

OFFICIAL MICROSOFT LEARNING PRODUCT

20464D

**Developing Microsoft SQL Server 2014
Databases**

Information in this document, including URL and other Internet Web site references, is subject to change without notice. Unless otherwise noted, the example companies, organizations, products, domain names, e-mail addresses, logos, people, places, and events depicted herein are fictitious, and no association with any real company, organization, product, domain name, e-mail address, logo, person, place or event is intended or should be inferred. Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

The names of manufacturers, products, or URLs are provided for informational purposes only and Microsoft makes no representations and warranties, either expressed, implied, or statutory, regarding these manufacturers or the use of the products with any Microsoft technologies. The inclusion of a manufacturer or product does not imply endorsement of Microsoft of the manufacturer or product. Links may be provided to third party sites. Such sites are not under the control of Microsoft and Microsoft is not responsible for the contents of any linked site or any link contained in a linked site, or any changes or updates to such sites. Microsoft is not responsible for webcasting or any other form of transmission received from any linked site. Microsoft is providing these links to you only as a convenience, and the inclusion of any link does not imply endorsement of Microsoft of the site or the products contained therein.

© 2016 Microsoft Corporation. All rights reserved.

Microsoft and the trademarks listed at <http://www.microsoft.com/about/legal/en/us/IntellectualProperty/Trademarks/EN-US.aspx> are trademarks of the Microsoft group of companies. All other trademarks are property of their respective owners

Product Number: 20464D

Part Number (if applicable):

Released: 02/2016

Module 1

An Introduction to Database Development

Contents:

Lesson 1: Introduction to the SQL Server Platform	2
Lesson 2: Working with SQL Server Tools	4
Lesson 3: Configuring SQL Server Services	8
Module Review and Takeaways	10
Lab Review Questions and Answers	11

Lesson 1

Introduction to the SQL Server Platform

Contents:

Question and Answers

3

Question and Answers

Question: Which component of SQL Server includes an OLAP engine?

- ☐ Integration Services
- ☐ Reporting Services
- ☐ Analysis Services
- ☐ Master Data Services
- ☐ Data Quality Services

Answer:

- ☐ Integration Services
- ☐ Reporting Services
- ☒ Analysis Services
- ☐ Master Data Services
- ☐ Data Quality Services

Lesson 2

Working with SQL Server Tools

Contents:

Question and Answers	5
Demonstration: Using SQL Server Management Studio	6

Question and Answers

Categorize each item into the appropriate category. Indicate your answer by writing the category number to the right of each item.

Items	
1	TCP/IP
2	JDBC
3	Named Pipes
4	OLEDB
5	Relational Engine
6	SQL OS
7	Storage Engine
8	ODBC

Category 1		Category 2		Category 3
Client Libraries		Network Libraries		SQL Server Software Layers

Answer:

Category 1		Category 2		Category 3
Client Libraries		Network Libraries		SQL Server Software Layers
ODBC JDBC OLEDB		TCP/IP Named Pipes		Relational Engine Storage Engine SQL OS

Demonstration: Using SQL Server Management Studio**Demonstration Steps**

Use SSMS to connect to an on-premises instance of SQL Server 2014

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **AdventureWorks\Student** with the password **Pa\$\$w0rd**.
2. Run D:\Demofiles\Mod01\Setup.cmd as an administrator to revert any changes.
3. In the virtual machine, on the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, ensure that **Server type** is set to **Database Engine**.
5. In the **Server name** text box, type **(local)**.
6. In the **Authentication** drop-down list, select **Windows Authentication**, and then click **Connect**.

Run a T-SQL script

1. If required, on the **View** menu, click **Object Explorer**.
2. In Object Explorer, expand **Databases**, expand **AdventureWorks**, and then expand **Tables**. Review the database objects.
3. Right-click the **AdventureWorks** database, and then click **New Query**.
4. Type the query shown in the snippet below.

```
SELECT * FROM Production.Product ORDER BY ProductID;
```

5. Note the use of IntelliSense while you are typing this query, and then on the toolbar, click **Execute**. Note how the results can be returned.
6. On the **File** menu, click **Save SQLQuery1.sql**. Note that this saves the query to a file. In the **Save File As** window, click **Cancel**.
7. On the **Results** tab, right-click the cell for **ProductID 1** (first row and first cell), and then click **Save Results As**. In the **FileName** text box, type **Demonstration2AResults** and then click **Save**. Note that this saves the query results to a file.

8. On the **Query** menu, click **Display Estimated Execution Plan**. Note that SQL Server Management Studio can do more than simply execute queries.
9. On the **Tools** menu, click **Options**.
10. In the **Options** pane, expand **Query Results**, expand **SQL Server**, and then click **General**. Review the available configuration options and then click **Cancel**.
11. On the **File** menu, click **Close**. In the **Microsoft SQL Server Management Studio** window, click **No**.

Open a SQL Server Management Studio project

1. On the **File** menu, click **Open**, and then click **Project/Solution**.
2. In the **Open Project** window, open the **D:\Demofiles\Mod01\Demo01.ssmssl** project.
3. On the **View** menu, click **Solution Explorer**. Note the contents of Solution Explorer.
4. In Solution Explorer, click **Close**.

Connect to servers and databases

1. In Object Explorer, from the **Connect** toolbar icon, note the other SQL Server components to which connections can be made.
2. On the **File** menu, click **New**, and then click **Database Engine Query** to open a new connection.
3. In the **Connect to Database Engine** window, in the **Server name** box, type **(local)**.
4. In the **Authentication** drop-down list, select **Windows Authentication**, and then click **Connect**.
5. In the **Available Databases** drop-down list on the toolbar, click **tempdb**. Note that this will change the database against which the query is executed.
6. Right-click in the query window, click **Connection**, and then click **Change Connection**. This will reconnect the query to another instance of SQL Server.
7. In the **Connect to Database Engine** window, click **Cancel**.

Register servers

1. On the **View** menu, click **Registered Servers**.
2. In the **Registered Servers** window, expand **Database Engine**, right-click **Local Server Groups**, and then click **New Server Group**.
3. In the **New Server Group Properties** window, in the **Group name** box, type **Dev Servers** and then click **OK**.
4. Right-click **Dev Servers**, and then click **New Server Registration**.
5. In the **New Server Registration** window, click the **Server name** drop-down list, select **(local)** and then click **Save**.
6. Right-click **Dev Servers**, and then click **New Server Registration**.
7. In the **New Server Registration** window, in the **Server name** text box, type **MIA-SQL\SQL2**, and then click **Save**.
8. In the **Registered Servers** window, right-click the **Dev Servers** group, and then click **New Query**.
9. Type the query as shown in the snippet below, and then click **Execute**.

```
SELECT @@version;
```

10. Close SQL Server Management Studio.
11. In the **Microsoft SQL Server Management Studio** window, click **No**.

Lesson 3

Configuring SQL Server Services

Contents:

Question and Answers	9
Demonstration: Using SQL Server Profiler	9

Question and Answers

Question: You want to ensure that the clients do not need to include the IP address and port number of the SQL Server in their connection string. You also want to ensure that the SQL Server's IP address can change without reconfiguring all clients. What should you create?

Answer: Create a server alias

Demonstration: Using SQL Server Profiler

Demonstration Steps

Start a SQL Server Profiler trace

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. If you have not completed the previous demonstration, run D:\Demofiles\Mod01\Setup.cmd as an administrator to revert any changes.
3. On the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, ensure that **Server type** is set to **Database Engine**.
5. In the **Server name** text box, type **(local)**.
6. In the **Authentication** drop-down list, select **Windows Authentication**, and then click **Connect**.
7. On the **Tools** menu, click **SQL Server Profiler**.
8. In the **Connect to Server** window, ensure that **Server type** is set to **Database Engine**.
9. In the **Server name** text box, type **(local)**.
10. In the **Authentication** drop-down list, select **Windows Authentication**, and then click **Connect**.
11. In the **Trace Properties** window, in **Trace name**, type **Demonstration**.
12. Click **Run**. Note that this will start a new trace with the default options.

View a SQL Server Profiler trace

1. Switch to SQL Server Management Studio, and then click **New Query**.
2. In the query window, type the query as shown below, and then click **Execute**.

```
USE AdventureWorks;  
GO  
SELECT * FROM Person.Person  
ORDER BY FirstName;  
GO
```

3. Switch to SQL Server Profiler. Note the statement trace occurring in SQL Server Profiler.
4. On the **File** menu, click **Stop Trace**.
5. In the Results grid, click individual statements to see the detail shown in the lower pane.
6. Close SQL Server Management Studio and SQL Server Profiler without saving any changes.

Module Review and Takeaways

Review Question(s)

Question: What is the difference between a version of SQL Server and an edition of SQL Server?

Answer: Versions are releases of the product. Editions are levels of the product that have differing capabilities.

Question: What is the purpose of SQL Server Data Tools?

Answer: SQL Server Data Tools adds templates to Visual Studio for constructing and testing Business Intelligence projects.

Lab Review Questions and Answers

Lab: Introduction to Database Development

Question and Answers

Lab Review

Question: How can you configure SQL Server to use a different IP port?

Answer: Network configuration in SQL Server Configuration Manager enables you to configure ports for protocols.

Module 2

Designing and Implementing Tables

Contents:

Lesson 1: Using Data Types	2
Lesson 2: Working with Character Data	4
Lesson 3: Designing Tables	6
Lesson 4: Working with Schemas	8
Lesson 5: Creating and Altering Tables	10
Module Review and Takeaways	13
Lab Review Questions and Answers	14

Lesson 1

Using Data Types

Contents:

Question and Answers	3
Demonstration: Working with Numeric Data Types	3

Question and Answers

Question: When should a value be nullable?

Answer: A value should be nullable if it is possible for the value to be unknown.

Introducing Data Types

Question: Why would it be faster to compare two integer variables that are holding the values 3,240 and 19,704 than two **varchar(10)** variables that are holding the values "3240" and "19704"?

Answer: Because the number of bytes that need to be compared is much lower. Also, no collation or sorting rules need to be considered, as is needed for most string values.

Exact Numeric Data Types

Question: What would be a suitable data type for storing the value of a check box that can be 0 for cleared, 1 for selected, or -1 for disabled?

Answer: **smallint** (note that **tinyint** cannot be negative).

Date and Time Data Types

Question: Why is the specification of a date range from the year 0000 to the year 9999 based on the Gregorian calendar not entirely meaningful?

Answer: Because the Gregorian calendar was introduced by Pope Gregory XIII in 1582.

Unique Identifiers

Question: The slide mentions that a common error is to store GUIDs as strings. What would be wrong with this?

Answer: Storing GUIDs as strings leads to a greater draw on storage and much slower speeds.

Demonstration: Working with Numeric Data Types

Demonstration Steps

Work with NULL and insert GUIDs into a table

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. Run D:\Demofiles\Mod02\Setup.cmd as an administrator to revert any changes.
3. On the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod02\Demo02.ssmssln**, and then click **Open**.
6. If Solution Explorer is not visible, click the **View** menu and click **Solution Explorer**.
7. Expand the **Queries** folder and then double-click **11 - Demonstration 1A.sql**.
8. Follow the instructions contained within the comments of the script file.

Lesson 2

Working with Character Data

Contents:

Question and Answers

5

Question and Answers

Question: What are the code page and sensitivity values for the collation SQL_Scandinavian_Cp850_CI_AS?

Answer: Code page 850, case insensitive and accent sensitive.

Character Data Types

Question: Why would you use the **sysname** data type rather than the **nvarchar(128)** data type?

Answer: To minimize the amount of change required to scripts should the **sysname** data type ever change in SQL Server.

Lesson 3

Designing Tables

Contents:

Question and Answers

7

Question and Answers

Question: You have a database table that stores one entry for each book in your catalog. The primary key for the table is the International Standard Book Number (ISBN) issued by a registration agency to each published book. Is this a natural or surrogate key?

Answer: This is an example of a natural key

Lesson 4

Working with Schemas

Contents:

Question and Answers	9
Demonstration: Working with Schemas	9

Question and Answers

Question: Which of the following values is not used in the multipart naming convention for a table in a SQL Server 2014 database?

- () The name of the server.
- () The name of the database.
- () The name of the schema.
- () The name of the table owner.
- () The name of the table.

Answer:

- () The name of the server.
- () The name of the database.
- () The name of the schema.
- (√) The name of the table owner.
- () The name of the table.

Demonstration: Working with Schemas

Demonstration Steps

Create a schema, create a schema with an included object, and drop a schema

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. Ensure that you have completed the previous demonstrations in this module.
3. In the virtual machine, on the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod02\Demo02.ssmssln**, and then click **Open**.
6. Open the **21 - Demonstration 2A.sql** script file.
7. Follow the instructions contained within the comments of the script file.

Lesson 5

Creating and Altering Tables

Contents:

Question and Answers	11
Demonstration: Working with Tables	11
Demonstration: Working with Temporary Tables	11
Demonstration: Working with Computed Columns	11

Question and Answers

Question: You have a table that stores stock market information for many companies. You want to add a computed column that returns the total value of the company, as calculated from the NumberOfShares and ShareValue columns. This value is requested occasionally by board members but the ShareValue changes many times in a day. Should you use a persisted or a nonpersisted computed column?

Answer: Use a nonpersisted computed column.

Demonstration: Working with Tables

Demonstration Steps

Create tables and alter tables

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. Ensure that you have completed the previous demonstrations in this module.
3. In the virtual machine, on the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod02\Demo02.ssmssln**, and then click **Open**.
6. Open the **31 - Demonstration 3A.sql** script file.
7. Follow the instructions contained within the comments of the script file.

Demonstration: Working with Temporary Tables

Demonstration Steps

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **AdventureWorks\Student** with the password **Pa\$\$w0rd**.
2. Ensure that you have completed the previous demonstrations in this module.
3. In the virtual machine, on the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod02\Demo02.ssmssln**, and then click **Open**.
6. Open the **32 - Demonstration 3B.sql** script file.
7. Follow the instructions contained within the comments of the script file.

Demonstration: Working with Computed Columns

Demonstration Steps

Work with computed columns

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. Ensure that you have completed the previous demonstrations in this module.
3. In the virtual machine, on the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.

5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod02\Demo02.ssmssln**, and then click **Open**.
6. Open the **33 - Demonstration 3C.sql** script file.
7. Follow the instructions contained within the comments of the script file.
8. Close SQL Server Management Studio without saving any changes.

Module Review and Takeaways

Best Practice

All tables should have primary keys.

Foreign keys should be declared within the database in almost all circumstances. Developers often suggest that the application will ensure referential integrity, but experience shows that this is a poor option. Databases are often accessed by multiple applications, and bugs are also easy to miss when they first start to occur.

Review Question(s)

Question: What is a primary key?

Answer: A primary key is a key that consists of one or more columns that can uniquely identify a row in a table. It cannot be NULL.

Question: What is a foreign key?

Answer: A foreign key is a key in one table that references a candidate key (normally a primary key) from another table.

Question: What is meant by the term “referential integrity”?

Answer: Referential integrity means ensuring that foreign key relationships are enforced.

Lab Review Questions and Answers

Lab: Designing and Implementing Tables

Question and Answers

Lab Review

Question: When should a column be declared as nullable?

Answer: A column should be declared as nullable when the value can be unknown.

Question: Would it be reasonable to have columns called, for example, **AddressLine1**, **AddressLine2**, and **AddressLine3** in a normalized design?

Answer: Yes, these columns might represent distinct attributes of an object such as a customer. As an example, the addresses might represent different lines on a form.

Module 3

Ensuring Data Integrity through Constraints

Contents:

Lesson 1: Enforcing Data Integrity	2
Lesson 2: Implementing Domain Integrity	4
Lesson 3: Implementing Entity and Referential Integrity	6
Module Review and Takeaways	9
Lab Review Questions and Answers	10

Lesson 1

Enforcing Data Integrity

Contents:

Question and Answers

3

Question and Answers

Categorize each item into the appropriate type of data integrity. Indicate your answer by writing the category number to the right of each item.

Items	
1	A default constraint
2	A CHECK that restricts the range of values
3	A foreign key constraint
4	A primary key constraint
5	A tinyint data type
6	A unique constraint

Category 1		Category 2		Category 3
Domain Integrity		Entity Integrity		Referential Integrity

Answer:

Category 1		Category 2		Category 3
Domain Integrity		Entity Integrity		Referential Integrity
A tinyint data type A CHECK that restricts the range of values A default constraint		A primary key constraint A unique constraint		A foreign key constraint

Lesson 2

Implementing Domain Integrity

Contents:

Question and Answers	5
Demonstration: Data and Domain Integrity	5

Question and Answers

Question: You have a table that stores employees. The Bonus column stores the value of an employee's end-of-year bonus. You want to ensure that, when a new Employee is created, if the user does not specify a bonus, the Bonus column stores a zero. What type of constraint should you use?

Answer: Use a DEFAULT constraint

Demonstration: Data and Domain Integrity

Demonstration Steps

Enforce data and domain integrity

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. Run D:\Demofiles\Mod03\Setup.cmd as an administrator to revert any changes.
3. In the virtual machine, on the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod03\Demo03.ssmssln**, and then click **Open**.
6. If Solution Explorer is not visible, click the **View** menu and click **Solution Explorer**.
7. Expand the **Queries** folder and double-click **21 - Demonstration 2A.sql**.
8. Follow the instructions contained within the comments of the script file.

Lesson 3

Implementing Entity and Referential Integrity

Contents:

Question and Answers	7
Demonstration: Entity and Referential Integrity	7
Demonstration: Working with Identity and Sequences	7

Question and Answers

Question: You have a table that stores employees and their details. When an employee joins the company, they are given an employee number, which they keep for their entire employment. The EmployeeNumber column has already been created but you want to make sure that the same number can never be used twice. What kind of object should you use?

- ☐ A PRIMARY KEY constraint
- ☐ A UNIQUE constraint
- ☐ A FOREIGN KEY constraint
- ☐ An IDENTITY property
- ☐ A sequence

Answer:

- ☐ A PRIMARY KEY constraint
- ☒ A UNIQUE constraint
- ☐ A FOREIGN KEY constraint
- ☐ An IDENTITY property
- ☐ A sequence

Demonstration: Entity and Referential Integrity

Demonstration Steps

Define entity integrity for table, define referential integrity for tables, and define cascading referential integrity constraints

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. Run D:\Demofiles\Mod03\Setup.cmd as an administrator to revert any changes
3. In the virtual machine, on the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod03\Demo03.ssmssln**, and then click **Open**.
6. Open the **31 - Demonstration 3A.sql** script file.
7. Follow the instructions contained within the comments of the script file.

Demonstration: Working with Identity and Sequences

Demonstration Steps

Work with identity constraints, create a sequence, and use a sequence to provide key values for two tables

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. Run D:\Demofiles\Mod03\Setup.cmd as an administrator to revert any changes.
3. In the virtual machine, on the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.

5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod03\Demo03.ssmssln**, and then click **Open**.
6. Open the **32 - Demonstration 3B.sql** script file.
7. Follow the instructions contained within the comments of the script file.
8. Close SQL Server Management Studio and SQL Server Profiler without saving any changes.

Module Review and Takeaways

Best Practice

When you create a constraint on a column, if you do not specify a name for the constraint, SQL Server will generate a unique name for the constraint. However, you should always name constraints to adhere to your naming conventions.

Review Question(s)

Question: Why implement CHECK constraints if an application is already checking the input data?

Answer: Even if an application checks that data conforms in the user interface or back-end code, error conditions may arise that cause fields to become corrupt or null. Also, procedures for archiving, backing up, and triggers may attempt to copy bad data into the table, which can then cause an application to fail. Multiple applications may be accessing the same data.

Question: What are some scenarios in which you may want to temporarily disable constraint checking?

Answer: Constraint checking can affect performance, so you might want to disable it when performing large inserts, such as in a restore procedure or copying large number of records for an archive. In addition, you may know that duplicate or invalid data exists in your source or destination and have a plan to deal with cleaning up the data afterwards, such as by using a script or other procedure.

Lab Review Questions and Answers

Lab: Ensuring Data Integrity Through Constraints

Question and Answers

Lab Review

Question: In SQL Server Management Studio, you have successfully run a script that created a table, but you don't see the table in Object Explorer. What do you need to do?

Answer: In Object Explorer, expand the database, right-click **Tables**, and then click **Refresh**.

Question: What does the **DEFAULT** option do when you create a column?

Answer: The **DEFAULT** option specifies a default value for the column for new records.

Question: What requirement does a PRIMARY KEY constraint have that a UNIQUE constraint does not?

Answer: Primary keys cannot be NULL.

Module 4

Introduction to Indexes

Contents:

Lesson 1: Core Indexing Concepts	2
Lesson 2: Single-Column and Composite Indexes	4
Lesson 3: Table Structures in SQL Server	6
Lesson 4: Working with Clustered Indexes	8
Lesson 5: Working with Nonclustered Indexes	10
Module Review and Takeaways	12
Lab Review Questions and Answers	13

Lesson 1

Core Indexing Concepts

Contents:

Question and Answers	3
Demonstration: Viewing Index Fragmentation	3

Question and Answers

Question: You have a table in a Human Resources database that stores details of the company's employees. You want to build an index to perform as fast as possible. Which of the following columns should you use?

- () The employee's family name.
- () The employee's GUID.
- () The employee's start date.
- () The employee's age.
- () The employee's department name.

Answer:

- () The employee's family name.
- () The employee's GUID.
- () The employee's start date.
- (v) The employee's age.
- () The employee's department name.

Demonstration: Viewing Index Fragmentation

Demonstration Steps

Identify fragmented indexes

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. Run D:\Demofiles\Mod04\Setup.cmd as an administrator to revert any changes.
3. In the virtual machine, on the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod04\Demo04.ssmssln**, and then click **Open**.
6. On the **View** menu, click **Solution Explorer**.
7. Expand the **Queries** folder.
8. Open the **11 – Demonstration 1Asql** script file.
9. Follow the instructions contained within the comments of the script file.

Lesson 2

Single-Column and Composite Indexes

Contents:

Question and Answers

5

Question and Answers

Question: A recent bulk import has changed most of the data in the Employees table. You have reason to believe that SQL Server has not updated the statistics since the import happened. What are the likely consequences for queries executed when statistics are out-of-date?

Answer: Queries are likely to return results more slowly whilst statistics are out-of-date.

Lesson 3

Table Structures in SQL Server

Contents:

Question and Answers	7
Demonstration: Rebuilding Heaps	7

Question and Answers

Question: In the Employees table, you have created a clustered index on the FamilyName column. In this situation, what does SQL Server use to distinguish between employees called Archibald Smith and Sarah Smith?

- ☐ A uniqueifier
- ☐ A forwarding pointer
- ☐ A composite index
- ☐ A descending index
- ☐ A rebuilt heap

Answer:

- ☒ A uniqueifier
- ☐ A forwarding pointer
- ☐ A composite index
- ☐ A descending index
- ☐ A rebuilt heap

Demonstration: Rebuilding Heaps

Demonstration Steps

Create a table as a heap, check the fragmentation and forwarding pointers for a heap, and rebuild a heap

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. Run D:\Demofiles\Mod04\Setup.cmd as an administrator to revert any changes
3. In the virtual machine, on the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod04\Demo04.ssmssln**, and then click **Open**.
6. On the **View** menu, click **Solution Explorer**.
7. Open the **31 – Demonstration 3A.sql** script file.
8. Follow the instructions contained within the comments of the script file.

Lesson 4

Working with Clustered Indexes

Contents:

Question and Answers	9
Resources	9
Demonstration: Working with Clustered Indexes	9

Question and Answers

Question: You have created a clustered index in the Employees table based on the SocialSecurityNumber column. You want to reduce that potential for fragmentation in the index. What should you do?

Answer: Set the FILLFACTOR to a value less than 100, such as 70

Resources

Persisting Data by Using Indexes



Reference Links: For a full list of the requirements for creating indexed views, see the *Creating Indexed Views* topic in SQL Server Books Online.



Reference Links: For information about the requirements for creating indexes on computed columns, see the *Indexes on Computed Columns* topic in SQL Server Books Online.

Demonstration: Working with Clustered Indexes

Demonstration Steps

Create a table that has a clustered index, detect fragmentation in a clustered index, and correct fragmentation in a clustered index

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. Run D:\Demofiles\Mod04\Setup.cmd as an administrator to revert any changes.
3. In the virtual machine, on the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod04\Demo04.ssmssln**, and then click **Open**.
6. On the **View** menu, click **Solution Explorer**.
7. Open the **41 – Demonstration 4A.sql** script file.
8. Follow the instructions contained within the comments of the script file.

Lesson 5

Working with Nonclustered Indexes

Contents:

Question and Answers	11
Demonstration: Working with Nonclustered Indexes	11

Question and Answers

Question: You have created a nonclustered index on a column in a heap. What value does the index return to enable SQL Server to locate the right rows to return for a query?

Answer: The index returns the Row IDs of the right rows.

Demonstration: Working with Nonclustered Indexes

Demonstration Steps

Create covering indexes

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. Run D:\Demofiles\Mod04\Setup.cmd as an administrator to revert any changes.
3. Start SQL Server Management Studio.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod04\Demo04.ssmssln**, and then click **Open**.
6. On the **View** menu, click **Solution Explorer**.
7. Open the **51 – Demonstration 5A.sql** script file.
8. Follow the instructions contained within the comments of the script file.
9. Close SQL Server Management Studio and SQL Server Profiler without saving any changes.

Module Review and Takeaways

Best Practice

Unless specific circumstances arise, most tables should have a clustered index.

The clustered index may or may not be placed on the table's primary key.

When using GUID primary keys in the logical data model, consider avoiding their use throughout the physical implementation of the data model.

Review Question(s)

Question: What is the main problem with using unique identifiers as primary keys?

Answer: The main problem with using unique identifiers as primary keys is the random order of their values.

Question: Where are newly inserted rows placed when a table is structured as a heap?

Answer: When a table is structured as a heap, newly inserted rows are placed in any page that has sufficient space available.

Lab Review Questions and Answers

Lab: Creating Indexes

Question and Answers

Lab Review

Question: When is it important that a clustered index has an increasing key?

Answer: It is important that a clustered index has an increasing key when significant **INSERT** operations are expected in the order of the key.

Question: Which table structure is automatically assigned when a table is assigned a primary key during table creation and no structure is specified?

Answer: A clustered index structure is automatically assigned when a table is assigned a primary key during table creation and no structure is specified.

Module 5

Advanced Indexing

Contents:

Lesson 1: Core Concepts of Execution Plans	2
Lesson 2: Common Execution Plan Elements	5
Lesson 3: Working with Execution Plans	7
Lesson 4: Designing Effective Nonclustered Indexes	9
Lesson 5: Performance Monitoring	11
Module Review and Takeaways	13
Lab Review Questions and Answers	14

Lesson 1

Core Concepts of Execution Plans

Contents:

Question and Answers	3
Resources	3
Demonstration: Viewing Execution Plans in SQL Server Management Studio	4

Question and Answers

Put the following phases in the creation of a query plan in order by numbering each to indicate the correct order.

	Steps
	Query Plan Caching
	Object Name Resolution
	Query Plan Execution
	Transact-SQL Parsing
	Query Optimization

Answer:

	Steps
5	Query Plan Caching
2	Object Name Resolution
4	Query Plan Execution
1	Transact-SQL Parsing
3	Query Optimization

Resources

SET Statements



Reference Links: For more information about the **SET** commands, see the *Displaying Execution Plans by Using the Showplan SET Options* topic in the Microsoft Developer Network (MSDN) library.

Demonstration: Viewing Execution Plans in SQL Server Management Studio

Demonstration Steps

Use execution plans

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. Run D:\Demofiles\Mod05\Setup.cmd as an administrator to revert any changes.
3. In the virtual machine, on the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod05\Demo05.ssmssl**, and then click **Open**.
6. On the **View** menu, click **Solution Explorer**.
7. Expand the **Queries** folder.
8. Open the **11 – Demonstration 1A.sql** script file.
9. Follow the instructions contained within the comments of the script file.

Lesson 2

Common Execution Plan Elements

Contents:

Question and Answers	6
Demonstration: Working with Common Execution Plan Elements	6

Question and Answers

Question: A user has submitted a query that joins the SalesRep table with the Customers table. The SalesRep table has a clustered index based on the SalesRep.EmployeeID column. The Customers table has a clustered index based on the Customers.CustomerID column. What common execution plan element is the query optimizer likely to add to the query plan to implement this join?

- () A table scan
- () A clustered index seek
- () A nested lookup
- () A merge join
- () A hash match

Answer:

- () A table scan
- () A clustered index seek
- () A nested lookup
- () A merge join
- (√) A hash match

Demonstration: Working with Common Execution Plan Elements

Demonstration Steps

Run queries that demonstrate the most common execution plan elements

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. If you have not completed the previous demonstrations in this module, run D:\Demofiles\Mod05\Setup.cmd as an administrator to revert any changes.
3. In the virtual machine, on the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod05\Demo05.ssmssln**, and then click **Open**.
6. On the **View** menu, click **Solution Explorer**.
7. Open the **21 – Demonstration 2A.sql** script file.
8. Follow the instructions contained within the comments of the script file.

Lesson 3

Working with Execution Plans

Contents:

Question and Answers	8
Demonstration: Viewing Cached Plans	8

Question and Answers

Question: You have identified a poorly performing request and its plan handle by using the **sys.dm_exec_requests** DMV. Which DMV can you use to obtain the query plan in XML format?

- () sys.dm_exec_sql_text
- () sys.dm_exec_query_plan
- () sys.dm_exec_cached_plans
- () sys.dm_exec_requests
- () sys.dm_exec_query_stats

Answer:

- () sys.dm_exec_sql_text
- (√) sys.dm_exec_query_plan
- () sys.dm_exec_cached_plans
- () sys.dm_exec_requests
- () sys.dm_exec_query_stats

Demonstration: Viewing Cached Plans

Demonstration Steps

Viewing cached execution plans

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. Run D:\Demofiles\Mod05\Setup.cmd as an administrator to revert any changes.
3. In the virtual machine, on the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod05\Demo05.ssmssln**, and then click **Open**.
6. On the **View** menu, click **Solution Explorer**.
7. Open the **31 – Demonstration 3A.sql** script file.
8. Follow the instructions contained within the comments of the script file.

Lesson 4

Designing Effective Nonclustered Indexes

Contents:

Question and Answers	10
Demonstration: Obtaining Index Information	10

Question and Answers

Question: You have recently added a new nonclustered index to the Customers table. You are evaluating its performance and you would like to know how often the new index has been used. What DMV could you use to obtain this information?

Answer: sys.dm_db_index_usage_stats

Demonstration: Obtaining Index Information

Demonstration Steps

Viewing index information

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **AdventureWorks\Student** with the password **Pa\$\$w0rd**.
2. If you have not completed the previous demonstrations in this module, run D:\Demofiles\Mod05\Setup.cmd as an administrator to revert any changes.
3. In the virtual machine, on the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod05\Demo05.ssmssl**, and then click **Open**.
6. On the **View** menu, click **Solution Explorer**.
7. Open the **41 - Demonstration 4A.sql** script file.
8. Follow the instructions contained within the comments of the script file.
9. Close SQL Server Management Studio without saving any changes.

Lesson 5

Performance Monitoring

Contents:

Question and Answers

12

Question and Answers

Question: In the baseline you established one month after database deployment, the **System:Processor Queue Length** counter showed an average value of 2.3. Your recent monitoring shows that this counter now averages 1.8. What action should you take?

Answer: Take no action because the counter shows an improvement

Module Review and Takeaways

Best Practice

Avoid capturing execution plans for large numbers of statements when you use SQL Server Profiler.

If you need to capture plans by using SQL Server Profiler, make sure that the trace is filtered to reduce the number of events that are captured.

Review Question(s)

Question: What is the difference between a graphical execution plan and an XML execution plan?

Answer: A graphical plan is a rendering of an XML plan. It contains the most important information from the plan in an easy-to-read format.

Question: Why might a Transact-SQL **DELETE** statement have a complex execution plan?

Answer: There might be referential integrity checks to be done (that is, foreign keys).

Lab Review Questions and Answers

Lab: Advanced Indexing

Question and Answers

Lab Review

Question: Can two different queries end up with the same execution plan?

Answer: Yes, two different queries can end up with the same execution plan. The execution plan details the steps to retrieve the required results. There are many ways to express the requirements of a single query.

Module 6

In-Memory Database Capabilities

Contents:

Lesson 1: The Buffer Pool Extension	2
Lesson 2: Columnstore Indexes	4
Module Review and Takeaways	6

Lesson 1

The Buffer Pool Extension

Contents:

Question and Answers	3
Demonstration: Using the Buffer Pool Extension	3

Question and Answers

Question: You have a database running on 64 bit SQL Server 2014 Web Edition. You have a database that often experiences a high volume of read operations and occasional write operations. You have 28 GB of physical memory and a spare slot where you can install a new SSD disk. Should you enable the buffer pool extension?

Answer: In this scenario you cannot enable the buffer pool extension.

Demonstration: Using the Buffer Pool Extension

Demonstration Steps

Configure the buffer pool extension

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. Run D:\Demofiles\Mod06\Setup.cmd as an administrator to revert any changes.
3. In the virtual machine, on the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod06\Demo06.ssmssln**, and then click **Open**.
6. On the **View** menu, click **Solution Explorer**.
7. Expand the **Queries** folder.
8. Open the **11 – Demonstration 1A.sql** script file.
9. Follow the instructions contained within the comments of the script file.

Lesson 2

Columnstore Indexes

Contents:

Question and Answers	5
Resources	5
Demonstration: Using a Columnstore Index	5

Question and Answers

Question: You have recently created a nonclustered columnstore index on a table in your data warehouse. You have seen a large improvement in the performance of data mining operations. However, a fellow administrator contacts you to say that he can no longer import data to the table. How can you solve this problem?

Answer: You must either drop the index or use table partitioning to enable updates to data in the table.

Resources

Columnstore Index Scenarios



Reference Links: For a full list of the limitations of using columnstore indexes, see the *Columnstore Indexes* topic in SQL Server Books Online.

Demonstration: Using a Columnstore Index

Demonstration Steps

Create a columnstore index

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **AdventureWorks\Student** with the password **Pa\$\$w0rd**.
2. Run D:\Demofiles\Mod06\Setup.cmd as an administrator to revert any changes.
3. In the virtual machine, on the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod06\Demo06.ssmssln**, and then click **Open**.
6. On the **View** menu, click **Solution Explorer**.
7. Open the **21 – Demonstration 2A.sql** script file.
8. Follow the instructions contained within the comments of the script file.
9. Close SQL Server Management Studio without saving any changes.

Module Review and Takeaways

In this module, you have learned about some of the ways in which SQL Server 2014 takes advantage of the increasing amount of RAM in modern servers, and provides in-memory optimizations for database workloads.

Module 7

Designing and Implementing Views

Contents:

Lesson 1: Introduction to Views	2
Lesson 2: Creating and Managing Views	4
Lesson 3: Performance Considerations for Views	6
Module Review and Takeaways	8
Lab Review Questions and Answers	9

Lesson 1

Introduction to Views

Contents:

Question and Answers	3
Demonstration: Querying System Views and Dynamic Management Views	3

Question and Answers

Question: You have a table named Employees in the HumanResources database that contains sensitive personal data. You want to enable members of the TeamLeaders group to read the FirstName, LastName, and Skills columns but to prevent them from accessing the Age, Ethnicity, and Spouse columns. How can you configure this?

Answer: Ensure that TeamLeaders do not have access to the Employees table. Create a View that includes FirstName, LastName, and Skills columns but not the Age, Ethnicity, and Spouse columns. Grant TeamLeaders access to the view.

Demonstration: Querying System Views and Dynamic Management Views

Demonstration Steps

Query system views and dynamic management views

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **AdventureWorks\Student** with the password **Pa\$\$w0rd**.
2. Run D:\Demofiles\Mod07\Setup.cmd as an administrator to revert any changes.
3. In the virtual machine, on the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod07\Demo07.ssmssln**, and then click **Open**.
6. Expand the **Queries** folder.
7. Open the **11 - Demonstration 1A.sql** script file.
8. Follow the instructions contained within the comments of the script file.

Lesson 2

Creating and Managing Views

Contents:

Question and Answers	5
Demonstration: Implementing Views	5

Question and Answers

Question: You are creating a new view and you want to ensure that changes to the underlying tables do not render the view meaningless. Which T-SQL element should you use on the CREATE VIEW command?

- () The DISTINCT
- () The WITH ENCRYPTION option
- () The HAVING keyword
- () The WITH SCHEMABINDING option
- () The WITH CHECK option

Answer:

- () The DISTINCT
- () The WITH ENCRYPTION option
- () The HAVING keyword
- (√) The WITH SCHEMABINDING option
- () The WITH CHECK option

Demonstration: Implementing Views

Demonstration Steps

Create, query and drop views

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. Run D:\Demofiles\Mod07\Setup.cmd as an administrator to revert any changes.
3. In the virtual machine, on the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod07\Demo07.ssmssln**, and then click **Open**.
6. Open the **21 - Demonstration 2A.sql** script file.
7. Follow the instructions contained within the comments of the script file.

Lesson 3

Performance Considerations for Views

Contents:

Question and Answers	7
Demonstration: Investigating Views and Performance	7

Question and Answers

Question: A user submits a query against a view. The view definition joins one table and a second view. The second view definition is a union of two other tables. How many query execution plans will the optimizer create for this query?

Answer: One

Demonstration: Investigating Views and Performance

Demonstration Steps

Investigate how views can affect query performance

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. If you have not completed the previous demonstrations, run D:\Demofiles\Mod07\Setup.cmd as an administrator to revert any changes.
3. In the virtual machine, on the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod07\Demo07.ssmssln**, and then click **Open**.
6. Open the **31 - Demonstration 3A.sql** script file.
7. Follow the instructions contained within the comments of the script file.
8. Close SQL Server Management Studio and SQL Server Profiler without saving any changes.

Module Review and Takeaways

Best Practice

Use views to focus data for users.

Avoid nesting many layers within views.

Avoid ownership chain problems within views.

Ensure consistent connection **SET** options when intending to index views.

Review Question(s)

Question: How does SQL Server store the view in the database?

Answer: The **SELECT** statement is what SQL Server stores in the database.

Question: What is a standard, nonindexed view?

Answer: Standard views combine data from one or more base tables (or views) into a new virtual table. They are materialized at run time.

Question: What is an unbroken ownership chain?

Answer: When the same user owns the source object, view, stored procedure, or user-defined function, and all target objects (underlying tables, views, or other objects), the ownership chain is said to be unbroken.

Lab Review Questions and Answers

Lab: Designing and Implementing Views

Question and Answers

Lab Review

Question: What considerations are there for views that involve multiple tables?

Answer: If the view is updatable, only data from a single table can be updated in any **UPDATE** statement.

Question: What is required for columns in views that are created from expressions?

Answer: Columns in views that are based on expressions need to be aliased.

Module 8

Designing and Implementing Stored Procedures

Contents:

Lesson 1: Introduction to Stored Procedures	2
Lesson 2: Working with Stored Procedures	4
Lesson 3: Implementing Parameterized Stored Procedures	6
Lesson 4: Controlling Execution Context	8
Module Review and Takeaways	10
Lab Review Questions and Answers	11

Lesson 1

Introduction to Stored Procedures

Contents:

Question and Answers	3
Demonstration: Working with System Stored Procedures and System Extended Stored Procedures	3

Question and Answers

Question: Which of the following statements can you use in a stored procedure?

- ☐ CREATE DEFAULT
- ☐ ALTER PROCEDURE
- ☐ CREATE RULE
- ☐ USE HumanResources
- ☐ CREATE VIEW

Answer:

- ☐ CREATE DEFAULT
- ☐ ALTER PROCEDURE
- ☐ CREATE RULE
- ☐ USE HumanResources
- ☒ CREATE VIEW

Demonstration: Working with System Stored Procedures and System Extended Stored Procedures

Demonstration Steps

Execute system stored procedures

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. Run D:\Demofiles\Mod08\Setup.cmd as an administrator to revert any changes.
3. In the virtual machine, on the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod08\Demo08.ssmssln**, and then click **Open**.
6. Expand the **Queries** folder.
7. Open the **11 - Demonstration 1A.sql** script file.
8. Follow the instructions contained within the comments of the script file.

Lesson 2

Working with Stored Procedures

Contents:

Question and Answers	5
Demonstration: Implementing Stored Procedures	5

Question and Answers

Question: You have many stored procedures in your database and many of them call other stored procedures. You want to return the number of nested levels as an output parameter to the calling user so that developers can use this information for optimization. How can you determine the current number of nested stored procedures?

Answer: Use the @@nestlevel to determine the current number of nested levels.

Demonstration: Implementing Stored Procedures

Demonstration Steps

Create, execute, and alter a stored procedure

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **AdventureWorks\Student** with the password **Pa\$\$w0rd**.
2. Ensure that you have run the previous demos in this module.
3. In the virtual machine, on the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod08\Demo08.ssmssln**, and then click **Open**.
6. Expand the **Queries** folder.
7. Open the **21 - Demonstration 2A.sql** script file.
8. Follow the instructions contained within the comments of the script file.

Lesson 3

Implementing Parameterized Stored Procedures

Contents:

Question and Answers	7
Demonstration: Passing Parameters to Stored Procedures	7

Question and Answers

Question: You have created a stored procedure by using the following code:

```
CREATE PROC HumanResources.GetVacationDaysByEmployeeName
@FamilyName NVARCHAR(50), @VacationDays int OUTPUT
AS
BEGIN
    SELECT @FamilyName = VacationDays
    FROM HumanResources.Employees AS e
    WHERE e.FamilyName = @FamilyName
END;
```

You call the procedure with the following code:

```
DECLARE @FamilyName NVARCHAR(50) = 'Andrews';
DECLARE @VacationDays int;
EXEC HumanResources.GetVacationDaysByEmployeeName @FamilyName,
@VacationDays;
SELECT @VacationDays;
```

This code always returns NULL even though all rows have an integer value in the VacationDays column. What can you do to fix this code?

Answer: In the EXEC statement, add the OUTPUT keyword after “@VacationDays”

Demonstration: Passing Parameters to Stored Procedures

Demonstration Steps

Pass parameters to stored procedures

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. Ensure that you have run the previous demos in this module
3. In the virtual machine, on the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod08\Demo08.ssmssl**, and then click **Open**.
6. Expand the **Queries** folder.
7. Open the **31 - Demonstration 3A.sql** script file.
8. Follow the instructions contained within the comments of the script file.

Lesson 4

Controlling Execution Context

Contents:

Question and Answers	9
Demonstration: Viewing Execution Context	9

Question and Answers

Question: A stored procedure is returning an access denied error. It uses the EXECUTE AS clause to specify a login named "MaintenanceLogin". You want to find out which server-level roles this login is a member of. Which system view can you use to find this information?

Answer: Use the sys.login_token system view.

Demonstration: Viewing Execution Context

Demonstration Steps

View and change the execution context

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. Ensure that you have run the previous demos in this module.
3. In the virtual machine, on the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod08\Demo08.ssmssln**, and then click **Open**.
6. Expand the **Queries** folder.
7. Open the **41 - Demonstration 4A.sql** script file.
8. Follow the instructions contained within the comments of the script file.
9. Close SQL Server Management Studio without saving any changes.

Module Review and Takeaways

Best Practice

Use the EXECUTE AS clause to override the execution context of stored procedures that use dynamic SQL, rather than granting permissions on the underlying tables to users.

Design procedures to perform individual tasks. Avoid designing procedures that perform a large number of tasks, unless those tasks are performed by executing other stored procedures.

Keep consistent ownership of stored procedures, views, tables, and other objects within databases.

Review Question(s)

Question: What happens to the **WITH RECOMPILE** option when you use it with a **CREATE PROC** statement?

Answer: When you use the **WITH RECOMPILE** option with a **CREATE PROC** statement, it causes a new execution plan to be generated every time the procedure is executed.

Question: What happens to the **WITH RECOMPILE** option when you use it with an **EXECUTE** statement?

Answer: When you use the **WITH RECOMPILE** option with an **EXECUTE** statement, it causes a new execution plan to be generated for this particular execution of the procedure and causes the plan to be discarded after execution.

Lab Review Questions and Answers

Lab: Designing and Implementing Stored Procedures

Question and Answers

Lab Review

Question: When do you need the **OUTPUT** keyword for output parameters when you are working with stored procedures?

Answer: You need the **OUTPUT** keyword when you are declaring the parameters in the stored procedure and when you are calling the stored procedure in the **EXEC** statement.

Module 9

Designing and Implementing User-Defined Functions

Contents:

Lesson 1: Overview of Functions	2
Lesson 2: Designing and Implementing Scalar Functions	4
Lesson 3: Designing and Implementing Table-Valued Functions	6
Lesson 4: Considerations for Implementing Functions	8
Lesson 5: Alternatives to Functions	10
Module Review and Takeaways	12

Lesson 1

Overview of Functions

Contents:

Question and Answers

3

Question and Answers

Question: You want to add a function to look up the total value of all orders for each customer in a month. You want to be able to pass the relevant month and receive a table of data with the values per customer. What kind of object should you create?

- ☐ A system function
- ☐ A scalar function
- ☐ A table-valued function
- ☐ A view
- ☐

Answer:

- ☐ A system function
- ☐ A scalar function
- ☒ A table-valued function
- ☐ A view
- ☐

Lesson 2

Designing and Implementing Scalar Functions

Contents:

Question and Answers	5
Demonstration: Working with Scalar Functions	5

Question and Answers

Question: You want to create a database object that will accept an ISBN number and return the number of times the corresponding book has been sold in the current year. You also want the object to store this number in the Books table. Should you create a scalar function or a stored procedure?

Answer: Create a stored procedure.

Demonstration: Working with Scalar Functions

Demonstration Steps

Work with scalar functions

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. Run D:\Demofiles\Mod09\Setup.cmd as an administrator to revert any changes.
3. In the virtual machine, on the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod09\Demo09.ssmssln**, and then click **Open**.
6. Expand the **Queries** folder.
7. Open the **21 - Demonstration 2A.sql** script file.
8. Follow the instructions contained within the comments of the script file.

Lesson 3

Designing and Implementing Table-Valued Functions

Contents:

Question and Answers	7
Demonstration: Implementing Table-Valued Functions	7

Question and Answers

Question: You have defined a table-valued function with the following code:

```
CREATE FUNCTION Sales.GetLastOrdersForCustomer
    (@CustomerID int, @NumberOfOrders int)
AS
RETURN (SELECT TOP(@NumberOfOrders), soh.SalesorderID,
        soh.OrderDate, soh.PurchaseOrderNumber
        FROM Sales.SalesOrderHeader AS soh
        WHERE soh.CustomerID = @CustomerID
        ORDER BY soh.OrderDate DESC, soh.SalesOrderID DESC);
```

You receive an error. What is missing from the code?

Answer: Add "RETURNS TABLE" before the word "AS".

Demonstration: Implementing Table-Valued Functions

Demonstration Steps

1. In the virtual machine, on the taskbar, click **SQL Server 2014 Management Studio**.
2. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
3. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod09\Demo09.ssmssln**, and then click **Open**.
4. Expand the **Queries** folder.
5. Open the **31 - Demonstration 3A.sql** script file.
6. Follow the instructions contained within the comments of the script file.

Lesson 4

Considerations for Implementing Functions

Contents:

Demonstration: Controlling the Execution Context

9

Demonstration: Controlling the Execution Context

Demonstration Steps

Alter the execution context of a function

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. Run D:\Demofiles\Mod09\Setup.cmd as an administrator to revert any changes.
3. In the virtual machine, on the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod09\Demo09.ssmssln**, and then click **Open**.
6. Expand the **Queries** folder.
7. Open the **41 - Demonstration 4A.sql** script file.
8. Follow the instructions contained within the comments of the script file.
9. Close SQL Server Management Studio without saving any changes

Lesson 5

Alternatives to Functions

Contents:

Question and Answers

11

Question and Answers

Question: You want to create a database object that accepts a price and returns the number of orders placed in the current month that are greater than that price. You want to include error-handling code to ensure that, when there are no orders in the month, the error is handled properly and not returned to the caller. What type of object should you create?

Answer: Create a stored procedure

Module Review and Takeaways

Best Practice

Avoid calling multistatement TVFs for each row of a query. In many cases, you can dramatically improve performance by extracting the code from the query into the surrounding query.

Use the WITH EXECUTE AS clause to override the security context of code that needs to perform actions that the user who is executing the code does not have.

Review Question(s)

Question: When you are using the EXECUTE AS clause, what privileges should you grant to the login or user that is being impersonated?

Answer: For the login or user that is being impersonated, specify a login or user that has the least privileges required to perform the operations that are required in the session. For example, do not specify a login name with server-level permissions if only database-level permissions are required. Also, do not specify a database owner account unless those permissions are required.

Question: When you are using the EXECUTE AS clause, what privileges should you grant to the login or user that is creating the code?

Answer: You should grant IMPERSONATE permission for the login or user that is creating the code.

Module 10

Responding to Data Manipulation via Triggers

Contents:

Lesson 1: Designing DML Triggers	2
Lesson 2: Implementing DML Triggers	4
Lesson 3: Advanced Trigger Concepts	7
Module Review and Takeaways	9

Lesson 1

Designing DML Triggers

Contents:

Question and Answers

3

Question and Answers

Question: You have a custom desktop application that manipulates data in the Human Resources database. You have recently added a trigger to the Employee table to implement auditing by recording the details of update operations in the UserActions table. Since you added the trigger, the application has been displaying an error message that says "An unexpected number of rows was returned". What can you do to remove this error?

Answer: Use the SET NOCOUNT ON option for the trigger.

Lesson 2

Implementing DML Triggers

Contents:

Question and Answers	5
Demonstration: Working with AFTER INSERT Triggers	5
Demonstration: Working with AFTER DELETE Triggers	5
Demonstration: Working with AFTER UPDATE Triggers	6

Question and Answers

Question: You are testing a trigger on a table in your staging database. You want to remove all the testing entries from the table without firing the DELETE AFTER trigger you have defined, which would add many uninteresting rows to the AuditEvents table. How can you perform this operation without firing the trigger?

Answer: Use a TRUNCATE TABLE statement

Demonstration: Working with AFTER INSERT Triggers

Demonstration Steps

Create an AFTER INSERT trigger

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. Run D:\Demofiles\Mod10\Setup.cmd as an administrator to revert any changes.
3. On the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod10\Demo10.ssmssln**, and then click **Open**.
6. Open the **Queries** folder.
7. Open the **21 - Demonstration 2A.sql** script file.
8. Follow the instructions contained within the comments of the script file.

Demonstration: Working with AFTER DELETE Triggers

Demonstration Steps

Create and test AFTER DELETE triggers

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **AdventureWorks\Student** with the password **Pa\$\$w0rd**.
2. If you have not completed the previous demonstration in this module, then run D:\Demofiles\Mod10\Setup.cmd as an administrator to revert any changes
3. On the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod10\Demo10.ssmssln**, and then click **Open**.
6. Open the **Queries** folder.
7. Open the **22 - Demonstration 2B.sql** script file.
8. Follow the instructions contained within the comments of the script file.

Demonstration: Working with AFTER UPDATE Triggers

Demonstration Steps

Create and test AFTER UPDATE triggers

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **AdventureWorks\Student** with the password **Pa\$\$w0rd**.
2. If you have not completed the previous demonstrations in this module, then run D:\Demofiles\Mod10\Setup.cmd as an administrator to revert any changes. On the taskbar, click **SQL Server 2014 Management Studio**.
3. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
4. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod10\Demo10.ssmssl**, and then click **Open**.
5. Open the **Queries** folder.
6. Open the **23 - Demonstration 2C.sql** script file.
7. Follow the instructions contained within the comments of the script file.

Lesson 3

Advanced Trigger Concepts**Contents:**

Question and Answers	8
Demonstration: Working with INSTEAD OF Triggers	8
Demonstration: Replacing Triggers with Computed Columns	8

Question and Answers

Question: You have two AFTER INSERT triggers defined on the Customers table. You want to ensure that the CustomerNameCheck triggers, which ensures that the new customer does not have the same name as any existing customer, executes before the CreateCustomerID trigger, which creates a unique customer ID. How can you ensure the correct order?

Answer: Use the **sp_settriggerorder** system stored procedure.

Demonstration: Working with INSTEAD OF Triggers

Demonstration Steps

Create and test an INSTEAD OF DELETE trigger

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **AdventureWorks\Student** with the password **Pa\$\$w0rd**.
2. If you have not completed the previous demonstrations in this module, then run D:\Demofiles\Mod10\Setup.cmd as an administrator to revert any changes.
3. On the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod10\Demo10.ssmssl**, and then click **Open**.
6. Open the **31 - Demonstration 3A.sql** script file.
7. Follow the instructions contained within the comments of the script file.

Demonstration: Replacing Triggers with Computed Columns

Demonstration Steps

Replace a trigger with a computed column

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **AdventureWorks\Student** with the password **Pa\$\$w0rd**.
2. If you have not completed the previous demonstrations in this module, then run D:\Demofiles\Mod10\Setup.cmd as an administrator to revert any changes.
3. On the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod10\Demo10.ssmssl**, and then click **Open**.
6. Open the **Queries** folder.
7. Open the **32 - Demonstration 3B.sql** script file.
8. Follow the instructions contained within the comments of the script file.

Module Review and Takeaways

Best Practice

In many business scenarios, it makes sense to mark records as deleted with a status column and use a trigger or stored procedure to update an audit trail table. The changes can then be audited, the data is not lost, and the IT staff can perform purges or archival of the deleted records.

Avoid using triggers in situations where constraints could be used instead.

Review Question(s)

Question: How do constraints and triggers differ regarding timing of execution?

Answer: Constraints fire before the data manipulation has occurred. AFTER triggers fire after the data manipulation has occurred. INSTEAD OF triggers fire instead of the data manipulation.

Module 11

Using In-Memory Tables

Contents:

Lesson 1: Memory-Optimized Tables	2
Lesson 2: Natively Compiled Stored Procedures	5
Module Review and Takeaways	7

Lesson 1

Memory-Optimized Tables

Contents:

Question and Answers	3
Demonstration: Using Memory-Optimized Tables	3

Question and Answers

Put the following steps from the Memory Optimization Advisor in order by numbering each to indicate the correct order.

	Steps
	Primary Key Migration
	Migration Warnings
	Migration Validation
	Index Migration
	Summary
	Migration Options

Answer:

	Steps
4	Primary Key Migration
2	Migration Warnings
1	Migration Validation
5	Index Migration
6	Summary
3	Migration Options

Demonstration: Using Memory-Optimized Tables

Demonstration Steps

Using memory-optimized tables

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. Run D:\Demofiles\Mod11\Setup.cmd as an administrator to revert any changes.
3. On the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod11\Demo11.ssmssln**, and then click **Open**.

6. On the **View** menu, click **Solution Explorer**.
7. Expand the **Queries** folder.
8. Open the **11 – Demonstration 1A.sql** script file.
9. Follow the instructions contained within the comments of the script file.

Lesson 2

Natively Compiled Stored Procedures

Contents:

Question and Answers	6
Demonstration: Creating a Natively Compiled Stored Procedure	6

Question and Answers

Question: Which of the following elements is not required when you create a natively compiled stored procedure?

- () The BEGIN ATOMIC clause
- () The NATIVE_COMPILATION option
- () The EXECUTE AS option.
- () The LANGUAGE option
- () The SCHEMABINDING option

Answer:

- () The BEGIN ATOMIC clause
- () The NATIVE_COMPILATION option
- () The EXECUTE AS option.
- (√) The LANGUAGE option
- () The SCHEMABINDING option

Demonstration: Creating a Natively Compiled Stored Procedure

Demonstration Steps

Create a natively compiled stored procedure

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. Ensure that you have run the previous demonstration.
3. On the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod11\Demo11.ssmssl**, and then click **Open**.
6. On the **View** menu, click **Solution Explorer**.
7. Expand the **Queries** folder.
8. Open the **21 – Demonstration 2A.sql** script file.
9. Follow the instructions contained within the comments of the script file.
10. Close SQL Server Management Studio without saving any changes

Module Review and Takeaways

In this module, you have learned how to store tables in memory and how to natively compile stored procedures to access the memory-optimized tables.

Module 12

Implementing Managed Code in SQL Server

Contents:

Lesson 1: Introduction to CLR Integration in SQL Server	2
Lesson 2: Importing and Cataloging Assemblies	4
Lesson 3: Implementing CLR Integration in SQL Server	6
Module Review and Takeaways	8
Lab Review Questions and Answers	9

Lesson 1

Introduction to CLR Integration in SQL Server

Contents:

Question and Answers

3

Question and Answers

Question: Which of the following objects cannot be created by using Transact-SQL code?

- ☐ () A user-defined datatype
- ☐ () A scalar user-defined function
- ☐ () A table-valued user-defined function
- ☐ () A DML trigger
- ☐ () A DDL trigger

Answer:

- ☒ (√) A user-defined datatype
- ☐ () A scalar user-defined function
- ☐ () A table-valued user-defined function
- ☐ () A DML trigger
- ☐ () A DDL trigger

Lesson 2

Importing and Cataloging Assemblies

Contents:

Question and Answers	5
Demonstration: Importing and Cataloging an Assembly	5

Question and Answers

Question: You have successfully imported and cataloged a .dll assembly that contains several custom data types. When you try and use the datatypes in a CREATE TABLE statement, you receive an error. What could be the problem?

Answer: Answers may vary. However administrators should certainly check that the "clr enabled" server option has been set to 1 in these circumstances.

Demonstration: Importing and Cataloging an Assembly

Demonstration Steps

Import and catalog an assembly

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. Run D:\Demofiles\Mod12\Setup.cmd as an administrator to revert any changes.
3. On the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod12\Demo12.ssmssln**, and then click **Open**.
6. Expand the **Queries** folder.
7. Open the **21 - Demonstration 2A.sql** script file.
8. Follow the instructions contained within the comments of the script file.

Lesson 3

Implementing CLR Integration in SQL Server

Contents:

Question and Answers	7
Demonstration: Creating Aggregates and User-Defined Data Types	7

Question and Answers

Question: You are implementing a user-defined aggregate that a developer has created in a .NET framework .dll. The aggregate calculates a weighted mean of the values in the input column. Which of the following properties of the **SqlUserDefinedAggregate** attribute can the developer use to indicate that the mean is calculated only over non-null values?

- ☐ **IsInvariantToDuplicates**
- ☐ **IsInvariantToNulls**
- ☐ **IsNullIfEmpty**
- ☐ **Name**
- ☐ **Format.Native**

Answer:

- ☐ **IsInvariantToDuplicates**
- ☒ **IsInvariantToNulls**
- ☐ **IsNullIfEmpty**
- ☐ **Name**
- ☐ **Format.Native**

Demonstration: Creating Aggregates and User-Defined Data Types

Demonstration Steps

Create aggregates and user-defined data types

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **AdventureWorks\Student** with the password **Pa\$\$w0rd**.
2. Run D:\Demofiles\Mod12\Setup.cmd as an administrator to revert any changes.
3. On the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod12\Demo12.ssmssln**, and then click **Open**.
6. If the previous demonstration was not performed, open the **21 - Demonstration 2A.sql** script file and execute steps 1 to 3.
7. Expand the **Queries** folder.
8. Open the **31 - Demonstration 3A.sql** script file.
9. Follow the instructions contained within the comments of the script file.
10. Close SQL Server Management Studio without saving any changes

Module Review and Takeaways

Best Practice

When you are deciding between using Transact-SQL and managed code, the biggest mistake is to assume that either of them is the correct answer for every situation. Each has its benefits and limitations and should be used for the appropriate tasks.

Developers should avoid using SQL CLR to implement code that would be better placed on another application tier (such as on a client system).

Database administrators should avoid refusing to allow SQL CLR code without consideration. As you have seen in this module, there is code that should be implemented in managed code rather than in Transact-SQL.

Database administrators should set boundaries for developers:

No row-based code that should be set-based Transact-SQL operations.

Limited use of `EXTERNAL_ACCESS` permissions and only after justification.

Rare use of `UNSAFE` permissions and only after very serious justifications and testing.

Review Question(s)

Question: Which types of database objects can you implement by using managed code?

Answer: You can implement user-defined functions (scalar and table-valued), stored procedures, triggers (DML and DDL), user-defined aggregates, and user-defined data types by using managed code.

Question: What purpose do attributes have in CLR managed code?

Answer: In CLR managed code, attributes can relate to deployment, performance, or correctness.

Lab Review Questions and Answers

Lab: Implementing Managed Code in SQL Server

Question and Answers

Lab Review

Question: Suggest some possible uses for user-defined data types.

Answer: Answers will vary, but interesting options would be **jpeg** data types, **chemical** data types, **audio** data types, and so on.

Module 13

Optional – Storing and Querying XML Data in SQL Server

Contents:

Lesson 1: Introduction to XML and XML Schemas	2
Lesson 2: Storing XML Data and XML Schemas in SQL Server	4
Lesson 3: Implementing XML Indexes	6
Lesson 4: Using the Transact-SQL FOR XML Statement	8
Lesson 5: Getting Started with XQuery	10
Lesson 6: Shredding XML	12
Module Review and Takeaways	15
Lab Review Questions and Answers	16

Lesson 1

Introduction to XML and XML Schemas

Contents:

Question and Answers	3
Demonstration: Using XML and XML Schemas	3

Question and Answers

Question: Is the following XML fragment well formed?

```
<Employee Name="Germain Arenas" Department="HR" EmployeeID="254613">
<Employee Name="Joshua Murphy" Department="Sales" EmployeeID="254689">
<Employee Name="Daisuke Nakayama" Department="Engineering" EmployeeID="254636">
```

Answer: This is not a well-formed fragment.

Demonstration: Using XML and XML Schemas

Demonstration Steps

Structure XML and structure XML schemas

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. Run D:\Demofiles\Mod13\Setup.cmd as an administrator to revert any changes
3. On the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod13\Demo13.ssmssln**, and then click **Open**.
6. Expand the **Queries** folder.
7. Open the **11 - Demonstration 1A.sql** script file.
8. Follow the instructions contained within the comments of the script file.

Lesson 2

Storing XML Data and XML Schemas in SQL Server

Contents:

Question and Answers	5
Demonstration: Working with Typed vs. Untyped XML	5

Question and Answers

Question: A partner organization is providing lists of marketing leads to you in XML format but they are not providing their XML schema. You want to import the XML documents into SQL Server. How can you ensure the XML is well-formed?

Answer: Use the xml data format for the column you import the XML documents into.

Demonstration: Working with Typed vs. Untyped XML

Demonstration Steps

Work with typed and untyped XML

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **AdventureWorks\Student** with the password **Pa\$\$w0rd**.
2. Ensure that you have run all previous demonstrations
3. On the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod13\Demo13.ssmssln**, and then click **Open**.
6. Open the **21 - Demonstration 2A.sql** script file.
7. Follow the instructions contained within the comments of the script file.

Lesson 3

Implementing XML Indexes

Contents:

Question and Answers	7
Demonstration: Implementing XML Indexes	7

Question and Answers

Question: You have an xml format column in a SQL Server table. A popular application is coded to use PATH expressions to retrieve one or more values from elements in the XML. What kind of index should you create to optimize these queries?

- ☐ A PROPERTY secondary XML index
- ☐ A VALUE secondary XML index
- ☐ A clustered index
- ☐ A non-clustered index
- ☐ A PATH secondary XML index

Answer:

- ☒ A PROPERTY secondary XML index
- ☐ A VALUE secondary XML index
- ☐ A clustered index
- ☐ A non-clustered index
- ☐ A PATH secondary XML index

Demonstration: Implementing XML Indexes

Demonstration Steps

Implement XML indexes

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **AdventureWorks\Student** with the password **Pa\$\$w0rd**.
2. Ensure that you have run all previous demonstrations.
3. On the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod13\Demo13.ssmssln**, and then click **Open**.
6. Open the **31 - Demonstration 3A.sql** script file.
7. Follow the instructions contained within the comments of the script file.

Lesson 4

Using the Transact-SQL FOR XML Statement

Contents:

Question and Answers	9
Demonstration: Using FOR XML Queries	9

Question and Answers

Question: You are generating XML to export from the Human Resources database. You want to ensure that columns in the Employees table are rendered as attributes of the <Employee> element. What property must you add to your FOR XML clause?

Answer: No additional property is required because this is the default behavior.

Demonstration: Using FOR XML Queries

Demonstration Steps

Use FOR XML queries

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **AdventureWorks\Student** with the password **Pa\$\$w0rd**.
2. Ensure that you have run all previous demonstrations.
3. On the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod13\Demo13.ssmssln**, and then click **Open**.
6. Open the **41 - Demonstration 4A.sql** script file.
7. Follow the instructions contained within the comments of the script file.

Lesson 5

Getting Started with XQuery

Contents:

Demonstration: Using XQuery Methods in DDL Triggers	11
---	----

Demonstration: Using XQuery Methods in DDL Triggers

Demonstration Steps

Use XQuery in DDL triggers

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **AdventureWorks\Student** with the password **Pa\$\$w0rd**.
2. Ensure that you have run all previous demonstrations
3. On the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod13\Demo13.ssmssln**, and then click **Open**.
6. Open the **51 - Demonstration 5A.sql** script file.
7. Follow the instructions contained within the comments of the script file.

Lesson 6

Shredding XML

Contents:

Question and Answers	13
Demonstration: Shredding XML	14

Question and Answers

Put the following XML shredding steps in order by numbering each to indicate the correct order.

	Steps
	Calling <code>sp_xml_removedocument</code> removes the node tree from memory.
	The relational data is combined with other relational data as part of standard Transact-SQL queries.
	The <code>sp_xml_preparedocument</code> stored procedure creates an in-memory node tree, based on the input XML.
	An XML document is received from a client application.
	The OPENXML table-valued function queries the in-memory node tree and extracts relational data.

Answer:

	Steps
5	Calling <code>sp_xml_removedocument</code> removes the node tree from memory.
4	The relational data is combined with other relational data as part of standard Transact-SQL queries.
2	The <code>sp_xml_preparedocument</code> stored procedure creates an in-memory node tree, based on the input XML.
1	An XML document is received from a client application.
3	The OPENXML table-valued function queries the in-memory node tree and extracts relational data.

Demonstration: Shredding XML

Demonstration Steps

Shred XML data by using the nodes() method

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **AdventureWorks\Student** with the password **Pa\$\$w0rd**.
2. Ensure that you have run all previous demonstrations.
3. On the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod13\Demo13.ssmssln**, and then click **Open**.
6. Open the **61 - Demonstration 6A.sql** script file.
7. Follow the instructions contained within the comments of the script file.

Module Review and Takeaways

Best Practice

Use appropriate data types for your database columns. Do not store all of your data in XML columns.

Use XML schemas only when they are required. Validating data against schemas incurs substantial processing overhead.

Ensure that you have at least basic XML proficiency when you are working with SQL Server, even if you will be working primarily in database administration.

Index XML data that is stored in database columns. Use the appropriate type of index for the types of queries that you expect.

Review Question(s)

Question: What is XML?

Answer: XML is a plain-text, unicode-based metalanguage. It is a language for defining markup languages. It is not tied to any particular programming language, operating system, or software vendor. XML provides access to a wide range of technologies for manipulating, structuring, transforming, and querying data.

Question: What is the difference between an element and an attribute?

Answer: An attribute is a property of an element.

Question: What is AUTO mode?

Answer: AUTO mode gives you more control over the returned XML than RAW mode. It generates nesting in the resulting XML where necessary, based on the **SELECT** statement that is supplied.

Question: What is PATH mode?

Answer: PATH mode is a simpler way to introduce additional nesting for representing complex properties.

Lab Review Questions and Answers

Lab: Storing and Querying XML Data in SQL Server

Question and Answers

Lab Review

Question: What is the purpose of an XML schema?

Answer: The purpose of an XML schema is to define the permitted structure of an XML document.

Question: When would you use untyped XML?

Answer: You would use untyped XML if you did not have a schema for your XML data or you did not want the server to validate the data against the schema.

Question: You could pass XML data to a stored procedure by using either the **xml** data type or the **nvarchar** data type. What advantage does the **xml** data type provide over the **nvarchar** data type for this purpose?

Answer: You already know that the data is well-formed XML. Using the **xml** data type simplifies the level of error-checking that you would otherwise need to perform.

Question: Which XML query mode did you use for implementing the WebStock.GetAvailableModelsAsXML stored procedure?

Answer: The XML query mode that was used for implementing the WebStock.GetAvailableModelsAsXML stored procedure was RAW mode.

Module 14

Working with Spatial Data in SQL Server

Contents:

Lesson 1: Introduction to Spatial Data	2
Lesson 2: Working with Spatial Data Types in SQL Server	4
Lesson 3: Using Spatial Data in Applications	6
Module Review and Takeaways	8
Lab Review Questions and Answers	9

Lesson 1

Introduction to Spatial Data

Contents:

Question and Answers	3
Demonstration: Viewing Available Spatial Reference Systems	3

Question and Answers

Types of Spatial Data

Question: Which existing SQL Server data type could you use to store (but not directly process) raster data?

Answer: You could use the **varbinary(max)** data type to store (but not directly process) raster data.

Planar vs. Geodetic Data Types

Question: What is the difference between an ellipsoid and a sphere?

Answer: A sphere is one type of ellipsoid. Other ellipsoids are more like squashed or elongated spheres.

Viewing Available Spatial Reference Systems

Question: Do you currently use GPS data in any existing applications within your organization?

Answer: Answers will vary depending upon the experience of the students, but it is not uncommon to find students who have done this by using existing data types.

Demonstration: Viewing Available Spatial Reference Systems

Demonstration Steps

View the available special reference systems

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **AdventureWorks\Student** with the password **Pa\$\$w0rd**.
2. Run D:\Demofiles\Mod14\Demo14\Setup.cmd as an administrator to revert any changes.
3. On the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod14\Demo14\Demo14.ssmssln**, and then click **Open**.
6. Expand the **Queries** folder.
7. Open the **11 - Demonstration 1A.sql** script file.
8. Follow the instructions contained within the comments of the script file.

Lesson 2

Working with Spatial Data Types in SQL Server

Contents:

Question and Answers	5
Demonstration: Working with Spatial Data Types	5

Question and Answers

Question: You have defined and stored a polygon in a SQL Server table that describes a sales representative's area of responsibility. The polygon approximates the limits of the city of Seattle, WA. When you use the STArea method, the value returned is over 500 million km². What has caused this problem?

Answer: You have defined the polygon by listing points in clockwise order.

Spatial Data in SQL Server

Question: You may have used a web service to calculate the coordinates of an address. What is this process commonly called?

Answer: The process of calculating the coordinates of an address is called geocoding. Bing provides a geocoding service.

Spatial Data Formats

Question: Why is there a need to represent spatial data types as strings within SQL Server?

Answer: There is a need to represent spatial data types as strings within SQL Server because there are no literal geometry or geography formats.

Demonstration: Working with Spatial Data Types

Demonstration Steps

Work with spatial data types in SQL Server

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. Run D:\Demofiles\Mod14\Demo14\Setup.cmd as an administrator to revert any changes.
3. On the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod14\Demo14\Demo14.ssmssln**, and then click **Open**.
6. Expand the **Queries** folder.
7. Open the **21 - Demonstration 2A.sql** script file.
8. Follow the instructions contained within the comments of the script file.

Lesson 3

Using Spatial Data in Applications

Contents:

Question and Answers	7
Demonstration: Using Spatial Data in Applications	7

Question and Answers

Performance Issues in Spatial Queries

Question: What would be the challenge in locating the intersecting streets in your suburb or region?

Answer: The list of streets might be very large.

Question: Which streets would you need to check?

Answer: You would need to check all of the streets unless there was some way of narrowing them down.

Question: How could you minimize the problem of needing to check all of the streets?

Answer: You could minimize the problem of needing to check all of the streets by eliminating any street that does not come near the target area, and then only checking those that are in the vicinity.

Demonstration: Using Spatial Data in Applications

Demonstration Steps

Use spatial data in SQL Server to solve some business questions

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. Run D:\Demofiles\Mod14\Demo14\Setup.cmd as an administrator to revert any changes.
3. On the taskbar, click **SQL Server 2014 Management Studio**.
4. In the **Connect to Server** window, in **Server name**, type **MIA-SQL** and then click **Connect**.
5. On the **File** menu, click **Open**, click **Project/Solution**, navigate to **D:\Demofiles\Mod14\Demo14\Demo14.ssmssln**, and then click **Open**.
6. Expand the **Queries** folder.
7. Open the **31 - Demonstration 3A.sql** script file.
8. Follow the instructions contained within the comments of the script file.
9. Close SQL Server Management Studio without saving any changes.

Module Review and Takeaways

Best Practice

Set the SRID for **geometry** objects to 0 to ensure that it is always possible to perform operations on multiple **geometry** objects.

Use a CHECK constraint to ensure that the SRID values for a column are consistent across all rows.

Before you create spatial indexes, make sure that the queries that need to be executed against the data use predicate forms that are supported by the types of index that you are creating.

Review Question(s)

Question: What is the main difference between the **geometry** and **geography** data types?

Answer: The **geometry** data type is for flat-earth calculations; the **geography** data type is for round-earth calculations.

Question: When you are defining a polygon, why does it matter how you specify the order of the points?

Answer: When you are defining a polygon, you need to consider how you specify the order of the points because a counterclockwise order returns the inner area, whereas a clockwise order returns the outside area.

Lab Review Questions and Answers

Lab: Working with Spatial Data in SQL Server

Question and Answers

Lab Review

Question: Where would you imagine you might use spatial data in your own business applications?

Answer: Spatial data will be a completely new concept for many students. Answers will vary depending upon their backgrounds and experience levels.

Module 15

Incorporating Data Files into Databases

Contents:

Lesson 1: Considerations for Working with Data Files in SQL Server 2014	2
Lesson 2: Implementing FILESTREAM and FileTables	4
Lesson 3: Searching Data Files	7
Module Review and Takeaways	10
Lab Review Questions and Answers	11

Lesson 1

Considerations for Working with Data Files in SQL Server 2014

Contents:

Question and Answers	3
Resources	3

Question and Answers

Question: You want to configure data file storage for the Human Resources database. You want to be able to backup all data, including data files in the database backups. Data files are very often larger than 1 MB. Developers have written code in the database application that accesses files through T-SQL queries but you also want users to be able to edit Word files with their Microsoft Office desktop applications. How should you configure data file storage.

- () Store data files on a separate hard disk from the database. Store a path to each file in each relevant row in the database.
- () Store data files as BLOBs within the database. Use varchar(max) columns wherever data files are used.
- () Enable FILESTREAM on the database so that data files can be stored on a separate hard disk.
- () Enable FileTable on the database so that data files can be stored on a separate hard disk.

Answer:

- () Store data files on a separate hard disk from the database. Store a path to each file in each relevant row in the database.
- () Store data files as BLOBs within the database. Use varchar(max) columns wherever data files are used.
- () Enable FILESTREAM on the database so that data files can be stored on a separate hard disk.
- (v) Enable FileTable on the database so that data files can be stored on a separate hard disk.

Resources

Options for Storing Data Files in SQL Server 2014



Reference Links: For more information about RBS, see the *Remote Blob Store (RBS) (SQL Server)* topic in SQL Server Books Online.

Lesson 2

Implementing FILESTREAM and FileTables

Contents:

Question and Answers	5
Resources	5
Demonstration: Configuring FILESTREAM and FileTable	6

Question and Answers

Question: You are migrating data files that were stored on a separate hard drive into a FileTable. Which function can you use to replace the old UNC path that was stored for each file with the new FileTable UNC path?

- ☐ DIRECTORY_NAME
- ☐ FileTableRootPath
- ☐ NON_TRANSACTED_ACCESS
- ☐ GetFileNameSpacePath
- ☐ GetPathLocator

Answer:

- ☒ DIRECTORY_NAME
- ☐ FileTableRootPath
- ☐ NON_TRANSACTED_ACCESS
- ☐ GetFileNameSpacePath
- ☐ GetPathLocator

Resources

Considerations for Implementing FILESTREAM



Reference Links: For more information about the best practices for using FILESTREAM, see the *FILESTREAM Best Practices* topic in SQL Server Books Online.

Considerations for Implementing FileTables



Reference Links: For more information about compatibility with other SQL Server features, see the *FileTable Compatibility with Other SQL Server Features* topic in SQL Server Books Online. For more information about using FileTable with AlwaysOn Availability Groups, see the *FILESTREAM and FileTable with AlwaysOn Availability Groups* topic in SQL Server Books Online.

Accessing FILESTREAM Data and FileTable Data



Reference Links: For more information about using the **FileTableRootPath** and **GetFileNameSpacePath** functions, see the *Work with Directories and Paths in FileTables* article in the MSDN library.



Reference Links: For more information about using **GetPathLocator** when you are migrating files from the file system to a FileTable, see the *Load Files into FileTables* article in the MSDN library.

Demonstration: Configuring FILESTREAM and FileTable

Demonstration Steps

View FILESTREAM configuration

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **AdventureWorks\Student** with the password **Pa\$\$w0rd**.
2. Run D:\Demofiles\Mod15\Setup.cmd as an administrator to revert any changes.
3. On the taskbar, click **Start** to open the Start page.
4. Click the down arrow to view all apps.
5. Scroll across to the **Microsoft SQL Server 2014** group.
6. Click **SQL Server 2014 Configuration Manager**.
7. In the **User Account Control** dialog box, click **Yes**.
8. In SQL Server Configuration Manager, click **SQL Server Services**, right-click **SQL Server (MSSQLSERVER)**, and then click **Properties**.
9. In the **SQL Server (MSSQLSERVER) Properties** dialog box, click the **FILESTREAM** tab.
10. Note that FILESTREAM is enabled for Transact-SQL access, File I/O access, and remote access, and that the FILESTREAM share is named MSSQLSERVER, and then click **Cancel**.

Configure FILESTREAM access and create a FileTable

1. On the taskbar, click **SQL Server 2014 Management Studio**.
2. In the **Connect to Server** dialog box, in the **Server Name** field, type **MIA-SQL** and then click **Connect**.
3. In SQL Server Management Studio, click **File**, point to **Open**, and then click **File**.
4. In the **Open File** dialog box, browse to **D:\Demofiles\Mod15**, and then double-click **FilesDemo.sql**.
5. In the query window, under the **Enable filestream** comment, highlight the Transact-SQL statement, and then click **Execute**.
6. Under the **Create filestream database** comment, highlight the Transact-SQL statement, and then click **Execute**.
7. Under the **Create filetable** comment, highlight the Transact-SQL statement, and then click **Execute**.
8. In File Explorer, in the D:\Demofiles\Mod15 folder, click **Document1**, press Shift, click **Document3**, right-click **Document3**, and then click **Copy**.
9. On the taskbar, click **Start**, type **Run** and then press Enter.
10. In the **Run** dialog box, type **\\localhost\MSSQLSERVER\FilestreamData\Documents** and then click **OK**.
11. Right-click in the empty folder, and then click **Paste**.
12. In SQL Server Management Studio, in the query window, under the **Query FileTable** comment, highlight the Transact-SQL statement, and then click **Execute**.
13. Leave SQL Server Management Studio open for the next demonstration.

Lesson 3

Searching Data Files

Contents:

Question and Answers	8
Resources	8
Demonstration: Creating a Full-Text Index	8



Question and Answers

Question: You are searching a full-text index. You want to return all entries that contain the word "bicycle". You want entries that contain synonyms such as "bike" and "cycle". What kind of query should you perform?


Answer: Execute a FREETEXT query.

Resources

Considerations for Implementing Full-Text Indexes

-  **Reference Links:** For more information about populating full-text indexes, see the *Populate Full-Text Indexes* topic in SQL Server Books Online.
-  **Reference Links:** For more information about improving performance for full-text indexes, see the *Improve the Performance of Full-Text Queries* topic in SQL Server Books Online.

Semantic Search

-  **Reference Links:** For more information about installing and configuring the Semantic Language Statistics database, see the *Install and Configure Semantic Search* topic in SQL Server Books Online.

Demonstration: Creating a Full-Text Index

Demonstration Steps

Create and query a full-text index

1. Ensure that the 20464C-MIA-DC and 20464C-MIA-SQL virtual machines are running and then log on to 20464C-MIA-SQL as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. Ensure that you have run the previous demonstrations in this module.
3. On the taskbar, click **SQL Server 2014 Management Studio**.
4. In SQL Server Management Studio, in the **FilesDemo.sql** query window, under the **Create full-text catalog** comment, highlight the Transact-SQL statement, and then click **Execute**.
5. Under the **Get index ID for FileTable PK** comment, highlight the Transact-SQL statement, and then click **Execute**.
6. In the **Results** pane, in the **name** column, right-click the value in row 1, which begins with **PK_FileStor_**, and then click **Copy**.
7. In the **FilesDemo.sql** query window, under the **Create full-text index** comment, after **KEY INDEX**, highlight **PK_FileStor_REPLACE_WITH_INDEX_ID**, right-click **PK_FileStor_REPLACE_WITH_INDEX_ID**, and then click **Paste**.
8. Under the **Create full-text index** comment, highlight the Transact-SQL statement, and then click **Execute**.
9. Under the **Find documents containing "impe"diet" nea" "viva"us" (wi"hin 15 search terms)** comment, highlight the Transact-SQL statement, click **Execute**, and then review the results.
10. Close the **FilesDemo.sql** query window, and do not save any changes.

11. Close SQL Server Management Studio.

Module Review and Takeaways

Question: How have you enabled the storage of data files in your places of work? How could you use the features of SQL Server 2014 to improve the storage of data files?

Answer: Answers will vary, depending on students' opinions and experiences.

Lab Review Questions and Answers

Lab: Implementing a Solution for Storing Data Files

Question and Answers

Lab Review

Question: If the lab scenario were modified as described in the list below, how might this influence your choice of storage solution for the data files?

Administrators want to be able to perform point-in-time restores for all database data, including data files.

Most of the files have an average size of 0.5 MB.

Answer: The solution that students implemented has the **NON_TRANSACTED_ACCESS** option set to FULL. However, this setting prevents you from performing point-in-time restores, so you would need to consider which factor is the most important.

If the average file size was smaller, that is, 0.5 MB, you could consider using **varbinary(max)** to store the files, although it would not fulfill all of the other criteria.

