Amazon-Web-Services

Exam Questions DBS-C01

AWS Certified Database - Specialty

About Exambible

Your Partner of IT Exam

Found in 1998

Exambible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, Exambible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1
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 2
An ecommerce company uses Amazon DynamoDB as the backend for its payments system. A new regulation requires the company to log all data access requests for financial audits. For this purpose, the company plans to use AWS logging and save logs to Amazon S3
How can a database specialist activate logging on the database?

A. Use AWS CloudTrail to monitor DynamoDB control-plane operation
B. Create a DynamoDB stream to monitor data-plane operation
C. Pass the stream to Amazon Kinesis Data Stream
D. Use that stream as a source for Amazon Kinesis Data Firehose to store the data in an Amazon S3 bucket.
E. Use AWS CloudTrail to monitor DynamoDB data-plane operation
F. Create a DynamoDB stream to monitor control-plane operation
G. Pass the stream to Amazon Kinesis Data Stream
H. Use that stream as a source for Amazon Kinesis Data Firehose to store the data in an Amazon S3 bucket.
I. Create two trails in AWS CloudTrai
J. Use Trail1 to monitor DynamoDB control-plane operation
K. Use Trail2 to monitor DynamoDB data-plane operations.
L. Use AWS CloudTrail to monitor DynamoDB data-plane and control-plane operations.

Answer: D

Explanation:
https://aws.amazon.com/about-aws/whats-new/2021/04/you-now-can-use-aws-cloudtrail-to-log-amazon-dynamo

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

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

Answer: D

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

NEW QUESTION 4
A single MySQL database was moved to Amazon Aurora by a business. The production data is stored in a database cluster in VPC PROD, whereas 12 testing environments are hosted in VPC TEST with the same AWS account. Testing has a negligible effect on the test data. The development team requires that each environment be updated nightly to ensure that each test database has daily production data.
Which migration strategy will be the quickest and least expensive to implement?

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

Answer: A

Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Managing.Clone.html

NEW QUESTION 5
A 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 loggin
B. Direct the logging to a specified log group in Amazon CloudWatch Logs.
C. Edit the database configuration of the cluster by enabling audit loggin
D. Direct the logging to a specified Amazon S3 bucket
E. Modify the cluster by enabling continuous delivery of AWS CloudTrail logs to Amazon S3.
F. Create a new parameter group with the enable_user_activity_logging parameter set to tru
G. Configure the cluster to use the new parameter group.
H. Modify the system table to enable logging for each user.

Answer: AD

Explanation:
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

NEW QUESTION 6
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 7
A company has a web-based survey application that uses Amazon DynamoDB. During peak usage, when survey responses are being collected, a Database Specialist sees the ProvisionedThroughputExceededException error.
What can the Database Specialist do to resolve this error? (Choose two.)

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

Answer: CD

Explanation:
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/switching.capacitymode.html

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

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

Answer: B

Explanation:
Aurora Serverless is not compatible to all Aurora provisioned engine version. However, you can do clone with most engine version. Meanwhile, I also consider the performance while restoring snapshot to Aurora Serverless.
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-serverless.how-it-works.html#aurora
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-serverless.html#aurora-serverless.us

NEW QUESTION 9
AWS CloudFormation stack including an Amazon RDS database instance was mistakenly removed, resulting in the loss of recent data. A Database Specialist must apply RDS parameters to the CloudFormation template in order to minimize the possibility of future inadvertent instance data loss.
Which settings will satisfy this criterion? (Select three.)

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

Answer: ADF

Explanation:
A https://aws.amazon.com/about-aws/whats-new/2018/09/amazon-rds-now-provides-database-deletion-protection/
D https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_WorkingWithAutomatedBackups.html
F - https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute-deletionpolicy.html

NEW QUESTION 10
A database specialist has been entrusted by an ecommerce firm with designing a reporting dashboard that visualizes crucial business KPIs derived from the company's primary production database running on Amazon Aurora. The dashboard should be able to read data within 100 milliseconds after an update.
The Database Specialist must conduct an audit of the Aurora DB cluster's present setup and provide a
cost-effective alternative. The solution must support the unexpected read demand generated by the reporting dashboard without impairing the DB cluster's write availability and performance.
Which solution satisfies these criteria?

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

Answer: D

NEW QUESTION 10
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 15
Recently, an ecommerce business transferred one of its SQL Server databases to an Amazon RDS for SQL Server Enterprise Edition database instance. The corporation anticipates an increase in read traffic as a result of an approaching sale. To accommodate the projected read load, a database professional must establish a read replica of the database instance.
Which procedures should the database professional do prior to establishing the read replica? (Select two.)

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

Answer: BC

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

NEW QUESTION 18
A business is transferring a database from one AWS Region to another using an Amazon RDS for SQL Server DB instance. The organization wishes to keep database downtime to a minimum throughout the transfer.
Which migration strategy should the organization use for this cross-regional move?

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

Answer: C

Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.XRgn.html
With Amazon RDS, you can create a MariaDB, MySQL, Oracle, or PostgreSQL read replica in a different AWS Region from the source DB instance. Creating a cross-Region read replica isn't supported for SQL Server on Amazon RDS.

NEW QUESTION 23
A corporation is transitioning from an IBM Informix database to an Amazon RDS for SQL Server Multi-AZ implementation with Always On Availability Groups (AGs). SQL Server Agent tasks are scheduled to execute at 5-minute intervals on the Always On AG listener to synchronize data between the Informix and SQL Server databases. After a successful failover to the backup node with minimum delay, users endure hours of stale data.
How can a database professional guarantee that consumers view the most current data after a failover?

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

Answer: D

Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_SQLServerMultiAZ.html
If you have SQL Server Agent jobs, recreate them on the secondary. You do so because these jobs are stored in the msdb database, and you can't replicate this database by using Database Mirroring (DBM) or Always On Availability Groups (AGs). Create the jobs first in the original primary, then fail over, and create the same jobs in the new primary.

NEW QUESTION 28
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 29
A company uses Microsoft SQL Server on Amazon RDS in a Multi-AZ deployment as the database engine for its application. The company was recently acquired by another company. A database specialist must rename the database to follow a new naming standard.
Which combination of steps should the database specialist take to rename the database? (Choose two.)

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

Answer: BD

Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Appendix.SQLServer.CommonDBATasks.Renami

NEW QUESTION 30
A database professional is tasked with the task of migrating 25 GB of data files from an on-premises storage system to an Amazon Neptune database.
Which method of data loading is the FASTEST?

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

Answer: A

Explanation:
* 1. Copy the data files to an Amazon Simple Storage Service (Amazon S3) bucket.
* 2. Create an IAM role with Read and List access to the bucket.
* 3. Create an Amazon S3 VPC endpoint.
* 4. Start the Neptune loader by sending a request via HTTP to the Neptune DB instance.
* 5. The Neptune DB instance assumes the IAM role to load the data from the bucket.

NEW QUESTION 34
A large company has a variety of Amazon DB clusters. Each of these clusters has various configurations that adhere to various requirements. Depending on the team and use case, these configurations can be organized into broader categories.
A database administrator wants to make the process of storing and modifying these parameters more systematic. The database administrator also wants to ensure that changes to individual categories of configurations are automatically applied to all instances when required.
Which AWS service or feature will help automate and achieve this objective?

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

Answer: B

NEW QUESTION 38
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 41
A database specialist is designing an enterprise application for a large company. The application uses Amazon DynamoDB with DynamoDB Accelerator (DAX). The database specialist observes that most of the queries are not found in the DAX cache and that they still require DynamoDB table reads.
What should the database specialist review first to improve the utility of DAX?

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

Answer: C

Explanation:
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/DAX.cluster-management.html#DAX.clu

NEW QUESTION 46
A bike rental company operates an application to track its bikes. The application receives location and condition data from bike sensors. The application also receives rental transaction data from the associated mobile app.
The application uses Amazon DynamoDB as its database layer. The company has configured DynamoDB with provisioned capacity set to 20% above the expected peak load of the application. On an average day, DynamoDB used 22 billion read capacity units (RCUs) and 60 billion write capacity units (WCUs). The application is running well. Usage changes smoothly over the course of the day and is generally shaped like a bell curve. The timing and magnitude of peaks vary based on the weather and season, but the general shape is consistent.
Which solution will provide the MOST cost optimization of the DynamoDB database layer?

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

Answer: C

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

NEW QUESTION 48
A database specialist was alerted that a production Amazon RDS MariaDB instance with 100 GB of storage was out of space. In response, the database specialist modified the DB instance and added 50 GB of storage capacity. Three hours later, a new alert is generated due to a lack of free space on the same DB instance. The database specialist decides to modify the instance immediately to increase its storage capacity by 20 GB.
What will happen when the modification is submitted?

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

Answer: D

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

NEW QUESTION 51
An online gaming company is using an Amazon DynamoDB table in on-demand mode to store game scores. After an intensive advertisement campaign in South America, the average number of concurrent users rapidly increases from 100,000 to 500,000 in less than 10 minutes every day around 5 PM.
The on-call software reliability engineer has observed that the application logs contain a high number of DynamoDB throttling exceptions caused by game score insertions around 5 PM. Customer service has also reported that several users are complaining about their scores not being registered.
How should the database administrator remediate this issue at the lowest cost?

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

Answer: B

NEW QUESTION 54
A company's Security department established new requirements that state internal users must connect to an existing Amazon RDS for SQL Server DB instance using their corporate Active Directory (AD) credentials. A Database Specialist must make the modifications needed to fulfill this requirement.
Which combination of actions should the Database Specialist take? (Choose three.)

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

Answer: BCF

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

NEW QUESTION 55
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 59
A company is deploying a solution in Amazon Aurora by migrating from an on-premises system. The IT department has established an AWS Direct Connect link from the company's data center. The company's Database Specialist has selected the option to require SSL/TLS for connectivity to prevent plaintext data from being set over the network. The migration appears to be working successfully, and the data can be queried from a desktop machine.
Two Data Analysts have been asked to query and validate the data in the new Aurora DB cluster. Both Analysts are unable to connect to Aurora. Their user names and passwords have been verified as valid and the Database Specialist can connect to the DB cluster using their accounts. The Database Specialist also verified that the security group configuration allows network from all corporate IP addresses.

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

A. Restart the DB cluster to apply the SSL change.
B. Instruct the Data Analysts to download the root certificate and use the SSL certificate on the connection string to connect.
C. Add explicit mappings between the Data Analysts' IP addresses and the instance in the security group assigned to the DB cluster.
D. Modify the Data Analysts' local client firewall to allow network traffic to AWS.

Answer: B

Explanation:
• To connect using SSL:
• Provide the SSLTrust certificate (can be downloaded from AWS)
• Provide SSL options when connecting to database
• Not using SSL on a DB that enforces SSL would result in error https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/ssl-certificate-rotation-aurora-postgresql.ht

NEW QUESTION 64
A business's production database is hosted on a single-node Amazon RDS for MySQL DB instance. The database instance is hosted in a United States AWS Region.
A week before a significant sales event, a fresh database maintenance update is released. The maintenance update has been designated as necessary. The firm want to minimize the database instance's downtime and requests that a database expert make the database instance highly accessible until the sales event concludes.
Which solution will satisfy these criteria?

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

Answer: D

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/rds-required-maintenance/

NEW QUESTION 68
A company's ecommerce website uses Amazon DynamoDB for purchase orders. Each order is made up of a Customer ID and an Order ID. The DynamoDB table uses the Customer ID as the partition key and the Order ID as the sort key.
To meet a new requirement, the company also wants the ability to query the table by using a third attribute named Invoice ID. Queries using the Invoice ID must be strongly consistent. A database specialist must provide this capability with optimal performance and minimal overhead.
What should the database administrator do to meet these requirements?

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

Answer: C

Explanation:
as Local secondary index can only be created while creating the Dynamodb table. and query needs to use third attribute on top of primary and sort key, so Local Secondary index has primary and sort key as well as the third attribute. Global secondary index can be created without primary and sort key

NEW QUESTION 72
A company developed an AWS CloudFormation template used to create all new Amazon DynamoDB tables in its AWS account. The template configures provisioned throughput capacity using hard-coded values. The company wants to change the template so that the tables it creates in the future have independently configurable read and write capacity units assigned.
Which solution will enable this change?

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

Answer: B

Explanation:
Input parameter and FindInMap You can use an input parameter with the Fn::FindInMap function to refer to a specific value in a map. For example, suppose you have a list of regions and environment types that map to a specific AMI ID. You can select the AMI ID that your stack uses by using an input parameter (EnvironmentType). To determine the region, use the AWS::Region pseudo parameter, which gets the AWS Region in which you create the stack.
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/parameters-section-structure.html

NEW QUESTION 75

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 77
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 80
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 83
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 84
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 87
A financial services company runs an on-premises MySQL database for a critical application. The company is dissatisfied with its current database disaster recovery (DR) solution. The application experiences a significant amount of downtime whenever the database fails over to its DR facility. The application also experiences slower response times when reports are processed on the same database. To minimize the downtime in DR situations, the company has decided to migrate the database to AWS. The company requires a solution that is highly available and the most cost-effective.
Which solution meets these requirements?

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

Answer: C

Explanation:
https://aws.amazon.com/about-aws/whats-new/2016/09/reader-end-point-for-amazon-aurora/

NEW QUESTION 89
A company is about to launch a new product, and test databases must be re-created from production data. The company runs its production databases on an Amazon Aurora MySQL DB cluster. A Database Specialist needs to deploy a solution to create these test databases as quickly as possible with the least amount of administrative effort.
What should the Database Specialist do to meet these requirements?

A. Restore a snapshot from the production cluster into test clusters
B. Create logical dumps of the production cluster and restore them into new test clusters
C. Use database cloning to create clones of the production cluster
D. Add an additional read replica to the production cluster and use that node for testing

Answer: C

Explanation:
https://aws.amazon.com/getting-started/hands-on/aurora-cloning-backtracking/
"Cloning an Aurora cluster is extremely useful if you want to assess the impact of changes to your database, or if you need to perform workload-intensive operations—such as exporting data or running analytical queries, or simply if you want to use a copy of your production database in a development or testing environment. You can make multiple clones of your Aurora DB cluster. You can even create additional clones from other clones, with the constraint that the clone databases must be created in the same region as the source databases.

NEW QUESTION 92
Application developers have reported that an application is running slower as more users are added. The application database is running on an Amazon Aurora DB cluster with an Aurora Replica. The application is written to take advantage of read scaling through reader endpoints. A database specialist looks at the performance metrics of the database and determines that, as new users were added to the database, the primary instance CPU utilization steadily increased while the Aurora Replica CPU utilization remained steady.
How can the database specialist improve database performance while ensuring minimal downtime?

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

Answer: C

NEW QUESTION 97
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 DBinstance 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 98
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 103
Amazon DynamoDB global tables are being used by a business to power an online gaming game. The game is played by gamers from all around the globe. As the game became popularity, the amount of queries to DynamoDB substantially rose. Recently, gamers have complained about the game's condition being inconsistent between nations. A database professional notices that the ReplicationLatency metric for many replica tables is set to an abnormally high value.
Which strategy will resolve the issue?

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

Answer: A

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

NEW QUESTION 105
A company is going to use an Amazon Aurora PostgreSQL DB cluster for an application backend. The DB cluster contains some tables with sensitive data. A Database Specialist needs to control the access privileges at the table level.
How can the Database Specialist meet these requirements?

A. Use AWS IAM database authentication and restrict access to the tables using an IAM policy.
B. Configure the rules in a NACL to restrict outbound traffic from the Aurora DB cluster.
C. Execute GRANT and REVOKE commands that restrict access to the tables containing sensitive data.
D. Define access privileges to the tables containing sensitive data in the pg_hba.conf file.

Answer: C

NEW QUESTION 110
A database specialist wants to ensure that an Amazon Aurora DB cluster is always automatically upgraded to the most recent minor version available. Noticing that there is a new minor version available, the database specialist has issues an AWS CLI command to enable automatic minor version updates. The command runs successfully, but checking the Aurora DB cluster indicates that no update to the Aurora version has been made.
What might account for this? (Choose two.)

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

Answer: AD

Explanation:
"When Amazon RDS designates a minor engine version as the preferred minor engine version, each database that meets both of the following conditions is upgraded to the minor engine version automatically"
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_UpgradeDBInstance.Upgrading.html
Call 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.
https://docs.amazonaws.cn/en_us/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Updates.Patching.html#

NEW QUESTION 114
A database specialist needs to delete user data and sensor data 1 year after it was loaded in an Amazon DynamoDB table. TTL is enabled on one of the attributes. The database specialist monitors TTL rates on the Amazon CloudWatch metrics for the table and observes that items are not being deleted as

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

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

Answer: B

Explanation:
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/TTL.html#time-to-live-ttl-before-you-sta

NEW QUESTION 119
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 121
A software company is conducting a security audit of its three-node Amazon Aurora MySQL DB cluster. Which finding is a security concern that needs to be addressed?

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

Answer: D

Explanation:
Rotate your IAM credentials regularly.

NEW QUESTION 123
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 125
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 a Secret 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 128
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 133
A manufacturing company has an. inventory system that stores information in an Amazon Aurora MySQL DB cluster. The database tables are partitioned. The database size has grown to 3 TB. Users run one-time queries by using a SQL client. Queries that use an equijoin to join large tables are taking a long time to run.
Which action will improve query performance with the LEAST operational effort?

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

Answer: B

Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.BestPractices.html

NEW QUESTION 137
A 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 encryptio
B. When the read replica is available, promote the read replica and update the endpoint that is used by the applicatio
C. Delete the unencrypted DB instance.
D. Take a snapshot of the DB instanc
E. Make an encrypted copy of the snapsho
F. Restore the encrypted snapsho
G. When the new DB instance is available, update the endpoint that is used by the applicatio
H. Delete the unencrypted DB instance.
I. Create a new encrypted DB instanc
J. 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 applicatio
K. Delete the unencrypted DB instance.
L. Convert the DB instance to an Amazon Aurora DB cluster, and enable encryptio
M. When the DB cluster is available, update the endpoint that is used by the application to the cluster endpoin
N. Delete the unencrypted DB instance.

Answer: C

Explanation:
https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/encrypt-an-existing-amazon-rds-for-postgresq 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 the downtime. This pattern uses the AWS Database Migration Service (AWS DMS) to migrate and continuously replicate the data so that the cutover to the new, encrypted database can be done with minimal downtime.

NEW QUESTION 140

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

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

Answer: A

Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.AuroraMySQL.Monitoring.Metrics. Unlike for other DB engines, for Aurora DB instances this metric reports the amount of storage available to each DB instance. This value depends on the DB instance class (for pricing information, see the Amazon RDS product page). You can increase the amount of free storage space for an instance by choosing a larger DB instance class for your instance."

NEW QUESTION 145
A company is using an Amazon ElastiCache for Redis cluster to host its online shopping website. Shoppers receive the following error when the website's application queries the cluster:

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

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

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

Answer: ABC

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/oom-command-not-allowed-redis/

NEW QUESTION 148
A company with 500,000 employees needs to supply its employee list to an application used by human resources. Every 30 minutes, the data is exported using the LDAP service to load into a new Amazon DynamoDB table. The data model has a base table with Employee ID for the partition key and a global secondary index with Organization ID as the partition key.
While importing the data, a database specialist receives ProvisionedThroughputExceededException errors.
After increasing the provisioned write capacity units
(WCUs) to 50,000, the specialist receives the same errors. Amazon CloudWatch metrics show a consumption of 1,500 WCUs.
What should the database specialist do to address the issue?

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

Answer: A

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/dynamodb-table-throttled/

NEW QUESTION 152
A business that specializes in internet advertising is developing an application that will show adverts to its customers. The program stores data in an Amazon DynamoDB database. Additionally, the application caches its reads using a DynamoDB Accelerator (DAX) cluster. The majority of reads come via the GetItem and BatchGetItem queries. The application does not need consistency of readings.
The application cache does not behave as intended after deployment. Specific extremely consistent queries to the DAX cluster are responding in several milliseconds rather than microseconds.
How can the business optimize cache behavior in order to boost application performance?

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

Answer: C

NEW QUESTION 157
A ride-hailing application uses an Amazon RDS for MySQL DB instance as persistent storage for bookings. This application is very popular and the company expects a tenfold increase in the user base in next few months. The application experiences more traffic during the morning and evening hours.
This application has two parts:

- An in-house booking component that accepts online bookings that directly correspond to simultaneous requests from users.
- A third-party customer relationship management (CRM) component used by customer care representatives. The CRM uses queries to access booking data.

A database specialist needs to design a cost-effective database solution to handle this workload. Which solution meets these requirements?

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

NEW QUESTION 161
A company is concerned about the cost of a large-scale, transactional application using Amazon DynamoDB that only needs to store data for 2 days before it is deleted. In looking at the tables, a Database Specialist notices that much of the data is months old, and goes back to when the application was first deployed.
What can the Database Specialist do to reduce the overall cost?

A. Create a new attribute in each table to track the expiration time and create an AWS Glue transformation to delete entries more than 2 days old.
B. Create a new attribute in each table to track the expiration time and enable DynamoDB Streams on each table.
C. Create a new attribute in each table to track the expiration time and enable time to live (TTL) on each table.
D. Create an Amazon CloudWatch Events event to export the data to Amazon S3 daily using AWS Data Pipeline and then truncate the Amazon DynamoDB table.

Answer: C

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

NEW QUESTION 164
A gaming company is designing a mobile gaming app that will be accessed by many users across the globe. The company wants to have replication and full support for multi-master writes. The company also wants to ensure low latency and consistent performance for app users.
Which solution meets these requirements?

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

Answer: A

NEW QUESTION 167
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 168
A database specialist manages a critical Amazon RDS for MySQL DB instance for a company. The data stored daily could vary from .01% to 10% of the current database size. The database specialist needs to ensure that the DB instance storage grows as needed.
What is the MOST operationally efficient and cost-effective solution?

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

Answer: A

Explanation:
If your workload is unpredictable, you can enable storage autoscaling for an Amazon RDS DB instance. With storage autoscaling enabled, when Amazon RDS detects that you are running out of free database space it automatically scales up your storage.
https://aws.amazon.com/about-aws/whats-new/2019/06/rds-storage-auto-scaling/
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_PIOPS.StorageTypes.html#USER_PIOPS.

NEW QUESTION 169

A 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 item
B. Use the Query API operation to read multiple items that have a specific partition key and sort ke
C. Use the GetItem API operation to read a single item.
D. Use the Scan API operation with a filter expression that allows multiple items to be rea
E. Use the Query API operation to read multiple items that have a specific partition key and sort ke
F. Use the GetItem API operation to read a single item.
G. Use the Scan API operation with a filter expression that allows multiple items to be rea
H. Use the Query API operation to read a single item that has a specific primary ke
I. Use the BatchGetItem API operation to read multiple items.
J. Use the Scan API operation in parallel with many workers to read all the item
K. 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.

Answer: B

Explanation:
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/SQLtoNoSQL.ReadData.html

NEW QUESTION 174
An ecommerce company is using Amazon DynamoDB as the backend for its order-processing application.
The steady increase in the number of orders is resulting in increased DynamoDB costs. Order verification and reporting perform many repeated GetItem functions that pull similar datasets, and this read activity is contributing to the increased costs. The company wants to control these costs without significant development efforts.
How should a Database Specialist address these requirements?

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

Answer: D

Explanation:
https://docs.amazonaws.cn/en_us/amazondynamodb/latest/developerguide/DAX.html
"Applications that are read-intensive, but are also cost-sensitive. With DynamoDB, you provision the number of reads per second that your application requires. If read activity increases, you can increase your tables' provisioned read throughput (at an additional cost). Or, you can offload the activity from your application to a DAX cluster, and reduce the number of read capacity units that you need to purchase otherwise."

NEW QUESTION 177
A company is planning to use Amazon RDS for SQL Server for one of its critical applications. The company's security team requires that the users of the RDS for SQL Server DB instance are authenticated with on-premises Microsoft Active Directory credentials. Which combination of steps should a database specialist take to meet this requirement? (Choose three.)

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

Answer: CEF

NEW QUESTION 178
A company has an on-premises system that tracks various database operations that occur over the lifetime of a database, including database shutdown, deletion, creation, and backup.
The company recently moved two databases to Amazon RDS and is looking at a solution that would satisfy these requirements. The data could be used by other systems within the company.
Which solution will meet these requirements with minimal effort?

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

Answer: C

NEW QUESTION 180
A company is using Amazon Aurora MySQL as the database for its retail application on AWS. The company receives a notification of a pending database upgrade and wants to ensure upgrades do not occur before or during the most critical time of year. Company leadership is concerned that an Amazon RDS maintenance window will cause an outage during data ingestion.
Which step can be taken to ensure that the application is not interrupted?

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

Answer: C

NEW QUESTION 182
A database specialist is working on an Amazon RDS for PostgreSQL DB instance that is experiencing application performance issues due to the addition of new workloads. The database has 5 ¢' of storage space with Provisioned IOPS. Amazon CloudWatch metrics show that the average disk queue depth is greater than 200 and that the disk I/O response time is significantly higher than usual.
What should the database specialist do to improve the performance of the application immediately?

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

Answer: A

NEW QUESTION 186
A Database Specialist is constructing a new Amazon Neptune DB cluster and tries to load data from Amazon S3 using the Neptune bulk loader API. The Database Specialist is confronted with the following error message:
€Unable to establish a connection to the s3 endpoint. The source URL is s3:/mybucket/graphdata/ and the region code is us-east-1. Kindly confirm your Configuration S3.
Which of the following activities should the Database Specialist take to resolve the issue? (Select two.)

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

Answer: BE

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

NEW QUESTION 187
A company is using an Amazon Aurora MySQL database with Performance Insights enabled. A database specialist is checking Performance Insights and observes an alert message that starts with the following phrase:
`Performance Insights is unable to collect SQL Digest statistics on new queries`¦`
Which action will resolve this alert message?

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

Answer: A

Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_PerfInsights.UsingDashboard.AnalyzeDBL

NEW QUESTION 191
A 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 PostgreSQL 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: B

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/dms-mapping-oracle-postgresql/

NEW QUESTION 196
A database professional maintains a fleet of Amazon RDS database instances that are configured to utilize the default database parameter group. A database expert must connect a custom parameter group with certain database instances.
When will the instances be allocated to this new parameter group once the database specialist performs this change?

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

Answer: C

Explanation:
When you associate a new DB parameter group with a DB instance, the modified static and dynamic parameters are applied only after the DB instance is rebooted.

NEW QUESTION 197
A company is developing a multi-tier web application hosted on AWS using Amazon Aurora as the database. The application needs to be deployed to production and other non-production environments. A Database Specialist needs to specify different MasterUsername and MasterUserPassword properties in the AWS CloudFormation templates used for automated deployment. The CloudFormation templates are version controlled in the company's code repository. The company also needs to meet compliance requirement by routinely rotating its database master password for production.
What is most secure solution to store the master password?

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

Answer: C

Explanation:
"By using the secure string support in CloudFormation with dynamic references you can better maintain your infrastructure as code. You'll be able to avoid hard coding passwords into your templates and you can keep these runtime configuration parameters separated from your code. Moreover, when properly used, secure strings will help keep your development and production code as similar as possible, while continuing to make your infrastructure code suitable for continuous deployment pipelines."
https://aws.amazon.com/blogs/mt/using-aws-systems-manager-parameter-store-secure-string-parameters-in-aws https://aws.amazon.com/blogs/security/how-to-use-aws-secrets-manager-rotate-credentials-amazon-rds-database

NEW QUESTION 202
A development team asks a database specialist to create a copy of a production Amazon RDS for MySQL DB instance every morning. The development team will use the copied DB instance as a testing environment for development. The original DB instance and the copy will be hosted in different VPCs of the same AWS account. The development team wants the copy to be available by 6 AM each day and wants to use the same endpoint address each day.
Which combination of steps should the database specialist take to meet these requirements MOST cost-effectively? (Choose three.)

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

Answer: ABC

NEW QUESTION 204
A database specialist must create nightly backups of an Amazon DynamoDB table in a mission-critical workload as part of a disaster recovery strategy.
Which backup methodology should the database specialist use to MINIMIZE management overhead?

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

Answer: C

Explanation:
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/CreateBackup.html#:~:text=If%20you%2
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/backuprestore_HowItWorks.html

NEW QUESTION 206
A gaming company uses Amazon Aurora Serverless for one of its internal applications. The company's developers use Amazon RDS Data API to work with the Aurora Serverless DB cluster. After a recent security review, the company is mandating security enhancements. A database specialist must ensure that access to RDS Data API is private and never passes through the public internet. What should the database specialist do to meet this requirement?

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

Answer: C

Explanation:
https://aws.amazon.com/about-aws/whats-new/2020/02/amazon-rds-data-api-now-supports-aws-privatelink/

NEW QUESTION 209
A company has an on-premises production Microsoft SQL Server with 250 GB of data in one database. A database specialist needs to migrate this on-premises SQL Server to Amazon RDS for SQL Server. The nightly native SQL Server backup file is approximately 120 GB in size. The application can be down for an extended period of time to complete the migration. Connectivity between the on-premises environment and AWS can be initiated from on-premises only.
How can the database be migrated from on-premises to Amazon RDS with the LEAST amount of effort?

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

Answer: B

Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/SQLServer.Procedural.Importing.html

NEW QUESTION 212
Recently, a gaming firm purchased a popular iOS game that is especially popular during the Christmas season. The business has opted to include a leaderboard into the game, which will be powered by Amazon DynamoDB. The application's load is likely to increase significantly throughout the Christmas season.
Which solution satisfies these criteria at the lowest possible cost?

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

Answer: D

Explanation:
"On-demand is ideal for bursty, new, or unpredictable workloads whose traffic can spike in seconds or minutes"
vs.
'DynamoDB released auto scaling to make it easier for you to manage capacity efficiently, and auto scaling continues to help DynamoDB users lower the cost of workloads that have a predictable traffic pattern."
https://aws.amazon.com/blogs/database/amazon-dynamodb-auto-scaling-performance-and-cost-optimization-at

NEW QUESTION 215
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 217
A company has a 20 TB production Amazon Aurora DB cluster. The company runs a large batch job overnight to load data into the Aurora DB cluster. To ensure the company’s development team has the most up-to-date data for testing, a copy of the DB cluster must be available in the shortest possible time after the batch job completes.
How should this be accomplished?

A. Use the AWS CLI to schedule a manual snapshot of the DB cluste
B. Restore the snapshot to a new DB cluster using the AWS CLI.
C. Create a dump file from the DB cluste
D. Load the dump file into a new DB cluster.

E. Schedule a job to create a clone of the DB cluster at the end of the overnight batch process.
F. Set up a new daily AWS DMS task that will use cloning and change data capture (CDC) on the DB cluster to copy the data to a new DB cluste
G. Set up a time for the AWS DMS stream to stop when the new cluster is current.

Answer: C

NEW QUESTION 220
A 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 replicas
B. Amazon DynamoDB global tables
C. Amazon Aurora global database
D. Amazon Athena and Amazon S3 with S3 Cross Region Replication

Answer: B

NEW QUESTION 224
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 229
A company is launching a new Amazon RDS for MySQL Multi-AZ DB instance to be used as a data store for a custom-built application. After a series of tests with point-in-time recovery disabled, the company decides that it must have point-in-time recovery reenabled before using the DB instance to store production data.
What should a database specialist do so that point-in-time recovery can be successful?

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

Answer: C

Explanation:
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-m "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

NEW QUESTION 232
A company is writing a new survey application to be used with a weekly televised game show. The application will be available for 2 hours each week. The company expects to receive over 500,000 entries every week, with each survey asking 2-3 multiple choice questions of each user. A Database Specialist needs to select a platform that is highly scalable for a large number of concurrent writes to handle he anticipated volume.
Which AWS services should the Database Specialist consider? (Choose two.)

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

Answer: AE

Explanation:
https://docs.aws.amazon.com/AmazonElastiCache/latest/mem-ug/Strategies.html#Strategies.WriteThrough https://aws.amazon.com/products/databases/real-time-apps-elasticache-for-redis/

NEW QUESTION 234
A stock market analysis firm maintains two locations: one in the us-east-1 Region and another in the eu-west-2 Region. The business want to build an AWS

database solution capable of providing rapid and accurate updates.
Dashboards with advanced analytical queries are used to present data in the eu-west-2 office. Because the corporation will use these dashboards to make purchasing choices, they must have less than a second to obtain application data.
Which solution satisfies these criteria and gives the MOST CURRENT dashboard?

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

Answer: C

Explanation:
Amazon Aurora global databases span multiple AWS Regions, enabling low latency global reads and providing fast recovery from the rare outage that might affect an entire AWS Region. An Aurora global database has a primary DB cluster in one Region, and up to five secondary DB clusters in different Regions.
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-global-database.html

NEW QUESTION 236
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 and use 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 240
A company is running Amazon RDS for MySQL for its workloads. There is downtime when AWS operating system patches are applied during the Amazon RDS-specified maintenance window.
What is the MOST cost-effective action that should be taken to avoid downtime?

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

Answer: D

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/rds-required-maintenance/
To minimize downtime, modify the Amazon RDS DB instance to a Multi-AZ deployment. For Multi-AZ deployments, OS maintenance is applied to the secondary instance first, then the instance fails over, and then the primary instance is updated. The downtime is during failover. For more information, see Maintenance for Multi-AZ Deployments. https://aws.amazon.com/rds/faqs/ The availability benefits of Multi-AZ also extend to planned maintenance. For example, with automated backups, I/O activity is no longer suspended on your primary during your preferred backup window, since backups are taken from the standby. In the case of patching or DB instance class scaling, these operations occur first on the standby, prior to automatic fail over. As a result, your availability impact is limited to the time required for automatic failover to complete.

NEW QUESTION 244
A 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: AD

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/rds-snapshots-share-account/

NEW QUESTION 246
A business's production databases are housed on a 3 TB Amazon Aurora MySQL DB cluster. The database cluster is installed in the region us-east-1. For disaster recovery (DR) requirements, the company's database expert needs to fast deploy the DB cluster in another AWS Region to handle the production load with an RTO of less than two hours.
Which approach is the MOST OPERATIONALLY EFFECTIVE in meeting these requirements?

A. Implement an AWS Lambda function to take a snapshot of the production DB cluster every 2 hours, and copy that snapshot to an Amazon S3 bucket in the DR Regio
B. Restore the snapshot to an appropriately sized DB cluster in the DR Region.
C. Add a cross-Region read replica in the DR Region with the same instance type as the current primary instanc
D. If the read replica in the DR Region needs to be used for production, promote the read replica to become a standalone DB cluster.
E. Create a smaller DB cluster in the DR Regio
F. Configure an AWS Database Migration Service (AWS DMS) task with change data capture (CDC) enabled to replicate data from the current production DB cluster to the DB cluster in the DR Region.
G. Create an Aurora global database that spans two Region
H. Use AWS Database Migration Service (AWS DMS) to migrate the existing database to the new global database.

Answer: B

Explanation:
RTO is 2 hours. With 3 TB database, cross-region replica is a better option

NEW QUESTION 251
A database specialist is managing an application in the us-west-1 Region and wants to set up disaster recovery in the us-east-1 Region. The Amazon Aurora MySQL DB cluster needs an RPO of 1 minute and an RTO of 2 minutes.
Which approach meets these requirements with no negative performance impact?

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

Answer: C

Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-global-database-disaster-recovery.ht

NEW QUESTION 255
A retail company uses Amazon Redshift Spectrum to run complex analytical queries on objects that are stored in an Amazon S3 bucket. The objects are joined with multiple dimension tables that are stored in an Amazon Redshift database. The company uses the database to create monthly and quarterly aggregated reports. Users who attempt to run queries are reporting the following error message: error: Spectrum Scan Error: Access throttled
Which solution will resolve this error?

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

Answer: C

Explanation:
https://docs.aws.amazon.com/redshift/latest/dg/c-spectrum-troubleshooting.html#spectrum-troubleshooting-acce https://docs.aws.amazon.com/redshift/latest/dg/c-spectrum-troubleshooting.html https://docs.aws.amazon.com/redshift/latest/dg/c-spectrum-troubleshooting.html#spectrum-troubleshooting-acce

NEW QUESTION 260
A company runs online transaction processing (OLTP) workloads on an Amazon RDS for PostgreSQL Multi- AZ DB instance. Tests were run on the database after work hours, which generated additional database logs. The free storage of the RDS DB instance is low due to these additional logs.
What should the company do to address this space constraint issue?

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

Answer: B

Explanation:
To set the retention period for system logs, use the rds.log_retention_period parameter. You can find rds.log_retention_period in the DB parameter group associated with your DB instance. The unit for this parameter is minutes. For example, a setting of 1,440 retains logs for one day. The default value is 4,320 (three days). The maximum value is 10,080 (seven days).
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/USER_LogAccess.Concepts.PostgreSQL.ht

NEW QUESTION 263
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 265
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 267
A 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 instance
E. Use Amazon Route 53 DNS with failover health checks configured.

Answer: B

NEW QUESTION 269
A gaming company has recently acquired a successful iOS game, which is particularly popular during the holiday season. The company has decided to add a leaderboard to the game that uses Amazon DynamoDB. The application load is expected to ramp up over the holiday season.
Which solution will meet these requirements at the lowest cost?

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

Answer: C

NEW QUESTION 270
A Database Specialist is designing a new database infrastructure for a ride hailing application. The application data includes a ride tracking system that stores GPS coordinates for all rides. Real-time statistics and metadata lookups must be performed with high throughput and microsecond latency. The database should be fault tolerant with minimal operational overhead and development effort.
Which solution meets these requirements in the MOST efficient way?

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

Answer: B

Explanation:
https://aws.amazon.com/dynamodb/dax/#:~:text=Amazon%20DynamoDB%20Accelerator%20(DAX)%20is,mil "Amazon DynamoDB Accelerator (DAX) is a fully managed, highly available, in-memory cache for DynamoDB that delivers up to a 10x performance improvement – from milliseconds to microseconds – even at millions of requests per second. "

NEW QUESTION 271
The Security team for a finance company was notified of an internal security breach that happened 3 weeks ago. A Database Specialist must start producing audit logs out of the production Amazon Aurora PostgreSQL cluster for the Security team to use for monitoring and alerting. The Security team is required to perform

real- time alerting and monitoring outside the Aurora DB cluster and wants to have the cluster push encrypted files to the chosen solution.
Which approach will meet these requirements?

A. Use pg_audit to generate audit logs and send the logs to the Security team.
B. Use AWS CloudTrail to audit the DB cluster and the Security team will get data from Amazon S3.
C. Set up database activity streams and connect the data stream from Amazon Kinesis to consumer applications.
D. Turn on verbose logging and set up a schedule for the logs to be dumped out for the Security team.

Answer: C

Explanation:
https://aws.amazon.com/about-aws/whats-new/2019/05/amazon-aurora-with-postgresql-compatibility-supports- "Database Activity Streams for Amazon Aurora with PostgreSQL compatibility provides a near real-time data stream of the database activity in your relational database to help you monitor activity. When integrated with third party database activity monitoring tools, Database Activity Streams can monitor and audit database activity to provide safeguards for your database and help meet compliance and regulatory requirements."
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Overview.LoggingAndMonitoring.html

NEW QUESTION 272
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 276
A company is running an on-premises application comprised of a web tier, an application tier, and a MySQL database tier. The database is used primarily during business hours with random activity peaks throughout the day. A database specialist needs to improve the availability and reduce the cost of the MySQL database tier as part of the company's migration to AWS.
Which MySQL database option would meet these requirements?

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

Answer: C

NEW QUESTION 277
An online retail company is planning a multi-day flash sale that must support processing of up to 5,000 orders per second. The number of orders and exact schedule for the sale will vary each day. During the sale, approximately 10,000 concurrent users will look at the deals before buying items. Outside of the sale, the traffic volume is very low. The acceptable performance for read/write queries should be under 25 ms. Order items are about 2 KB in size and have a unique identifier. The company requires the most cost-effective solution that will automatically scale and is highly available.
Which solution meets these requirements?

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

Answer: A

Explanation:
The number of orders and exact schedule for the sale will vary each day. During the sale, approximately 10,000 concurrent users will look at the deals before buying items. Outside of the sale, the traffic volume is very low ==> Setting provisioning DynamoDB fix read 5000/write 10000 with will waste the resource when the traffic is low. It is not cost-effective.

NEW QUESTION 281
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 284
Amazon Neptune is being used by a corporation as the graph database for one of its products. During an ETL procedure, the company's data science team produced enormous volumes of temporary data by unintentionally. The Neptune DB cluster extended its storage capacity automatically to handle the added data, but the data science team erased the superfluous data.
What should a database professional do to prevent incurring extra expenditures for cluster volume space that is not being used?

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

Answer: C

Explanation:
The only way to shrink the storage space used by your DB cluster when you have a large amount of unused allocated space is to export all the data in your graph and then reload it into a new DB cluster. Creating and restoring a snapshot does not reduce the amount of storage allocated for your DB cluster, because a snapshot retains the original image of the cluster's underlying storage.

NEW QUESTION 288
A company hosts an on-premises Microsoft SQL Server Enterprise edition database with Transparent Data Encryption (TDE) enabled. The database is 20 TB in size and includes sparse tables. The company needs to migrate the database to Amazon RDS for SQL Server during a maintenance window that is scheduled for an upcoming weekend. Data-at-rest encryption must be enabled for the target DB instance.
Which combination of steps should the company take to migrate the database to AWS in the MOST operationally efficient manner? (Choose two.)

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

Answer: BC

Explanation:
https://aws.amazon.com/blogs/database/migrate-tde-enabled-sql-server-databases-to-amazon-rds-for-sql-server/

NEW QUESTION 293
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 294
A financial services company is developing a shared data service that supports different applications from throughout the company. A Database Specialist designed a solution to leverage Amazon ElastiCache for Redis with cluster mode enabled to enhance performance and scalability. The cluster is configured to listen on port 6379.
Which combination of steps should the Database Specialist take to secure the cache data and protect it from unauthorized access? (Choose three.)

A. Enable in-transit and at-rest encryption on the ElastiCache cluster.
B. Ensure that Amazon CloudWatch metrics are configured in the ElastiCache cluster.
C. Ensure the security group for the ElastiCache cluster allows all inbound traffic from itself and inbound traffic on TCP port 6379 from trusted clients only.
D. Create an IAM policy to allow the application service roles to access all ElastiCache API actions.
E. Ensure the security group for the ElastiCache clients authorize inbound TCP port 6379 and port 22 traffic from the trusted ElastiCache cluster’s security group.
F. Ensure the cluster is created with the auth-token parameter and that the parameter is used in all subsequent commands.

Answer: ACF

Explanation:
https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/encryption.html

NEW QUESTION 298
A company wants to build a new invoicing service for its cloud-native application on AWS. The company has a small development team and wants to focus on service feature development and minimize operations and maintenance as much as possible. The company expects the service to handle billions of requests and millions of new records every day. The service feature requirements, including data access patterns are well-defined. The service has an availability target of 99.99 % with a milliseconds latency requirement. The database for the service will be the system of record for invoicing data.
Which database solution meets these requirements at the LOWEST cost?

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

Answer: D

Explanation:
Known patterns, minimum maintenance, miliseconds latency

NEW QUESTION 299
A 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 instanc
B. Promote the read replica at the time of failover.
C. Set up SQL replication from the DB instance to an Amazon EC2 instance in the disaster recovery Regio
D. Promote the EC2 instance as the primary server.
E. Use AWS Database Migration Service (AWS KMS) for ongoing replication of the DB instance in the disaster recovery Region.
F. Take manual snapshots of the DB instance in the primary Regio
G. Copy the snapshots to the disaster recovery Region.

Answer: C

Explanation:
https://aws.amazon.com/blogs/database/cross-region-disaster-recovery-of-amazon-rds-for-sql-server/

NEW QUESTION 303
.....

Relate Links

100% Pass Your DBS-C01 Exam with Exambible Prep Materials

https://www.exambible.com/DBS-C01-exam/

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - https://www.exambible.com/