

MOC 20762 C: Developing SQL Databases

Course Summary

Description

This five-day instructor-led course provides students with the knowledge and skills to develop a Microsoft SQL Server database. The course focuses on teaching individuals how to use SQL Server product features and tools related to developing a database.

Objectives

By the end of this course, students will be able to:

- Design and Implement Tables.
- Describe advanced table designs
- Ensure Data Integrity through Constraints.
- Describe indexes, including Optimized and Columnstore indexes
- Design and Implement Views.
- Design and Implement Stored Procedures.
- Design and Implement User Defined Functions.
- Respond to data manipulation using triggers.
- Design and Implement In-Memory Tables.
- Implement Managed Code in SQL Server.
- Store and Query XML Data.
- Work with Spatial Data.
- Store and Query Blobs and Text Documents.

Topics

- Introduction to Database Development
- Designing and Implementing Tables
- Advanced Table Designs
- Ensuring Data Integrity through Constraints
- Introduction to Indexes
- Designing Optimized Index Strategies
- Columnstore Indexes
- Designing and Implementing Views
- Designing and Implementing Stored Procedures
- Designing and Implementing User-Defined Functions
- Responding to Data Manipulation via Triggers
- Using In-Memory Tables
- Implementing Managed Code in SQL Server
- Storing and Querying XML Data in SQL Server
- Storing and Querying Spatial Data in SQL Server
- Storing and Querying BLOBs and Text Documents in SQL Server
- SQL Server Concurrency
- Performance and Monitoring

Audience

The primary audience for this course is IT Professionals who want to become skilled on SQL Server product features and technologies for implementing a database.

The secondary audience for this course is individuals who are developers from other product platforms looking to become skilled in the implementation of a SQL Server database

Prerequisite

Before taking this course, students should have basic knowledge of the Microsoft Windows operating system and its core functionality. Students should also have a working knowledge of Transact-SQL and relational databases.

Duration

Five Days

MOC 20762 C: Developing SQL Databases

Course Outline

Course Outline

I. *Introduction to Database Development*

Before beginning to work with Microsoft SQL Server in either a development or an administration role, it is important to understand the scope of the SQL Server platform. In particular, it is useful to understand that SQL Server is not just a database engine—it is a complete platform for managing enterprise data. SQL Server provides a strong data platform for all sizes of organizations, in addition to a comprehensive set of tools to make development easier, and more robust.

A. Introduction to the SQL Server Platform
B. SQL Server Database Development Tasks

II. *Designing and Implementing Tables*

In a relational database management system (RDBMS), user and system data is stored in tables. Each table consists of a set of rows that describe entities and a set of columns that hold the attributes of an entity. For example, a Customer table might have columns such as CustomerName and CreditLimit, and a row for each customer. In Microsoft SQL Server data management software tables are contained within schemas that are very similar in concept to folders that contain files in the operating system. Designing tables is one of the most important tasks that a database developer undertakes, because incorrect table design leads to the inability to query the data efficiently. After an appropriate design has been created, it is important to know how to correctly implement the design.

- A. Designing Tables
- B. Data Types
- C. Working with Schemas
- D. Creating and Altering Tables

Lab: Designing and Implementing Tables

- Designing Tables
- Creating Schemas
- Creating Tables

III. *Advanced Table Designs*

The physical design of a database can have a significant impact on the ability of the database to meet the storage and performance requirements set out by the stakeholders. Designing a physical database implementation includes planning the filegroups, how to use partitioning to manage large tables, and using compression to improve storage and performance. Temporal tables are a new feature in SQL Server and offer a

straightforward solution to collecting changes to your data.

- A. Partitioning Data
- B. Compressing Data
- C. Temporal Tables

Lab: Using Advanced Table Designs

- Partitioning Data
- Compressing Data

IV. *Ensuring Data Integrity through Constraints*

The quality of data in your database largely determines the usefulness and effectiveness of applications that rely on it—the success or failure of an organization or a business venture could depend on it. Ensuring data integrity is a critical step in maintaining high-quality data. You should enforce data integrity at all levels of an application from first entry or collection through storage. Microsoft SQL Server data management software provides a range of features to simplify the job.

- A. Enforcing Data Integrity
- B. Implementing Data Domain Integrity
- C. Implementing Entity and Referential Integrity

Lab: Using Data Integrity Through Constraints

- Add Constraints
- Test the Constraints

V. *Introduction to Indexes*

An index is a collection of pages associated with a table. Indexes are used to improve the performance of queries or enforce uniqueness. Before learning to implement indexes, it is helpful to understand how they work, how effective different data types are when used within indexes, and how indexes can be constructed from multiple columns. This module discusses table structures that do not have indexes, and the different index types available in Microsoft SQL Server.

- A. Core Indexing Concepts
- B. Data Types and Indexes
- C. Heaps, Clustered, and Nonclustered Indexes
- D. Single Column and Composite Indexes

Lab: Implementing Indexes

- Creating a Heap
- Creating a Clustered Index
- Creating a Covered Index

MOC 20762 C: Developing SQL Databases

Course Outline (cont.)

VI. *Designing Optimized Index Strategies*

Indexes play an important role in enabling SQL Server to retrieve data from a database quickly and efficiently. This module discusses advanced index topics including covering indexes, the INCLUDE clause, query hints, padding and fill factor, statistics, using DMOs, the Database Tuning Advisor, and Query Store.

- A. Index Strategies
- B. Managing Indexes
- C. Execution Plans
- D. The Database Engine Tuning Advisor
- E. Query Store

Lab: Optimizing Indexes

- Using Query Store
- Heaps and Clustered Indexes
- Creating a Covered Index

VII. *Columnstore Indexes*

Introduced in Microsoft SQL Server 2012, columnstore indexes are used in large data warehouse solutions by many organizations. This module highlights the benefits of using these indexes on large datasets; the improvements made to columnstore indexes in SQL Server 2016; and the considerations needed to use columnstore indexes effectively in your solutions.

- A. Introduction to Columnstore Indexes
- B. Creating Columnstore Indexes
- C. Working with Columnstore Indexes

Lab: Using Columnstore Indexes

- Creating a Columnstore Index
- Create a Memory Optimized Columnstore Table

VIII. *Designing and Implementing Views*

This module describes the design and implementation of views. A view is a special type of query—one that is stored and can be used in other queries—just like a table. With a view, only the query definition is stored on disk; not the result set. The only exception to this is indexed views, when the result set is also stored on disk, just like a table. Views simplify the design of a database by providing a layer of abstraction, and hiding the complexity of table joins. Views are also a way of securing your data by giving users permissions to use a view, without giving them permissions to the underlying objects. This means data can be

kept private, and can only be viewed by appropriate users.

- A. Introduction to Views
- B. Creating and Managing Views
- C. Performance Considerations for Views

Lab: Designing and Implementing Views

- Creating Standard Views
- Creating an Updateable view

IX. *Designing and Implementing Stored Procedures*

This module describes the design and implementation of stored procedures.

- A. Introduction to Stored Procedures
- B. Working with Stored Procedures
- C. Implementing Parameterized Stored Procedures
- D. Controlling Execution Context

Lab: Designing and Implementing Stored Procedures

- Create Stored procedures
- Create Parameterized Stored procedures
- Changes Stored Procedure Execution Context

X. *Designing and Implementing User-Defined Functions*

Functions are routines that you use to encapsulate frequently performed logic. Rather than having to repeat the function logic in many places, code can call the function. This makes code more maintainable, and easier to debug. In this module, you will learn to design and implement user-defined functions (UDFs) that enforce business rules or data consistency. You will also learn how to modify and maintain existing functions.

- A. Overview of Functions
- B. Designing and Implementing Scalar Functions
- C. Designing and Implementing Table-Valued Functions
- D. Considerations for Implementing Functions
- E. Alternatives to Functions

Lab: Designing and Implementing User-Defined Functions

- Format Phone numbers
- Modify an Existing Function

MOC 20762 C: Developing SQL Databases

Course Outline (cont.)

XI. *Responding to Data Manipulation via Triggers*

Data Manipulation Language (DML) triggers are powerful tools that you can use to enforce domain, entity, referential data integrity and business logic. The enforcement of integrity helps you to build reliable applications. In this module, you will learn what DML triggers are, how they enforce data integrity, the different types of trigger that are available to you, and how to define them in your database.

- A. Designing DML Triggers
- B. Implementing DML Triggers
- C. Advanced Trigger Concepts

Lab: Responding to Data Manipulation by Using Triggers

- Create and Test the Audit Trigger
- Improve the Audit Trigger

XII. *Using In-Memory Tables*

Microsoft SQL Server 2014 data management software introduced in-memory online transaction processing (OLTP) functionality features to improve the performance of OLTP workloads. SQL Server adds several enhancements, such as the ability to alter a memory-optimized table without recreating it. Memory-optimized tables are primarily stored in memory, which provides the improved performance by reducing hard disk access. Natively compiled stored procedures further improve performance over traditional interpreted Transact-SQL.

- A. Memory-Optimized Tables
- B. Natively Compiled Stored Procedures

Lab: Using In-Memory Database Capabilities

- Using Memory-Optimized Tables
- Using Natively Compiled Stored procedures

XIII. *Implementing Managed Code in SQL Server*

As a SQL Server professional, you are likely to be asked to create databases that meet business needs. Most requirements can be met using Transact-SQL. However, occasionally you may need additional capabilities that can only be met by using common language runtime (CLR) code. As functionality is added to SQL Server with each new release, the necessity to use managed code decreases. However, there are times when you might need to create aggregates,

stored procedures, triggers, user-defined functions, or user-defined types. You can use any .NET Framework language to develop these objects. In this module, you will learn how to use CLR managed code to create user-defined database objects for SQL Server.

- A. Introduction to CLR Integration in SQL Server
- B. Implementing and Publishing CLR Assemblies

Lab: Implementing Managed Code in SQL Server

- Assessing Proposed CLR Code
- Creating a Scalar-Valued CLR Function
- Creating a Table Valued CLR Function

XIV. *Storing and Querying XML Data in SQL Server*

XML provides rules for encoding documents in a machine-readable form. It has become a widely adopted standard for representing data structures, rather than sending unstructured documents. Servers that are running Microsoft SQL Server data management software often need to use XML to interchange data with other systems; many SQL Server tools provide an XML-based interface. SQL Server offers extensive handling of XML, both for storage and querying. This module introduces XML, shows how to store XML data within SQL Server, and shows how to query the XML data. The ability to query XML data directly avoids the need to extract data into a relational format before executing Structured Query Language (SQL) queries. To effectively process XML, you need to be able to query XML data in several ways: returning existing relational data as XML, and querying data that is already XML.

- A. Introduction to XML and XML Schemas
- B. Storing XML Data and Schemas in SQL Server
- C. Implementing the XML Data Type
- D. Using the Transact-SQL FOR XML Statement
- E. Getting Started with XQuery
- F. Shredding XML

Lab: Storing and Querying XML Data in SQL Server

- Determining when to use XML
- Testing XML Data Storage in Variables
- Using XML Schemas

MOC 20762 C: Developing SQL Databases

Course Outline (cont.)

- Using FOR XML Queries
- Creating a Stored Procedure to Return XML

XV. *Storing and Querying Spatial Data in SQL Server*

This module describes spatial data and how this data can be implemented within SQL Server.

- A. Introduction to Spatial Data
- B. Working with SQL Server Spatial Data Types
- C. Using Spatial Data in Applications

Lab: Working with SQL Server Spatial Data

- Become Familiar with the Geometry Data Type
- Add Spatial Data to an Existing Table
- Find Nearby Locations

XVI. *Storing and Querying BLOBs and Text Documents in SQL Server*

Traditionally, databases have been used to store information in the form of simple values—such as integers, dates, and strings—that contrast with more complex data formats, such as documents, spreadsheets, image files, and video files. As the systems that databases support have become more complex, administrators have found it necessary to integrate this more complex file data with the structured data in database tables. For example, in a product database, it can be helpful to associate a product record with the service manual or instructional videos for that product. SQL Server provides several ways to integrate these files—that are often known as Binary Large Objects (BLOBs)—and enable their content to be indexed and included in search results. In this module, you will learn how to design and optimize a database that includes BLOBs.

- A. Considerations for BLOB Data
- B. Working with FILESTREAM
- C. Using Full-Text Search

Lab: Storing and Querying BLOBs and Text Documents in SQL Server

- Enabling and Using FILESTREAM Columns
- Enabling and Using File Tables
- Using a Full-Text Index

XVII. *SQL Server Concurrency*

This module explains how to name, declare, assign values to, and use variables. It also describes how to store data in an array. Concurrency control is a critical feature of multiuser database systems; it allows data to remain consistent when many users are modifying data at the same time. This module covers the implementation of concurrency in Microsoft SQL Server. You will learn about how SQL Server implements concurrency controls, and the different ways you can configure and work with concurrency settings.

- A. Concurrency and Transactions
- B. Locking Internals

Lab: SQL Server Concurrency

- Implement Snapshot Isolation
- Implement Partition Level Locking

XVIII. *Performance and Monitoring*

This module explains how to name, declare, assign values to, and use variables. It also describes how to store data in an array. This module looks at how to measure and monitor the performance of your SQL Server databases. The first two lessons look at SQL Server Extended Events, a flexible, lightweight event-handling system built into the Microsoft SQL Server Database Engine. These lessons focus on the architectural concepts, troubleshooting strategies and usage scenarios.

- A. Extended Events
- B. Working with extended Events
- C. Live Query Statistics
- D. Optimize Database File Configuration
- E. Metrics

Lab: Monitoring, Tracing, and Baselineing

- Collecting and Analyzing Data Using Extended Events
- Implementing Baseline Methodology