

MOC 40074 A Microsoft SQL Server 2014 for Oracle DBAs

Course Summary

Description

This four-day instructor-led course provides students with the knowledge and skills to capitalize on their skills and experience as an Oracle DBA to manage a Microsoft SQL Server system. This course provides training for Oracle DBA to compare and contrast Oracle database management to SQL Server database management.

Objectives

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

- Extend their existing competencies as Oracle DBAs to SQL Server.
- Manage SQL Server using the same perspective of an Oracle DBA.
- Understand the underlying architecture of SQL Server.
- Manage the SQL Server system, databases, and users.
- Manage database files by backing up or migrating to other systems.
- Define and implement monitoring and tuning solutions to the SQL Server system.
- Express High Availability options to SQL Server.
- Explain the process and tool to migrate Oracle schemas to SQL Server databases.

Topics

- Database and instances
- Instance architecture
- Database architecture
- Data objects
- Data access
- Basic administration
- Managing schema objects
- Data protection and security
- Data transport
- Backup and recovery
- Monitoring and performance tuning
- Scalability and high availability
- Microsoft SQL Server Migration Assistant

Audience

This course is intended for experienced Oracle database administrators (DBAs) who work in an enterprise-level environment and require the skills to begin supporting and maintaining a SQL Server database.

Prerequisites

Before attending this course, students must have:

- Oracle DBA experience
- Familiarity with Microsoft Windows platforms
- Understanding of operating system fundamentals

Duration

Four days

MOC 40074 A Microsoft SQL Server 2014 for Oracle DBAs**Course Outline****I. Database and instances**

This module provides an understanding of the two major components of a database system. The database constitutes the files that store data, and the instance is the collection of server resources that provide a powerful, high performance interface to the data. It also illustrates how the two interact to provide data requested by the clients. Viewing the database and the instance as two separate interactive components of the RDBMS helps us to divide-and-conquer the vast set of topics covered in this workshop.

- A. Terminology concepts
- B. Client interaction with database and instance
- C. Understanding database limits

II. Instance architecture

This module discusses the memory and process architectures that are key to a database's performance. The module goes into the details of the hierarchy of memory areas of an instance and its configuration. This module also describes how the various functions of the RDBMS are accomplished by the different processes running in the background. Finally, in this module, we look at the changes Oracle has made in its internal architecture on Microsoft Windows platform between 9i and 12c to exploit the advantages offered by the operating system mechanisms and how they compare to SQL Server's implementation.

- A. Configuring a database server
- B. Memory architecture overview
- C. Understanding processes and threads in the database engine
- D. Background processes

Lab: Instance architecture

- Using the Resource Governor
- View multi-instance shared resources

III. Database architecture

This module goes in-depth into structure, components and contents of the files that constitute the database. To be able to manage hundreds of gigabytes, terabytes, or even petabytes of data, it is important to learn the techniques by which storage is viewed (physical and logical) and allocated. Databases use various hierarchies of storage structures such as blocks, extents, segments and table spaces to control storage allocation. The definition of schema and the objects that comprise the schema are introduced here. SQL Server uses similar techniques as Oracle; however the differentiation from Oracle is in the functionality. Also covered in this module is how SQL Server 2014 supports placing data and log files on Microsoft Azure storage.

- A. Schema and data storage
- B. Tablespaces and datafiles
- C. Logging and data dictionary

Lab: Database architecture

- Working with filegroups

IV. Data objects

This module examines the schema objects in both databases and introduces the new SQL Server 2014 In-Memory OLTP tables and clustered columnstore index objects. While all schema objects are mentioned, of particular interest are tables, the type of data they can hold, and their storage layout. A proper understanding of data types and storage architecture of tables and indexes is useful in many aspects of database design and administration, such as fragmentation, capacity planning, etc. A mapping of the native data types from Oracle to SQL Server provides the student with a very good reference on what data types are compatible and what are not.

- A. Database tables
- B. Schema objects
- C. Data and data types
- D. Non-native data types – Beyond relational

Lab: Data objects

- Create a partitioned table

V. Data access

This module focuses on how data is accessed and manipulated by the clients. Important concepts such as transaction, session, and so on are discussed here. This module describes the various commands available through SQL for manipulating data, metadata, transactions, sessions, and instances. An overview of procedural extensions to the SQL language available in Oracle (PL/SQL) and SQL Server (Transact-SQL) is given in this module. Cursors, which are data structures used to convey results of user transactions, are discussed as well to provide insight into what SQL Server supports and how they are used compared to Oracle.

- A. Comparing structured query language
- B. Control and procedural statements
- C. Developing robust queries

Lab: Data access

- Query designer
- Concatenation and SQL injection
- Stored procedures
- Cursors

MOC 40074 A Microsoft SQL Server 2014 for Oracle DBAs

Course Outline (cont'd)

VI. Basic administration

This module contains discussion on planning and installation of SQL Server. While basic duties such as creating, starting and shutting down a database are common to all databases, the options available for these functions are the key differences. True to the words "Knowledge is Power", familiarity with the data dictionary and the different ways a SQL Server DBA uses it compared to an Oracle DBA is an invaluable skill for any database administrator.

- A. Installing SQL Server
- B. Managing and configuring SQL Server
- C. Working with SQL Server databases

Lab: Basic administration

- Defining a user database
- Setting configuration parameters
- Filegroup maintenance

VII. Managing schema objects

This module provides the administrative aspect of schema objects described in Module 4. The discussion covers planning, creation and maintenance of many key schema objects. Choices in terms of table and index types, column types, and storage greatly influence the database growth, scalability, performance and maintainability.

- A. Managing tables, constraints, object identifiers, and naming
- B. Managing triggers
- C. Managing indexes and views

Lab: Managing schema objects

- Creating tables and associated objects
- Creating indexed views

VIII. Data protection and security

This module fulfills the twin tasks of protecting data against unauthorized access (database security) and also from the destructive interaction between authorized users working concurrently (concurrency control). Under security, the various features for securing and auditing the database are discussed. In addition, this module examines the various features available for providing security such as logins, roles, profiles, and privileges. This includes the new capabilities of SQL Server 2014 to support the separation of duties that can be used to provide the lowest level of privileges for server roles.

- A. Securing the database
- B. Managing users
- C. Understanding privileges
- D. Managing roles

Lab: Data protection and security

- Creating logins and users
- Grant permissions
- Revoke permissions
- Separation of duties

IX. Data transport

This module examines the non-transactional mechanisms for moving data into and out of a database. Included in this module is a discussion of how to copy databases to a Microsoft Azure VM using the SQL Server 2014 Deploy database to Azure VM wizard. Another key feature covered is SQL Server Integration Services (SSIS). SSIS provides users with the capabilities beyond complex ETL and high performance data movement from heterogeneous data sources. It also adds data mining capability to the process and more, all of which will be discussed in this module. Other approaches to bulk data movement will also be covered in this module and guidelines will be provided on which tools are appropriate for what scenarios.

- A. Getting data into and out of SQL Server
- B. Understanding SQL Server Integration Services
- C. Other transfer methods

Lab : Data transport

- Use SQL Server Integration Services to migrate data into a flat file
- Use SQL Server Integration Services to import data from a flat file

X. Backup and recovery

This module lists the types of errors encountered in a database and the various mechanisms that are available to safeguard against these errors. This module discusses various types of backups and recovery methods available. This module also covers Oracle's Recovery Manager (RMAN) and the equivalent functionality available in SQL Server Management Studio. In addition, this module covers the new capabilities with SQL Server 2014 for managed backups and manual backups to Microsoft Azure storage.

- A. Understanding database backups
- B. Data recovery
- C. Backup and recovery tools and solutions

Lab : Backup and recovery

- Create and execute a maintenance plan for backup
- Modify, backup, and restore a database

MOC 40074 A Microsoft SQL Server 2014 for Oracle DBAs

Course Outline (cont'd)

XI. Monitoring and performance tuning

This module discusses the different approaches to instance and application tuning. It then discusses the administration of various types of resources such as system resources (such as memory, processes, storage, and so on) as well as low-level database resources (such as lock, latches, queues, and so on). This session features new capabilities of SQL Server 2014 that improve performance including In-Memory OLTP tables, delayed durability for transactions, online operations and buffer pool extension. Given the significant difference in how resources are managed and utilized in SQL Server compared to Oracle, it is important for the Oracle DBA to get a firm understanding of what is under the SQL Server hood in order to appreciate and best leverage the technology. Finally, it shows the mechanisms by which the database can be monitored for availability, errors, and performance. Statistics that can be captured for both proactive and reactive administration of the databases are reviewed here.

- A. SQL Server performance tuning
- B. Managing memory and processes for SQL Server
- C. Managing database interactions
- D. Monitoring availability and errors
- E. Monitoring performance

Lab : Part 1: Monitoring and performance tuning

- Understanding threads

Lab : Part 2: Monitoring and performance tuning

- SQL Server Profiler
- Hints (T-SQL)
- Plan guides
- Setting alerts to automate monitoring

Lab : Part 3: Monitoring and performance tuning

- Capture a SQL trace for use with Distributed Replay
- Prepare trace data for use with Distributed Replay
- Replay trace data using Distributed Replay

XII. Scalability and high availability

This module provides a high-level overview of the scalability and high availability features available in each RDBMS. Oracle and SQL Server are both enterprise class RDBMS, therefore the topics of scalability and high availability are deemed important. However, due to the scope of the course the discussion here is at a conceptual level and does not go beyond the concepts. The features discussed include AlwaysOn, parallel query, replication, clustering, table partitioning, database mirroring, and database snapshots.

- A. Understanding availability
- B. Replicating databases
- C. Other methods to obtain availability

XIII. Microsoft SQL Server Migration Assistant

This module focuses on automating the process of migrating from Oracle to SQL Server. This module presents an overview of the SSMA and will cover schema conversion, data migration, business logic conversion, validation, integration, and performance analysis.

- A. Install and configure SQL Server Migration Assistant
- B. Migration projects
 - Migrating the data