

Databases

Candidates for this exam are seeking to prove introductory knowledge of how to design, create, and query relational databases, such as MySQL, Microsoft SQL Server, or Oracle.

Candidates should have at least 150 hours of instruction or hands-on experience with database design and object management, data retrieval and manipulation, and related troubleshooting. Candidates should be able to write ANSI SQL statements and troubleshoot syntax errors and runtime errors.

The syntax in the exam is provided in ANSI SQL.

1. Database Design

1.1 Given a scenario, design tables for storing data

- Identify entities, rows/records, columns/fields

1.2 Given a scenario, identify the appropriate primary key

- Primary key, composite/compound key

1.3 Given a scenario, choose data types to meet requirements

- Definition and importance of data types; how data types affect storage requirements; data types for storing text, numbers, dates and times, and Boolean values

1.4 Given a scenario, design relationships between tables

- How to establish relationships using primary and foreign keys, entity-relationship diagrams (ERDs), referential integrity

1.5 Normalize a database

- Reasons for normalization, how to normalize a database to third normal form (3NF)

1.6 Given a scenario, identify data protection measures

- Backups, restore, principle of least privilege, GRANT, WITH GRANT OPTION, REVOKE, purpose of roles, when to use encryption, hashing, and masking, purpose of auditing

2. Database Object Management using DDL

2.1 Construct and analyze queries that create, alter, and drop tables

- Create, alter, and drop tables by using proper ANSI SQL syntax; NULL and NOT NULL

2.2 Construct and analyze queries that create, alter, and drop views

- Create, alter, and drop views by using proper ANSI SQL syntax; purpose of views

2.3 Construct and analyze stored procedures and functions

- Input and output parameters, return values, purpose of stored procedures

2.4 Given a scenario, choose between clustered and non-clustered indexes

- When to use clustered vs. non-clustered indexes, syntax for creating indexes



3. Data Retrieval

3.1 Construct and analyze queries that select data

- INNER JOIN, LEFT JOIN, RIGHT JOIN, CROSS JOIN (Cartesian product), and FULL OUTER JOIN; self joins; combine result sets by using UNION and INTERSECT; DISTINCT; column alias; computed columns; concatenation using the + operator and CONCAT(); NULLIF()

3.2 Construct and analyze queries that sort and filter data

- ORDER BY, WHERE, LIKE, BETWEEN, AND, OR, NOT, TOP (LIMIT), IN, NOT IN, ANY, ALL, NULL, NOT NULL, comparison operators

3.3 Construct and analyze queries that aggregate data

- GROUP BY, HAVING, MIN, MAX, COUNT, AVG (AVERAGE), SUM

4. Data Manipulation using DML

4.1 Construct and analyze INSERT statements

- INSERT INTO SELECT, INSERT INTO VALUES

4.2 Construct and analyze UPDATE statements

- Update data in a single table

4.3 Construct and analyze DELETE statements

- Delete data from a single table

5. Troubleshooting

5.1 Troubleshoot data object management query failures

- Syntax and runtime errors

5.2 Troubleshoot data retrieval query failures

- Syntax and runtime errors

5.3 Troubleshoot data manipulation query failures

- Syntax and runtime errors

