

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.lead4pass.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 media company is using Amazon RDS for PostgreSQL to store user data. The RDS DB instance currently has a publicly accessible setting enabled and is hosted in a public subnet. Following a recent AWS Well-Architected Framework review, a Database Specialist was given new security requirements.

Only certain on-premises corporate network IPs should connect to the DB instance. Connectivity is allowed from the corporate network only.

Which combination of steps does the Database Specialist need to take to meet these new requirements? (Choose three.)

- A. Modify the `pg_hba.conf` file. Add the required corporate network IPs and remove the unwanted IPs.
- B. Modify the associated security group. Add the required corporate network IPs and remove the unwanted IPs.
- C. Move the DB instance to a private subnet using AWS DMS.
- D. Enable VPC peering between the application host running on the corporate network and the VPC associated with the DB instance.
- E. Disable the publicly accessible setting.
- F. Connect to the DB instance using private IPs and a VPN.

Correct Answer: BEF

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_VPC.WorkingWithRDSThroughVPC.html#USER_VPC.Hiding

QUESTION 2

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?

- 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.

Correct Answer: B

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 uswest-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 3

A company is developing an application that performs intensive in-memory operations on advanced data structures such as sorted sets. The application requires sub-millisecond latency for reads and writes. The application occasionally must run a group of commands as an ACID-compliant operation. A database specialist is setting up the database for this application. The database specialist needs the ability to create a new database cluster from the latest backup of the production cluster.

Which type of cluster should the database specialist create to meet these requirements?

- A. Amazon ElastiCache for Memcached
- B. Amazon Neptune
- C. Amazon ElastiCache for Redis
- D. Amazon DynamoDB Accelerator (DAX)

Correct Answer: C

Explanation: <https://aws.amazon.com/elasticache/redis-vs-memcached/>
<https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/elasticache-use-cases.html#elasticache-for-redis-use-cases-gaming>

QUESTION 4

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

Correct Answer: ABC

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

QUESTION 5

A database specialist is responsible for an Amazon RDS for MySQL DB instance with one read replica. The DB instance and the read replica are assigned to the default parameter group. The database team currently runs test queries against a read replica. The database team wants to create additional tables in the read replica that will only be accessible from the read replica to benefit the tests.

Which should the database specialist do to allow the database team to create the test tables?

- A. Contact AWS Support to disable read-only mode on the read replica. Reboot the read replica. Connect to the read replica and create the tables.
- B. Change the `read_only` parameter to false (`read_only=0`) in the default parameter group of the read replica. Perform a reboot without failover. Connect to the read replica and create the tables using the `local_only` MySQL option.
- C. Change the `read_only` parameter to false (`read_only=0`) in the default parameter group. Reboot the read replica. Connect to the read replica and create the tables.
- D. Create a new DB parameter group. Change the `read_only` parameter to false (`read_only=0`). Associate the read replica with the new group. Reboot the read replica. Connect to the read replica and create the tables.

Correct Answer: D

Explanation: <https://aws.amazon.com/premiumsupport/knowledge-center/rds-read-replica/>

QUESTION 6

A business's production database is hosted on a single-node Amazon RDS for MySQL DB instance. The database instance is hosted in a United States AWS Region.

A week before a significant sales event, a fresh database maintenance update is released. The maintenance update has been designated as necessary. The firm wants to minimize the database instance's downtime and requests that a

database expert make the database instance highly accessible until the sales event concludes.

Which solution will satisfy these criteria?

- A. Defer the maintenance update until the sales event is over.
- B. Create a read replica with the latest update. Initiate a failover before the sales event.
- C. Create a read replica with the latest update. Transfer all read-only traffic to the read replica during the sales event.
- D. Convert the DB instance into a Multi-AZ deployment. Apply the maintenance update.

Correct Answer: D

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

QUESTION 7

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.

Correct Answer: ACD

QUESTION 8

A database professional is tasked with the task of migrating 25 GB of data files from an on-premises storage system to an Amazon Neptune database.

Which method of data loading is the FASTEST?

- A. Upload the data to Amazon S3 and use the Loader command to load the data from Amazon S3 into the Neptune database.
- B. Write a utility to read the data from the on-premises storage and run INSERT statements in a loop to load the data into the Neptune database.
- C. Use the AWS CLI to load the data directly from the on-premises storage into the Neptune database.
- D. Use AWS DataSync to load the data directly from the on-premises storage into the Neptune database.

Correct Answer: A

1. Copy the data files to an Amazon Simple Storage Service (Amazon S3) bucket.

2.

Create an IAM role with Read and List access to the bucket.

3.

Create an Amazon S3 VPC endpoint.

4.

Start the Neptune loader by sending a request via HTTP to the Neptune DB instance.

5.

The Neptune DB instance assumes the IAM role to load the data from the bucket.

QUESTION 9

A small startup company is looking to migrate a 4 TB on-premises MySQL database to AWS using an Amazon RDS for MySQL DB instance.

Which strategy would allow for a successful migration with the LEAST amount of downtime?

- A. Deploy a new RDS for MySQL DB instance and configure it for access from the on-premises data center. Use the mysqldump utility to create an initial snapshot from the on-premises MySQL server, and copy it to an Amazon S3 bucket. Import the snapshot into the DB instance utilizing the MySQL utilities running on an Amazon EC2 instance. Immediately point the application to the DB instance.
- B. Deploy a new Amazon EC2 instance, install the MySQL software on the EC2 instance, and configure networking for access from the on-premises data center. Use the mysqldump utility to create a snapshot of the on-premises MySQL server. Copy the snapshot into the EC2 instance and restore it into the EC2 MySQL instance. Use AWS DMS to migrate data into a new RDS for MySQL DB instance. Point the application to the DB instance.
- C. Deploy a new Amazon EC2 instance, install the MySQL software on the EC2 instance, and configure networking for access from the on-premises data center. Use the mysqldump utility to create a snapshot of the on-premises MySQL server. Copy the snapshot into an Amazon S3 bucket and import the snapshot into a new RDS for MySQL DB instance using the MySQL utilities running on an EC2 instance. Point the application to the DB instance.
- D. Deploy a new RDS for MySQL DB instance and configure it for access from the on-premises data center. Use the mysqldump utility to create an initial snapshot from the on-premises MySQL server, and copy it to an Amazon S3 bucket. Import the snapshot into the DB instance using the MySQL utilities running on an Amazon EC2 instance. Establish replication into the new DB instance using MySQL replication. Stop application access to the on-premises MySQL server and let the remaining transactions replicate over. Point the application to the DB instance.

Correct Answer: B

QUESTION 10

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

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

- A. Use AWS CloudFormation templates. Deploy a stack with the DB cluster for each development group. Delete the stack at the end of the development cycle.
- B. Use the Aurora DB cloning feature. Deploy a single development and test Aurora DB instance, and create clone instances for the development groups. Delete the clones at the end of the development cycle.
- C. Use Aurora Replicas. From the master automatic pause compute capacity option, create replicas for each development group, and promote each replica to master. Delete the replicas at the end of the development cycle.

D. Use Aurora Serverless. Restore current Aurora snapshot and deploy to a serverless cluster for each development group. Enable the option to pause the compute capacity on the cluster and set an appropriate timeout.

Correct Answer: B

Explanation: Aurora Serverless is not compatible to all Aurora provisioned engine version. However, you can do clone with most engine version. Meanwhile, I also consider the performance while restoring snapshot to Aurora Serverless. <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-serverless-how-it-works.html#aurora-serverless-how-it-works-pause-resume> <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-serverless.html#aurora-serverless-use-cases>

QUESTION 11

A company is load testing its three-tier production web application deployed with an AWS CloudFormation template on AWS. The Application team is making changes to deploy additional Amazon EC2 and AWS Lambda resources to expand the load testing capacity. A Database Specialist wants to ensure that the changes made by the Application team will not change the Amazon RDS database resources already deployed.

Which combination of steps would allow the Database Specialist to accomplish this? (Choose two.)

- A. Review the stack drift before modifying the template
- B. Create and review a change set before applying it
- C. Export the database resources as stack outputs
- D. Define the database resources in a nested stack
- E. Set a stack policy for the database resources

Correct Answer: BE

https://docs.amazonaws.cn/en_us/AWSCloudFormation/latest/UserGuide/best-practices.html#cf-best-practices-changesets

QUESTION 12

A global company is developing an application across multiple AWS Regions. The company needs a database solution with low latency in each Region and automatic disaster recovery. The database must be deployed in an active-active configuration with automatic data synchronization between Regions.

Which solution will meet these requirements with the LOWEST latency?

- A. Amazon RDS with cross-Region read replicas
- B. Amazon DynamoDB global tables
- C. Amazon Aurora global database
- D. Amazon Athena and Amazon S3 with S3 Cross Region Replication

Correct Answer: B

QUESTION 13

An online retailer uses Amazon DynamoDB for its product catalog and order data. Some popular items have led to frequently accessed keys in the data, and the company is using DynamoDB Accelerator (DAX) as the caching solution to cater to the frequently accessed keys. As the number of popular products is growing, the company realizes that more items need to be cached. The company observes a high cache miss rate and needs a solution to address this issue.

What should a database specialist do to accommodate the changing requirements for DAX?

- A. Increase the number of nodes in the existing DAX cluster.
- B. Create a new DAX cluster with more nodes. Change the DAX endpoint in the application to point to the new cluster.
- C. Create a new DAX cluster using a larger node type. Change the DAX endpoint in the application to point to the new cluster.
- D. Modify the node type in the existing DAX cluster.

Correct Answer: C

Explanation:

QUESTION 14

A retail company uses Amazon Redshift Spectrum to run complex analytical queries on objects that are stored in an Amazon S3 bucket. The objects are joined with multiple dimension tables that are stored in an Amazon Redshift database. The company uses the database to create monthly and quarterly aggregated reports. Users who attempt to run queries are reporting the following error message: error: Spectrum Scan Error: Access throttled.

Which solution will resolve this error?

- A. Check file sizes of fact tables in Amazon S3, and look for large files. Break up large files into smaller files of equal size between 100 MB and 1 GB
- B. Reduce the number of queries that users can run in parallel.
- C. Check file sizes of fact tables in Amazon S3, and look for small files. Merge the small files into larger files of at least 64 MB in size.
- D. Review and optimize queries that submit a large aggregation step to Redshift Spectrum.

Correct Answer: C

Explanation: <https://docs.aws.amazon.com/redshift/latest/dg/c-spectrum-troubleshooting.html#spectrum-troubleshooting-access-throttled> <https://docs.aws.amazon.com/redshift/latest/dg/c-spectrum-troubleshooting.html> <https://docs.aws.amazon.com/redshift/latest/dg/c-spectrum-troubleshooting.html#spectrum-troubleshooting-access-throttled>

QUESTION 15

A company has an application that uses an Amazon DynamoDB table as its data store. During normal business days, the throughput requirements from the application are uniform and consist of 5 standard write calls per second to the

DynamoDB table. Each write call has 2 KB of data.

For 1 hour each day, the company runs an additional automated job on the DynamoDB table that makes 20 write requests per second. No other application writes to the DynamoDB table. The DynamoDB table does not have to meet any additional capacity requirements.

How should a database specialist configure the DynamoDB table's capacity to meet these requirements MOST cost-effectively?

- A. Use DynamoDB provisioned capacity with 5 WCUs and auto scaling.
- B. Use DynamoDB provisioned capacity with 5 WCUs and a write-through cache that DynamoDB Accelerator (DAX) provides.
- C. Use DynamoDB provisioned capacity with 10 WCUs and auto scaling.
- D. Use DynamoDB provisioned capacity with 10 WCUs and no auto scaling.

Correct Answer: C

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

[DBS-C01 PDF Dumps](#)

[DBS-C01 Practice Test](#)

[DBS-C01 Study Guide](#)