

T-SQL Querying



Professional

 SolidQ

Itzik Ben-Gan
Dejan Sarka
Adam Machanic
Kevin Farlee

T-SQL Querying

Itzik Ben-Gan

Dejan Sarka

Adam Machanic

Kevin Farlee

T-SQL Querying

Table of Contents

Contents

Introduction

Foreword

Chapter 1 Logical query processing

- Logical query-processing phases

 - Logical query-processing phases in brief

- Sample query based on customers/orders scenario

- Logical query-processing phase details

 - Step 1: The FROM phase

 - Step 2: The WHERE phase

 - Step 3: The GROUP BY phase

 - Step 4: The HAVING phase

 - Step 5: The SELECT phase

 - Step 6: The ORDER BY phase

 - Step 7: Apply the TOP or OFFSET-FETCH filter

- Further aspects of logical query processing

 - Table operators

 - Window functions

 - The UNION, EXCEPT, and INTERSECT operators

- Conclusion

Chapter 2 Query tuning

- Internals

 - Pages and extents

Table of Contents

Table organization

Tools to measure query performance

Access methods

Table scan/unordered clustered index scan

Unordered covering nonclustered index scan

Ordered clustered index scan

Ordered covering nonclustered index scan

The storage engines treatment of scans

Nonclustered index seek + range scan + lookups

Unordered nonclustered index scan + lookups

Clustered index seek + range scan

Covering nonclustered index seek + range scan

Cardinality estimates

Legacy estimator vs. 2014 cardinality estimator

Implications of underestimations and overestimations

Statistics

Estimates for multiple predicates

Ascending key problem

Unknowns

Indexing features

Descending indexes

Included non-key columns

Filtered indexes and statistics

Columnstore indexes

Inline index definition

Prioritizing queries for tuning with extended events

Index and query information and statistics

Temporary objects

Table of Contents

Set-based vs. iterative solutions

Query tuning with query revisions

Parallel query execution

- How intraquery parallelism works

- Parallelism and query optimization

- The parallel APPLY query pattern

Conclusion

Chapter 3 Multi-table queries

Subqueries

- Self-contained subqueries

- Correlated subqueries

- The EXISTS predicate

- Misbehaving subqueries

Table expressions

- Derived tables

- CTEs

- Views

- Inline table-valued functions

- Generating numbers

The APPLY operator

- The CROSS APPLY operator

- The OUTER APPLY operator

- Implicit APPLY

- Reuse of column aliases

Joins

- Cross join

- Inner join

- Outer join

Table of Contents

Self join

Equi and non-equi joins

Multi-join queries

Semi and anti semi joins

Join algorithms

Separating elements

The UNION, EXCEPT, and INTERSECT operators

The UNION ALL and UNION operators

The INTERSECT operator

The EXCEPT operator

Conclusion

Chapter 4 Grouping, pivoting, and windowing

Window functions

Aggregate window functions

Ranking window functions

Offset window functions

Statistical window functions

Gaps and islands

Pivoting

One-to-one pivot

Many-to-one pivot

Unpivoting

Unpivoting with CROSS JOIN and VALUES

Unpivoting with CROSS APPLY and VALUES

Using the UNPIVOT operator

Custom aggregations

Using a cursor

Using pivoting

Table of Contents

Specialized solutions

Grouping sets

GROUPING SETS subclause

CUBE and ROLLUP clauses

Grouping sets algebra

Materializing grouping sets

Sorting

Conclusion

Chapter 5 TOP and OFFSET-FETCH

The TOP and OFFSET-FETCH filters

The TOP filter

The OFFSET-FETCH filter

Optimization of filters demonstrated through paging

Optimization of TOP

Optimization of OFFSET-FETCH

Optimization of ROW_NUMBER

Using the TOP option with modifications

TOP with modifications

Modifying in chunks

Top N per group

Solution using ROW_NUMBER

Solution using TOP and APPLY

Solution using concatenation (a carry-along sort)

Median

Solution using PERCENTILE_CONT

Solution using ROW_NUMBER

Solution using OFFSET-FETCH and APPLY

Conclusion

Table of Contents

Chapter 6 Data modification

Inserting data

- SELECT INTO

- Bulk import

- Measuring the amount of logging

- BULK rowset provider

Sequences

- Characteristics and inflexibilities of the identity property

- The sequence object

- Performance considerations.

- Summarizing the comparison of identity with sequence

Deleting data

- TRUNCATE TABLE

- Deleting duplicates

Updating data

- Update using table expressions

- Update using variables

Merging data

- MERGE examples

- Preventing MERGE conflicts

- ON isn't a filter

- USING is similar to FROM

The OUTPUT clause

- Example with INSERT and identity

- Example for archiving deleted data

- Example with the MERGE statement

- Composable DML

Conclusion

Table of Contents

Chapter 7 Working with date and time

- Date and time data types

- Date and time functions

- Challenges working with date and time

 - Literals

 - Identifying weekdays

 - Handling date-only or time-only data with DATETIME and SMALLDATETIME

 - First, last, previous, and next date calculations

 - Search argument

 - Rounding issues

- Querying date and time data

 - Grouping by the week

 - Intervals

- Conclusion

Chapter 8 T-SQL for BI practitioners

- Data preparation

 - Sales analysis view

- Frequencies

 - Frequencies without window functions

 - Frequencies with window functions

- Descriptive statistics for continuous variables

 - Centers of a distribution

 - Spread of a distribution

 - Higher population moments

- Linear dependencies

 - Two continuous variables

 - Contingency tables and chi-squared

 - Analysis of variance

Table of Contents

Definite integration

Moving averages and entropy

Moving averages

Entropy

Conclusion

Chapter 9 Programmable objects

Dynamic SQL

Using the EXEC command

Using the sp_executesql procedure

Dynamic pivot

Dynamic search conditions.

Dynamic sorting

User-defined functions

Scalar UDFs

Multistatement TVFs

Stored procedures

Compilations, recompilations, and reuse of execution plans

Table type and table-valued parameters

EXECUTE WITH RESULT SETS

Triggers

Trigger types and uses

Efficient trigger programming

SQLCLR programming

SQLCLR architecture

CLR scalar functions and creating your first assembly

Streaming table-valued functions

SQLCLR stored procedures and triggers

SQLCLR user-defined types

Table of Contents

SQLCLR user-defined aggregates

Transaction and concurrency

Transactions described

Locks and blocking

Lock escalation

Delayed durability

Isolation levels

Deadlocks

Error handling

The TRY-CATCH construct

Errors in transactions

Retry logic

Conclusion

Chapter 10 In-Memory OLTP

In-Memory OLTP overview

Data is always in memory

Native compilation

Lock and latch-free architecture

SQL Server integration

Creating memory-optimized tables

Creating indexes in memory-optimized tables

Clustered vs. nonclustered indexes

Nonclustered indexes

Hash indexes

Execution environments

Query interop

Natively compiled procedures

Surface-area restrictions

Table of Contents

Table DDL

DML

Conclusion

Chapter 11 Graphs and recursive queries

Terminology

Graphs

Trees

Hierarchies

Scenarios

Employee organizational chart

Bill of materials (BOM)

Road system

Iteration/recursion

Subgraph/descendants

Ancestors/path

Subgraph/descendants with path enumeration

Sorting

Cycles

Materialized path

Maintaining data

Querying

Materialized path with the HIERARCHYID data type

Maintaining data

Querying

Further aspects of working with HIERARCHYID

Nested sets

Assigning left and right values

Querying

Table of Contents

Transitive closure

Directed acyclic graph

Conclusion

Index