DBS-C01 Q&As

AWS Certified Database - Specialty (DBS-C01)

Pass Amazon DBS-C01 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.certbus.com/aws-certified-database-specialty.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Amazon Official Exam Center

- **Instant Download** After Purchase
- **100% Money Back** Guarantee
- **365 Days** Free Update
- **800,000+** Satisfied Customers

QUESTION 1

A financial services company has an application deployed on AWS that uses an Amazon Aurora PostgreSQL DB cluster. A recent audit showed that no log files contained database administrator activity. A database specialist needs to recommend a solution to provide database access and activity logs. The solution should use the least amount of effort and have a minimal impact on performance.

Which solution should the database specialist recommend?

A. Enable Aurora Database Activity Streams on the database in synchronous mode. Connect the Amazon Kinesis data stream to Kinesis Data Firehose. Set the Kinesis Data Firehose destination to an Amazon S3 bucket.

B. Create an AWS CloudTrail trail in the Region where the database runs. Associate the database activity logs with the trail.

C. Enable Aurora Database Activity Streams on the database in asynchronous mode. Connect the Amazon Kinesis data stream to Kinesis Data Firehose. Set the Firehose destination to an Amazon S3 bucket.

D. Allow connections to the DB cluster through a bastion host only. Restrict database access to the bastion host and application servers. Push the bastion host logs to Amazon CloudWatch Logs using the CloudWatch Logs agent.

Correct Answer: C

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/DBActivityStreams.Ov erview.html

QUESTION 2

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.

Correct Answer: D

QUESTION 3

A company requires near-real-time notifications when changes are made to Amazon RDS DB security groups.

Which solution will meet this requirement with the LEAST operational overhead?

A. Configure an RDS event notification subscription for DB security group events.

B. Create an AWS Lambda function that monitors DB security group changes. Create an Amazon Simple Notification Service (Amazon SNS) topic for notification.

C. Turn on AWS CloudTrail. Configure notifications for the detection of changes to DB security groups.

D. Configure an Amazon CloudWatch alarm for RDS metrics about changes to DB security groups.

Correct Answer: A

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_Events.Messages.html#USER_Events.Messages.security-group

QUESTION 4

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. Set it to the capacity levels required for peak daytime throughput.

B. Use provisioned capacity. Set it to the capacity levels required for peak daytime throughput.

C. Use provisioned capacity. Create an AWS Application Auto Scaling policy to update capacity based on consumption.

D. Use on-demand capacity.

Correct Answer: C

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#HowItWorks.ProvisionedThroughput.Manual 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

QUESTION 5

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. Restore a new DB instance from the encrypted snapshot.

B. Modify the DB instance and enable encryption.

C. Restore a DB instance from the most recent automated snapshot and enable encryption.

D. Create an encrypted read replica of the DB instance. Promote the read replica to a standalone instance.

Correct Answer: A

https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/encrypt-an-existing- amazon-rds-for-postgresql-db-instance.html 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.

QUESTION 6

A company is using AWS CloudFormation to provision and manage infrastructure resources, including a production database. During a recent CloudFormation stack update, a database specialist observed that changes were made to a database resource that is named ProductionDatabase. The company wants to prevent changes to only ProductionDatabase during future stack updates.

Which stack policy will meet this requirement?

A.

```
{
    "Statement" : [
        {
            "Effect" : "Allow",
            "Action" : "Update:*",
            "Principal": "*",
            "Resource" : "*"
        },
        {
            "Effect" : "Deny",
            "Action" : "Update:*",
            "Principal": "*",
            "Resource" : "LogicalResourceId/ProductionDatabase"
        }
    ]
}
```

B.

```
{
    "Statement" : [
        {
            "Effect" : "Deny",
            "Action" : "Update:*",
            "Principal": "*",
            "Resource" : "LogicalResourceId/ProductionDatabase"
        }
    ]
}
```

C.

```
{
    "Statement" : [
        {
            "Effect" : "Deny",
            "Action" : "Update:*",
            "Principal": "*",
            "Resource" : "*"        },
        {
            "Effect" : "Deny",
            "Action" : "Update:*",
            "Principal": "*",
            "Resource" : "LogicalResourceId/ProductionDatabase"
        }
    ]
}
```

D.

```
{
    "Statement" : [
        {
            "Effect" : "Allow",
            "Action" : "Update:*",
            "Principal": "*",
            "Resource" : "*"
        },
        {
            "Effect" : "Deny",
            "Action" : "Delete:*",
            "Principal": "*",
            "Resource" : "LogicalResourceId/ProductionDatabase"
        }
    ]
}
```

A. Option A

B. Option B

C. Option C

D. Option D

Correct Answer: A

QUESTION 7

A financial company is hosting its web application on AWS. The application\\'s database is hosted on Amazon RDS for MySQL with automated backups enabled.

The application has caused a logical corruption of the database, which is causing the application to become unresponsive. The specific time of the corruption has been identified, and it was within the backup retention period.

How should a database specialist recover the database to the most recent point before corruption?

A. Use the point-in-time restore capability to restore the DB instance to the specified time. No changes to the application connection string are required.

B. Use the point-in-time restore capability to restore the DB instance to the specified time. Change the application connection string to the new, restored DB instance.

C. Restore using the latest automated backup. Change the application connection string to the new, restored DB instance.

D. Restore using the appropriate automated backup. No changes to the application connection string are required.

Correct Answer: B

The point-in-time restore capability of Amazon RDS for MySQL allows you to create a new DB instance with the same configuration as the original one, but with data restored to a specific time within your backup retention period. You can specify any time within your backup retention period, up to the last five minutes of your DB instance\\'s usage1. This feature is useful for recovering from logical corruption or user errors that affect your database. However, when you use the point-in-time restore capability, you are creating a new DB instance with a different endpoint. Therefore, you need to change the application connection string to point to the new, restored DB instance. You can also delete or rename the original DB instance if you no longer need it1

QUESTION 8

A database specialist is launching a test graph database using Amazon Neptune for the first time. The database specialist needs to insert millions of rows of test observations from a .csv file that is stored in Amazon S3. The database specialist has been using a series of API calls to upload the data to the Neptune DB instance.

Which combination of steps would allow the database specialist to upload the data faster? (Choose three.)

A. Ensure Amazon Cognito returns the proper AWS STS tokens to authenticate the Neptune DB instance to the S3 bucket hosting the CSV file.

B. Ensure the vertices and edges are specified in different .csv files with proper header column formatting.

C. Use AWS DMS to move data from Amazon S3 to the Neptune Loader.

D. Curl the S3 URI while inside the Neptune DB instance and then run the addVertex or addEdge commands.

E. Ensure an IAM role for the Neptune DB instance is configured with the appropriate permissions to allow access to the file in the S3 bucket.

F. Create an S3 VPC endpoint and issue an HTTP POST to the database\\'s loader endpoint.

Correct Answer: BEF

Explanation:

QUESTION 9

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

Which solution satisfies these criteria?

A. Leverage an Amazon Redshift data warehouse solution using a dense storage instance type while keeping all the data on local Amazon Redshift storage. Provision enough instances to support high demand.

B. Leverage an Amazon Redshift data warehouse solution using a dense storage instance to store the most recent data. Keep historical data on Amazon S3 and access it using the Amazon Redshift Spectrum layer. Provision enough instances to support high demand.

C. Leverage an Amazon Redshift data warehouse solution using a dense storage instance to store the most recent data. Keep historical data on Amazon S3 and access it using the Amazon Redshift Spectrum layer. Enable Amazon Redshift Concurrency Scaling.

D. Leverage an Amazon Redshift data warehouse solution using a dense storage instance to store the most recent data. Keep historical data on Amazon S3 and access it using the Amazon Redshift Spectrum layer. Leverage Amazon Redshift elastic resize.

Correct Answer: C

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

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.

Correct Answer: C

https://aws.amazon.com/about-aws/whats-new/2019/05/amazon-aurora-with-postgresql- compatibility-supports-database-activity-streams/

"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.LoggingAnd Monitoring.html

---

QUESTION 11

An online shopping company has a large inflow of shopping requests daily. As a result, there is a consistent load on the company\\'s Amazon RDS database. A database specialist needs to ensure the database is up and running at all times. The database specialist wants an automatic notification system for issues that may cause database downtime or for configuration changes made to the database.

What should the database specialist do to achieve this? (Choose two.)

A. Create an Amazon CloudWatch Events event to send a notification using Amazon SNS on every API call logged in AWS CloudTrail.

B. Subscribe to an RDS event subscription and configure it to use an Amazon SNS topic to send notifications.

C. Use Amazon SES to send notifications based on configured Amazon CloudWatch Events events.

D. Configure Amazon CloudWatch alarms on various metrics, such as FreeStorageSpace for the RDS instance.

E. Enable email notifications for AWS Trusted Advisor.

Correct Answer: BD

QUESTION 12

A company has a database monitoring solution that uses Amazon CloudWatch for its Amazon RDS for SQL Server

environment. The cause of a recent spike in CPU utilization was not determined using the standard metrics that were collected. The CPU spike caused the application to perform poorly, impacting users. A Database Specialist needs to determine what caused the CPU spike.

Which combination of steps should be taken to provide more visibility into the processes and queries running during an increase in CPU load? (Choose two.)

A. Enable Amazon CloudWatch Events and view the incoming T-SQL statements causing the CPU to spike.

B. Enable Enhanced Monitoring metrics to view CPU utilization at the RDS SQL Server DB instance level.

C. Implement a caching layer to help with repeated queries on the RDS SQL Server DB instance.

D. Use Amazon QuickSight to view the SQL statement being run.

E. Enable Amazon RDS Performance Insights to view the database load and filter the load by waits, SQL statements, hosts, or users.

Correct Answer: BE

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

QUESTION 13

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

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

A. Run the master in Amazon Aurora MySQL. Create 12 clones in VPC_TEST, and script the clones to be deleted and re-created nightly.

B. Run the master in Amazon Aurora MySQL. Take a nightly snapshot, and restore it into 12 databases in VPC_TEST using Aurora Serverless.

C. Run the master in Amazon Aurora MySQL. Create 12 Aurora Replicas in VPC_TEST, and script the replicas to be deleted and re-created nightly.

D. Run the master in Amazon Aurora MySQL using Aurora Serverless. Create 12 clones in VPC_TEST, and script the clones to be deleted and re-created nightly.

Correct Answer: A

QUESTION 14

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. This must be enabled expressly using the AWS Management Console.

C. Applying minor version upgrades requires sufficient free space.

D. The AWS CLI command did not include an apply-immediately parameter.

E. Aurora has detected a breaking change in the new minor version and has automatically rejected the upgrade.

Correct 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#AuroraMySQL.Updates.AMVU

QUESTION 15

Recently, an ecommerce business transferred one of its SQL Server databases to an Amazon RDS for SQL Server Enterprise Edition database instance. The corporation anticipates an increase in read traffic as a result of an approaching sale. To accommodate the projected read load, a database professional must establish a read replica of the database instance.

Which procedures should the database professional do prior to establishing the read replica? (Select two.)

A. Identify a potential downtime window and stop the application calls to the source DB instance.

B. Ensure that automatic backups are enabled for the source DB instance.

C. Ensure that the source DB instance is a Multi-AZ deployment with Always ON Availability Groups.

D. Ensure that the source DB instance is a Multi-AZ deployment with SQL Server Database Mirroring (DBM).

E. Modify the read replica parameter group setting and set the value to 1.

Correct Answer: BC

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/SQLServer.ReadReplicas.ht ml