

20762B: DEVELOPING SQL DATABASES

ABOUT THIS COURSE

This five day instructor-led course provides students with the knowledge and skills to develop a Microsoft SQL Server 2016 database. The course focuses on teaching individuals how to use SQL Server 2016 product features and tools related to developing a database.

AUDIENCE PROFILE

The primary audience for this course is IT Professionals who want to become skilled on SQL Server 2016 product features and technologies for implementing a database.

The secondary audiences for this course are individuals who are developers from other product platforms looking to become skilled in the implementation of a SQL Server 2016 database.

AT COURSE COMPLETION

After completing 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.

PREREQUISITES

Before attending this course, students must have:

- Basic knowledge of the Microsoft Windows operating system and its core functionality.
- Working knowledge of Transact-SQL.
- Working knowledge of relational databases.

COURSE OUTLINE

MODULE 1: 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.

20762B: DEVELOPING SQL DATABASES

LESSONS

- Introduction to the SQL Server Platform
 - SQL Server Database Development Tasks
-

AFTER COMPLETING THIS MODULE, YOU WILL BE ABLE TO:

- Describe the SQL Server platform.
- Use SQL Server administration tools.

MODULE 2: 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.

LESSONS

- Designing Tables
 - Data Types
 - Working with Schemas
 - Creating and Altering Tables
-

LAB: DESIGNING AND IMPLEMENTING TABLES

- Designing Tables
 - Creating Schemas
 - Creating Tables
-

AFTER COMPLETING THIS MODULE, YOU WILL BE ABLE TO:

- Design Tables using normalization, primary and foreign keys
- Work with identity columns.
- Understand built-in and user data types.
- Use schemas in your database designs to organize data, and manage object security.
- Work with computed columns and temporary tables.

MODULE 3: 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 2016 and offer a straightforward solution to collecting changes to your data.

LESSONS

- Partitioning Data
 - Compressing Data
 - Temporal Tables
-

LAB: USING ADVANCED TABLE DESIGNS

- Portioning Data
- Compressing Data

20762B: DEVELOPING SQL DATABASES

AFTER COMPLETING THIS MODULE, YOU WILL BE ABLE TO:

- Describe the considerations for using partitioned tables in a SQL Server database.
- Plan for using data compression in a SQL Server database.
- Use temporal tables to store and query changes to your data.

MODULE 4: 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.

LESSONS

- Enforcing Data Integrity
- Implementing Data Domain Integrity
- Implementing Entity and Referential Integrity

LAB: ENSURING DATA INTEGRITY THROUGH CONSTRAINTS

- Add Constraints
- Test the Constraints

AFTER COMPLETING THIS MODULE, YOU WILL BE ABLE TO:

- Describe the options for enforcing data integrity and the levels at which they should be applied
- Implement domain integrity through options such as check, unique, and default constraints
- Implement referential integrity through primary and foreign key constraints.

MODULE 5: 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.

LESSONS

- Core Indexing Concepts
- Data Types and Indexes
- Heaps, Clustered, and Nonclustered Indexes
- Single Column and Composite Indexes

LAB: IMPLEMENTING INDEXES

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

AFTER COMPLETING THIS MODULE, YOU WILL BE ABLE TO:

- Explain core indexing concepts
- Evaluate which index to use for different data types
- Describe the difference between single and composite column indexes

20762B: DEVELOPING SQL DATABASES

MODULE 6: 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.

LESSONS

- Index Strategies
- Managing Indexes
- Execution Plans
- The Database Engine Tuning Advisor
- Query Store

LAB: OPTIMIZING INDEXES

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

AFTER COMPLETING THIS MODULE, YOU WILL BE ABLE TO:

- What a covering index is and when to use one.
- The issues involved in managing indexes.
- Actual and estimated execution plans.
- How to use Database Tuning Advisor to improve the performance of queries.
- How to use Query Store to improve query performance.

MODULE 7: 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.

LESSONS

- Introduction to Columnstore Indexes
- Creating Columnstore Indexes
- Working with Columnstore Indexes

LAB: USING COLUMNSTORE INDEXES

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

AFTER COMPLETING THIS MODULE, YOU WILL BE ABLE TO:

- Describe columnstore indexes and identify suitable scenarios for their use.
- Create clustered and nonclustered columnstore indexes.
- Describe considerations for using columnstore indexes.

MODULE 8: 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

20762B: DEVELOPING SQL DATABASES

them permissions to the underlying objects. This means data can be kept private, and can only be viewed by appropriate users

LESSONS

- Introduction to Views
 - Creating and Managing Views
 - Performance Considerations for Views
-

LAB: DESIGNING AND IMPLEMENTING VIEWS

- Creating Standard Views
 - Creating an Updateable view
-

AFTER COMPLETING THIS MODULE, YOU WILL BE ABLE TO:

- Understand the role of views in database design.
 - Create and manage views.
 - Understand the performance considerations with views
-

MODULE 9: DESIGNING AND IMPLEMENTING STORED PROCEDURES

This module describes the design and implementation of stored procedures.

LESSONS

- Introduction to Stored Procedures
 - Working with Stored Procedures
 - Implementing Parameterized Stored Procedures
 - Controlling Execution Context
-

LAB: DESIGNING AND IMPLEMENTING STORED PROCEDURES

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

AFTER COMPLETING THIS MODULE, YOU WILL BE ABLE TO:

- Understand what stored procedures are, and what benefits they have
 - Design, Create and Alter Stored Procedures
 - Control the Execution Context of Stored Procedures
 - Implement Stored Procedures that use Parameters
-

MODULE 10: DESIGNING AND IMPLEMENTING USER-DEFINED FUNCTIONS

This module describes the design and implementation of functions, both scalar and table-valued. (Also discusses where they can lead to performance issues).

LESSONS

- Overview of Functions
 - Designing and Implementing Scalar Functions
 - Designing and Implementing Table-Valued Functions
 - Considerations for Implementing Functions
 - Alternatives to Functions
-

LAB: DESIGNING AND IMPLEMENTING USER-DEFINED FUNCTIONS

- Format Phone
 - Modify an Existing Function
-

20762B: DEVELOPING SQL DATABASES

AFTER COMPLETING THIS MODULE, YOU WILL BE ABLE TO:

- Design, create and alter scalar functions
- Design, create and alter table-valued functions
- Describe the performance impacts of functions

MODULE 11: 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

LESSONS

- Designing DML Triggers
- Implementing DML Triggers
- Advanced Trigger Concepts

LAB: RESPONDING TO DATA MANIPULATION BY USING TRIGGERS

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

AFTER COMPLETING THIS MODULE, YOU WILL BE ABLE TO:

- Design DML triggers
- Implement DML triggers
- Explain advanced DLM trigger concepts, such as nesting and recursion

MODULE 12: 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 2016 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.

LESSONS

- Memory-Optimized Tables
- Natively Compiled Stored Procedures

LAB: USING IN-MEMORY DATABASE CAPABILITIES

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

AFTER COMPLETING THIS MODULE, YOU WILL BE ABLE TO:

- Use memory optimized tables to improve performance for latch-bound workloads
- Use natively compiled stored procedures

MODULE 13: 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

20762B: DEVELOPING SQL DATABASES

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.

LESSONS

- Introduction to CLR Integration in SQL Server
 - 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
-

AFTER COMPLETING THIS MODULE, YOU WILL BE ABLE TO:

- Explain the importance of CLR integration in SQL Server
 - Implement and publish CLR assemblies using SQL Server Data Tools (SSDT)
-

MODULE 14: 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.

LESSONS

- Introduction to XML and XML Schemas
 - Storing XML Data and Schemas in SQL Server
 - Implementing the XML Data Type
 - Using the Transact-SQL FOR XML Statement
 - Getting Started with XQuery
 - 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
 - Using FOR XML Queries
 - Creating a Stored Procedure to Return XML
-

AFTER COMPLETING THIS MODULE, YOU WILL BE ABLE TO:

- Describe XML and XML schemas
 - Store XML data and associated XML schemas in SQL Server
 - Implement XML indexes with SQL Server
 - Use the Transact-SQL for XML statement
 - Work with basic XQuery queries
-

MODULE 15: STORING AND QUERYING SPATIAL DATA IN SQL SERVER

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

20762B: DEVELOPING SQL DATABASES

LESSONS

- Introduction to Spatial Data
 - Working with SQL Server Spatial Data Types
 - 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
-

AFTER COMPLETING THIS MODULE, YOU WILL BE ABLE TO:

- Describe how spatial data can be stored in SQL Server
 - Use basic methods of the GEOMETRY and GEOGRAPHY data types
 - Query databases containing spatial data
-

MODULE 16: 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.

LESSONS

- Considerations for BLOB Data
 - Working with FILESTREAM
 - 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
-

AFTER COMPLETING THIS MODULE, YOU WILL BE ABLE TO:

- Describe the considerations for designing databases that incorporate BLOB data.
 - Describe the benefits and design considerations for using FILESTREAM to store BLOB data on a Windows file system.
 - Describe the benefits of using full-text indexing and Semantic Search, and explain how to use these features to search SQL Server data, including unstructured data
-

MODULE 17: 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.

20762B: DEVELOPING SQL DATABASES

LESSONS

- Concurrency and Transactions
- Locking Internals

LAB: SQL SERVER CONCURRENCY

- Implement Snapshot Isolation
- Implement Partition Level Locking

AFTER COMPLETING THIS MODULE, YOU WILL BE ABLE TO:

- Describe concurrency and transactions in SQL Server.
- Describe SQL Server locking.

MODULE 18: 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.

LESSONS

- Extended Events
- Working with extended Events
- Live Query Statistics
- Optimize Database File Configuration
- Metrics

LAB: MONITORING, TRACING, AND BASELINING

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

AFTER COMPLETING THIS MODULE, YOU WILL BE ABLE TO:

- Understand Extended Events and how to use them.
- Work with Extended Events.
- Understand Live Query Statistics.
- Optimize the file configuration of your databases.
- Use DMVs and Performance Monitor to create baselines and gather performance metrics.