AWS-Certified-Database-Specialty Dumps

AWS Certified Database - Specialty

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

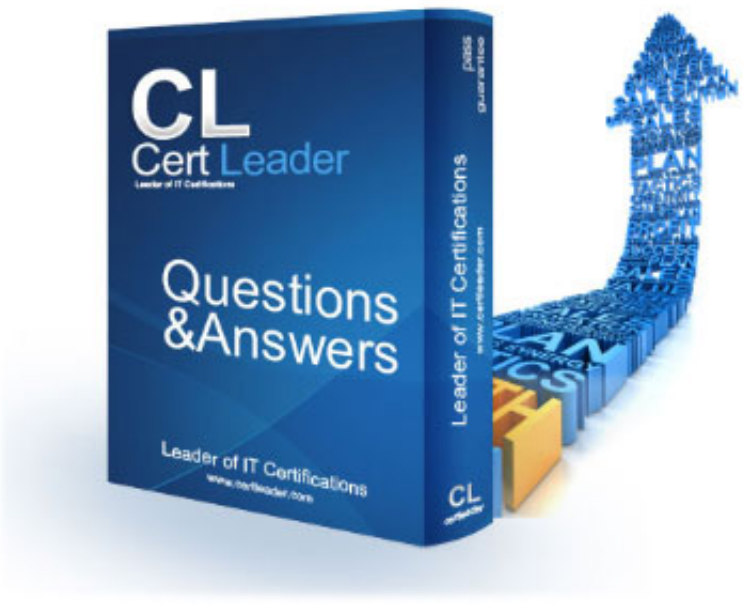

NEW QUESTION 1
A database specialist manages a critical Amazon RDS for MySQL DB instance for a company. The data stored daily could vary from .01% to 10% of the current database size. The database specialist needs to ensure that the DB instance storage grows as needed.
What is the MOST operationally 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 2
A company uses Amazon Aurora for secure financial transactions. The data must always be encrypted at rest and in transit to meet compliance requirements.
Which combination of actions should a database specialist take to meet these requirements? (Choose two.)

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

Answer: AB

Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-replicas-adding.html

NEW QUESTION 3
A financial services organization employs an Amazon Aurora PostgreSQL DB cluster to host an application on AWS. No log files detailing database administrator activity were discovered during a recent examination. A database professional must suggest a solution that enables access to the database and maintains activity logs. The solution should be simple to implement and have a negligible effect on performance.
Which database specialist solution should be recommended?

A. Enable Aurora Database Activity Streams on the database in synchronous mod
B. Connect the Amazon Kinesis data stream to Kinesis Data Firehos
C. Set the Kinesis Data Firehose destination to an Amazon S3 bucket.
D. Create an AWS CloudTrail trail in the Region where the database run
E. Associate the database activity logs with the trail.
F. Enable Aurora Database Activity Streams on the database in asynchronous mod
G. Connect the Amazon Kinesis data stream to Kinesis Data Firehos
H. Set the Firehose destination to an Amazon S3 bucket.
I. Allow connections to the DB cluster through a bastion host onl
J. Restrict database access to the bastion host and application server
K. Push the bastion host logs to Amazon CloudWatch Logs using the CloudWatch Logs agent.

Answer: C

Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/DBActivityStreams.Overview.html

NEW QUESTION 4
A company is closing one of its remote data centers. This site runs a 100 TB on-premises data warehouse solution. The company plans to use the AWS Schema Conversion Tool (AWS SCT) and AWS DMS for the migration to AWS. The site network bandwidth is 500 Mbps. A Database Specialist wants to migrate the on-premises data using Amazon S3 as the data lake and Amazon Redshift as the data warehouse. This move must take place during a 2-week period when source systems are shut down for maintenance. The data should stay encrypted at rest and in transit.
Which approach has the least risk and the highest likelihood of a successful data transfer?

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

M. Once complete, load the data to Amazon Redshift using AWS Glue.

Answer: B

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

NEW QUESTION 5
A Database Specialist is setting up a new Amazon Aurora DB cluster with one primary instance and three Aurora Replicas for a highly intensive, business-critical application. The Aurora DB cluster has one medium- sized primary instance, one large-sized replica, and two medium sized replicas. The Database Specialist did not assign a promotion tier to the replicas.
In the event of a primary failure, what will occur?

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

Answer: C

Explanation:
Priority: If you don't select a value, the default is tier-1. This priority determines the order in which Aurora
https://docs.amazonaws.cn/en_us/AmazonRDS/latest/AuroraUserGuide/aurora-replicas-adding.html
More than one Aurora Replica can share the same priority, resulting in promotion tiers. If two or more Aurora Replicas share the same priority, then Amazon RDS promotes the replica that is largest in size. If two or more Aurora Replicas share the same priority and size, then Amazon RDS promotes an arbitrary replica in the same promotion tier.
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Managing.Backups.html#Aurora.M If two or more Aurora Replicas share the same priority, then Amazon RDS promotes the replica that is largest in size. If two or more Aurora Replicas share the same priority and size, then Amazon RDS promotes an arbitrary replica in the same promotion tier. https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Concepts.AuroraHighAvailability.html

NEW QUESTION 6
A company has an ecommerce web application with an Amazon RDS for MySQL DB instance. The marketing team has noticed some unexpected updates to the product and pricing information on the website, which is impacting sales targets. The marketing team wants a database specialist to audit future database activity to help identify how and when the changes are being made.
What should the database specialist do to meet these requirements? (Choose two.)

A. Create an RDS event subscription to the audit event type.
B. Enable auditing of CONNECT and QUERY_DML events.
C. SSH to the DB instance and review the database logs.
D. Publish the database logs to Amazon CloudWatch Logs.
E. Enable Enhanced Monitoring on the DB instance.

Answer: BD

Explanation:
https://aws.amazon.com/blogs/database/configuring-an-audit-log-to-capture-database-activities-for-amazon-rds

NEW QUESTION 7
A company maintains several databases using Amazon RDS for MySQL and PostgreSQL. Each RDS database generates log files with retention periods set to their default values. The company has now mandated that database logs be maintained for up to 90 days in a centralized repository to facilitate real-time and afterthe-fact analyses.
What should a Database Specialist do to meet these requirements with minimal effort?

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

Answer: B

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

NEW QUESTION 8
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 9
A gaming company has implemented a leaderboard in AWS using a Sorted Set data structure within Amazon ElastiCache for Redis. The ElastiCache cluster has been deployed with cluster mode disabled and has a replication group deployed with two additional replicas. The company is planning for a worldwide gaming event and is anticipating a higher write load than what the current cluster can handle.
Which method should a Database Specialist use to scale the ElastiCache cluster ahead of the upcoming event?

A. Enable cluster mode on the existing ElastiCache cluster and configure separate shards for the Sorted Set across all nodes in the cluster.
B. Increase the size of the ElastiCache cluster nodes to a larger instance size.
C. Create an additional ElastiCache cluster and load-balance traffic between the two clusters.
D. Use the EXPIRE command and set a higher time to live (TTL) after each call to increment a given key.

Answer: B

NEW QUESTION 10
A huge gaming firm is developing a centralized method for storing the status of various online games' user sessions. The workload requires low-latency key-value storage and will consist of an equal number of reads and writes. Across the games' geographically dispersed user base, data should be written to the AWS Region nearest to the user. The design should reduce the burden associated with managing data replication across Regions.
Which solution satisfies these criteria?

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

Answer: D

Explanation:
https://aws.amazon.com/dynamodb/?nc1=h_ls

NEW QUESTION 10
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 14
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 15
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 17
A database specialist is responsible for an Amazon RDS for MySQL DB instance with one read replica. The DB instance and the read replica are assigned to the default parameter group. The database team currently runs test queries against a read replica. The database team wants to create additional tables in the read replica that will only be accessible from the read replica to benefit the tests.
Which should the database specialist do to allow the database team to create the test tables?

A. Contact AWS Support to disable read-only mode on the read replic
B. Reboot the read replic
C. Connect to the read replica and create the tables.
D. Change the read_only parameter to false (read_only=0) in the default parameter group of the read replic
E. Perform a reboot without failove
F. Connect to the read replica and create the tables using the local_only MySQL option.
G. Change the read_only parameter to false (read_only=0) in the default parameter grou
H. Reboot the read replic
I. Connect to the read replica and create the tables.
J. Create a new DB parameter grou
K. Change the read_only parameter to false (read_only=0). Associate the read replica with the new grou
L. Reboot the read replic
M. Connect to the read replica and create the tables.

Answer: D

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/rds-read-replica/

NEW QUESTION 19
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 20
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 24
A company is hosting critical business data in an Amazon Redshift cluster. Due to the sensitive nature of the data, the cluster is encrypted at rest using AWS KMS. As a part of disaster recovery requirements, the company needs to copy the Amazon Redshift snapshots to another Region.
Which steps should be taken in the AWS Management Console to meet the disaster recovery requirements?

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

Answer: C

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

NEW QUESTION 27
A marketing company is using Amazon DocumentDB and requires that database audit logs be enabled. A Database Specialist needs to configure monitoring so that all data definition language (DDL) statements performed are visible to the Administrator. The Database Specialist has set the audit_logs parameter to enabled in the cluster parameter group.
What should the Database Specialist do to automatically collect the database logs for the Administrator?

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

Answer: C

Explanation:
https://docs.aws.amazon.com/documentdb/latest/developerguide/event-auditing.html Auditing Amazon DocumentDB Events
PDF
Kindle RSS
With Amazon DocumentDB (with MongoDB compatibility), you can audit events that were performed in your cluster. Examples of logged events include successful and failed authentication attempts, dropping a collection in a database, or creating an index. By default, auditing is disabled on Amazon DocumentDB and requires that you opt in to use this feature.
When auditing is enabled, Amazon DocumentDB records Data Definition Language (DDL), authentication, authorization, and user management events to Amazon CloudWatch Logs. When auditing is enabled, Amazon DocumentDB exports your cluster's auditing records (JSON documents) to Amazon CloudWatch Logs. You can use Amazon CloudWatch Logs to analyze, monitor, and archive your Amazon DocumentDB auditing events.

NEW QUESTION 30
A company wants to migrate its on-premises MySQL databases to Amazon RDS for MySQL. To comply with the company's security policy, all databases must be encrypted at rest. RDS DB instance snapshots must also be shared across various accounts to provision testing and staging environments.
Which solution meets these requirements?

A. Create an RDS for MySQL DB instance with an AWS Key Management Service (AWS KMS) customer managed CM
B. Update the key policy to include the Amazon Resource Name (ARN) of the other AWS accounts as a principal, and then allow the kms:CreateGrant action.
C. Create an RDS for MySQL DB instance with an AWS managed CM
D. Create a new key policy to include the Amazon Resource Name (ARN) of the other AWS accounts as a principal, and then allow the kms:CreateGrant action.
E. Create an RDS for MySQL DB instance with an AWS owned CM
F. Create a new key policy to include the administrator user name of the other AWS accounts as a principal, and then allow the kms:CreateGrant action.
G. Create an RDS for MySQL DB instance with an AWS CloudHSM ke
H. Update the key policy to include the Amazon Resource Name (ARN) of the other AWS accounts as a principal, and then allow the kms:CreateGrant action.

Answer: A

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

NEW QUESTION 35
A company has a database monitoring solution that uses Amazon CloudWatch for its Amazon RDS for SQL Server environment. The cause of a recent spike in CPU utilization was not determined using the standard metrics that were collected. The CPU spike caused the application to perform poorly, impacting users. A Database Specialist needs to determine what caused the CPU spike.
Which combination of steps should be taken to provide more visibility into the processes and queries running during an increase in CPU load? (Choose two.)

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

Answer: BE

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

NEW QUESTION 40
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 43
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 44
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 46
A database specialist is building a system that uses a static vendor dataset of postal codes and related territory information that is less than 1 GB in size. The dataset is loaded into the application's cache at start up. The company needs to store this data in a way that provides the lowest cost with a low application startup time.
Which approach will meet these requirements?

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

Answer: D

Explanation:
https://www.sumologic.com/insight/s3-cost-optimization/
For example, for 1 GB file stored on S3 with 1 TB of storage provisioned, you are billed for 1 GB only. In a lot of other services such as Amazon EC2, Amazon Elastic Block Storage (Amazon EBS) and Amazon DynamoDB you pay for provisioned capacity. For example, in the case of Amazon EBS disk you pay for the size of 1 TB of disk even if you just save 1 GB file. This makes managing S3 cost easier than many other services including Amazon EBS and Amazon EC2. On S3 there is no risk of over-provisioning and no need to manage disk utilization.

NEW QUESTION 48
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 50
A large financial services company requires that all data be encrypted in transit. A Developer is attempting to connect to an Amazon RDS DB instance using the company VPC for the first time with credentials provided by a Database Specialist. Other members of the Development team can connect, but this user is consistently receiving an error indicating a communications link failure. The Developer asked the Database Specialist to reset the password a number of times, but the error persists.
Which step should be taken to troubleshoot this issue?

A. Ensure that the database option group for the RDS DB instance allows ingress from the Developer machine's IP address
B. Ensure that the RDS DB instance's subnet group includes a public subnet to allow the Developer to connect
C. Ensure that the RDS DB instance has not reached its maximum connections limit
D. Ensure that the connection is using SSL and is addressing the port where the RDS DB instance is listening for encrypted connections

Answer: D

Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/SQLServer.Concepts.General.SSL.Using.html

NEW QUESTION 53
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 57
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 60
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 62
A Database Specialist migrated an existing production MySQL database from on-premises to an Amazon RDS for MySQL DB instance. However, after the migration, the database needed to be encrypted at rest using AWS KMS. Due to the size of the database, reloading, the data into an encrypted database would be too time- consuming, so it is not an option.
How should the Database Specialist satisfy this new requirement?

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

Answer: A

Explanation:
"However, because you can encrypt a copy of an unencrypted DB snapshot, you can effectively add encryption to an unencrypted DB instance. That is, you can create a snapshot of your DB instance, and then create an encrypted copy of that snapshot. You can then restore a DB instance from the encrypted snapshot, and thus you have an encrypted copy of your original DB instance. For more information, see Copying a Snapshot."
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.Encryption.html

NEW QUESTION 66
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 70
A company needs a data warehouse solution that keeps data in a consistent, highly structured format. The company requires fast responses for end-user queries when looking at data from the current year, and users must have access to the full 15-year dataset, when needed. This solution also needs to handle a fluctuating number incoming queries. Storage costs for the 100 TB of data must be kept low.
Which solution meets these requirements?

A. Leverage an Amazon Redshift data warehouse solution using a dense storage instance type while keeping all the data on local Amazon Redshift storag
B. Provision enough instances to support high demand.
C. Leverage an Amazon Redshift data warehouse solution using a dense storage instance to store the most recent dat
D. Keep historical data on Amazon S3 and access it using the Amazon Redshift Spectrum laye
E. Provision enough instances to support high demand.
F. Leverage an Amazon Redshift data warehouse solution using a dense storage instance to store the most recent dat
G. Keep historical data on Amazon S3 and access it using the Amazon Redshift Spectrum laye
H. Enable Amazon Redshift Concurrency Scaling.
I. Leverage an Amazon Redshift data warehouse solution using a dense storage instance to store the most recent dat
J. Keep historical data on Amazon S3 and access it using the Amazon Redshift Spectrum laye
K. Leverage Amazon Redshift elastic resize.

Answer: C

Explanation:
https://docs.aws.amazon.com/redshift/latest/dg/concurrency-scaling.html
"With the Concurrency Scaling feature, you can support virtually unlimited concurrent users and concurrent queries, with consistently fast query performance. When concurrency scaling is enabled, Amazon Redshift automatically adds additional cluster capacity when you need it to process an increase in concurrent read queries. Write operations continue as normal on your main cluster. Users always see the most current data, whether the queries run on the main cluster or on a concurrency scaling cluster. You're charged for concurrency scaling clusters only for the time they're in use. For more information about pricing, see Amazon Redshift pricing. You manage which queries are sent to the concurrency scaling cluster by configuring WLM queues. When you enable concurrency scaling for a queue, eligible queries are sent to the concurrency scaling cluster instead of waiting in line."

NEW QUESTION 74
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 78
A database expert is responsible for building a highly available online transaction processing (OLTP) solution that makes use of Amazon RDS for MySQL production databases. Disaster recovery criteria include a cross-regional deployment and an RPO and RTO of 5 and 30 minutes, respectively.
What should the database professional do to ensure that the database meets the criteria for high availability and disaster recovery?

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

Answer: D

NEW QUESTION 81
A financial institution uses AWS to host its online application. Amazon RDS for MySQL is used to host the application's database, which includes automatic backups.
The program has corrupted the database logically, resulting in the application being unresponsive. The exact moment the corruption occurred has been determined, and it occurred within the backup retention period.
How should a database professional restore a database to its previous state prior to corruption?

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

Answer: B

Explanation:
When you perform a restore operation to a point in time or from a DB Snapshot, a new DB Instance is created with a new endpoint (the old DB Instance can be deleted if so desired). This is done to enable you to create multiple DB Instances from a specific DB Snapshot or point in time."

NEW QUESTION 86
The Development team recently executed a database script containing several data definition language (DDL) and data manipulation language (DML) statements on an Amazon Aurora MySQL DB cluster. The release accidentally deleted thousands of rows from an important table and broke some application functionality. This was discovered 4 hours after the release. Upon investigation, a Database Specialist tracked the issue to a DELETE command in the script with an incorrect WHERE clause filtering the wrong set of rows.
The Aurora DB cluster has Backtrack enabled with an 8-hour backtrack window. The Database Administrator also took a manual snapshot of the DB cluster before the release started. The database needs to be returned to the correct state as quickly as possible to resume full application functionality. Data loss must be minimal. How can the Database Specialist accomplish this?

A. Quickly rewind the DB cluster to a point in time before the release using Backtrack.
B. Perform a point-in-time recovery (PITR) of the DB cluster to a time before the release and copy the deleted rows from the restored database to the original database.
C. Restore the DB cluster using the manual backup snapshot created before the release and change the application configuration settings to point to the new DB cluster.
D. Create a clone of the DB cluster with Backtrack enable
E. Rewind the cloned cluster to a point in time before the releas
F. Copy deleted rows from the clone to the original database.

Answer: A

NEW QUESTION 89
A company is running a two-tier ecommerce application in one AWS account. The web server is deployed using an Amazon RDS for MySQL Multi-AZ DB instance. A Developer mistakenly deleted the database in the production environment. The database has been restored, but this resulted in hours of downtime and lost revenue.
Which combination of changes in existing IAM policies should a Database Specialist make to prevent an error like this from happening in the future? (Choose three.)

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

Answer: ACD

Explanation:

https://aws.amazon.com/blogs/database/using-iam-multifactor-authentication-with-amazon-rds/
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/security_iam_id-based-policy-htmlhttps://docs.aws

NEW QUESTION 93
In North America, a business launched a mobile game that swiftly expanded to 10 million daily active players. The game's backend is hosted on AWS and makes considerable use of a TTL-configured Amazon DynamoDB table.
When an item is added or changed, its TTL is set to 600 seconds plus the current epoch time. The game logic is reliant on the purging of outdated data in order to compute rewards points properly. At times, items from the table are read that are many hours beyond their TTL expiration.
How should a database administrator resolve this issue?

A. Use a client library that supports the TTL functionality for DynamoDB.
B. Include a query filter expression to ignore items with an expired TTL.
C. Set the ConsistentRead parameter to true when querying the table.
D. Create a local secondary index on the TTL attribute.

Answer: B

Explanation:
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/howitworks-ttl.html

NEW QUESTION 95
A company is running an Amazon RDS for PostgeSQL DB instance and wants to migrate it to an Amazon Aurora PostgreSQL DB cluster. The current database is 1 TB in size. The migration needs to have minimal downtime.
What is the FASTEST way to accomplish this?

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

Answer: D

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

NEW QUESTION 99
The website of a manufacturing firm makes use of an Amazon Aurora PostgreSQL database cluster. Which settings will result in the LEAST amount of downtime for the application during failover? (Select
three.)

A. Use the provided read and write Aurora endpoints to establish a connection to the Aurora DB cluster.
B. Create an Amazon CloudWatch alert triggering a restore in another Availability Zone when the primary Aurora DB cluster is unreachable.
C. Edit and enable Aurora DB cluster cache management in parameter groups.
D. Set TCP keepalive parameters to a high value.
E. Set JDBC connection string timeout variables to a low value.
F. Set Java DNS caching timeouts to a high value.

Answer: ACE

Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.BestPractices.html
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.cluster-cache-mgmt.htm
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.BestPractices.html#Aur

NEW QUESTION 101
A company is releasing a new mobile game featuring a team play mode. As a group of mobile device users play together, an item containing their statuses is updated in an Amazon DynamoDB table. Periodically, the other users' devices read the latest statuses of their teammates from the table using the BatchGetItemn operation.
Prior to launch, some testers submitted bug reports claiming that the status data they were seeing in the game was not up-to-date. The developers are unable to replicate this issue and have asked a database specialist for a recommendation.
Which recommendation would resolve this issue?

A. Ensure the DynamoDB table is configured to be always consistent.
B. Ensure the BatchGetItem operation is called with the ConsistentRead parameter set to false.
C. Enable a stream on the DynamoDB table and subscribe each device to the stream to ensure all devices receive up-to-date status information.
D. Ensure the BatchGetItem operation is called with the ConsistentRead parameter set to true.

Answer: D

Explanation:
https://docs.aws.amazon.com/ja_jp/amazondynamodb/latest/developerguide/API_BatchGetItem_v20111205.htm By default, BatchGetItem performs eventually consistent reads on every table in the request. If you want strongly consistent reads instead, you can set ConsistentRead to true for any or all tables.

NEW QUESTION 104
A major organization maintains a number of Amazon DB clusters. Each of these clusters is configured differently to meet certain needs. These configurations may be classified into wider groups based on the team and use case.
A database administrator wishes to streamline the process of storing and updating these settings. Additionally, the database administrator want to guarantee that changes to certain configuration categories are automatically implemented to all instances as necessary.
Which AWS service or functionality will assist in automating and achieving this goal?

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

Answer: B

Explanation:
Database parameters specify how the database is configured. For example, database parameters can specify the amount of resources, such as memory, to allocate to a database.

NEW QUESTION 105
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 107
An online advertising website uses an Amazon DynamoDB table with on-demand capacity mode as its data store. The website also has a DynamoDB Accelerator (DAX) cluster in the same VPC as its web application server. The application needs to perform infrequent writes and many strongly consistent reads from the data store by querying the DAX cluster.
During a performance audit, a systems administrator notices that the application can look up items by using the DAX cluster. However, the QueryCacheHits metric for the DAX cluster consistently shows 0 while the QueryCacheMisses metric continuously keeps growing in Amazon CloudWatch.
What is the MOST likely reason for this occurrence?

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

Answer: B

Explanation:
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/DAX.concepts.html
"If the request specifies strongly consistent reads, DAX passes the request through to DynamoDB. The results from DynamoDB are not cached in DAX. Instead, they are simply returned to the application."

NEW QUESTION 110
A Database Specialist modified an existing parameter group currently associated with a production Amazon RDS for SQL Server Multi-AZ DB instance. The change is associated with a static parameter type, which controls the number of user connections allowed on the most critical RDS SQL Server DB instance for the company. This change has been approved for a specific maintenance window to help minimize the impact on users.
How should the Database Specialist apply the parameter group change for the DB instance?

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

Answer: C

Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_WorkingWithParamGroups.html#USER_W

NEW QUESTION 113
An ecommerce company has tasked a Database Specialist with creating a reporting dashboard that visualizes critical business metrics that will be pulled from the core production database running on Amazon Aurora. Data that is read by the dashboard should be available within 100 milliseconds of an update.
The Database Specialist needs to review the current configuration of the Aurora DB cluster and develop a cost-effective solution. The solution needs to accommodate the unpredictable read workload from the
reporting dashboard without any impact on the write availability and performance of the DB cluster.
Which solution meets these requirements?

A. Turn on the serverless option in the DB cluster so it can automatically scale based on demand.
B. Provision a clone of the existing DB cluster for the new Application team.

C. Create a separate DB cluster for the new workload, refresh from the source DB cluster, and set up ongoing replication using AWS DMS change data capture (CDC).
D. Add an automatic scaling policy to the DB cluster to add Aurora Replicas to the cluster based on CPU consumption.

Answer: A

NEW QUESTION 118
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 120
A small startup firm wishes to move a 4 TB MySQL database from on-premises to AWS through an Amazon RDS for MySQL DB instance.
Which migration approach would result in the LEAST amount of downtime?

A. Deploy a new RDS for MySQL DB instance and configure it for access from the on-premises data cente
B. Use the mysqldump utility to create an initial snapshot from the on-premises MySQL server, and copy it to an Amazon S3 bucke
C. Import the snapshot into the DB instance utilizing the MySQL utilities running on an Amazon EC2 instanc
D. Immediately point the application to the DB instance.
E. Deploy a new Amazon EC2 instance, install the MySQL software on the EC2 instance, and configure networking for access from the on-premises data cente
F. Use the mysqldump utility to create a snapshot of the on-premises MySQL serve
G. Copy the snapshot into the EC2 instance and restore it into the EC2 MySQL instanc
H. Use AWS DMS to migrate data into a new RDS for MySQL DB instanc
I. Point the application to the DB instance.
J. Deploy a new Amazon EC2 instance, install the MySQL software on the EC2 instance, and configure networking for access from the on-premises data cente
K. Use the mysqldump utility to create a snapshot of the on-premises MySQL serve
L. Copy the snapshot into an Amazon S3 bucket and import the snapshot into a new RDS for MySQL DB instance using the MySQL utilities running on an EC2 instanc
M. Point the application to the DB instance.
N. Deploy a new RDS for MySQL DB instance and configure it for access from the on-premises data cente
O. Use the mysqldump utility to create an initial snapshot from the on-premises MySQL server, and copy it to an Amazon S3 bucke
P. Import the snapshot into the DB instance using the MySQL utilities running on an Amazon EC2 instanc
Q. Establish replication into the new DB instance using MySQL replicatio
R. Stop application access to the on-premises MySQL server and let the remaining transactions replicate ove
S. Point the application to the DB instance.

Answer: D

Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/MySQL.Procedural.Importing.NonRDSRepl.html
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/MySQL.Procedural.Importing.External.Repl.html

NEW QUESTION 121
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 122
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 124
A company's database specialist disabled TLS on an Amazon DocumentDB cluster to perform benchmarking tests. A few days after this change was implemented, a database specialist trainee accidentally deleted multiple tables. The database specialist restored the database from available snapshots. An hour after restoring the cluster, the database specialist is still unable to connect to the new cluster endpoint.
What should the database specialist do to connect to the new, restored Amazon DocumentDB cluster?

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

Answer: A

Explanation:
You can't modify the parameter settings of the default parameter groups. You can use a DB parameter group to act as a container for engine configuration values that are applied to one or more DB instances. If you create a DB instance without specifying a DB parameter group, the DB instance uses a default DB parameter group. Each default DB parameter group contains database engine defaults and Amazon RDS system defaults. You can't modify the parameter settings of a default parameter group. Instead, you create your own parameter group where you choose your own parameter settings. Not all DB engine parameters can be changed in a parameter group that you create.

NEW QUESTION 127
A company is using an Amazon RDS for MySQL DB instance for its internal applications. A security audit shows that the DB instance is not encrypted at rest. The company's application team needs to encrypt the DB instance.
What should the team do to meet this requirement?

A. Stop the DB instance and modify it to enable encryptio
B. Apply this setting immediately without waiting for the next scheduled RDS maintenance window.
C. Stop the DB instance and create an encrypted snapsho
D. Restore the encrypted snapshot to a new encrypted DB instanc
E. Delete the original DB instance, and update the applications to point to the new encrypted DB instance.
F. Stop the DB instance and create a snapsho
G. Copy the snapshot into another encrypted snapsho
H. Restore the encrypted snapshot to a new encrypted DB instanc
I. Delete the original DB instance, and update the applications to point to the new encrypted DB instance.
J. Create an encrypted read replica of the DB instanc
K. Promote the read replica to maste
L. Delete the original DB instance, and update the applications to point to the new encrypted DB instance.

Answer: C

NEW QUESTION 130
A company is using Amazon Aurora PostgreSQL for the backend of its application. The system users are complaining that the responses are slow. A database specialist has determined that the queries to Aurora take longer during peak times. With the Amazon RDS Performance Insights dashboard, the load in the chart for average active sessions is often above the line that denotes maximum CPU usage and the wait state shows that most wait events are IO:XactSync.
What should the company do to resolve these performance issues?

A. Add an Aurora Replica to scale the read traffic.
B. Scale up the DB instance class.
C. Modify applications to commit transactions in batches.
D. Modify applications to avoid conflicts by taking locks.

Answer: C

Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.Reference.html https://blog.dbi-services.com/aws-aurora-xactsync-batch-commit/

NEW QUESTION 131
A database specialist is constructing an AWS CloudFormation stack using AWS CloudFormation. The database expert wishes to avoid the stack's Amazon RDS ProductionDatabase resource being accidentally deleted.
Which solution will satisfy this criterion?

A. Create a stack policy to prevent update
B. Include €Effect€ : €ProductionDatabase€ and €Resource€€Deny€ in the policy.
C. Create an AWS CloudFormation stack in XML forma
D. Set xAttribute as false.
E. Create an RDS DB instance without the DeletionPolicy attribut
F. Disable termination protection.
G. Create a stack policy to prevent update
H. Include Effect, Deny, and Resource :ProductionDatabase in the policy.

Answer: D

Explanation:
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/protect-stack-resources.html "When you set a stack policy, all resources are protected by default. To allow updates on all resources, we add an Allow statement that allows all actions on all resources. Although the Allow statement specifies all resources, the explicit Deny statement overrides it for the resource with the ProductionDatabase logical ID. This Deny statement prevents all update actions, such as replacement or deletion, on the ProductionDatabase resource."

NEW QUESTION 134
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 136
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 140
A Database Specialist is migrating an on-premises Microsoft SQL Server application database to Amazon RDS for PostgreSQL using AWS DMS. The application requires minimal downtime when the RDS DB instance goes live.
What change should the Database Specialist make to enable the migration?

A. Configure the on-premises application database to act as a source for an AWS DMS full load with ongoing change data capture (CDC)
B. Configure the AWS DMS replication instance to allow both full load and ongoing change data capture (CDC)
C. Configure the AWS DMS task to generate full logs to allow for ongoing change data capture (CDC)
D. Configure the AWS DMS connections to allow two-way communication to allow for ongoing change data capture (CDC)

Answer: A

Explanation:
"requires minimal downtime when the RDS DB instance goes live" in order to do CDC: "you must first ensure that ARCHIVELOG MODE is on to provide information to LogMiner. AWS DMS uses LogMiner to read information from the archive logs so that AWS DMS can capture changes"
https://docs.aws.amazon.com/dms/latest/sbs/chap-oracle2postgresql.steps.configureoracle.html "If you want to capture and apply changes (CDC), then you also need the following privileges."

NEW QUESTION 141
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 142
An internet advertising firm stores its data in an Amazon DynamoDb table. Amazon DynamoDB Streams are enabled on the table, and one of the keys has a global secondary index. The table is encrypted using a customer-managed AWS Key Management Service (AWS KMS) key.
The firm has chosen to grow worldwide and want to duplicate the database using DynamoDB global tables in a new AWS Region.
An administrator observes the following upon review:
No role with the dynamodb: CreateGlobalTable permission exists in the account.
An empty table with the same name exists in the new Region where replication is desired.
A global secondary index with the same partition key but a different sort key exists in the new Region where replication is desired.

Which settings will prevent you from creating a global table or replica in the new Region? (Select two.)

A. A global secondary index with the same partition key but a different sort key exists in the new Region where replication is desired.
B. An empty table with the same name exists in the Region where replication is desired.
C. No role with the dynamodb:CreateGlobalTable permission exists in the account.
D. DynamoDB Streams is enabled for the table.
E. The table is encrypted using a KMS customer managed key.

Answer: AB

NEW QUESTION 144
A ride-hailing application stores bookings in a persistent Amazon RDS for MySQL DB instance. This program is very popular, and the corporation anticipates a tenfold rise in the application's user base over the next several months. The application receives a higher volume of traffic in the morning and evening.
This application is divided into two sections:
An internal booking component that takes online reservations in response to concurrent user queries. A component of a third-party customer relationship management (CRM) system that customer service
professionals utilize. Booking data is accessed using queries in the CRM.
To manage this workload effectively, a database professional must create a cost-effective database system. Which solution satisfies these criteria?

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

Answer: B

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

NEW QUESTION 146
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 150
A company has a production Amazon Aurora Db cluster that serves both online transaction processing (OLTP) transactions and compute-intensive reports. The reports run for 10% of the total cluster uptime while the OLTP transactions run all the time. The company has benchmarked its workload and determined that a six-node Aurora DB cluster is appropriate for the peak workload.
The company is now looking at cutting costs for this DB cluster, but needs to have a sufficient number of nodes in the cluster to support the workload at different times. The workload has not changed since the previous benchmarking exercise.
How can a Database Specialist address these requirements with minimal user involvement?

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

Answer: D

NEW QUESTION 154
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 157
A retail company with its main office in New York and another office in Tokyo plans to build a database solution on AWS. The company's main workload consists of a mission-critical application that updates its application data in a data store. The team at the Tokyo office is building dashboards with complex analytical queries using the application data. The dashboards will be used to make buying decisions, so they need to have access to the application data in less than 1 second.
Which solution meets these requirements?

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

Answer: D

Explanation:
https://aws.amazon.com/blogs/database/aurora-postgresql-disaster-recovery-solutions-using-amazon-aurora-glob

NEW QUESTION 158
A company just migrated to Amazon Aurora PostgreSQL from an on-premises Oracle database. After the migration, the company discovered there is a period of time every day around 3:00 PM where the response time of the application is noticeably slower. The company has narrowed down the cause of this issue to the database and not the application.
Which set of steps should the Database Specialist take to most efficiently find the problematic PostgreSQL query?

A. Create an Amazon CloudWatch dashboard to show the number of connections, CPU usage, and disk space 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

NEW QUESTION 160
A small startup company is looking to migrate a 4 TB on-premises MySQL database to AWS using an Amazon RDS for MySQL DB instance.
Which strategy would allow for a successful migration with the LEAST amount of downtime?

A. Deploy a new RDS for MySQL DB instance and configure it for access from the on-premises data cente
B. Use the mysqldump utility to create an initial snapshot from the on-premises MySQL server, and copy it to an Amazon S3 bucke
C. Import the snapshot into the DB instance utilizing the MySQL utilities running on an Amazon EC2 instanc
D. Immediately point the application to the DB instance.
E. Deploy a new Amazon EC2 instance, install the MySQL software on the EC2 instance, and configure networking for access from the on-premises data cente
F. Use the mysqldump utility to create a snapshot of the on-premises MySQL serve
G. Copy the snapshot into the EC2 instance and restore it into the EC2 MySQL instanc
H. Use AWS DMS to migrate data into a new RDS for MySQL DB instanc
I. Point the application to the DB instance.
J. Deploy a new Amazon EC2 instance, install the MySQL software on the EC2 instance, and configure networking for access from the on-premises data cente
K. Use the mysqldump utility to create a snapshot of the on-premises MySQL serve
L. Copy the snapshot into an Amazon S3 bucket and import the snapshot into a new RDS for MySQL DB instance using the MySQL utilities running on an EC2 instanc
M. Point the application to the DB instance.
N. Deploy a new RDS for MySQL DB instance and configure it for access from the on-premises data cente
O. Use the mysqldump utility to create an initial snapshot from the on-premises MySQL server, and copy it to an Amazon S3 bucke
P. Import the snapshot into the DB instance using the MySQL utilities running on an Amazon EC2 instanc
Q. Establish replication into the new DB instance using MySQL replicatio
R. Stop application access to the on-premises MySQL server and let the remaining transactions replicate ove

S. Point the application to the DB instance.

Answer: B

NEW QUESTION 163
To meet new data compliance requirements, a company needs to keep critical data durably stored and readily accessible for 7 years. Data that is more than 1 year old is considered archival data and must automatically be moved out of the Amazon Aurora MySQL DB cluster every week. On average, around 10 GB of new data is added to the database every month. A database specialist must choose the most operationally efficient solution to migrate the archival data to Amazon S3. Which solution meets these requirements?

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

Answer: B

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

NEW QUESTION 167
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 170
A company is running a finance application on an Amazon RDS for MySQL DB instance. The application is governed by multiple financial regulatory agencies. The RDS DB instance is set up with security groups to allow access to certain Amazon EC2 servers only. AWS KMS is used for encryption at rest.
Which step will provide additional security?

A. Set up NACLs that allow the entire EC2 subnet to access the DB instance
B. Disable the master user account
C. Set up a security group that blocks SSH to the DB instance
D. Set up RDS to use SSL for data in transit

Answer: D

NEW QUESTION 173
A company recently acquired a new business. A database specialist must migrate an unencrypted 12 TB Amazon RDS for MySQL DB instance to a new AWS account. The database specialist needs to minimize the amount of time required to migrate the database.
Which solution meets these requirements?

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

Answer: A

Explanation:
Sharing an unencrypted manual DB snapshot enables authorized AWS accounts to directly restore a DB instance from the snapshot instead of taking a copy of it and restoring from that. https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ShareSnapshot.html However Resource Access Manager could not share non-Aurora cluster. https://docs.aws.amazon.com/ram/latest/userguide/shareable.html

NEW QUESTION 177

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

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

Answer: C

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

NEW QUESTION 178
A Database Specialist has migrated an on-premises Oracle database to Amazon Aurora PostgreSQL. The schema and the data have been migrated successfully. The on-premises database server was also being used to run database maintenance cron jobs written in Python to perform tasks including data purging and generating data exports. The logs for these jobs show that, most of the time, the jobs completed within 5 minutes, but a few jobs took up to 10 minutes to complete. These maintenance jobs need to be set up for Aurora PostgreSQL.
How can the Database Specialist schedule these jobs so the setup requires minimal maintenance and provides high availability?

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

Answer: C

Explanation:
https://docs.aws.amazon.com/AmazonCloudWatch/latest/events/Create-CloudWatch-Events-Scheduled-Rule.ht https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/schedule-jobs-for-amazon-rds-and-aurora-pos a job for data extraction or a job for data purging can easily be scheduled using cron. For these jobs, database credentials are typically either hard-coded or stored in a properties file. However, when you migrate to Amazon Relational Database Service (Amazon RDS) or Amazon Aurora PostgreSQL, you lose the ability to log in to the host instance to schedule cron jobs. This pattern describes how to use AWS Lambda and AWS Secrets Manager to schedule jobs for Amazon RDS and Aurora PostgreSQL databases after migration.
https://docs.aws.amazon.com/AmazonCloudWatch/latest/events/RunLambdaSchedule.html

NEW QUESTION 183
A company uses the Amazon DynamoDB table contractDB in us-east-1 for its contract system with the following schema:
orderID (primary key) timestamp (sort key) contract (map) createdBy (string) customerEmail (string)
After a problem in production, the operations team has asked a database specialist to provide an IAM policy to read items from the database to debug the application. In addition, the developer is not allowed to access the value of the customerEmail field to stay compliant.
Which IAM policy should the database specialist use to achieve these requirements?
A)

B)

C)

D)

A. Option A
B. Option B
C. Option C
D. Option C

Answer: A

NEW QUESTION 187
A company has an application that uses an Amazon DynamoDB table to store user data. Every morning, a single-threaded process calls the DynamoDB API Scan operation to scan the entire table and generate a critical start-of-day report for management. A successful marketing campaign recently doubled the number of items in the table, and now the process takes too long to run and the report is not generated in time.
A database specialist needs to improve the performance of the process. The database specialist notes that, when the process is running, 15% of the table's provisioned read capacity units (RCUs) are being used.
What should the database specialist do?

A. Enable auto scaling for the DynamoDB table.

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

Answer: B

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

NEW QUESTION 189
An online shopping company has a large inflow of shopping requests daily. As a result, there is a consistent load on the company's Amazon RDS database. A database specialist needs to ensure the database is up and running at all times. The database specialist wants an automatic notification system for issues that may cause database downtime or for configuration changes made to the database.
What should the database specialist do to achieve this? (Choose two.)

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

Answer: BD

NEW QUESTION 193
A corporation wishes to move a 1 TB Oracle database from its current location to an Amazon Aurora PostgreSQL DB cluster. The database specialist at the firm noticed that the Oracle database stores 100 GB of large binary objects (LOBs) across many tables. The Oracle database supports LOBs up to 500 MB in size and an average of 350 MB. AWS DMS was picked by the Database Specialist to transfer the data with the most replication instances.
How should the database specialist improve the transfer of the database to AWS DMS?

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

Answer: C

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

NEW QUESTION 196
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 201
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 204
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 209
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 the DynamoDB 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 211
A company has applications running on Amazon EC2 instances in a private subnet with no internet connectivity. The company deployed a new application that uses Amazon DynamoDB, but the application cannot connect to the DynamoDB tables. A developer already checked that all permissions are set correctly.
What should a database specialist do to resolve this issue while minimizing access to external resources?

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

Answer: D

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

NEW QUESTION 212
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 214
A bank intends to utilize Amazon RDS to host a MySQL database instance. The database should be able to handle high-volume read requests with extremely few repeated queries.
Which solution satisfies these criteria?

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

Answer: D

NEW QUESTION 216
A company is using Amazon with Aurora Replicas for read-only workload scaling. A Database Specialist needs to split up two read-only applications so each application always connects to a dedicated replica. The Database Specialist wants to implement load balancing and high availability for the read-only applications.
Which solution meets these requirements?

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

Answer: D

Explanation:
https://aws.amazon.com/about-aws/whats-new/2018/11/amazon-aurora-simplifies-workload-management-with-c

NEW QUESTION 218
A business's mission-critical production workload is being operated on a 500 GB Amazon Aurora MySQL DB cluster. A database engineer must migrate the workload without causing data loss to a new Amazon Aurora Serverless MySQL DB cluster.
Which approach will result in the LEAST amount of downtime and the LEAST amount of application impact?

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

Answer: D

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

NEW QUESTION 223
A Database Specialist is planning to create a read replica of an existing Amazon RDS for MySQL Multi-AZ DB instance. When using the AWS Management Console to conduct this task, the Database Specialist discovers that the source RDS DB instance does not appear in the read replica source selection box, so the read replica cannot be created.
What is the most likely reason for this?

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

Answer: D

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

NEW QUESTION 225
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 227
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 232
A company uses an Amazon RDS for PostgreSQL DB instance for its customer relationship management (CRM) system. New compliance requirements specify that the database must be encrypted at rest.
Which action will meet these requirements?

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

Answer: A

Explanation:
https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/encrypt-an-existing-amazon-rds-for-postgresq You can enable encryption for an Amazon RDS DB instance when you create it, but not after it's created.
However, you can add encryption to an unencrypted DB instance by creating a snapshot of your DB instance, and then creating an encrypted copy of that snapshot. You can then restore a DB instance from the encrypted snapshot to get an encrypted copy of your original DB instance. The pattern uses AWS Database Migration Service (AWS DMS) to migrate data and AWS Key Management Service (AWS KMS) for encryption.

NEW QUESTION 233
A financial organization must ensure that the most current 90 days of MySQL database backups are accessible. Amazon RDS for MySQL DB instances are used to host all MySQL databases. A database expert must create
a solution that satisfies the criteria for backup retention with the least amount of development work feasible. Which strategy should the database administrator take?

A. Use AWS Backup to build a backup plan for the required retention perio
B. Assign the DB instances to the backup plan.
C. Modify the DB instances to enable the automated backup optio
D. Select the required backup retention period.
E. Automate a daily cron job on an Amazon EC2 instance to create MySQL dumps, transfer to Amazon S3, and implement an S3 Lifecycle policy to meet the retention requirement.
F. Use AWS Lambda to schedule a daily manual snapshot of the DB instance
G. Delete snapshots that exceed the retention requirement.

Answer: A

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

NEW QUESTION 238
A clothing company uses a custom ecommerce application and a PostgreSQL database to sell clothes to thousands of users from multiple countries. The company is migrating its application and database from its on- premises data center to the AWS Cloud. The company has selected Amazon EC2 for the application and Amazon RDS for PostgreSQL for the database. The company requires database passwords to be changed every 60 days. A Database Specialist needs to ensure that the credentials used by the web application to connect to the database are managed securely.
Which approach should the Database Specialist take to securely manage the database credentials?

A. Store the credentials in a text file in an Amazon S3 bucke
B. Restrict permissions on the bucket to the IAM role associated with the instance profile onl
C. Modify the application to download the text file and retrieve the credentials on start u
D. Update the text file every 60 days.
E. Configure IAM database authentication for the application to connect to the databas
F. Create an IAM user and map it to a separate database user for each ecommerce use
G. Require users to update their passwords every 60 days.

H. Store the credentials in AWS Secrets Manage
I. Restrict permissions on the secret to only the IAM role associated with the instance profil
J. Modify the application to retrieve the credentials from Secrets Manager on start u
K. Configure the rotation interval to 60 days.
L. Store the credentials in an encrypted text file in the application AM
M. Use AWS KMS to store the key for decrypting the text fil
N. Modify the application to decrypt the text file and retrieve the credentials on start u
O. Update the text file and publish a new AMI every 60 days.

Answer: C

NEW QUESTION 243
.....

Thank You for Trying Our Product

* 100% Pass or Money Back

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

* One year free update

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

* Trusted by Millions

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

* Shop Securely

All transactions are protected by VeriSign!

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

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