

Amazon-Web-Services

Exam Questions DBS-C01

AWS Certified Database - Specialty



NEW QUESTION 1

A database specialist is constructing an AWS CloudFormation stack using AWS CloudFormation. The database expert wishes to avoid the stack's Amazon RDS ProductionDatabase resource being accidentally deleted.

Which solution will satisfy this criterion?

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

Answer: D

Explanation:

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

NEW QUESTION 2

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

How should a Database Specialist ensure DynamoDB can handle the increased traffic?

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

Answer: D

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/bp-partition-key-design.html#bp-partition>

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

NEW QUESTION 3

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

Which migration strategy will be the quickest and least expensive to implement?

- A. Run the master in Amazon Aurora 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

Explanation:

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

NEW QUESTION 4

A company's database specialist implements an AWS Database Migration Service (AWS DMS) task for change data capture (CDC) to replicate data from an on-premises Oracle database to Amazon S3. When usage of the company's application increases, the database specialist notices multiple hours of latency with the CDC.

Which solutions will reduce this latency? (Choose two.)

- A. Configure the DMS task to run in full large binary object (LOB) mode.
- B. Configure the DMS task to run in limited large binary object (LOB) mode.
- C. Create a Multi-AZ replication instance.
- D. Load tables in parallel by creating multiple replication instances for sets of tables that participate in common transactions.
- E. Replicate tables in parallel by creating multiple DMS tasks for sets of tables that do not participate in common transactions.

Answer: BE

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

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

Answer: AD

Explanation:

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

NEW QUESTION 7

A company is using an Amazon Aurora PostgreSQL DB cluster with an xlarge primary instance master and two large Aurora Replicas for high availability and read-only workload scaling. A failover event occurs and application performance is poor for several minutes. During this time, application servers in all Availability Zones are healthy and responding normally.

What should the company do to eliminate this application performance issue?

- A. Configure both of the Aurora Replicas to the same instance class as the primary DB instance.
- B. Enable cache coherence on the DB cluster, set the primary DB instance failover priority to tier-0, and assign a failover priority of tier-1 to the replicas.
- C. Deploy an AWS Lambda function that calls the DescribeDBInstances action to establish which instance has failed, and then use the PromoteReadReplica operation to promote one Aurora Replica to be the primary DB instance.
- D. Configure an Amazon RDS event subscription to send a notification to an Amazon SNS topic to which the Lambda function is subscribed.
- E. Configure one Aurora Replica to have the same instance class as the primary DB instance.
- F. Implement Aurora PostgreSQL DB cluster cache management.
- G. Set the failover priority to tier-0 for the primary DB instance and one replica with the same instance class.
- H. Set the failover priority to tier-1 for the other replicas.
- I. Configure both Aurora Replicas to have the same instance class as the primary DB instance.
- J. Implement Aurora PostgreSQL DB cluster cache management.
- K. Set the failover priority to tier-0 for the primary DB instance and to tier-1 for the replicas.

Answer: C

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.cluster-cache-mgmt.htm>

<https://aws.amazon.com/blogs/database/introduction-to-aurora-postgresql-cluster-cache-management/>

"You can customize the order in which your Aurora Replicas are promoted to the primary instance after a failure by assigning each replica a priority. Priorities range from 0 for the first priority to 15 for the last priority. If the primary instance fails, Amazon RDS promotes the Aurora Replica with the better priority to the new primary instance. You can modify the priority of an Aurora Replica at any time. Modifying the priority doesn't trigger a failover. More than one Aurora Replica can share the same priority, resulting in promotion tiers. If two or more Aurora Replicas share the same priority, then Amazon RDS promotes the replica that is largest in size. If two or more Aurora Replicas share the same priority and size, then Amazon RDS promotes an arbitrary replica in the same promotion tier. "

Amazon Aurora with PostgreSQL compatibility now supports cluster cache management, providing a faster path to full performance if there's a failover. With cluster cache management, you designate a specific reader DB instance in your Aurora PostgreSQL cluster as the failover target. Cluster cache management keeps the data in the designated reader's cache synchronized with the data in the read-write instance's cache. If a failover occurs, the designated reader is promoted to be the new read-write instance, and workloads benefit immediately from the data in its cache.

NEW QUESTION 8

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 9

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

Which of the following provides the MOST cost-effective solution?

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

Answer: B

Explanation:

Aurora Serverless is not compatible to all Aurora provisioned engine version. However, you can do clone with most engine version. Meanwhile, I also consider the performance while restoring snapshot to Aurora Serverless.

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-serverless-how-it-works.html#aurora>

[https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-serverless.html#aurora-serverless.us](https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-serverless.html#aurora-serverless-us)

NEW QUESTION 10

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 10

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 VPC
- D. Use the VPC endpoint to access the table.
- E. Create an AWS Lambda function that has access to DynamoDB
- 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 11

A business is transferring a database from one AWS Region to another using an Amazon RDS for SQL Server DB instance. The organization wishes to keep database downtime to a minimum throughout the transfer.

Which migration strategy should the organization use for this cross-regional move?

- A. Back up the source database using native backup to an Amazon S3 bucket in the same Region
- B. Then restore the backup in the target Region.
- C. Back up the source database using native backup to an Amazon S3 bucket in the same Region
- D. Use Amazon S3 Cross-Region Replication to copy the backup to an S3 bucket in the target Region
- E. Then restore the backup in the target Region.
- F. Configure AWS Database Migration Service (AWS DMS) to replicate data between the source and the target database
- G. Once the replication is in sync, terminate the DMS task.

- H. Add an RDS for SQL Server cross-Region read replica in the target Region
- I. Once the replication is in sync, promote the read replica to master.

Answer: C

Explanation:

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

With Amazon RDS, you can create a MariaDB, MySQL, Oracle, or PostgreSQL read replica in a different AWS Region from the source DB instance. Creating a cross-Region read replica isn't supported for SQL Server on Amazon RDS.

NEW QUESTION 12

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 full
- F. The Database Specialist needs to scale up the instance.
- G. The local storage used to store temporary tables is full
- H. The Database Specialist needs to enable local storage scaling.

Answer: C

NEW QUESTION 16

A Database Specialist is troubleshooting an application connection failure on an Amazon Aurora DB cluster with multiple Aurora Replicas that had been running with no issues for the past 2 months. The connection failure lasted for 5 minutes and corrected itself after that. The Database Specialist reviewed the Amazon RDS events and determined a failover event occurred at that time. The failover process took around 15 seconds to complete.

What is the MOST likely cause of the 5-minute connection outage?

- A. After a database crash, Aurora needed to replay the redo log from the last database checkpoint
- B. The client-side application is caching the DNS data and its TTL is set too high
- C. After failover, the Aurora DB cluster needs time to warm up before accepting client connections
- D. There were no active Aurora Replicas in the Aurora DB cluster

Answer: B

Explanation:

When your application tries to establish a connection after a failover, the new Aurora PostgreSQL writer will be a previous reader, which can be found using the Aurora read only endpoint before DNS updates have fully propagated. Setting the java DNS TTL to a low value helps cycle between reader nodes on subsequent connection attempts.

Amazon Aurora is designed to recover from a crash almost instantaneously and continue to serve your application data. Unlike other databases, after a crash Amazon Aurora does not need to replay the redo log from the last database checkpoint before making the database available for operations. Amazon Aurora performs crash recovery asynchronously on parallel threads, so your database is open and available immediately after a crash. Because the storage is organized in many small segments, each with its own redo log, the underlying storage can replay redo records on demand in parallel and asynchronously as part of a disk read after a crash. This approach reduces database restart times to less than 60 seconds in most cases

NEW QUESTION 20

A gaming company is developing a new mobile game and decides to store the data for each user in Amazon DynamoDB. To make the registration process as easy as possible, users can log in with their existing Facebook or Amazon accounts. The company expects more than 10,000 users.

How should a database specialist implement access control with the LEAST operational effort?

- A. Use web identity federation on the mobile app and AWS STS with an attached IAM role to get temporary credentials to access DynamoDB.
- B. Use web identity federation on the mobile app and create individual IAM users with credentials to access DynamoDB.
- C. Use a self-developed user management system on the mobile app that lets users access the data from DynamoDB through an API.
- D. Use a single IAM user on the mobile app to access DynamoDB.

Answer: A

NEW QUESTION 24

A Database Specialist is designing a disaster recovery strategy for a production Amazon DynamoDB table. The table uses provisioned read/write capacity mode, global secondary indexes, and time to live (TTL). The Database Specialist has restored the latest backup to a new table.

To prepare the new table with identical settings, which steps should be performed? (Choose two.)

- A. Re-create global secondary indexes in the new table
- B. Define IAM policies for access to the new table
- C. Define the TTL settings
- D. Encrypt the table from the AWS Management Console or use the update-table command
- E. Set the provisioned read and write capacity

Answer: BC

Explanation:

The following items need to be reconfigured after restoring the DynamoDB table.

- AutoScaling policy
- IAM policy

--CloudWatch settings
--Tags
--Stream settings
--TTL
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/backuprestore_HowItWorks.html

NEW QUESTION 25

A company is using Amazon RDS for MySQL to redesign its business application. A Database Specialist has noticed that the Development team is restoring their MySQL database multiple times a day when Developers make mistakes in their schema updates. The Developers sometimes need to wait hours to the restores to complete.

Multiple team members are working on the project, making it difficult to find the correct restore point for each mistake. Which approach should the Database Specialist take to reduce downtime?

- A. Deploy multiple read replicas and have the team members make changes to separate replica instances
- B. Migrate to Amazon RDS for SQL Server, take a snapshot, and restore from the snapshot
- C. Migrate to Amazon Aurora MySQL and enable the Aurora Backtrack feature
- D. Enable the Amazon RDS for MySQL Backtrack feature

Answer: C

Explanation:

"Amazon Aurora, a fully-managed relational database service in AWS, is now offering a backtrack feature. With Amazon Aurora with MySQL compatibility, users can backtrack, or "rewind", a database cluster to a specific point in time, without restoring data from a backup. The backtrack process allows a point in time to be specified with one second resolution, and the rewind process typically takes minutes. This new feature facilitates developers in undoing mistakes like deleting data inappropriately or dropping the wrong table."

NEW QUESTION 30

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 33

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

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

Answer: BCF

Explanation:

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

NEW QUESTION 36

A company is deploying a solution in Amazon Aurora by migrating from an on-premises system. The IT department has established an AWS Direct Connect link from the company's data center. The company's Database Specialist has selected the option to require SSL/TLS for connectivity to prevent plaintext data from being set over the network. The migration appears to be working successfully, and the data can be queried from a desktop machine.

Two Data Analysts have been asked to query and validate the data in the new Aurora DB cluster. Both Analysts are unable to connect to Aurora. Their user names and passwords have been verified as valid and the Database Specialist can connect to the DB cluster using their accounts. The Database Specialist also verified that the security group configuration allows network from all corporate IP addresses.

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

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

Answer: B

Explanation:

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

NEW QUESTION 37

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

Which solution will enable this change?

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

Answer: B

Explanation:

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

NEW QUESTION 38

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 40

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 43

A Database Specialist is creating a new Amazon Neptune DB cluster, and is attempting to load data from Amazon S3 into the Neptune DB cluster using the Neptune bulk loader API. The Database Specialist receives the following error:

"Unable to connect to s3 endpoint. Provided source = s3://mybucket/graphdata/ and region = us-east-1. Please verify your S3 configuration."

Which combination of actions should the Database Specialist take to troubleshoot the problem? (Choose two.)

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

Answer: BD

NEW QUESTION 48

A company is releasing a new mobile game featuring a team play mode. As a group of mobile device users play together, an item containing their statuses is updated in an Amazon DynamoDB table. Periodically, the other users' devices read the latest statuses of their teammates from the table using the BatchGetItem operation.

Prior to launch, some testers submitted bug reports claiming that the status data they were seeing in the game was not up-to-date. The developers are unable to replicate this issue and have asked a database specialist for a recommendation.

Which recommendation would resolve this issue?

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

Answer: D

Explanation:

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

NEW QUESTION 52

A software-as-a-service (SaaS) company is using an Amazon Aurora Serverless DB cluster for its production MySQL database. The DB cluster has general logs and slow query logs enabled. A database engineer must use the most operationally efficient solution with minimal resource utilization to retain the logs and facilitate interactive search and analysis.

Which solution meets these requirements?

- A. Use an AWS Lambda function to ship database logs to an Amazon S3 bucket
- B. Use Amazon Athena and Amazon QuickSight to search and analyze the logs.
- C. Download the logs from the DB cluster and store them in Amazon S3 by using manual script
- D. Use Amazon Athena and Amazon QuickSight to search and analyze the logs.
- E. Use an AWS Lambda function to ship database logs to an Amazon S3 bucket
- F. Use Amazon Elasticsearch Service (Amazon ES) and Kibana to search and analyze the logs.
- G. Use Amazon CloudWatch Logs Insights to search and analyze the logs when the logs are automatically uploaded by the DB cluster.

Answer: D

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/aurora-serverless-logs-enable-view/>
<https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/AnalyzingLogData.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 54

A company is running a two-tier ecommerce application in one AWS account. The web server is deployed using an Amazon RDS for MySQL Multi-AZ DB instance. A Developer mistakenly deleted the database in the production environment. The database has been restored, but this resulted in hours of downtime and lost revenue.

Which combination of changes in existing IAM policies should a Database Specialist make to prevent an error like this from happening in the future? (Choose three.)

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

Answer: ACD

Explanation:

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

NEW QUESTION 59

A company is loading sensitive data into an Amazon Aurora MySQL database. To meet compliance requirements, the company needs to enable audit logging on the Aurora MySQL DB cluster to audit database activity. This logging will include events such as connections, disconnections, queries, and tables queried. The company also needs to publish the DB logs to Amazon CloudWatch to perform real-time data

analysis.

Which solution meets these requirements?

- A. Modify the default option group parameters to enable Advanced Auditin
- B. Restart the database for the changes to take effect.
- C. Create a custom DB cluster parameter grou
- D. Modify the parameters for Advanced Auditin
- E. Modify the cluster to associate the new custom DB parameter group with the Aurora MySQL DB cluster.
- F. Take a snapshot of the databas
- G. Create a new DB instance, and enable custom auditing and logging to CloudWate
- H. Deactivate the DB instance that has no logging.
- I. Enable AWS CloudTrail for the DB instanc
- J. Create a filter that provides only connections, disconnections, queries, and tables queried.

Answer: B

Explanation:

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

NEW QUESTION 64

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

Which solution would meet these requirements and deploy the DynamoDB tables?

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

Answer: C

Explanation:

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

NEW QUESTION 69

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

What should the company do to achieve this in the shortest amount of time?

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

Answer: B

Explanation:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_RebootInstance.html <https://exain.wordpress.com/2017/07/12/amazon-rds-multi-az-setup-failover-simulation/>

"Rebooting with failover is beneficial when you want to simulate a failure of a DB instance for testing, or restore operations to the original AZ after a failover occurs."

NEW QUESTION 71

A company is using Amazon Redshift as its data warehouse solution. The Redshift cluster handles the following types of workloads:

*Real-time inserts through Amazon Kinesis Data Firehose

*Bulk inserts through COPY commands from Amazon S3

*Analytics through SQL queries

Recently, the cluster has started to experience performance issues.

Which combination of actions should a database specialist take to improve the cluster's performance? (Choose three.)

- A. Modify the Kinesis Data Firehose delivery stream to stream the data to Amazon S3 with a high buffer size and to load the data into Amazon Redshift by using the COPY command.
- B. Stream real-time data into Redshift temporary tables before loading the data into permanent tables.
- C. For bulk inserts, split input files on Amazon S3 into multiple files to match the number of slices on Amazon Redshif
- D. Then use the COPY command to load data into Amazon Redshift.
- E. For bulk inserts, use the parallel parameter in the COPY command to enable multi-threading.
- F. Optimize analytics SQL queries to use sort keys.
- G. Avoid using temporary tables in analytics SQL queries.

Answer: BCE

Explanation:

<https://aws.amazon.com/blogs/big-data/top-10-performance-tuning-techniques-for-amazon-redshift/> Tip #6: Improving the efficiency of temporary tables

Tip #9: Maintaining efficient data loads

Amazon Redshift best practices suggest using the COPY command to perform data loads of file-based data. Tip #3: Sort key recommendation

Sorting a table on an appropriate sort key can accelerate query performance, especially queries with range-restricted predicates, by requiring fewer table blocks to be read from disk.

NEW QUESTION 76

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

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

Answer: A

NEW QUESTION 77

A company uses the Amazon DynamoDB table contractDB in us-east-1 for its contract system with the following schema: orderID (primary key) timestamp (sort key) contract (map) createdBy (string) customerEmail (string)

After a problem in production, the operations team has asked a database specialist to provide an IAM policy to read items from the database to debug the application. In addition, the developer is not allowed to access the value of the customerEmail field to stay compliant.

Which IAM policy should the database specialist use to achieve these requirements?

A)

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "IAMPolicy",
      "Effect": "Allow",
      "Action": [
        "dynamodb: Query"
      ],
      "Resource": [
        "arn:aws:dynamodb:us-east-1:123456789012:table/contractDB"
      ],
      "Condition": {
        "ForAllValues:StringLike": {
          "dynamodb:Attributes": [
            "orderID",
            "timestamp",
            "contract",
            "createdBy"
          ]
        },
        "StringEquals": {
          "dynamodb:Select": "SPECIFIC_ATTRIBUTES"
        }
      }
    }
  ]
}
```

B)

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "IAMPolicy",
      "Effect": "Allow",
      "Action": [
        "dynamodb: Query"
      ],
      "Resource": [
        "arn:aws:dynamodb:us-east-1:123456789012:table/contractDB"
      ],
      "Condition": {
        "ForAllValues:StringLike": {
          "dynamodb:Attributes": [
            "customerEmail"
          ]
        },
        "StringEquals": {
          "dynamodb:Select": "SPECIFIC_ATTRIBUTES"
        }
      }
    }
  ]
}
```

C)

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "IAMPolicy",
      "Effect": "Deny",
      "Action": [
        "dynamodb: Query"
      ],
      "Resource": [
        "arn:aws:dynamodb:us-east-1:123456789012:table/contractDB"
      ],
      "Condition": {
        "ForAllValues:StringLike": {
          "dynamodb:Attributes": [
            "customerEmail"
          ]
        },
        "StringEquals": {
          "dynamodb:Select": "SPECIFIC_ATTRIBUTES"
        }
      }
    }
  ]
}
```

D)

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "IAMPolicy",
      "Effect": "Deny",
      "Action": [
        "dynamodb:Query"
      ],
      "Resource": [
        "arn:aws:dynamodb:us-east-1:123456789012:table/contractDB"
      ],
      "Condition": {
        "ForAllValues:StringLike": {
          "dynamodb:Attributes": [
            "orderId",
            "timestamp",
            "contract",
            "createdBy"
          ]
        },
        "StringEquals": {
          "dynamodb:Select": "SPECIFIC_ATTRIBUTES"
        }
      }
    }
  ]
}
```

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

Answer: A

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 80

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

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

Answer: C

Explanation:

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

NEW QUESTION 83

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

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

Answer: A

Explanation:

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

NEW QUESTION 86

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

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

Answer: C

NEW QUESTION 89

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

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

Answer: AD

Explanation:

"When Amazon RDS designates a minor engine version as the preferred minor engine version, each database that meets both of the following conditions is upgraded to the minor engine version automatically"

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

Call the modify-db-instance Amazon CLI command. Specify the name of your DB instance for the --db-instance-identifier option and true for the --auto-minor-version-upgrade option. Optionally, specify the --apply-immediately option to immediately enable this setting for your DB instance. Run a separate modify-db-instance command for each DB instance in the cluster.

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

NEW QUESTION 90

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

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

Answer: D

NEW QUESTION 93

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

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

Answer: D

Explanation:

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

NEW QUESTION 97

A business uses Amazon EC2 instances in VPC A to serve an internal file-sharing application. This application is supported by an Amazon ElastiCache cluster in VPC B that is peering with VPC A. The corporation migrates the instances of its applications from VPC A to VPC B. The file-sharing application is no longer able to connect to the ElastiCache cluster, as shown by the logs.

What is the best course of action for a database professional to take in order to remedy this issue?

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

Answer: D

Explanation:

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

NEW QUESTION 102

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

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

Answer: A

Explanation:

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

NEW QUESTION 104

A Database Specialist is setting up a new Amazon Aurora DB cluster with one primary instance and three Aurora Replicas for a highly intensive, business-critical application. The Aurora DB cluster has one medium-sized primary instance, one large-sized replica, and two medium-sized replicas. The Database Specialist did not assign a promotion tier to the replicas.

In the event of a primary failure, what will occur?

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

Answer: C

Explanation:

Priority: If you don't select a value, the default is tier-1. This priority determines the order in which Aurora

https://docs.amazonaws.cn/en_us/AmazonRDS/latest/AuroraUserGuide/aurora-replicas-adding.html

More than one Aurora Replica can share the same priority, resulting in promotion tiers. If two or more Aurora Replicas share the same priority, then Amazon RDS promotes the replica that is largest in size. If two or more Aurora Replicas share the same priority and size, then Amazon RDS promotes an arbitrary replica in the same promotion tier.

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Managing.Backups.html#Aurora.M> If two or more Aurora Replicas share the same priority, then Amazon RDS promotes the replica that is largest in size. If two or more Aurora Replicas share the same priority and size, then Amazon RDS promotes an

arbitrary replica in the same promotion tier. <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Concepts.AuroraHighAvailability.html>

NEW QUESTION 106

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

addressed?

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

Answer: D

Explanation:

Rotate your IAM credentials regularly.

NEW QUESTION 108

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

How could a database professional approach these needs while using DynamoDB?

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

Answer: D

NEW QUESTION 109

A retail company manages a web application that stores data in an Amazon DynamoDB table. The company is undergoing account consolidation efforts. A database engineer needs to migrate the DynamoDB table from the current AWS account to a new AWS account.

Which strategy meets these requirements with the LEAST amount of administrative work?

- A. Use AWS Glue to crawl the data in the DynamoDB tabl
- B. Create a job using an available blueprint to export the data to Amazon S3. Import the data from the S3 file to a DynamoDB table in the new account.
- C. Create an AWS Lambda function to scan the items of the DynamoDB table in the current account and write to a file in Amazon S3. Create another Lambda function to read the S3 file and restore the items of a DynamoDB table in the new account.
- D. Use AWS Data Pipeline in the current account to export the data from the DynamoDB table to a file in Amazon S3. Use Data Pipeline to import the data from the S3 file to a DynamoDB table in the new account.
- E. Configure Amazon DynamoDB Streams for the DynamoDB table in the current accoun
- F. Create an AWS Lambda function to read from the stream and write to a file in Amazon S3. Create another Lambda function to read the S3 file and restore the items to a DynamoDB table in the new account.

Answer: C

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/dynamodb-cross-account-migration/> <https://aws.amazon.com/premiumsupport/knowledge-center/data-pipeline-account-access-dynamodb-s3/>

NEW QUESTION 111

A large IT hardware manufacturing company wants to deploy a MySQL database solution in the AWS Cloud. The solution should quickly create copies of the company's production databases for test purposes. The solution must deploy the test databases in minutes, and the test data should match the latest production data as closely as possible. Developers must also be able to make changes in the test database and delete the instances afterward.

Which solution meets these requirements?

- A. Leverage Amazon RDS for MySQL with write-enabled replicas running on Amazon EC2. Create the test copies using a mysqldump backup from the RDS for MySQL DB instances and importing them into the new EC2 instances.
- B. Leverage Amazon Aurora MySQL
- C. Use database cloning to create multiple test copies of the production DB clusters.
- D. Leverage Amazon Aurora MySQL
- E. Restore previous production DB instance snapshots into new test copies of Aurora MySQL DB clusters to allow them to make changes.
- F. Leverage Amazon RDS for MySQL
- G. Use database cloning to create multiple developer copies of the production DB instance.

Answer: B

NEW QUESTION 115

A company is running a blogging platform. A security audit determines that the Amazon RDS DB instance that is used by the platform is not configured to encrypt the data at rest. The company must encrypt the DB instance within 30 days.

What should a database specialist do to meet this requirement with the LEAST amount of downtime?

- A. Create a read replica of the DB instance, and enable encryptio
- B. When the read replica is available, promote the read replica and update the endpoint that is used by the applicatio
- C. Delete the unencrypted DB instance.

- D. Take a snapshot of the DB instance
- E. Make an encrypted copy of the snapshot
- F. Restore the encrypted snapshot
- G. When the new DB instance is available, update the endpoint that is used by the application
- H. Delete the unencrypted DB instance.
- I. Create a new encrypted DB instance
- J. Perform an initial data load, and set up logical replication between the two DB instances. When the new DB instance is in sync with the source DB instance, update the endpoint that is used by the application
- K. Delete the unencrypted DB instance.
- L. Convert the DB instance to an Amazon Aurora DB cluster, and enable encryption
- M. When the DB cluster is available, update the endpoint that is used by the application to the cluster endpoint
- N. Delete the unencrypted DB instance.

Answer: C

Explanation:

<https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/encrypt-an-existing-amazon-rds-for-postgresql> When the new, encrypted copy of the DB instance becomes available, you can point your applications to the new database. However, if your project doesn't allow for significant downtime for this activity, you need an alternate approach that helps minimize the downtime. This pattern uses the AWS Database Migration Service (AWS DMS) to migrate and continuously replicate the data so that the cutover to the new, encrypted database can be done with minimal downtime.

NEW QUESTION 120

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 123

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 125

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 130

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

OOM command not allowed when used memory > 'maxmemory'

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

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

Answer: ABC

Explanation:

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

NEW QUESTION 135

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:XactSync 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 140

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 instanc
- 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 144

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

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

Answer: D

NEW QUESTION 146

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

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

Answer: C

Explanation:

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

NEW QUESTION 148

A company has a database fleet that includes an Amazon RDS for MySQL DB instance. During an audit, the company discovered that the data that is stored on the DB instance is unencrypted.

A database specialist must enable encryption for the DB instance. The database specialist also must encrypt all connections to the DB instance. Which combination of actions should the database specialist take to meet these requirements? (Choose three.)

- A. In the RDS console, choose "Enable encryption" to encrypt the DB instance by using an AWS Key Management Service (AWS KMS) key.
- B. Encrypt the read replica of the unencrypted DB instance by using an AWS Key Management Service (AWS KMS) key.
- C. Fail over the read replica to the primary DB instance.
- D. Create a snapshot of the unencrypted DB instance.
- E. Encrypt the snapshot by using an AWS Key Management Service (AWS KMS) key.
- F. Restore the DB instance from the encrypted snapshot.
- G. Delete the original DB instance.
- H. Require SSL connections for applicable database user accounts.
- I. Use SSL/TLS from the application to encrypt a connection to the DB instance.
- J. Enable SSH encryption on the DB instance.

Answer: ACE

Explanation:

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

NEW QUESTION 150

A company uses Amazon DynamoDB as the data store for its ecommerce website. The website receives little to no traffic at night, and the majority of the traffic occurs during the day. The traffic growth during peak hours is gradual and predictable on a daily basis, but it can be orders of magnitude higher than during off-peak hours.

The company initially provisioned capacity based on its average volume during the day without accounting for the variability in traffic patterns. However, the website is experiencing a significant amount of throttling during peak hours. The company wants to reduce the amount of throttling while minimizing costs. What should a database specialist do to meet these requirements?

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

Answer: C

Explanation:

On-demand mode is a good option if any of the following are true: You create new tables with unknown workloads. You have unpredictable application traffic. You prefer the ease of paying for only what you use.

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/HowItWorks.ReadWriteCapacityMode.html>

Amazon DynamoDB auto scaling uses the AWS Application Auto Scaling service to dynamically adjust provisioned throughput capacity on your behalf
<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/AutoScaling.html>

NEW QUESTION 155

A social media company is using Amazon DynamoDB to store user profile data and user activity data. Developers are reading and writing the data, causing the size of the tables to grow significantly. Developers have started to face performance bottlenecks with the tables.

Which solution should a database specialist recommend to read items the FASTEST without consuming all the provisioned throughput for the tables?

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

Answer: B

Explanation:

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

NEW QUESTION 160

An ecommerce company is using Amazon DynamoDB as the backend for its order-processing application.

The steady increase in the number of orders is resulting in increased DynamoDB costs. Order verification and reporting perform many repeated GetItem functions that pull similar datasets, and this read activity is contributing to the increased costs. The company wants to control these costs without significant development efforts.

How should a Database Specialist address these requirements?

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

Answer: D

Explanation:

https://docs.amazonaws.cn/en_us/amazondynamodb/latest/developerguide/DAX.html

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

NEW QUESTION 163

A company has an on-premises system that tracks various database operations that occur over the lifetime of a database, including database shutdown, deletion, creation, and backup.

The company recently moved two databases to Amazon RDS and is looking at a solution that would satisfy these requirements. The data could be used by other systems within the company.

Which solution will meet these requirements with minimal effort?

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

Answer: C

NEW QUESTION 168

A company is migrating a mission-critical 2-TB Oracle database from on premises to Amazon Aurora. The cost for the database migration must be kept to a minimum, and both the on-premises Oracle database and the Aurora DB cluster must remain open for write traffic until the company is ready to completely cut over to Aurora.

Which combination of actions should a database specialist take to accomplish this migration as quickly as possible? (Choose two.)

- A. Use the AWS Schema Conversion Tool (AWS SCT) to convert the source database schem
- B. Then restore the converted schema to the target Aurora DB cluster.
- C. Use Oracle's Data Pump tool to export a copy of the source database schema and manually edit the schema in a text editor to make it compatible with Aurora.
- D. Create an AWS DMS task to migrate data from the Oracle database to the Aurora DB cluste
- E. Select the migration type to replicate ongoing changes to keep the source and target databases in sync until the company is ready to move all user traffic to the Aurora DB cluster.
- F. Create an AWS DMS task to migrate data from the Oracle database to the Aurora DB cluste
- G. Once the initial load is complete, create an AWS Kinesis Data Firehose stream to perform change data capture (CDC) until the company is ready to move all user traffic to the Aurora DB cluster.
- H. Create an AWS Glue job and related resources to migrate data from the Oracle database to the Aurora DB cluste
- I. Once the initial load is complete, create an AWS DMS task to perform change data capture (CDC) until the company is ready to move all user traffic to the Aurora DB cluster.

Answer: AC

NEW QUESTION 173

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 178

A company hosts a 2 TB Oracle database in its on-premises data center. A database specialist is migrating the database from on premises to an Amazon Aurora PostgreSQL database on AWS.

The database specialist identifies a problem that relates to compatibility Oracle stores metadata in its data dictionary in uppercase, but PostgreSQL stores the metadata in lowercase. The database specialist must resolve this problem to complete the migration.

What is the MOST operationally efficient solution that meets these requirements?

- A. Override the default uppercase format of Oracle schema by encasing object names in quotation marks during creation.
- B. Use AWS Database Migration Service (AWS DMS) mapping rules with rule-action as convert-lowercase.
- C. Use the AWS Schema Conversion Tool conversion agent to convert the metadata from uppercase to lowercase.

D. Use an AWS Glue job that is attached to an AWS Database Migration Service (AWS DMS) replication task to convert the metadata from uppercase to lowercase.

Answer: B

Explanation:

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

NEW QUESTION 181

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

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

Answer: ABC

NEW QUESTION 185

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%20>
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/backuprestore_HowItWorks.html

NEW QUESTION 187

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 190

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

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

Answer: ACE

Explanation:

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

NEW QUESTION 195

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 198

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

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

Answer: B

Explanation:

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

NEW QUESTION 203

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 204

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

How should this be accomplished?

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

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

Answer: C

NEW QUESTION 206

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 207

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

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

Answer: C

Explanation:

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

NEW QUESTION 212

A company just migrated to Amazon Aurora PostgreSQL from an on-premises Oracle database. After the migration, the company discovered there is a period of time every day around 3:00 PM where the response time of the application is noticeably slower. The company has narrowed down the cause of this issue to the database and not the application.

Which set of steps should the Database Specialist take to most efficiently find the problematic PostgreSQL query?

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

Answer: D

NEW QUESTION 213

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 cluster
- 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 217

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

Which authentication solution should a database professional provide?

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

Answer: D

Explanation:

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

NEW QUESTION 218

A company has an AWS CloudFormation stack that defines an Amazon RDS DB instance. The company accidentally deletes the stack and loses recent data from the DB instance. A database specialist must change the CloudFormation template for the RDS resource to reduce the chance of accidental data loss from the DB instance in the future.

Which combination of actions should the database specialist take to meet this requirement? (Choose three.)

- A. Set the DeletionProtection property to True.
- B. Set the MultiAZ property to True.
- C. Set the TerminationProtection property to True.
- D. Set the DeleteAutomatedBackups property to False.
- E. Set the DeletionPolicy attribute to No.
- F. Set the DeletionPolicy attribute to Retain.

Answer: ADF

NEW QUESTION 222

A financial institution uses AWS to host its online application. Amazon RDS for MySQL is used to host the application's database, which includes automatic backups.

The program has corrupted the database logically, resulting in the application being unresponsive. The exact moment the corruption occurred has been determined, and it occurred within the backup retention period.

How should a database professional restore a database to its previous state prior to corruption?

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

Answer: B

Explanation:

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

NEW QUESTION 226

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 228

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

What should the Database Specialist do to automatically collect the database logs for the Administrator?

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

Answer: C

Explanation:

<https://docs.aws.amazon.com/documentdb/latest/developerguide/event-auditing.html> Auditing Amazon DocumentDB Events
PDF

Kindle RSS

With Amazon DocumentDB (with MongoDB compatibility), you can audit events that were performed in your cluster. Examples of logged events include successful and failed authentication attempts, dropping a collection in a database, or creating an index. By default, auditing is disabled on Amazon DocumentDB and requires that you opt in to use this feature.

When auditing is enabled, Amazon DocumentDB records Data Definition Language (DDL), authentication, authorization, and user management events to Amazon CloudWatch Logs. When auditing is enabled, Amazon DocumentDB exports your cluster's auditing records (JSON documents) to Amazon CloudWatch Logs. You can use Amazon CloudWatch Logs to analyze, monitor, and archive your Amazon DocumentDB auditing events.

NEW QUESTION 233

A company is running its line of business application on AWS, which uses Amazon RDS for MySQL at the persistent data store. The company wants to minimize downtime when it migrates the database to Amazon Aurora.

Which migration method should a Database Specialist use?

- A. Take a snapshot of the RDS for MySQL DB instance and create a new Aurora DB cluster with the option to migrate snapshots.
- B. Make a backup of the RDS for MySQL DB instance using the `mysqldump` utility, create a new Aurora DB cluster, and restore the backup.
- C. Create an Aurora Replica from the RDS for MySQL DB instance and promote the Aurora DB cluster.
- D. Create a clone of the RDS for MySQL DB instance and promote the Aurora DB cluster.

Answer: C

Explanation:

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

NEW QUESTION 234

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 235

A company has deployed an e-commerce web application in a new AWS account. An Amazon RDS for MySQL Multi-AZ DB instance is part of this deployment with a `database-1.xxxxxxxxxx.us-east-1.rds.amazonaws.com` endpoint listening on port 3306. The company's Database Specialist is able to log in to MySQL and run queries from the bastion host using these details.

When users try to utilize the application hosted in the AWS account, they are presented with a generic error message. The application servers are logging a "could not connect to server: Connection times out" error message to Amazon CloudWatch Logs.

What is the cause of this error?

- A. The user name and password the application is using are incorrect.
- B. The security group assigned to the application servers does not have the necessary rules to allow inbound connections from the DB instance.
- C. The security group assigned to the DB instance does not have the necessary rules to allow inbound connections from the application servers.
- D. The user name and password are correct, but the user is not authorized to use the DB instance.

Answer: C

NEW QUESTION 240

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

Which solution meets these requirements?

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

Answer: A

Explanation:

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

NEW QUESTION 245

A company hosts an on-premises Microsoft SQL Server Enterprise edition database with Transparent Data Encryption (TDE) enabled. The database is 20 TB in size and includes sparse tables. The company needs to migrate the database to Amazon RDS for SQL Server during a maintenance window that is scheduled for an upcoming weekend. Data-at-rest encryption must be enabled for the target DB instance.

Which combination of steps should the company take to migrate the database to AWS in the MOST operationally efficient manner? (Choose two.)

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

Answer: BC

Explanation:

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

NEW QUESTION 249

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 253

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

Which database specialist solution should be recommended?

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

Answer: C

Explanation:

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

NEW QUESTION 254

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

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

Answer: A

Explanation:

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

NEW QUESTION 259

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

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

Answer: A

Explanation:

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

NEW QUESTION 262

An online advertising website uses an Amazon DynamoDB table with on-demand capacity mode as its data store. The website also has a DynamoDB Accelerator (DAX) cluster in the same VPC as its web application server. The application needs to perform infrequent writes and many strongly consistent reads from the data store by querying the DAX cluster.

During a performance audit, a systems administrator notices that the application can look up items by using the DAX cluster. However, the QueryCacheHits metric for the DAX cluster consistently shows 0 while the QueryCacheMisses metric continuously keeps growing in Amazon CloudWatch. What is the MOST likely reason for this occurrence?

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

Answer: B

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/DAX.concepts.html>

"If the request specifies strongly consistent reads, DAX passes the request through to DynamoDB. The results from DynamoDB are not cached in DAX. Instead, they are simply returned to the application."

NEW QUESTION 267

A company uses an on-premises Microsoft SQL Server database to host relational and JSON data and to run daily ETL and advanced analytics. The company wants to migrate the database to the AWS Cloud. Database specialist must choose one or more AWS services to run the company's workloads. Which solution will meet these requirements in the MOST operationally efficient manner?

- A. Use Amazon Redshift for relational data
- B. Use Amazon DynamoDB for JSON data
- C. Use Amazon Redshift for relational data and JSON data.
- D. Use Amazon RDS for relational data
- E. Use Amazon Neptune for JSON data
- F. Use Amazon Redshift for relational data
- G. Use Amazon S3 for JSON data.

Answer: B

Explanation:

<https://docs.aws.amazon.com/redshift/latest/dg/super-overview.htm>

NEW QUESTION 272

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 275

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 278

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 281

A company uses Amazon Aurora MySQL as the primary database engine for many of its applications. A database specialist must create a dashboard to provide the company with information about user connections to databases. According to compliance requirements, the company must retain all connection logs for at least 7 years.

Which solution will meet these requirements MOST cost-effectively?

- A. Enable advanced auditing on the Aurora cluster to log CONNECT event
- B. Export audit logs from Amazon CloudWatch to Amazon S3 by using an AWS Lambda function that is invoked by an Amazon EventBridge (Amazon CloudWatch Events) scheduled event
- C. Build a dashboard by using Amazon QuickSight.
- D. Capture connection attempts to the Aurora cluster with AWS Cloud Trail by using the DescribeEvents API operation
- E. Create a CloudTrail trail to export connection logs to Amazon S3. Build a dashboard by using Amazon QuickSight.
- F. Start a database activity stream for the Aurora cluster
- G. Push the activity records to an Amazon Kinesis data stream
- H. Build a dynamic dashboard by using AWS Lambda.
- I. Publish the DatabaseConnections metric for the Aurora DB instances to Amazon CloudWatch
- J. Build a dashboard by using CloudWatch dashboards.

Answer: A

Explanation:

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

NEW QUESTION 284

.....

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](#)