

AWS DBS-C01

AWS DATABASE SPECIALTY CERTIFICATION QUESTIONS & ANSWERS

Exam Summary – Syllabus – Questions

DBS-C01

[AWS Certified Database - Specialty](#)

65 Questions Exam – 750 / 1000 Cut Score – Duration of 180 minutes

www.VMExam.com

Table of Contents

Know Your DBS-C01 Certification Well:	2
AWS DBS-C01 Database Specialty Certification Details:	2
DBS-C01 Syllabus:.....	3
Workload-Specific Database Design - 26%	3
Deployment and Migration - 20%	3
Management and Operations - 18%	4
Monitoring and Troubleshooting - 18%	4
Database Security - 18%	5
AWS DBS-C01 Sample Questions:.....	6
Study Guide to Crack AWS Database Specialty DBS- C01 Exam:.....	11

Know Your DBS-C01 Certification Well:

The DBS-C01 is best suitable for candidates who want to gain knowledge in the AWS Specialty. Before you start your DBS-C01 preparation you may struggle to get all the crucial Database Specialty materials like DBS-C01 syllabus, sample questions, study guide.

But don't worry the DBS-C01 PDF is here to help you prepare in a stress free manner.

The PDF is a combination of all your queries like-

- What is in the DBS-C01 syllabus?
- How many questions are there in the DBS-C01 exam?
- Which Practice test would help me to pass the DBS-C01 exam at the first attempt?

Passing the DBS-C01 exam makes you AWS Certified Database - Specialty. Having the Database Specialty certification opens multiple opportunities for you. You can grab a new job, get a higher salary or simply get recognition within your current organization.

AWS DBS-C01 Database Specialty Certification Details:

Exam Name	AWS Certified Database - Specialty
Exam Code	DBS-C01
Exam Price	\$300 USD
Duration	180 minutes
Number of Questions	65
Passing Score	750 / 1000
Recommended Training / Books	AWS Database Offerings Planning and Designing Databases on AWS Exam Readiness - AWS Certified Database - Specialty
Schedule Exam	PEARSON VUE
Sample Questions	AWS DBS-C01 Sample Questions
Recommended Practice	AWS Certified Database - Specialty Practice Test

DBS-C01 Syllabus:

Section	Objectives
Workload-Specific Database Design - 26%	
Select appropriate database services for specific types of data and workloads.	<ul style="list-style-type: none"> - Differentiate between ACID vs. BASE workloads - Explain appropriate uses of types of databases (e.g., relational, key-value, document, in-memory, graph, time series, ledger) - Identify use cases for persisted data vs. ephemeral data
Determine strategies for disaster recovery and high availability.	<ul style="list-style-type: none"> - Select Region and Availability Zone placement to optimize database performance - Determine implications of Regions and Availability Zones on disaster recovery/high availability strategies - Differentiate use cases for read replicas and Multi-AZ deployments
Design database solutions for performance, compliance, and scalability.	<ul style="list-style-type: none"> - Recommend serverless vs. instance-based database architecture - Evaluate requirements for scaling read replicas - Define database caching solutions - Evaluate the implications of partitioning, sharding, and indexing - Determine appropriate instance types and storage options - Determine auto-scaling capabilities for relational and NoSQL databases - Determine the implications of Amazon DynamoDB adaptive capacity - Determine data locality based on compliance requirements
Compare the costs of database solutions.	<ul style="list-style-type: none"> - Determine cost implications of Amazon DynamoDB capacity units, including on-demand vs. provisioned capacity - Determine costs associated with instance types and automatic scaling - Design for costs including high availability, backups, multi-Region, Multi-AZ, and storage type options - Compare data access costs
Deployment and Migration - 20%	
Automate database solution deployments.	<ul style="list-style-type: none"> - Evaluate application requirements to determine components to deploy - Choose appropriate deployment tools and services (e.g., AWS CloudFormation, AWS CLI)

Section	Objectives
Determine data preparation and migration strategies.	<ul style="list-style-type: none"> - Determine the data migration method (e.g., snapshots, replication, restore) - Evaluate database migration tools and services (e.g., AWS DMS, native database tools) - Prepare data sources and targets - Determine schema conversion methods (e.g., AWS Schema Conversion Tool) - Determine heterogeneous vs. homogeneous migration strategies
Execute and validate data migration.	<ul style="list-style-type: none"> - Design and script data migration - Run data extraction and migration scripts - Verify the successful load of data
Management and Operations - 18%	
Determine maintenance tasks and processes.	<ul style="list-style-type: none"> - Account for the AWS shared responsibility model for database services - Determine appropriate maintenance window strategies - Differentiate between major and minor engine upgrades
Determine backup and restore strategies.	<ul style="list-style-type: none"> - Identify the need for automatic and manual backups/snapshots - Differentiate backup and restore strategies (e.g., full backup, point-in-time, encrypting backups cross-Region) - Define retention policies - Correlate the backup and restore to recovery point objective (RPO) and recovery time objective (RTO) requirements
Manage the operational environment of a database solution.	<ul style="list-style-type: none"> - Orchestrate the refresh of lower environments - Implement configuration changes (e.g., in Amazon RDS option/parameter groups or Amazon DynamoDB indexing changes) - Automate operational tasks - Take action based on AWS Trusted Advisor reports
Monitoring and Troubleshooting - 18%	
Determine monitoring and alerting strategies.	<ul style="list-style-type: none"> - Evaluate monitoring tools (e.g., Amazon CloudWatch, Amazon RDS Performance Insights, database native) - Determine appropriate parameters and thresholds for alert conditions - Use tools to notify users when thresholds are breached (e.g., Amazon SNS, Amazon SQS, Amazon CloudWatch dashboards)
Troubleshoot and resolve common database issues.	<ul style="list-style-type: none"> - Identify, evaluate, and respond to categories of failures (e.g., troubleshoot connectivity; instance, storage, and

Section	Objectives
	partitioning issues) - Automate responses when possible
Optimize database performance.	- Troubleshoot database performance issues - Identify appropriate AWS tools and services for database optimization - Evaluate the configuration, schema design, queries, and infrastructure to improve performance
Database Security - 18%	
Encrypt data at rest and in transit.	- Encrypt data in relational and NoSQL databases - Apply SSL connectivity to databases - Implement key management (e.g., AWS KMS, AWS CloudHSM)
Evaluate auditing solutions.	- Determine auditing strategies for structural/schema changes (e.g., DDL) - Determine auditing strategies for data changes (e.g., DML) - Determine auditing strategies for data access (e.g., queries) - Determine auditing strategies for infrastructure changes (e.g., AWS CloudTrail) - Enable the export of database logs to Amazon CloudWatch Logs
Determine access control and authentication mechanisms.	- Recommend authentication controls for users and roles (e.g., IAM, native credentials, Active Directory) - Recommend authorization controls for users (e.g., policies)
Recognize potential security vulnerabilities within database solutions.	- Determine security group rules and NACLs for database access - Identify relevant VPC configurations (e.g., VPC endpoints, public vs. private subnets, demilitarized zone) - Determine appropriate storage methods for sensitive data

AWS DBS-C01 Sample Questions:

Question: 1

A global company wants to run an application in several AWS Regions to support a global user base.

The application will need a database that can support a high volume of low-latency reads and writes that is expected to vary over time. The data must be shared across all of the Regions to support dynamic company-wide reports.

Which database meets these requirements?

- a) Use Amazon Aurora Serverless and configure endpoints in each Region.
- b) Use Amazon RDS for MySQL and deploy read replicas in an auto scaling group in each Region.
- c) Use Amazon DocumentDB (with MongoDB compatibility) and configure read replicas in an auto scaling group in each Region.
- d) Use Amazon DynamoDB global tables and configure DynamoDB auto scaling for the tables.

Answer: d

Question: 2

A media company is running a critical production application that uses Amazon RDS for PostgreSQL with Multi-AZ deployments. The database size is currently 25 TB.

The IT director wants to migrate the database to Amazon Aurora PostgreSQL with minimal effort and minimal disruption to the business.

What is the best migration strategy to meet these requirements?

- a) Use the AWS Schema Conversion Tool (AWS SCT) to copy the database schema from RDS for PostgreSQL to an Aurora PostgreSQL DB cluster. Create an AWS DMS task to copy the data.
- b) Create a script to continuously back up the RDS for PostgreSQL instance using `pg_dump`, and restore the backup to an Aurora PostgreSQL DB cluster using `pg_restore`.
- c) Create a read replica from the existing production RDS for PostgreSQL instance. Check that the replication lag is zero and then promote the read replica as a standalone Aurora PostgreSQL DB cluster.
- d) Create an Aurora Replica from the existing production RDS for PostgreSQL instance. Stop the writes on the master, check that the replication lag is zero, and then promote the Aurora Replica as a standalone Aurora PostgreSQL DB cluster.

Answer: d

Question: 3

A company's security department has mandated that their existing Amazon RDS for MySQL DB instance be encrypted at rest.

What should a database specialist do to meet this requirement?

- a) Modify the database to enable encryption. Apply this setting immediately without waiting for the next scheduled maintenance window.
- b) Export the database to an Amazon S3 bucket with encryption enabled. Create a new database and import the export file.
- c) Create a snapshot of the database. Create an encrypted copy of the snapshot. Create a new database from the encrypted snapshot.
- d) Create a snapshot of the database. Restore the snapshot into a new database with encryption enabled.

Answer: c

Question: 4

A company undergoing a security audit has determined that its database administrators are presently sharing an administrative database user account for the company's Amazon Aurora deployment.

To support proper traceability, governance, and compliance, each database administration team member must start using individual, named accounts. Furthermore, long-term database user credentials should not be used.

Which solution should a database specialist implement to meet these requirements?

- a) Use the AWS CLI to fetch the AWS IAM users and passwords for all team members. For each IAM user, create an Aurora user with the same password as the IAM user.
- b) Enable IAM database authentication on the Aurora cluster. Create a database user for each team member without a password. Attach an IAM policy to each administrator's IAM user account that grants the connect privilege using their database user account.
- c) Create a database user for each team member. Share the new database user credentials with the team members. Have users change the password on the first login to the same password as their IAM user.
- d) Create an IAM role and associate an IAM policy that grants the connect privilege using the shared account. Configure a trust policy that allows the administrator's IAM user account to assume the role.

Answer: b

Question: 5

An operations team in a large company wants to centrally manage resource provisioning for its development teams across multiple accounts.

When a new AWS account is created, the developers require full privileges for a database environment that uses the same configuration, data schema, and source data as the company's production Amazon RDS for MySQL DB instance.

How can the operations team achieve this?

- a) Enable the source DB instance to be shared with the new account so the development team may take a snapshot. Create an AWS CloudFormation template to launch the new DB instance from the snapshot.
- b) Create an AWS CLI script to launch the approved DB instance configuration in the new account. Create an AWS DMS task to copy the data from the source DB instance to the new DB instance.
- c) Take a manual snapshot of the source DB instance and share the snapshot privately with the new account. Specify the snapshot ARN in an RDS resource in an AWS CloudFormation template and use StackSets to deploy to the new account.
- d) Create a DB instance read replica of the source DB instance. Share the read replica with the new AWS account.

Answer: c

Question: 6

A company has a highly available production 10 TB SQL Server relational database running on Amazon EC2. Users have recently been reporting performance and connectivity issues.

A database specialist has been asked to configure a monitoring and alerting strategy that will provide metrics visibility and notifications to troubleshoot these issues.

Which solution will meet these requirements?

- a) Configure AWS CloudTrail logs to monitor and detect signs of potential problems. Create an AWS Lambda function that is triggered when specific API calls are made and send notifications to an Amazon SNS topic.
- b) Install an Amazon Inspector agent on the DB instance. Configure the agent to stream server and database activity to Amazon CloudWatch Logs. Configure metric filters and alarms to send notifications to an Amazon SNS topic.
- c) Migrate the database to Amazon RDS for SQL Server and use Performance Insights to monitor and detect signs of potential problems. Create a scheduled AWS Lambda function that retrieves metrics from the Performance Insights API and send notifications to an Amazon SNS topic.
- d) Configure Amazon CloudWatch Application Insights for .NET and SQL Server to monitor and detect signs of potential problems. Configure CloudWatch Events to send notifications to an Amazon SNS topic.

Answer: d

Question: 7

A company's ecommerce application stores order transactions in an Amazon RDS for MySQL database. The database has run out of available storage and the application is currently unable to take orders.

Which action should a database specialist take to resolve the issue in the shortest amount of time?

- a) Add more storage space to the DB instance using the ModifyDBInstance action.
- b) Create a new DB instance with more storage space from the latest backup.
- c) Change the DB instance status from STORAGE_FULL to AVAILABLE.
- d) Configure a read replica with more storage space.

Answer: a

Question: 8

A company's customer relationship management application uses an Amazon RDS for PostgreSQL Multi-AZ database. The database size is approximately 100 GB.

A database specialist has been tasked with developing a cost-effective disaster recovery plan that will restore the database in a different Region within 2 hours. The restored database should not be missing more than 8 hours of transactions.

What is the MOST cost-effective solution that meets the availability requirements?

- a) Create an RDS read replica in the second Region. For disaster recovery, promote the read replica to a standalone instance.
- b) Create an RDS read replica in the second Region using a smaller instance size. For disaster recovery, scale the read replica and promote it to a standalone instance.
- c) Schedule an AWS Lambda function to create an hourly snapshot of the DB instance and another Lambda function to copy the snapshot to the second Region. For disaster recovery, create a new RDS Multi-AZ DB instance from the last snapshot.
- d) Create a new RDS Multi-AZ DB instance in the second Region. Configure an AWS DMS task for ongoing replication.

Answer: c

Question: 9

A medical company is planning to migrate its on-premises PostgreSQL database, along with application and web servers, to AWS.

Amazon RDS for PostgreSQL is being considered as the target database engine. Access to the database should be limited to application servers and a bastion host in a VPC.

Which solution meets the security requirements?

- a) Launch the RDS for PostgreSQL database in a DB subnet group containing private subnets. Modify the `pg_hba.conf` file on the DB instance to allow connections from only the application servers and bastion host.
- b) Launch the RDS for PostgreSQL database in a DB subnet group containing public subnets. Create a new security group with inbound rules to allow connections from only the security groups of the application servers and bastion host. Attach the new security group to the DB instance.
- c) Launch the RDS for PostgreSQL database in a DB subnet group containing private subnets. Create a new security group with inbound rules to allow connections from only the security groups of the application servers and bastion host. Attach the new security group to the DB instance.
- d) Launch the RDS for PostgreSQL database in a DB subnet group containing private subnets. Create a NACL attached to the VPC and private subnets. Modify the inbound and outbound rules to allow connections to and from the application servers and bastion host.

Answer: c

Question: 10

A database specialist is troubleshooting complaints from an application's users who are experiencing performance issues when saving data in an Amazon ElastiCache for Redis cluster with cluster mode disabled.

The database specialist finds that the performance issues are occurring during the cluster's backup window. The cluster runs in a replication group containing three nodes. Memory on the nodes is fully utilized. Organizational policies prohibit the database specialist from changing the backup window time.

How could the database specialist address the performance concern?

(Select TWO.)

- a) Add an additional node to the cluster in the same Availability Zone as the primary.
- b) Configure the backup job to take a snapshot of a read replica.
- c) Increase the local instance storage size for the cluster nodes.
- d) Increase the reserved-memory-percent parameter value.
- e) Configure the backup process to flush the cache before taking the backup.

Answer: b, d

Study Guide to Crack AWS Database Specialty DBS-C01 Exam:

- Getting details of the DBS-C01 syllabus, is the first step of a study plan. This pdf is going to be of ultimate help. Completion of the syllabus is must to pass the DBS-C01 exam.
- Making a schedule is vital. A structured method of preparation leads to success. A candidate must plan his schedule and follow it rigorously to attain success.
- Joining the AWS provided training for DBS-C01 exam could be of much help. If there is specific training for the exam, you can discover it from the link above.
- Read from the DBS-C01 sample questions to gain your idea about the actual exam questions. In this PDF useful sample questions are provided to make your exam preparation easy.
- Practicing on DBS-C01 practice tests is must. Continuous practice will make you an expert in all syllabus areas.

Reliable Online Practice Test for DBS-C01 Certification

Make VMExam.com your best friend during your AWS Certified Database - Specialty exam preparation. We provide authentic practice tests for the DBS-C01 exam. Experts design these online practice tests, so we can offer you an exclusive experience of taking the actual DBS-C01 exam. We guarantee you 100% success in your first exam attempt if you continue practicing regularly. Don't bother if you don't get 100% marks in initial practice exam attempts. Just utilize the result section to know your strengths and weaknesses and prepare according to that until you get 100% with our practice tests. Our evaluation makes you confident, and you can score high in the DBS-C01 exam.

Start Online practice of DBS-C01 Exam by visiting URL

<https://www.vmexam.com/aws/dbs-c01-aws-certified-database-specialty>