

Amazon-Web-Services

Exam Questions DBS-C01

AWS Certified Database - Specialty



NEW QUESTION 1

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

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

Answer: A

NEW QUESTION 2

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

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

Answer: C

Explanation:

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

NEW QUESTION 3

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

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

Answer: D

Explanation:

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

NEW QUESTION 4

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

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

Answer: A

Explanation:

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

NEW QUESTION 5

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

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

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

Answer: CDE

NEW QUESTION 6

A financial company is running an Amazon Redshift cluster for one of its data warehouse solutions. The company needs to generate connection logs, user logs, and user activity logs. The company also must make these logs available for future analysis.

Which combination of steps should a database specialist take to meet these requirements? (Choose two.)

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

Answer: AD

Explanation:

AWS CloudWatch Logs are stored indefinitely and CloudWatch Log Insights is used to analyze the logs and query upon them.

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/AnalyzingLogData.html>

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/WhatIsCloudWatchLogs.html>

"Log retention – By default, logs are kept indefinitely and never expire. You can adjust the retention policy for each log group, keeping the indefinite retention, or choosing a retention period between 10 years and one day."

<https://docs.aws.amazon.com/redshift/latest/mgmt/db-auditing.html>

NEW QUESTION 7

A Database Specialist is migrating a 2 TB Amazon RDS for Oracle DB instance to an RDS for PostgreSQL DB instance using AWS DMS. The source RDS Oracle DB instance is in a VPC in the us-east-1 Region. The target RDS for PostgreSQL DB instance is in a VPC in the us-west-2 Region.

Where should the AWS DMS replication instance be placed for the MOST optimal performance?

- A. In the same Region and VPC of the source DB instance
- B. In the same Region and VPC as the target DB instance
- C. In the same VPC and Availability Zone as the target DB instance
- D. In the same VPC and Availability Zone as the source DB instance

Answer: C

Explanation:

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_ReplicationInstance.VPC.html#CHAP_ReplicationInstance In fact, all the configurations list on above url prefer the replication instance putting into target vpc region / subnet / az.

https://docs.aws.amazon.com/dms/latest/sbs/CHAP_SQLServer2Aurora.Steps.CreateReplicationInstance.html

NEW QUESTION 8

An Amazon RDS EBS-optimized instance with Provisioned IOPS (PIOPS) storage is using less than half of its allocated IOPS over the course of several hours under constant load. The RDS instance exhibits multi-second read and write latency, and uses all of its maximum bandwidth for read throughput, yet the instance uses less than half of its CPU and RAM resources.

What should a Database Specialist do in this situation to increase performance and return latency to sub-second levels?

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

Answer: D

Explanation:

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

NEW QUESTION 9

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

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

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

Answer: CD

Explanation:

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

NEW QUESTION 10

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

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

Answer: B

Explanation:

<https://docs.aws.amazon.com/prescriptive-guidance/latest/migration-sql-server/ec2-fci.html>

An FCI is generally preferable over an Always on availability group when: You're using SQL Server Standard edition instead of Enterprise edition.

NEW QUESTION 10

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

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

Answer: A

Explanation:

<https://aws.amazon.com/blogs/mt/integrating-aws-cloudformation-with-aws-systems-manager-parameter-store/>

NEW QUESTION 15

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

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

Answer: A

Explanation:

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Target.S3.html Using Amazon S3 as a target for AWS Database Migration Service

NEW QUESTION 17

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

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

Answer: C

NEW QUESTION 20

An ecommerce company uses a backend application that stores data in an Amazon DynamoDB table. The backend application runs in a private subnet in a VPC and must connect to this table.

The company must minimize any network latency that results from network connectivity issues, even during periods of heavy application usage. A database administrator also needs the ability to use a private connection to connect to the DynamoDB table from the application.

Which solution will meet these requirements?

- A. Use network ACLs to ensure that any outgoing or incoming connections to any port except DynamoDB are deactivate
- B. Encrypt API calls by using TLS.
- C. Create a VPC endpoint for DynamoDB in the application's VP
- D. Use the VPC endpoint to access the table.
- E. Create an AWS Lambda function that has access to DynamoD
- F. Restrict outgoing access only to this Lambda function from the application.
- G. Use a VPN to route all communication to DynamoDB through the company's own corporate network infrastructure.

Answer: B

Explanation:

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

NEW QUESTION 24

A Database Specialist is performing a proof of concept with Amazon Aurora using a small instance to confirm a simple database behavior. When loading a large dataset and creating the index, the Database Specialist encounters the following error message from Aurora:

ERROR: cloud not write block 7507718 of temporary file: No space left on device

What is the cause of this error and what should the Database Specialist do to resolve this issue?

- A. The scaling of Aurora storage cannot catch up with the data loadin
- B. The Database Specialist needs to modify the workload to load the data slowly.
- C. The scaling of Aurora storage cannot catch up with the data loadin
- D. The Database Specialist needs to enable Aurora storage scaling.
- E. The local storage used to store temporary tables is ful
- F. The Database Specialist needs to scale up the instance.
- G. The local storage used to store temporary tables is ful
- H. The Database Specialist needs to enable local storage scaling.

Answer: C

NEW QUESTION 26

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 30

An worldwide gaming company's development team is experimenting with using Amazon DynamoDB to store in-game events for three mobile titles. Maximum concurrent users for the most popular game is 500,000, while the least popular game is 10,000. The typical event is 20 KB in size, while the average user session generates one event each second. Each event is assigned a millisecond time stamp and a globally unique identification.

The lead developer generated a single DynamoDB database with the following structure for the events:

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

In a small-scale development setting, the tests were successful. When the application was deployed to production, however, new events were not being added to the database, and the logs indicated DynamoDB failures with the `ItemCollectionSizeLimitExceededException` issue code.

Which design modification should a database professional offer to the development team?

- A. Use the player identifier as the partition ke
- B. Use the event time as the sort ke

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

Answer: D

NEW QUESTION 32

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 35

A business's production database is hosted on a single-node Amazon RDS for MySQL DB instance. The database instance is hosted in a United States AWS Region.

A week before a significant sales event, a fresh database maintenance update is released. The maintenance update has been designated as necessary. The firm want to minimize the database instance's downtime and requests that a database expert make the database instance highly accessible until the sales event concludes.

Which solution will satisfy these criteria?

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

Answer: D

Explanation:

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

NEW QUESTION 39

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 tomove 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 44

A company is building a software as a service application. As part of the new user sign-on workflow, a Python script invokes the CreateTable operation using the

Amazon DynamoDB API. After the call returns, the script attempts to call PutItem.

Occasionally, the PutItem request fails with a ResourceNotFoundException error, which causes the workflow to fail. The development team has confirmed that the same table name is used in the two API calls.

How should a database specialist fix this issue?

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

Answer: C

Explanation:

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

NEW QUESTION 45

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 49

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 center
- B. Use the mysqldump utility to create an initial snapshot from the on-premises MySQL server, and copy it to an Amazon S3 bucket
- C. Import the snapshot into the DB instance utilizing the MySQL utilities running on an Amazon EC2 instance
- 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 center
- F. Use the mysqldump utility to create a snapshot of the on-premises MySQL server
- G. Copy the snapshot into the EC2 instance and restore it into the EC2 MySQL instance
- H. Use AWS DMS to migrate data into a new RDS for MySQL DB instance
- 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 center
- K. Use the mysqldump utility to create a snapshot of the on-premises MySQL server
- 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 instance
- 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 center
- O. Use the mysqldump utility to create an initial snapshot from the on-premises MySQL server, and copy it to an Amazon S3 bucket
- P. Import the snapshot into the DB instance using the MySQL utilities running on an Amazon EC2 instance
- Q. Establish replication into the new DB instance using MySQL replication
- R. Stop application access to the on-premises MySQL server and let the remaining transactions replicate over
- 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 53

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

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

- A. Perform a logical dump from the Db2 database and restore it to an Aurora DB cluster
- 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 cluster
- 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 55

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

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

Answer: D

Explanation:

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

NEW QUESTION 56

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

- A. Use an Amazon RDS DB instance deployed in the us-east-1 Region with a read replica instance in the ap-northeast-1 Region
- 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 Region
- 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 Region
- F. Have the dashboard application read from the read replica.
- G. Use an Amazon Aurora global database
- H. Deploy the writer instance in the us-east-1 Region and the replica in the ap-northeast-1 Region
- I. Have the dashboard application read from the replica in the ap-northeast-1 Region.

Answer: D

Explanation:

<https://aws.amazon.com/blogs/database/aurora-postgresql-disaster-recovery-solutions-using-amazon-aurora-global>

NEW QUESTION 59

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

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

Answer: C

Explanation:

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

NEW QUESTION 64

An information management services company is storing JSON documents on premises. The company is using a MongoDB 3.6 database but wants to migrate to AWS. The solution must be compatible, scalable, and fully managed. The solution also must result in as little downtime as possible during the migration. Which solution meets these requirements?

- A. Create an AWS Database Migration Service (AWS DMS) replication instance, a source endpoint for MongoDB, and a target endpoint of Amazon DocumentDB (with MongoDB compatibility).
- B. Create an AWS Database Migration Service (AWS DMS) replication instance, a source endpoint for MongoDB, and a target endpoint of a MongoDB image that is hosted on Amazon EC2
- C. Use the mongodump and mongorestore tools to migrate the data from the source MongoDB deployment to Amazon DocumentDB (with MongoDB compatibility).
- D. Use the mongodump and mongorestore tools to migrate the data from the source MongoDB deployment to a MongoDB image that is hosted on Amazon EC2.

Answer: A

Explanation:

<https://docs.aws.amazon.com/documentdb/latest/developerguide/docdb-migration.html#docdb-migration-approach>

NEW QUESTION 67

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 period.
- B. Assign the DB instances to the backup plan.
- C. Modify the DB instances to enable the automated backup option.
- 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 70

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

- A. Create an Aurora Replica with encryption enabled using AWS Key Management Service (AWS KMS). Then promote the replica to master.
- B. Use SSL/TLS to secure the in-transit connection between the financial application and the Aurora DB cluster.
- C. Modify the existing Aurora DB cluster and enable encryption using an AWS Key Management Service (AWS KMS) encryption key.
- 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 key.
- 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 73

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

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

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

Answer: C

NEW QUESTION 78

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 queue.
- 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 83

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

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

Answer: B

Explanation:

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

NEW QUESTION 86

A manufacturing company's website uses an Amazon Aurora PostgreSQL DB cluster. Which configurations will result in the LEAST application downtime during a failover? (Choose three.)

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

Answer: ABC

NEW QUESTION 87

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

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

Answer: D

Explanation:

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

NEW QUESTION 89

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

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

Answer: A

Explanation:

[https://aws.amazon.com/fr/blogs/database/best-practices-for-configuring-parameters-for-amazon-rds-for-mysql-](https://aws.amazon.com/fr/blogs/database/best-practices-for-configuring-parameters-for-amazon-rds-for-mysql/) "You can set parameters globally using a parameter group. Alternatively, you can set them for a particular session using the SET command."
<https://aws.amazon.com/blogs/database/best-practices-for-configuring-parameters-for-amazon-rds-for-mysql-pa>

NEW QUESTION 91

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

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

Answer: B

Explanation:

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

NEW QUESTION 92

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 97

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

Which approach should the Database Specialist take?

- A. Dump all the tables from the Oracle database into an Amazon S3 bucket using datapump (expdp). Run data transformations in AWS Glue.
- 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 appliance.
- D. Once the Snowball data is delivered to Amazon S3, create a new Aurora DB cluster.
- 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 migration.
- 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 instance.
- 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 98

A company is using a Single-AZ Amazon RDS for MySQL DB instance for development. The DB instance is experiencing slow performance when queries are executed. Amazon CloudWatch metrics indicate that the instance requires more I/O capacity.

Which actions can a database specialist perform to resolve this issue? (Choose two.)

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

Answer: BE

Explanation:

<https://aws.amazon.com/blogs/database/best-storage-practices-for-running-production-workloads-on-hosted-data>

"If you find the pattern of IOPS usage consistently going beyond more than 16,000, you should modify the DB instance and change the storage type from gp2 to io1.

NEW QUESTION 100

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 102

A large retail company recently migrated its three-tier ecommerce applications to AWS. The company's backend database is hosted on Amazon Aurora PostgreSQL. During peak times, users complain about longer page load times. A database specialist reviewed Amazon RDS Performance Insights and found a spike in IO:WaitSync wait events. The SQL attached to the wait events are all single INSERT statements.

How should this issue be resolved?

- A. Modify the application to commit transactions in batches
- B. Add a new Aurora Replica to the Aurora DB cluster.
- C. Add an Amazon ElastiCache for Redis cluster and change the application to write through.
- D. Change the Aurora DB cluster storage to Provisioned IOPS (PIOPS).

Answer: A

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.Reference.html> "This wait most often arises when there is a very high rate of commit activity on the system. You can sometimes alleviate this wait by modifying applications to commit transactions in batches. "
<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/apg-waits.xactsync.html>

NEW QUESTION 103

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

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

Answer: C

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/dms-disable-public-access/>

NEW QUESTION 104

A company has a production environment running on Amazon RDS for SQL Server with an in-house web application as the front end. During the last application maintenance window, new functionality was added to the web application to enhance the reporting capabilities for management. Since the update, the application is slow to respond to some reporting queries. How should the company identify the source of the problem?

- A. Install and configure Amazon CloudWatch Application Insights for Microsoft .NET and Microsoft SQL Serve
- B. Use a CloudWatch dashboard to identify the root cause.
- C. Enable RDS Performance Insights and determine which query is creating the problem
- D. Request changes to the query to address the problem.
- E. Use AWS X-Ray deployed with Amazon RDS to track query system traces.
- F. Create a support request and work with AWS Support to identify the source of the issue.

Answer: B

Explanation:

Amazon RDS Performance Insights is a database performance tuning and monitoring feature that helps you quickly assess the load on your database, and determine when and where to take action. Performance Insights allows non-experts to detect performance problems with an easy-to-understand dashboard that visualizes database load. <https://aws.amazon.com/rds/performance-insights/>

NEW QUESTION 105

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

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

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

Answer: ABC

Explanation:

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

NEW QUESTION 106

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

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

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

- C. Only allow incoming traffic from the subnet of the application servers on port 3306.
- D. Only allow incoming traffic from the subnet of the application servers on port 443.

Answer: A

Explanation:

most restrictive approach is to allow only incoming connections from SG of EC2 instance on port 3306

NEW QUESTION 110

A company requires near-real-time notifications when changes are made to Amazon RDS DB security groups. Which solution will meet this requirement with the LEAST operational overhead?

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

Answer: A

Explanation:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_Events.Messages.html#USER_Events.Mes

NEW QUESTION 113

A business is transferring its on-premises database workloads to the Amazon Web Services (AWS) Cloud. A database professional migrating an Oracle database with a huge table to Amazon RDS has picked AWS DMS. The database professional observes that AWS DMS is consuming considerable time migrating the data. Which activities would increase the pace of data migration? (Select three.)

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

Answer: ACD

Explanation:

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Tasks.LOBSupport.html

NEW QUESTION 116

A company is using an Amazon Aurora MySQL database with Performance Insights enabled. A database specialist is checking Performance Insights and observes an alert message that starts with the following phrase:

`Performance Insights is unable to collect SQL Digest statistics on new queries`!

Which action will resolve this alert message?

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

Answer: A

Explanation:

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

NEW QUESTION 117

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

- A. Install the AWS CLI on an Amazon EC2 instance
- B. Write a CLI command that creates a backup of the DynamoDB table
- 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 table
- 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 119

A Database Specialist needs to speed up any failover that might occur on an Amazon Aurora PostgreSQL DB cluster. The Aurora DB cluster currently includes the

primary instance and three Aurora Replicas.

How can the Database Specialist ensure that failovers occur with the least amount of downtime for the application?

- A. Set the TCP keepalive parameters low
- B. Call the AWS CLI failover-db-cluster command
- C. Enable Enhanced Monitoring on the DB cluster
- D. Start a database activity stream on the DB cluster

Answer: A

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.BestPractices.html#Aur>

NEW QUESTION 122

A company's database specialist is building an Amazon RDS for Microsoft SQL Server DB instance to store hundreds of records in CSV format. A customer service tool uploads the records to an Amazon S3 bucket.

An employee who previously worked at the company already created a custom stored procedure to map the necessary CSV fields to the database tables. The database specialist needs to implement a solution that reuses this previous work and minimizes operational overhead.

Which solution will meet these requirements?

- A. Create an Amazon S3 event to invoke an AWS Lambda function
- B. Configure the Lambda function to parse the .csv file and use a SQL client library to run INSERT statements to load the data into the tables.
- C. Write a custom .NET app that is hosted on Amazon EC2. Configure the .NET app to load the .csv file and call the custom stored procedure to insert the data into the tables.
- D. Download the .csv file from Amazon S3 to the RDS D drive by using an AWS msdb stored procedure. Call the custom stored procedure to insert the data from the RDS D drive into the tables.
- E. Create an Amazon S3 event to invoke AWS Step Functions to parse the .csv file and call the custom stored procedure to insert the data into the tables.

Answer: C

Explanation:

Step 1: Download S3 Files

Amazon RDS for SQL Server comes with several custom stored procedures and functions. These are located in the msdb database. The stored procedure to download files from S3 is "rds_download_from_s3". The syntax for this stored procedure is shown here:

```
exec msdb.dbo.rds_download_from_s3
@s3_arn_of_file='arn:aws:s3:::<bucket_name>/<file_name>',
@rds_file_path='D:\S3\<custom_folder_name>\<file_name>',
@overwrite_file=1;
```

NEW QUESTION 124

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

Which AWS solution meets these requirements?

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

Answer: D

Explanation:

RDS event subscriptions do not cover "data is inserted into a table" - see

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/USER_Events.Messages.html We can use stored procedure to invoke Lambda function -

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Integrating.Lambda.html>

NEW QUESTION 129

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

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

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

Answer: C

Explanation:

You can restore a DB instance to a specific point in time (PITR), creating a new DB instance. To support PITR, your DB instances must have backup retention set to a nonzero value. <https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/custom-backup-sqlserver.html>

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

<https://aws.amazon.com/blogs/storage/point-in-time-recovery-and-continuous-backup-for-amazon-rds-with-aws>

NEW QUESTION 131

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

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 135

A company's application development team wants to share an automated snapshot of its Amazon RDS database with another team. The database is encrypted with a custom AWS Key Management Service (AWS KMS) key under the "WeShare" AWS account. The application development team needs to share the DB snapshot under the "WeReceive" AWS account.

Which combination of actions must the application development team take to meet these requirements? (Choose two.)

- A. Add access from the "WeReceive" account to the custom AWS KMS key policy of the sharing team.
- B. Make a copy of the DB snapshot, and set the encryption option to disable.
- C. Share the DB snapshot by setting the DB snapshot visibility option to public.
- D. Make a copy of the DB snapshot, and set the encryption option to enable.
- E. Share the DB snapshot by using the default AWS KMS encryption key.

Answer: AD

Explanation:

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

NEW QUESTION 140

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

Which approach meets these requirements with no negative performance impact?

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

Answer: C

Explanation:

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

NEW QUESTION 145

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 148

A company is running an Amazon RDS for PostgreSQL DB instance and wants to migrate it to an Amazon Aurora PostgreSQL DB cluster. The current database is 1 TB in size. The migration needs to have minimal downtime.

What is the FASTEST way to accomplish this?

- A. Create an Aurora PostgreSQL DB 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 an Aurora 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 153

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 156

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

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

Answer: C

NEW QUESTION 158

A business is launching a new Amazon RDS for SQL Server database instance. The organization wishes to allow auditing of the SQL Server database. Which measures should a database professional perform in combination to achieve this requirement? (Select two.)

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

Answer: AE

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Appendix.SQLServer.Options.Audit.html>
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/security_iam_service-with-iam.html

NEW QUESTION 159

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 162

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 164

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 168

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 172

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

Which database solution meets these requirements at the LOWEST cost?

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

Answer: D

Explanation:

Known patterns, minimum maintenance, miliseconds latency

NEW QUESTION 175

A company has migrated a single MySQL database to Amazon Aurora. The production data is hosted in a DB cluster in VPC_PROD, and 12 testing environments are hosted in VPC_TEST using the same AWS account. Testing results in minimal changes to the test data. The Development team wants each environment refreshed nightly so each test database contains fresh production data every day.

Which migration approach will be the fastest and most cost-effective to implement?

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

Answer: A

NEW QUESTION 177

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

Which of the following will resolve this issue?

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

Answer: B

Explanation:

Max_connection is a formula in RDS parameter group:

$\text{GREATEST}(\{\log(\text{DBInstanceClassMemory}/805306368)*45\}, \{\log(\text{DBInstanceClassMemory}/8187281408)*100\})$

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Managing.Performance.htm> You can increase the maximum number of connections to your Aurora MySQL DB instance by scaling the instance up to a DB instance class with more memory, or by setting a larger value for the max_connections parameter in the DB parameter group for your instance, up to 16,000. You must change a larger value for the max_connections parameter in the DB parameter group, not edit my.cnf, it is not physical server hosting MySQL.

NEW QUESTION 180

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

Which solution addresses these requirements?

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

Answer: D

Explanation:

PostgreSQL: sslrootcert=rds-cert.pem sslmode=[verify-ca | verify-full]

NEW QUESTION 183

A company has a production Amazon Aurora Db cluster that serves both online transaction processing (OLTP) transactions and compute-intensive reports. The reports run for 10% of the total cluster uptime while the OLTP transactions run all the time. The company has benchmarked its workload and determined that a six-node Aurora DB cluster is appropriate for the peak workload.

The company is now looking at cutting costs for this DB cluster, but needs to have a sufficient number of nodes in the cluster to support the workload at different times. The workload has not changed since the previous benchmarking exercise.

How can a Database Specialist address these requirements with minimal user involvement?

- A. Split up the DB cluster into two different clusters: one for OLTP and the other for reporting
- B. Monitor and set up replication between the two clusters to keep data consistent.
- C. Review and evaluate the peak combined workload

- D. Ensure that utilization of the DB cluster node is at an acceptable level.
- 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 workload.
- G. The cluster can be restarted again depending on the workload at the time.
- H. Set up automatic scaling on the DB cluster.
- I. This will allow the number of reader nodes to adjust automatically to the reporting workload, when needed.

Answer: D

NEW QUESTION 185

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

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

Answer: C

NEW QUESTION 190

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

Which approach will meet these requirements?

- A. Use an Amazon RDS DB instance.
- B. Shut down the instance once the data has been read.
- C. Use Amazon Aurora Serverless.
- 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 193

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

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

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

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

Answer: B

NEW QUESTION 194

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

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

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

F. Choose a preferred maintenance window at the time of lowest usage to apply any required changes and mandatory updates.

Answer: AE

Explanation:

<https://aws.amazon.com/blogs/database/configuring-amazon-elasticache-for-redis-for-higher-availability/>

NEW QUESTION 195

A team of Database Specialists is currently investigating performance issues on an Amazon RDS for MySQL DB instance and is reviewing related metrics. The team wants to narrow the possibilities down to specific database wait events to better understand the situation.

How can the Database Specialists accomplish this?

- A. Enable the option to push all database logs to Amazon CloudWatch for advanced analysis
- B. Create appropriate Amazon CloudWatch dashboards to contain specific periods of time
- C. Enable Amazon RDS Performance Insights and review the appropriate dashboard
- D. Enable Enhanced Monitoring with the appropriate settings

Answer: C

Explanation:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_PerfInsights.Enabling.html <https://aws.amazon.com/rds/performance-insights/>

<https://aws.amazon.com/blogs/database/tuning-amazon-rds-for-mysql-with-performance-insights/>

NEW QUESTION 197

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

- A. Create a custom script that exports archival data from the DB cluster to Amazon S3 using a SQL view, then deletes the archival data from the DB cluster
- 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 cluster
- 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 cluster
- 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 199

On AWS, a business is developing a web application. The application needs that the database supports concurrent read and write activities in several AWS Regions. Additionally, the database must communicate data changes across Regions as they occur. The application must be highly available and have a latency of less than a few hundred milliseconds.

Which solution satisfies these criteria?

- A. Amazon DynamoDB global tables
- B. Amazon DynamoDB streams with AWS Lambda to replicate the data
- C. An Amazon ElastiCache for Redis cluster with cluster mode enabled and multiple shards
- D. An Amazon Aurora global database

Answer: A

Explanation:

Aurora Global Databases provides a writer and a reader endpoints in the primary region but only a reader endpoints in other region. Although strongly consistent, it does not fulfill the requirements that "there are plenty of read / write activities" in all regions.

NEW QUESTION 201

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 202

A company migrated one of its business-critical database workloads to an Amazon Aurora Multi-AZ DB cluster. The company requires a very low RTO and needs to improve the application recovery time after database failovers.

Which approach meets these requirements?

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

Answer: C

Explanation:

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

<https://aws.amazon.com/rds/proxy/>

NEW QUESTION 204

A business need a data warehouse system that stores data consistently and in a highly organized fashion. The organization demands rapid response times for end-user inquiries including current-year data, and users must have access to the whole 15-year dataset when necessary. Additionally, this solution must be able to manage a variable volume of incoming inquiries. Costs associated with storing the 100 TB of data must be maintained to a minimum.

Which solution satisfies these criteria?

- A. Leverage an Amazon Redshift data warehouse solution using a dense storage instance type while keeping all the data on local Amazon Redshift storage
- 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 data
- D. Keep historical data on Amazon S3 and access it using the Amazon Redshift Spectrum layer
- 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 data
- G. Keep historical data on Amazon S3 and access it using the Amazon Redshift Spectrum layer
- H. Enable Amazon Redshift Concurrency Scaling.
- I. Leverage an Amazon Redshift data warehouse solution using a dense storage instance to store the most recent data
- J. Keep historical data on Amazon S3 and access it using the Amazon Redshift Spectrum layer
- 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 208

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

Which solution will meet these requirements?

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

Answer: B

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Overview.Encryption.html>

NEW QUESTION 210

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

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

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

Answer: B

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Managing.FaultInjectionQu> "You can test the fault tolerance of your Amazon Aurora DB cluster by using fault injection queries. Fault injection queries are issued as SQL commands to an Amazon Aurora instance and they enable you to schedule a simulated occurrence of one of the following events: A crash of a writer or reader DB instance A failure of an Aurora Replica A disk failure Disk congestion When a fault injection query specifies a crash, it forces a crash of the Aurora DB instance. The other fault injection queries result in simulations of failure events, but don't cause the event to occur. When you submit a fault injection query, you also specify an amount of time for the failure event simulation to occur for."

NEW QUESTION 213

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 214

Developers have requested a new Amazon Redshift cluster so they can load new third-party marketing data. The new cluster is ready and the user credentials are given to the developers. The developers indicate that their copy jobs fail with the following error message:

"Amazon Invalid operation: S3ServiceException:Access Denied,Status 403,Error AccessDenied."

The developers need to load this data soon, so a database specialist must act quickly to solve this issue. What is the MOST secure solution?

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

Answer: B

Explanation:

<https://docs.aws.amazon.com/redshift/latest/gsg/rs-gsg-create-an-iam-role.html>

"Now that you have created the new role, your next step is to attach it to your cluster. You can attach the role when you launch a new cluster or you can attach it to an existing cluster. In the next step, you attach the role to a new cluster."

https://docs.aws.amazon.com/redshift/latest/dg/copy-usage_notes-access-permissions.html

NEW QUESTION 216

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 219

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

What should the database specialist do to resolve this?

- A. Disable automated backups in the DB instanc
- B. Share both the automated snapshot and the default KMS key with the staging accoun
- C. Restore the snapshot in the staging account and enable automated backups.

- D. Copy the automated snapshot specifying a custom KMS encryption ke
- E. Share both the copied snapshot and the custom KMS encryption key with the staging accoun
- F. Restore the snapshot to the staging account within the same Region.
- G. Modify the DB instance to use a custom KMS encryption ke
- H. Share both the automated snapshot and the custom KMS encryption key with the staging accoun
- I. Restore the snapshot in the staging account.
- J. Copy the automated snapshot while keeping the default KMS ke
- K. Share both the snapshot and the default KMS key with the staging accoun
- L. Restore the snapshot in the staging account.

Answer: B

Explanation:

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

NEW QUESTION 224

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

DBS-C01 Practice Exam Features:

- * DBS-C01 Questions and Answers Updated Frequently
- * DBS-C01 Practice Questions Verified by Expert Senior Certified Staff
- * DBS-C01 Most Realistic Questions that Guarantee you a Pass on Your First Try
- * DBS-C01 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The DBS-C01 Practice Test Here](#)