

Exam Questions AWS-Certified-Database-Specialty

AWS Certified Database - Specialty

<https://www.2passeasy.com/dumps/AWS-Certified-Database-Specialty/>



NEW QUESTION 1

A database specialist needs to review and optimize an Amazon DynamoDB table that is experiencing performance issues. A thorough investigation by the database specialist reveals that the partition key is causing hot partitions, so a new partition key is created. The database specialist must effectively apply this new partition key to all existing and new data.

How can this solution be implemented?

- A. Use Amazon EMR to export the data from the current DynamoDB table to Amazon S3. Then use Amazon EMR again to import the data from Amazon S3 into a new DynamoDB table with the new partition key.
- B. Use AWS DMS to copy the data from the current DynamoDB table to Amazon S3. Then import the DynamoDB table to create a new DynamoDB table with the new partition key.
- C. Use the AWS CLI to update the DynamoDB table and modify the partition key.
- D. Use the AWS CLI to back up the DynamoDB table.
- E. Then use the restore-table-from-backup command and modify the partition key.

Answer: A

Explanation:

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

NEW QUESTION 2

A company has an ecommerce web application with an Amazon RDS for MySQL DB instance. The marketing team has noticed some unexpected updates to the product and pricing information on the website, which is impacting sales targets. The marketing team wants a database specialist to audit future database activity to help identify how and when the changes are being made.

What should the database specialist do to meet these requirements? (Choose two.)

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

Answer: BD

Explanation:

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

NEW QUESTION 3

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 4

A gaming company has implemented a leaderboard in AWS using a Sorted Set data structure within Amazon ElastiCache for Redis. The ElastiCache cluster has been deployed with cluster mode disabled and has a replication group deployed with two additional replicas. The company is planning for a worldwide gaming event and is anticipating a higher write load than what the current cluster can handle.

Which method should a Database Specialist use to scale the ElastiCache cluster ahead of the upcoming event?

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

Answer: B

NEW QUESTION 5

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

Which solution addresses these requirements?

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

- G. Set the rds.force_ssl=1 parameter in DB parameter group
- H. Download and use the Amazon RDS certificate bundle and configure the PostgreSQL connection string with sslmode=verify-full.

Answer: D

Explanation:

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

NEW QUESTION 6

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

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

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

Answer: C

Explanation:

<https://aws.amazon.com/blogs/aws/use-cloudformation-stacksets-to-provision-resources-across-multiple-aws-ac>

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stacksets-concepts.html>

NEW QUESTION 7

A business that specializes in internet advertising is developing an application that will show adverts to its customers. The program stores data in an Amazon DynamoDB database. Additionally, the application caches its reads using a DynamoDB Accelerator (DAX) cluster. The majority of reads come via the GetItem and BatchGetItem queries. The application does not need consistency of readings.

The application cache does not behave as intended after deployment. Specific extremely consistent queries to the DAX cluster are responding in several milliseconds rather than microseconds.

How can the business optimize cache behavior in order to boost application performance?

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

Answer: C

NEW QUESTION 8

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 9

A company's ecommerce website uses Amazon DynamoDB for purchase orders. Each order is made up of a Customer ID and an Order ID. The DynamoDB table uses the Customer ID as the partition key and the Order ID as the sort key.

To meet a new requirement, the company also wants the ability to query the table by using a third attribute named Invoice ID. Queries using the Invoice ID must be strongly consistent. A database specialist must provide this capability with optimal performance and minimal overhead.

What should the database administrator do to meet these requirements?

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

Answer: C

Explanation:

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

NEW QUESTION 10

A gaming company has recently acquired a successful iOS game, which is particularly popular during the holiday season. The company has decided to add a leaderboard to the game that uses Amazon DynamoDB. The application load is expected to ramp up over the holiday season.

Which solution will meet these requirements at the lowest cost?

- A. DynamoDB Streams
- B. DynamoDB with DynamoDB Accelerator
- C. DynamoDB with on-demand capacity mode

D. DynamoDB with provisioned capacity mode with Auto Scaling

Answer: C

NEW QUESTION 10

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 11

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 16

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 17

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

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

Answer: D

Explanation:

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

NEW QUESTION 19

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

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

Answer: C

NEW QUESTION 21

A company is moving its fraud detection application from on premises to the AWS Cloud and is using Amazon Neptune for data storage. The company has set up a 1 Gbps AWS Direct Connect connection to migrate 25 TB of fraud detection data from the on-premises data center to a Neptune DB instance. The company already has an Amazon S3 bucket and an S3 VPC endpoint, and 80% of the company's network bandwidth is available. How should the company perform this data load?

- A. Use an AWS SDK with a multipart upload to transfer the data from on premises to the S3 bucket.
- B. Use the Copy command for Neptune to move the data in bulk from the S3 bucket to the Neptune DB instance.
- C. Use AWS Database Migration Service (AWS DMS) to transfer the data from on premises to the S3 bucket.
- D. Use the Loader command for Neptune to move the data in bulk from the S3 bucket to the Neptune DB instance.
- E. Use AWS DataSync to transfer the data from on premises to the S3 bucket.
- F. Use the Loader command for Neptune to move the data in bulk from the S3 bucket to the Neptune DB instance.
- G. Use the AWS CLI to transfer the data from on premises to the S3 bucket.
- H. Use the Copy command for Neptune to move the data in bulk from the S3 bucket to the Neptune DB instance.

Answer: C

Explanation:

"AWS DataSync is an online data transfer service that simplifies, automates, and accelerates moving data between on-premises storage systems and AWS storage services, and also between AWS storage services."

<https://docs.aws.amazon.com/neptune/latest/userguide/bulk-load.html>

NEW QUESTION 25

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

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

Answer: C

Explanation:

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

NEW QUESTION 29

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

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

Answer: D

NEW QUESTION 30

A company has two separate AWS accounts: one for the business unit and another for corporate analytics. The company wants to replicate the business unit data stored in Amazon RDS for MySQL in us-east-1 to its corporate analytics Amazon Redshift environment in us-west-1. The company wants to use AWS DMS with Amazon RDS as the source endpoint and Amazon Redshift as the target endpoint. Which action will allow AWS DMS to perform the replication?

- A. Configure the AWS DMS replication instance in the same account and Region as Amazon Redshift.
- B. Configure the AWS DMS replication instance in the same account as Amazon Redshift and in the same Region as Amazon RDS.
- C. Configure the AWS DMS replication instance in its own account and in the same Region as Amazon Redshift.
- D. Configure the AWS DMS replication instance in the same account and Region as Amazon RDS.

Answer: A

Explanation:

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Target.Redshift.html

NEW QUESTION 32

A major organization maintains a number of Amazon DB clusters. Each of these clusters is configured differently to meet certain needs. These configurations may be classified into wider groups based on the team and use case.

A database administrator wishes to streamline the process of storing and updating these settings. Additionally, the database administrator wants to guarantee that changes to certain configuration categories are automatically implemented to all instances as necessary.

Which AWS service or functionality will assist in automating and achieving this goal?

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

Answer: B

Explanation:

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

NEW QUESTION 37

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

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

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

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

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

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

Answer: D

NEW QUESTION 38

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 40

Recently, a gaming firm purchased a popular iOS game that is especially popular during the Christmas season. The business has opted to include a leaderboard into the game, which will be powered by Amazon DynamoDB. The application's load is likely to increase significantly throughout the Christmas season.

Which solution satisfies these criteria at the lowest possible cost?

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

Answer: D

Explanation:

"On-demand is ideal for bursty, new, or unpredictable workloads whose traffic can spike in seconds or minutes"

vs.

"DynamoDB released auto scaling to make it easier for you to manage capacity efficiently, and auto scaling continues to help DynamoDB users lower the cost of workloads that have a predictable traffic pattern."

<https://aws.amazon.com/blogs/database/amazon-dynamodb-auto-scaling-performance-and-cost-optimization-at>

NEW QUESTION 42

A financial company has allocated an Amazon RDS MariaDB DB instance with large storage capacity to accommodate migration efforts. Post-migration, the company purged unwanted data from the instance. The company now want to downsize storage to save money. The solution must have the least impact on production and near-zero downtime. Which solution would meet these requirements?

- A. Create a snapshot of the old databases and restore the snapshot with the required storage
- B. Create a new RDS DB instance with the required storage and move the databases from the old instances to the new instance using AWS DMS
- C. Create a new database using native backup and restore
- D. Create a new read replica and make it the primary by terminating the existing primary

Answer: B

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/rds-db-storage-size/> Use AWS Database Migration Service (AWS DMS) for minimal downtime.

NEW QUESTION 44

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

Which process should the Database Specialist recommend to meet these requirements?

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

Answer: A

Explanation:

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

NEW QUESTION 48

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 51

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

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

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 54

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

What should the Database Specialist do to meet these requirements?

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

Answer: C

Explanation:

<https://aws.amazon.com/getting-started/hands-on/aurora-cloning-backtracking/>

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

NEW QUESTION 58

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 59

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 63

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

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 64

An internet advertising firm stores its data in an Amazon DynamoDB table. Amazon DynamoDB Streams are enabled on the table, and one of the keys has a global secondary index. The table is encrypted using a customer-managed AWS Key Management Service (AWS KMS) key.

The firm has chosen to grow worldwide and want to duplicate the database using DynamoDB global tables in a new AWS Region.

An administrator observes the following upon review:

No role with the dynamodb: CreateGlobalTable permission exists in the account.
An empty table with the same name exists in the new Region where replication is desired.
A global secondary index with the same partition key but a different sort key exists in the new Region where replication is desired.
Which settings will prevent you from creating a global table or replica in the new Region? (Select two.)

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

Answer: AB

NEW QUESTION 66

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 69

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

ProvisionedThroughputExceededException error.

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

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

Answer: CD

Explanation:

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

NEW QUESTION 73

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

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

Answer: AE

Explanation:

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

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/security_iam_service-with-iam.html

NEW QUESTION 75

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

What is likely causing the timeouts?

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

Answer: C

Explanation:

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

NEW QUESTION 76

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 81

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 85

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

Which solution will enable this change?

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

Answer: B

Explanation:

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

NEW QUESTION 86

A Database Specialist is constructing a new Amazon Neptune DB cluster and tries to load data from Amazon S3 using the Neptune bulk loader API. The Database Specialist is confronted with the following error message:

€Unable to establish a connection to the s3 endpoint. The source URL is s3://mybucket/graphdata/ and the region code is us-east-1. Kindly confirm your Configuration S3.

Which of the following activities should the Database Specialist take to resolve the issue? (Select two.)

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

Answer: BE

Explanation:

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

NEW QUESTION 91

A business's mission-critical production workload is being operated on a 500 GB Amazon Aurora MySQL DB cluster. A database engineer must migrate the workload without causing data loss to a new Amazon Aurora Serverless MySQL DB cluster.

Which approach will result in the LEAST amount of downtime and the LEAST amount of application impact?

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

Answer: D

Explanation:

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

NEW QUESTION 94

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 96

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

What is most secure solution to store the master password?

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

Answer: C

Explanation:

"By using the secure string support in CloudFormation with dynamic references you can better maintain your infrastructure as code. You'll be able to avoid hard coding passwords into your templates and you can keep these runtime configuration parameters separated from your code. Moreover, when properly used, secure strings will help keep your development and production code as similar as possible, while continuing to make your infrastructure code suitable for continuous deployment pipelines."

<https://aws.amazon.com/blogs/mt/using-aws-systems-manager-parameter-store-secure-string-parameters-in-aws> <https://aws.amazon.com/blogs/security/how-to-use-aws-secrets-manager-rotate-credentials-amazon-rds-database>

NEW QUESTION 100

A company uses an Amazon RDS for PostgreSQL DB instance for its customer relationship management (CRM) system. New compliance requirements specify that the database must be encrypted at rest.

Which action will meet these requirements?

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

Answer: A

Explanation:

<https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/encrypt-an-existing-amazon-rds-for-postgresql> You can enable encryption for an Amazon RDS DB instance when you create it, but not after it's created.

However, you can add encryption to an unencrypted DB instance by creating a snapshot of your DB instance, and then creating an encrypted copy of that snapshot. You can then restore a DB instance from the encrypted snapshot to get an encrypted copy of your original DB instance. The pattern uses AWS Database Migration Service (AWS DMS) to migrate data and AWS Key Management Service (AWS KMS) for encryption.

NEW QUESTION 102

In one AWS account, a business runs a two-tier ecommerce application. An Amazon RDS for MySQL

Multi-AZ database instance serves as the application's backend. A developer removed the database instance in the production environment by accident. Although the organization recovers the database, the incident results in hours of outage and financial loss.

Which combination of adjustments would reduce the likelihood that this error will occur again in the future? (Select three.)

- A. Grant least privilege to groups, IAM users, and roles.
- B. Allow all users to restore a database from a backup.
- C. Enable deletion protection on existing production DB instances.
- D. Use an ACL policy to restrict users from DB instance deletion.
- E. Enable AWS CloudTrail logging and Enhanced Monitoring.

Answer: ACD

NEW QUESTION 104

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

- A. Use AWS Backup to build a backup plan for the required retention period
- B. Assign the DB instances to the backup plan.
- C. Modify the DB instances to enable the automated backup option

- D. Select the required backup retention period.
- E. Automate a daily cron job on an Amazon EC2 instance to create MySQL dumps, transfer to Amazon S3, and implement an S3 Lifecycle policy to meet the retention requirement.
- F. Use AWS Lambda to schedule a daily manual snapshot of the DB instance
- G. Delete snapshots that exceed the retention requirement.

Answer: A

Explanation:

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

NEW QUESTION 108

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

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

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

Answer: ABC

Explanation:

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

NEW QUESTION 111

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual AWS-Certified-Database-Specialty Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the AWS-Certified-Database-Specialty Product From:

<https://www.2passeasy.com/dumps/AWS-Certified-Database-Specialty/>

Money Back Guarantee

AWS-Certified-Database-Specialty Practice Exam Features:

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