Amazon

Exam Questions AWS-Certified-Database-Specialty

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

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

Answer: AB

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

NEW QUESTION 3
A 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 4
A company runs online transaction processing (OLTP) workloads on an Amazon RDS for PostgreSQL Multi- AZ DB instance. Tests were run on the database after work hours, which generated additional database logs. The free storage of the RDS DB instance is low due to these additional logs.
What should the company do to address this space constraint issue?

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

Answer: B

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

NEW QUESTION 5

After restoring an Amazon RDS snapshot from 3 days ago, a company's Development team cannot connect to the restored RDS DB instance. What is the likely cause of this problem?

A. The restored DB instance does not have Enhanced Monitoring enabled
B. The production DB instance is using a custom parameter group
C. The restored DB instance is using the default security group
D. The production DB instance is using a custom option group

Answer: C

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/rds-cannot-connect/
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_RestoreFromSnapshot.html

NEW QUESTION 6
A company has migrated a single MySQL database to Amazon Aurora. The production data is hosted in a DB cluster in VPC_PROD, and 12 testing environments are hosted in VPC_TEST using the same AWS account. Testing results in minimal changes to the test data. The Development team wants each environment refreshed nightly so each test database contains fresh production data every day.
Which migration approach will be the fastest and most cost-effective to implement?

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

Answer: A

NEW QUESTION 7
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 8
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 use-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 9
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 10
A company's database specialist disabled TLS on an Amazon DocumentDB cluster to perform benchmarking tests. A few days after this change was implemented, a database specialist trainee accidentally deleted multiple tables. The database specialist restored the database from available snapshots. An hour after restoring the cluster, the database specialist is still unable to connect to the new cluster endpoint.
What should the database specialist do to connect to the new, restored Amazon DocumentDB cluster?

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

Answer: A

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

NEW QUESTION 11
On AWS, a business is developing a web application. The application needs that the database supports concurrent read and write activities in several AWS Regions. Additionally, the database must communicate data changes across Regions as they occur. The application must be highly available and have a latency of less than a few hundred milliseconds.
Which solution satisfies these criteria?

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

Answer: A

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

NEW QUESTION 12
An electric utility company wants to store power plant sensor data in an Amazon DynamoDB table. The utility company has over 100 power plants and each power plant has over 200 sensors that send data every 2 seconds. The sensor data includes time with milliseconds precision, a value, and a fault attribute if the sensor is malfunctioning. Power plants are identified by a globally unique identifier. Sensors are identified by a unique identifier within each power plant. A database specialist needs to design the table to support an efficient method of finding all faulty sensors within a given power plant.
Which schema should the database specialist use when creating the DynamoDB table to achieve the fastest query time when looking for faulty sensors?

A. Use the plant identifier as the partition key and the measurement time as the sort ke
B. Create a global secondary index (GSI) with the plant identifier as the partition key and the fault attribute as the sort key.
C. Create a composite of the plant identifier and sensor identifier as the partition ke
D. Use the measurement time as the sort ke
E. Create a local secondary index (LSI) on the fault attribute.
F. Create a composite of the plant identifier and sensor identifier as the partition ke
G. Use the measurement time as the sort ke
H. Create a global secondary index (GSI) with the plant identifier as the partition key and the fault attribute as the sort key.
I. Use the plant identifier as the partition key and the sensor identifier as the sort ke
J. Create a local secondary index (LSI) on the fault attribute.

Answer: D

Explanation:
Plant id as partition key and Sensor id as a sort key. Fault can be identified quickly using the local secondary index and associated plant and sensor can be identified easily.

NEW QUESTION 13
A company is using Amazon RDS for PostgreSQL. The Security team wants all database connection requests to be logged and retained for 180 days. The RDS for PostgreSQL DB instance is currently using the default parameter group. A Database Specialist has identified that setting the log_connections parameter to 1 will enable connections logging.
Which combination of steps should the Database Specialist take to meet the logging and retention requirements? (Choose two.)

A. Update the log_connections parameter in the default parameter group
B. Create a custom parameter group, update the log_connections parameter, and associate the parameter with the DB instance
C. Enable publishing of database engine logs to Amazon CloudWatch Logs and set the event expiration to 180 days
D. Enable publishing of database engine logs to an Amazon S3 bucket and set the lifecycle policy to 180 days
E. Connect to the RDS PostgreSQL host and update the log_connections parameter in the postgresql.conf file

Answer: AE

NEW QUESTION 14
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 15
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 16
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 17
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 18
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 19

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 20

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 21
......

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons and Answers in PDF Format

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 FirstTry

* AWS-Certified-Database-Specialty Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The AWS-Certified-Database-Specialty Practice Test Here](https://www.surepassexam.com/AWS-Certified-Database-Specialty-exam-dumps.html)