



Free Questions for DBS-C01 by certscare

Shared by Spears on 18-01-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

A global company is creating an application. The application must be highly available. The company requires an RTO and an RPO of less than 5 minutes. The company needs a database that will provide the ability to set up an active-active configuration and near real-time synchronization of data across tables in multiple AWS Regions.

Which solution will meet these requirements?

Options:

- A) Amazon RDS for MariaDB with cross-Region read replicas
- B) Amazon RDS With a Multi-AZ deployment
- C) Amazon DynamoDB global tables
- D) Amazon DynamoDB With a global secondary index (GSI)

Answer:

C

Explanation:

Amazon DynamoDB global tables are a fully managed, serverless, multi-Region, and multi-active database option that delivers fast and localized read and write performance for massively scaled global applications. Global tables provide you 99.999% availability, increased application resiliency, and improved business continuity¹. As global tables replicate your Amazon DynamoDB tables automatically across your choice of AWS Regions, you can achieve fast, local read and write performance¹. In the rare event of a single Region failure, there is no need for a database failover as global tables' multi-active architecture allows customers to read and write to any replica table¹. Global tables also eliminate the difficult work of replicating data between Regions and resolving update conflicts for multi-active workloads¹.

Using Amazon DynamoDB global tables, the company can:

Set up an active-active configuration and near real-time synchronization of data across tables in multiple AWS Regions.

Meet the RTO and RPO of less than 5 minutes, as global tables provide automatic failover and zero data loss in case of a Region failure.

Achieve high availability, durability, and scalability for their application.

Question 2

Question Type: MultipleChoice

A company has an existing system that uses a single-instance Amazon DocumentDB (with MongoDB compatibility) cluster. Read requests account for 75% of the system queries. Write requests are expected to increase by 50% after an upcoming global release. A database specialist needs to design a solution that improves the overall database performance without creating additional application

overhead.

Which solution will meet these requirements?

Options:

- A) Recreate the cluster with a shared cluster volume. Add two instances to serve both read requests and write requests.
- B) Add one read replica instance. Activate a shared cluster volume. Route all read queries to the read replica instance.
- C) Add one read replica instance. Set the read preference to secondary preferred.
- D) Add one read replica instance. Update the application to route all read queries to the read replica instance.

Answer:

C

Explanation:

By default, an application directs its read operations to the primary member in a replica set (i.e. read preference mode 'primary'). But, clients can specify a read preference to send read operations to secondaries. <https://www.mongodb.com/docs/manual/core/read-preference/>

A read replica instance is an Amazon DocumentDB instance that supports only read operations. An Amazon DocumentDB cluster can have up to 15 replicas in addition to the primary instance. Having multiple replicas enables you to distribute read workloads and increase

the read throughput of your cluster¹.

The read preference option determines how your MongoDB client or driver routes read requests to instances in your Amazon DocumentDB cluster. The secondary preferred option instructs the client to route read queries to the replicas, unless none are available, in which case the queries are routed to the primary instance². This option can improve the overall database performance by offloading read requests from the primary instance, which handles all write requests, and balancing them across the replicas.

Using this solution, the company can:

Improve the overall database performance by adding one read replica instance and setting the read preference to secondary preferred.

Avoid creating additional application overhead by using the built-in read preference capabilities of the MongoDB client or driver, without having to update the application code or connection string.

Handle the expected increase in write requests by freeing up resources on the primary instance.

Question 3

Question Type: MultipleChoice

A company has more than 100 AWS accounts that need Amazon RDS instances. The company wants to build an automated solution to deploy the RDS instances with specific compliance parameters. The data does not need to be replicated. The company needs to create the databases within 1 day

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

Options:

- A) Create RDS resources by using AWS CloudFormation. Share the CloudFormation template with each account.
- B) Create an RDS snapshot. Share the snapshot With each account Deploy the snapshot into each account
- C) use AWS CloudFormation to create RDS instances in each account. Run AWS Database Migration Service (AWS DMS) replication to each of the created instances.
- D) Create a script by using the AWS CLI to copy the ROS instance into the other accounts from a template account.

Answer:

A

Explanation:

AWS CloudFormation is a service that helps you model and set up your AWS resources so that you can spend less time managing those resources and more time focusing on your applications that run in AWS. You create a template that describes all the AWS resources that you want (like Amazon RDS instances), and CloudFormation takes care of provisioning and configuring those resources for you¹.

Using AWS CloudFormation to create RDS resources and share the template with each account is a suitable solution for the company's requirements, as it can:

Automate the deployment of RDS instances with specific compliance parameters, such as security groups, encryption, backup settings, etc.

Reduce the operational overhead and human errors of manually creating RDS instances in each account.

Enable the company to create the databases within one day, as CloudFormation can provision resources in parallel and in a consistent manner.

Question 4

Question Type: MultipleChoice

A manufacturing company stores its inventory details in an Amazon DynamoDB table in the us-east-2 Region. According to new compliance and regulatory policies, the company is required to back up all of its tables nightly and store these backups in the us-west-2 Region for disaster recovery for 1 year

Which solution MOST cost-effectively meets these requirements?

Options:

A) Convert the existing DynamoDB table into a global table and create a global table replica in the us-west-2 Region.

- B)** Use AWS Backup to create a backup plan. Configure cross-Region replication in the plan and assign the DynamoDB table to this plan
- C)** Create an on-demand backup of the DynamoDB table and restore this backup in the us-west-2 Region.
- D)** Enable Amazon S3 Cross-Region Replication (CRR) on the S3 bucket where DynamoDB on-demand backups are stored.

Answer:

B

Explanation:

AWS Backup is a fully managed service that simplifies data protection across AWS services, in the cloud, and on premises. You can use AWS Backup to create backup plans that define how and when your backups are created, how long they are stored, and where they are replicated¹. You can also use AWS Backup to monitor and audit your backup activity.

Using AWS Backup to create a backup plan and configure cross-Region replication is a cost-effective solution for the company's requirements, as it can automate the nightly backup of the DynamoDB table and store the backups in the us-west-2 Region for one year. You can specify the source and destination Regions, the backup vault, and the retention period for your cross-Region replication rule in your backup plan². You can also assign your DynamoDB table to your backup plan by using a resource assignment³.

Question 5

Question Type: MultipleChoice

A company uses a large, growing, and high performance on-premises Microsoft SQL Server instance With an Always On availability group cluster size of 120 TIE. The company uses a third-party backup product that requires system-level access to the databases. The company will continue to use this third-party backup product in the future.

The company wants to move the DB cluster to AWS with the least possible downtime and data loss. The company needs a 2 Gbps connection to sustain Always On asynchronous data replication between the company's data center and AWS.

Which combination of actions should a database specialist take to meet these requirements? (Select THREE.)

Establish an AWS Direct Connect hosted connection between the companfs data center and AWS

Create an AWS Site-to-Site VPN connection between the companVs data center and AWS over the internet

Use AWS Database Migration Service (AWS DMS) to migrate the on-premises SQL Server databases to Amazon RDS for SQL Server
Configure Always On availability groups for SQL Server.

Deploy a new SQL Server Always On availability group DB cluster on Amazon EC2 Configure Always On distributed availability groups between the on-premises DB cluster and the AWS DB cluster_ Fail over to the AWS DB cluster when it is time to migrate.

Grant system-level access to the third-party backup product to perform backups of the Amazon RDS for SQL Server DB instance.

Configure the third-party backup product to perform backups of the DB cluster on Amazon EC2.

Options:

A) Establish an AWS Direct Connect hosted connection between the company's data center and AWS. This will provide a secure and

high-bandwidth connection for the Always On data replication and minimize the network latency and data loss.

D) Deploy a new SQL Server Always On availability group DB cluster on Amazon EC2. Configure Always On distributed availability groups between the on-premises DB cluster and the AWS DB cluster. Fail over to the AWS DB cluster when it is time to migrate. This will allow the company to use the same SQL Server version and edition as on-premises, and leverage the distributed availability group feature to span two separate availability groups across different locations. The failover process will be fast and seamless, with minimal downtime and data loss.

F) Configure the third-party backup product to perform backups of the DB cluster on Amazon EC2. This will enable the company to continue using their existing backup solution, which requires system-level access to the databases. Amazon RDS for SQL Server does not support system-level access, so it is not a suitable option for this requirement.

Answer:

A, D, F

Explanation:

The best combination of actions to meet the company's requirements are:

Question 6

Question Type: MultipleChoice

A database specialist needs to reduce the cost of an application's database. The database is running on a Multi-AZ deployment of an Amazon ROS for Microsoft SQL Server DB instance. The application requires the database to support stored procedures, SQL Server Wire Protocol (TDS), and T-SQC The database must also be highly available. The database specialist is using AWS Database Migration Service (AWS DMS) to migrate the database to a new data store.

Which solution will reduce the cost of the database with the LEAST effort?

Options:

- A)** Use AWS Database Migration Service (DMS) to migrate to an RDS for MySQL Multi-AZ database. Update the application code to use the features of MySQL that correspond to SQL Server. Update the application to use the MySQL port.
- B)** use AWS Database Migration Serve (OMS) to migrate to an RDS for PostgreSQL Multi-AZ database. Turn on the SQL_COMPAT optional extension within the database to allow the required features. Update the application to use the PostgreSQL port
- C)** Use AWS Database Migration Service (OMS) to migrate to an RDS for SQL Server Single-AZ database. Update the application to use the new database endpoint
- D)** Use AWS Database Migration Service (DMS) to migrate the database to Amazon Aurora PostgreSQL_ Turn on Babelfish for Aurora PostgreSQL_ Update the application to use the Babelfish TDS port.

Answer:

D

Explanation:

Amazon Aurora PostgreSQL is a fully managed, compatible, and scalable relational database service that supports the PostgreSQL open source database engine¹. Amazon Aurora PostgreSQL can reduce the cost of running a database compared to Amazon RDS for SQL Server, which is a commercial database engine that requires licensing fees².

Babelfish for Aurora PostgreSQL is a new capability for Amazon Aurora PostgreSQL-Compatible Edition that enables Aurora to understand commands from applications written for Microsoft SQL Server³. Babelfish allows Aurora PostgreSQL to support the SQL Server wire-level protocol (TDS) and commonly used T-SQL language and semantics, which reduces the amount of code changes required to migrate applications from SQL Server to Aurora³. Babelfish also provides high availability by replicating data across multiple Availability Zones in a single AWS Region⁴.

Using AWS Database Migration Service (DMS) to migrate the database to Amazon Aurora PostgreSQL and turning on Babelfish for Aurora PostgreSQL is a suitable solution for reducing the cost of the database with the least effort, as it can preserve the compatibility and availability of the database while minimizing the code changes in the application. The only change required in the application is to update the connection string to use the Babelfish TDS port, which is 1433 by default.

Question 7

Question Type: MultipleChoice

A marketing company is developing an application to track responses to email message campaigns. The company needs a database storage solution that is optimized to work with highly connected dat

a. The database needs to limit connections and programmatic access to the data by using IAM policies.

Which solution will meet these requirements?

Options:

- A) Amazon ElastiCache for Redis cluster
- B) Amazon Aurora MySQL DB cluster
- C) Amazon DynamoDB table
- D) Amazon Neptune DB cluster

Answer:

D

Explanation:

Amazon Neptune is a fast, reliable, fully managed graph database service that makes it easy to build and run applications that work with highly connected data sets¹. Graph databases are designed to store and query data that has complex relationships and interconnections, such as social networks, recommendation engines, fraud detection, and knowledge graphs². Amazon Neptune supports two popular graph models: Property Graph and Resource Description Framework (RDF), and their respective query languages: Apache TinkerPop Gremlin and SPARQL².

Amazon Neptune also supports IAM policies to control access to the database resources and operations. You can use IAM database authentication to authenticate users and applications that connect to a Neptune DB cluster. IAM database authentication works with MySQL and PostgreSQL database clients. You can also use IAM roles to manage access to Neptune from other AWS services, such as Amazon EC2, AWS Lambda, and Amazon SageMaker2.

Therefore, Amazon Neptune DB cluster is a suitable solution for the marketing company's requirements, as it can provide a graph database storage solution that is optimized for highly connected data and can limit connections and programmatic access by using IAM policies.

Question 8

Question Type: MultipleChoice

A financial services company is using AWS Database Migration Service (AWS DMS) to migrate its databases from on-premises to AWS. A database administrator is working on replicating a database to AWS from on-premises using full load and change data capture (CDC). During the CDC replication, the database administrator observed that the target latency was high and slowly increasing-

What could be the root causes for this high target latency? (Select TWO.)

Options:

- A) There was ongoing maintenance on the replication instance
- B) The source endpoint was changed by modifyng the task
- C) Loopback changes had affected the source and target instances-
- D) There was no primary key or index in the target database.
- E) There were resource bottlenecks in the replication instance

Answer:

D, E

Explanation:

Target latency is the amount of time that AWS DMS takes to apply changes from the source database to the target database¹. High target latency can indicate performance issues or replication errors in the AWS DMS task.

One possible cause of high target latency is the lack of a primary key or index in the target database. A primary key or index helps AWS DMS identify and apply changes to the corresponding rows in the target database. Without a primary key or index, AWS DMS has to scan the entire table to find the matching rows, which can increase the target latency and consume more CPU and memory resources².

Another possible cause of high target latency is the resource bottlenecks in the replication instance. The replication instance is the compute resource that runs the AWS DMS task and connects to the source and target endpoints. If the replication instance is under-provisioned or overloaded, it can affect the replication performance and cause high target latency. Some factors that can contribute to resource bottlenecks are insufficient network bandwidth, low disk space, high CPU utilization, or large transaction sizes³.

To Get Premium Files for DBS-C01 Visit

<https://www.p2pexams.com/products/dbs-c01>

For More Free Questions Visit

<https://www.p2pexams.com/amazon/pdf/dbs-c01>

