

70-464 – DEVELOPING MICROSOFT SQL SERVER DATABASES CERTIFICATION QUESTIONS AND STUDY GUIDE

Developing Microsoft SQL Server Databases (70-464)



Contents

Developing Microsoft SQL Server Databases Details.....	2
Developing Microsoft SQL Server Databases Syllabus for 70-464 Exam (Study Aid)	3
Developing Microsoft SQL Server Databases (70-464) Sample Questions	6

Developing Microsoft SQL Server Databases Certification Details

Exam Name	Developing Microsoft SQL Server Databases
Exam Code	70-464
Duration	120 Minutes
Passing Percentage	700 out of 1000
Negative Marking	No Negative Marking
Partial Credit	No Partial Credits
Reference Book	
Schedule Your exam	Developing Microsoft SQL Server Databases
Sample Questions	Developing Microsoft SQL Server Databases Certification Sample Question
Recommended Practice tool	Developing Microsoft SQL Server Databases Certification Practice Exam

Developing Microsoft SQL Server Databases Certification Syllabus for 70-464 (Study Aid)

IMPLEMENT DATABASE OBJECTS (30–35%)

1. Create and alter tables

Develop an optimal strategy for using temporary objects, including table variables and temporary tables; define alternatives to triggers; define data version control and management; implement @Table and #table appropriately; create calculated columns; implement partitioned tables, schemas and functions; implement column collation; implement online transaction processing (OLTP); implement columnstore and sparse columns

2. Design, implement and troubleshoot security

Implement data control language statements appropriately, troubleshoot connection issues, implement execute as statements, implement certificate-based security, create loginless users, define appropriate database roles and permissions, implement contained users, implement cross db ownership chaining, implement schema security, implement server roles

3. Design the locking granularity level

Choose the right lock mechanism for a given task; handle deadlocks; design index locking properties; fix locking and blocking issues; analyse a deadlock scenario; design appropriate isolation level, including Microsoft ActiveX data objects defaults; design for locks and lock escalation; design transactions that minimise locking; reduce locking contention; identify bottlenecks in data design; design appropriate concurrency control, such as pessimistic or optimistic

4. Implement indexes

Inspect physical characteristics of indexes and perform index maintenance; identify unused indexes; implement indexes; optimise indexes, including full, filter, statistics and force

5. Implement data types

Select appropriate data types, including BLOBs, GUIDs, XML and spatial data; develop a Common Language Runtime (CLR) data type; implement appropriate use of @Table and #table; determine values based on implicit and explicit conversions

6. Create and modify constraints

Create constraints on tables, define constraints, modify constraints according to performance implications, implement cascading deletes, configure constraints for bulk inserts

IMPLEMENT PROGRAMMING OBJECTS (15-20%)

1. Design and implement stored procedures

Create stored procedures and other programmatic objects; implement different types of stored procedure results; create a stored procedure for the data access layer; analyse and rewrite procedures and processes; program stored procedures by using T-SQL and CLR; implement parameters, including table valued, input and output; implement error handling, including TRY...CATCH; configure appropriate connection settings

2. Design T-SQL table-valued and scalar functions

Modify scripts that use cursors and loops into a SET-based operation, design deterministic and non-deterministic functions

3. Create, use and alter user-defined functions (UDFs)

Implement deterministic or non-deterministic functions; implement CROSS APPLY by using UDFs; implement CLR functions

4. Create and alter views

Set up and configure partitioned tables and partitioned views; create indexed views

DESIGN DATABASE OBJECTS (25–30%)

1. Design tables

Apply data design patterns; develop appropriately normalised and de-normalised SQL tables; design transactions; design views; implement GUID as a clustered index appropriately; design temp tables appropriately, including # vs. @; design an encryption strategy; design table partitioning; design a BLOB storage strategy, including filestream and filetable; design tables for In-Memory OLTP

2. Design for concurrency

Develop a strategy to maximise concurrency; define a locking and concurrency strategy; design a transaction isolation strategy, including server database and session; design triggers for concurrency

3. Design indexes

Design indexes and data structures; design filtered indexes; design an indexing strategy, including column store, semantic indexes and INCLUDE; design statistics; assess which indexes on a table are likely to be used, given different search arguments (SARG); design spatial and XML indexes

4. Design data integrity

Design a table data integrity policy, including checks, primary key, foreign key, uniqueness, XML schema and nullability; select a primary key

5. Design for implicit and explicit transactions

Manage transactions; ensure data integrity by using transactions; manage distributed transaction escalations; design savepoints; design error handling for transactions, including TRY, CATCH and THROW

OPTIMISE AND TROUBLESHOOT QUERIES (25–30%)

1. Optimise and tune queries

Tune a poorly performing query, including avoiding unnecessary data-type conversions; identify long-running queries; review and optimise code; analyse execution plans to optimise queries; tune queries using execution plans and Microsoft Database Tuning Advisor (DTA); optimise queries using pivots and common table expressions (CTE); design database layout to optimise queries; implement query hints; tune query workloads; implement recursive CTE; implement full text and semantic search; analyse execution plans; implement plan guides

2. Troubleshoot and resolve performance problems

Interpret performance monitor data; integrate performance monitor data with SQL Traces

3. Optimise indexes

Develop an optimal strategy for clustered indexes; analyse index usage; optimise indexes for workload, including data warehousing and OLTP; generate appropriate indexes and statistics by using INCLUDE columns; create filtered indexes; implement full-text indexing; implement columnstore indexes; optimise online index maintenance

4. Capture and analyse execution plans

Collect and read execution plans, create an index based on an execution plan, batch or split implicit transactions, split large queries, consolidate smaller queries, review and optimise parallel plans

5. Collect performance and system information

Monitor performance using Dynamic Management Views, collect output from the Database Engine Tuning Advisor, design Extended Events Sessions, review and interpret Extended Event logs; optimise Extended Event session settings, use Activity Monitor to minimise server impact and determine IO bottlenecks, monitor In-Memory OLTP resources

Developing Microsoft SQL Server Databases Exam (70-464) Sample Questions

- Below are the 10 sample questions which will help you be familiar with Developing Microsoft SQL Server Databases (70-464) exam style and Structure.
- These questions are just for demonstration purpose, there are many scenario based question are included in **Premium Developing Microsoft SQL Server Databases Practice Exam**
- Access to all 145+ questions is available only through premium practice exam available to members at www.analyticsexam.com

Q 1: You review a query that runs slowly. The query accesses data in a table named Schema1.Table1. The following is the relevant portion of the execution plan for the query:

```
<MissingIndexes>
  <MissingIndexGroup Impact="95.8296">
    <MissingIndex Database="DB1" Schema="Schema1" Table="Table1">
      <ColumnGroup Usage="EQUALITY">
        <Column Name="Column1" ColumnId="14" />
      </ColumnGroup>
      <ColumnGroup Usage="INEQUALITY">
        <Column Name="Column2" ColumnId="17" />
        <Column Name="Column3" ColumnId="21" />
      </ColumnGroup>
      <ColumnGroup Usage="INCLUDE">
        <Column Name="Column4" ColumnId="11" />
      </ColumnGroup>
    </MissingIndex>
  </MissingIndexGroup>
</MissingIndexes>
```

You need to create the missing index. Which code segment should you execute?

Options:

- CREATE NONCLUSTERED INDEX IX1 on schema1.Table1 (Column1) INCLUDE (Column4)
- CREATE NONCLUSTERED INDEX IX1 on Schema1.Table1 (Column1) INCLUDE (Column4) WHERE Column2 <> Column3
- CREATE NONCLUSTERED INDEX IX1 on Schema1.Table1 (Column1)
- CREATE NONCLUSTERED INDEX IX1 on Schema1.Table1 (Column1, Column2, Column3) INCLUDE (Column4)

Q 2: Your network contains a server that has SQL Server 2014 installed. You create a table by using the following script:

```
CREATE TABLE dbo.Products
(
    id int NOT NULL,
    ProductName nvarchar(50) NULL,
    ProductManufacturer nvarchar(50) NULL,
    ProductDescription nvarchar(200) NULL,
    CONSTRAINT PK_Products PRIMARY KEY CLUSTERED (id)
)
ON [PRIMARY]
GO
```

You need to recommend a solution to ensure that each combination of ProductName and Product Manufacturer is not duplicated. What should you recommend creating?

Options:

- A. A columnstore index
- B. A filtered index
- C. A UNIQUE constraint
- D. A CHECK constraint

Q 3: You use SQL Server 2012 to maintain the data used by the applications at your company. You plan to create a table named Table1 by using the following statement.

```
01 CREATE TABLE dbo.table1(
02     ID int IDENTITY(1,1) NOT NULL,
03
04     Email varchar(100) NULL,
05     CONSTRAINT PK_table1 PRIMARY KEY CLUSTERED(ID ASC)
06 )
```

You need to ensure that Table1 contains a column named UserName. The UserName column will:

- Store string values in any language.
- Accept a maximum of 200 characters.
- Be case-insensitive and accent-insensitive.

Which code segment should you add at line 03?

Options:

- A. UserName nvarchar(200) COLLATE Latin1_General_CS_AS NOT NULL,
- B. UserName varchar(200) COLLATE Latin 1_General_CS_AS NOT NULL,
- C. UserName nvarchar(200) COLLATE Latin1_General_CI_AI NOT NULL,
- D. UserName varchar(200) COLLATE Latin1_General_CI_AI NOT NULL,

Q 4: You plan to modify a stored procedure to use temporary data. The stored procedure must meet the following requirements:

- Favor physical memory when physical memory is available.
- Be able to roll back changes to the temporary data.

You need to recommend which object to add to the stored procedure. Which T-SQL command should you recommend?

Options:

- A. DECLARE TABLE @ Table...
- B. CREATE TABLE Table...
- C. CREATE VIEW Table...
- D. CREATE PARTITION SCHEME Table...
- E. CREATE TABLE ##Table...

Q 5: You plan to modify a procedure that contains hundreds of lines of code. The modification must support the following guidelines:

- Use only tables that are not persistent in the database.
- Minimize the amount of time required to execute and recompile procedures.

You need to identify which type of table must be used to support the planned modification. Which type of table should you identify?

Options:

- A. A system table
- B. A temporary table
- C. A partitioned table
- D. A table variable

Q 6: Your network contains a server named SQL1 that has SQL Server 2012 installed. SQL1 contains a database name DB1 and a table named Customers.

You add an additional server named SQL2 that runs SQL Server 2012. You need to create a distributed partitioned view. The solution must minimize the amount of network traffic. What should you do?

Options:

- A. Add the Customers table to SQL2.
- B. Create the view on SQL1.
- C. Remove the Customers table from SQL1.
- D. Create the view on SQL2.
- E. Add SQL2 as a linked server.
- F. Add SQL2 as a Distributor.

Q 7: Your company has a SQL Azure subscription. You implement a database named Database1. Database1 has two tables named Table1 and Table2. You create a stored procedure named sp1. Sp1 reads data from Table1 and inserts data into Table2.

A user named User1 informs you that he is unable to run sp1. You verify that User1 has the SELECT permission on Table1 and Table2. You need to ensure that User1 can run sp1. The solution must minimize the number of permissions assigned to User1.

What should you do?

Options:

- A. Change sp1 to run as the saUser.
- B. Add User1 to the db_datawriter role.
- C. Grant User1 the INSERT permission on Table2.
- D. Grant User1 the EXECUTE permission on sp1.

Q 8: Your company has a main office in London and a branch office in New York. Your network contains a server named Server5 that has SQL Server 2012 installed. Server5 contains a database name ContentDB and a table named ContentTable. You add an additional server named Server9 that runs SQL Server 2012.

You need to create a distributed partitioned view. The solution must minimize the amount of network traffic. What should you do?

Options:

- A. Add Server9 as a linked server.
- B. Add the Customers table to Server9.
- C. Create the view on Server5.
- D. Create the view on Server9.

Q 9: You use SQL Server 2012 to store data used by an e-commerce application. You develop a stored procedure named sp1. Sp1 is used to read the price of all the products sold on the ecommerce site.

You need to ensure that sp1 can read data even while another transaction is modifying the price of a product. Sp1 must only read committed data. Which transaction isolation level should you use in sp1?

Options:

- A. Read committed
- B. Snapshot
- C. Repeatable read
- D. Serializable

Q 10: You use SQL Azure to store data used by an e-commerce application. You develop a stored procedure named sp1. Sp1 is used to read and change the price of all the products sold on the e-commerce site. You need to ensure that other transactions are blocked from updating product data while sp1 is executing. Which transaction isolation level should you use in sp1?

Options:

- A. Read committed
- B. Serializable
- C. Snapshot
- D. Repeatable read

Answers:

Question: 1	Answer: D	Question: 2	Answer: C
Question: 3	Answer: C	Question: 4	Answer: E
Question: 5	Answer: D	Question: 6	Answer: B, D, E
Question: 7	Answer: D	Question: 8	Answer: ALL
Question: 9	Answer: B	Question: 10	Answer: B