

SQL Server 2014 Enhanced OLTP and Data Warehouse Performance



WorkshopPLUS

Learn how to develop, support, monitor and tune high performing applications on SQL Server 2014.

Target Audience:

To ensure the high-quality knowledge-transfer expected by attendees of this three day workshop, class size is limited to a maximum of 16 students who meet the following criteria:

- *Database Administrators*
- *Database Developers*
- *Database Support Engineers*
- *Database Architects*
- *ISV developers*
- *A good understanding of SQL Server Performance Tuning and Optimization.*
- *A recommended Pre-requisite is the SQL 2012 Performance Tuning and Optimization WorkshopPlus*

Overview

SQL Server 2014: Enhanced OLTP and Data Warehouse Performance is a three-day course that provides an in-depth study of the new high-performance capabilities built into SQL Server 2014 for both OLTP and data warehousing.

Through lecture, demo, discussions, and hands-on labs, students will discover the new performance-related features of SQL Server 2014. They will learn to use In-Memory OLTP and Columnstore indexes, examine memory management with Buffer Pool Extensions, explore the new cardinality engine, understand incremental statistics, and improve programming efficiency.

Key Features and Benefits

Major modules are organized by key mission-critical performance scenarios. Hands-on labs help attendees solidify concepts and gain experience using these new performance capabilities.

Technical Highlights

After completing this course, you will be able to:

- Design an effective In-Memory OLTP strategy
- Implement In-Memory Columnstore for Data warehousing
- Develop a migration strategy for both In-Memory OLTP and In-Memory Columnstore for DW
- Analyze use of Buffer Pool Extensions to SSDs
- Gain understanding of scenarios to use the New Cardinality Estimation, Incremental Statistics and Application Availability improvements such as Managed Lock Priority.

Syllabus

Hardware

Requirements:

Contact your TAM if the necessary hardware needs to be provided.

If you are attending an Open enrollment workshop, the hardware will be provided for you.

This workshop runs for a full three days. Students should anticipate consistent start and end times for each day. Early departure on any day is not recommended.

Module 1: In-Memory OLTP

In-Memory OLTP is a highly optimized engine that works with memory-optimized data and was designed to handle the most demanding OLTP workloads.

- **Section 1: Overview of In-Memory OLTP**

We discuss the high level architecture, the design decisions and key advantages in a OLTP environment.

- **Section 2: Implementing In-Memory OLTP**

We discuss the details of creating memory-optimized objects, the internals of native compilation, memory management, and storage and logging behind the scenes, as well as key changes in the concurrency model.

- **Section 3: In-Memory OLTP Performance**

We discuss the statistics and query indexing choices and coding constructs that are suited to extract maximum performance from the system. We also discuss hardware and configuration considerations.

- **Section 4: In-Memory OLTP Migration**

We discuss application patterns that can hugely benefit from this improvement and also look at a methodology to test and migrate into In-Memory OLTP with the associated tools.

Module 2: ColumnStore Indexes

ColumnStore indexes are designed for Data Warehouses and can improve the performance of Data Warehouse queries several folds.

- **Section 1: ColumnStore Overview**

An overview of the high-level architecture as well as the targeted intended workloads that would benefit.

- **Section 2: Implementing ColumnStore Indexes**

We cover creation of ColumnStore indexes, data loading practices and optimizations, concurrency implications as well as the index build process and memory implications of ColumnStore on the system.

Syllabus

Module 2: ColumnStore Indexes Continued

- **Section 3: Performance Considerations**

We look at query design considerations as well as Query plans and performance implications of Batch mode v/s row mode, Segment Elimination and statistics.

Module 3: Additional Performance Features

In this module we cover other key performance features introduced in SQL 2014.

- **Section 1 : Online Application Enhancements**

We discuss enhancements such as Managed Lock Priority and Online index rebuild changes that enable scenarios to reduce downtime due to concurrency.

- **Section 2 : Buffer Pool Extension to SSD**

We discuss enabling integration of SSDs to extended Server Memory for the buffer pool, which enables performance gains on read-heavy OLTP workloads.

- **Section 3 : Incremental Statistics**

Statistics are a key to the optimizer producing good query plans. Incremental Statistics solve some big problems in the Data Warehouse realm. We discuss their use and scenarios where they can be implemented.

- **Section 4 : New Cardinality Estimation**

The Cardinality estimation model has been changed in SQL 2014. We discuss the changes, how they affect query performance and plan choice, areas that will benefit performance, and how to enable and disable cardinality estimation in certain scenarios.