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/backuprestore_HowItWorks.html

NEW QUESTION 120
A business's mission-critical production workload is being operated on a 500 GB Amazon Aurora MySQL DB cluster. A database engineer must migrate the workload without causing data loss to a new Amazon Aurora Serverless MySQL DB cluster.
Which approach will result in the LEAST amount of downtime and the LEAST amount of application impact?

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

Answer: D

Explanation:
https://medium.com/@souri29/how-to-migrate-from-amazon-rds-aurora-or-mysql-to-amazon-aurora-serverless

NEW QUESTION 123
A Database Specialist is planning to create a read replica of an existing Amazon RDS for MySQL Multi-AZ DB instance. When using the AWS Management Console to conduct this task, the Database Specialist discovers that the source RDS DB instance does not appear in the read replica source selection box, so the read replica cannot be created.
What is the most likely reason for this?

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

Answer: D

Explanation:
>Your source DB instance must have backup retention enabled.
https://docs.aws.amazon.com/AmazonRDS/latest/APIReference/API_CreateDBInstanceReadReplica.html

NEW QUESTION 125
A large company is using an Amazon RDS for Oracle Multi-AZ DB instance with a Java application. As a part of its disaster recovery annual testing, the company would like to simulate an Availability Zone failure and record how the application reacts during the DB instance failover activity. The company does not want to make any code changes for this activity.
What should the company do to achieve this in the shortest amount of time?

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

Answer: B

Explanation:
 https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_RebootInstance.html https://exain.wordpress.com/2017/07/12/amazon-rds-multi-az-setup-failover-simulation/
"Rebooting with failover is beneficial when you want to simulate a failure of a DB instance for testing, or restore operations to the original AZ after a failover occurs."

NEW QUESTION 129
A business just transitioned from an on-premises Oracle database to Amazon Aurora PostgreSQL. Following the move, the organization observed that every day around 3:00 PM, the application's response time is substantially slower. The firm has determined that the problem is with the database, not the application.
Which set of procedures should the Database Specialist do to locate the erroneous PostgreSQL query most efficiently?

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

Answer: D

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

NEW QUESTION 134
In one AWS account, a business runs a two-tier ecommerce application. An Amazon RDS for MySQL
Multi-AZ database instance serves as the application's backend. A developer removed the database instance in the production environment by accident. Although the organization recovers the database, the incident results in hours of outage and financial loss.
Which combination of adjustments would reduce the likelihood that this error will occur again in the future? (Select 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.

Answer: ACD

NEW QUESTION 138
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

Answer: A

NEW QUESTION 139
A company is running a website on Amazon EC2 instances deployed in multiple Availability Zones (AZs). The site performs a high number of repetitive reads and writes each second on an Amazon RDS for MySQL Multi- AZ DB instance with General Purpose SSD (gp2) storage. After comprehensive testing and analysis, a database specialist discovers that there is high read latency and high CPU utilization on the DB instance.
Which approach should the database specialist to take to resolve this issue without changing the application?

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

Answer: D

NEW QUESTION 140
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

Answer: ABC

Explanation:
Connect - logins / DCL - authorizations (grant,revoke), DDL - schema updates

NEW QUESTION 143
......

Thank You for Trying Our Product

* 100% Pass or Money Back

   All our products come with a 90-day Money Back Guarantee.

* One year free update

   You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

   We currently serve more than 30,000,000 customers.

* Shop Securely

   All transactions are protected by VeriSign!

100% Pass Your AWS-Certified-Database-Specialty Exam with Our Prep Materials Via below:

https://www.certleader.com/AWS-Certified-Database-Specialty-dumps.html