

# SQL Server Query Tuning



**Klaus Aschenbrenner**


Independent SQL Server Consultant

SQLpassion.at

Twitter: @Aschenbrenner

# About me

- Independent SQL Server Consultant
- International Speaker, Author
- „Pro SQL Server 2008 Service Broker“
- SQLpassion.at
- Twitter: @Aschenbrenner



# SQL Server Performance & Troubleshooting Workshop

- 19. – 21. Juni in Bern
  - Gurten
- Agenda
  - SQL Server Indexing
  - SQL Server Locking & Blocking
  - SQL Server Performance Monitoring & Troubleshooting
- Weitere Informationen
  - <http://www.SQLpassion.at/events.html>

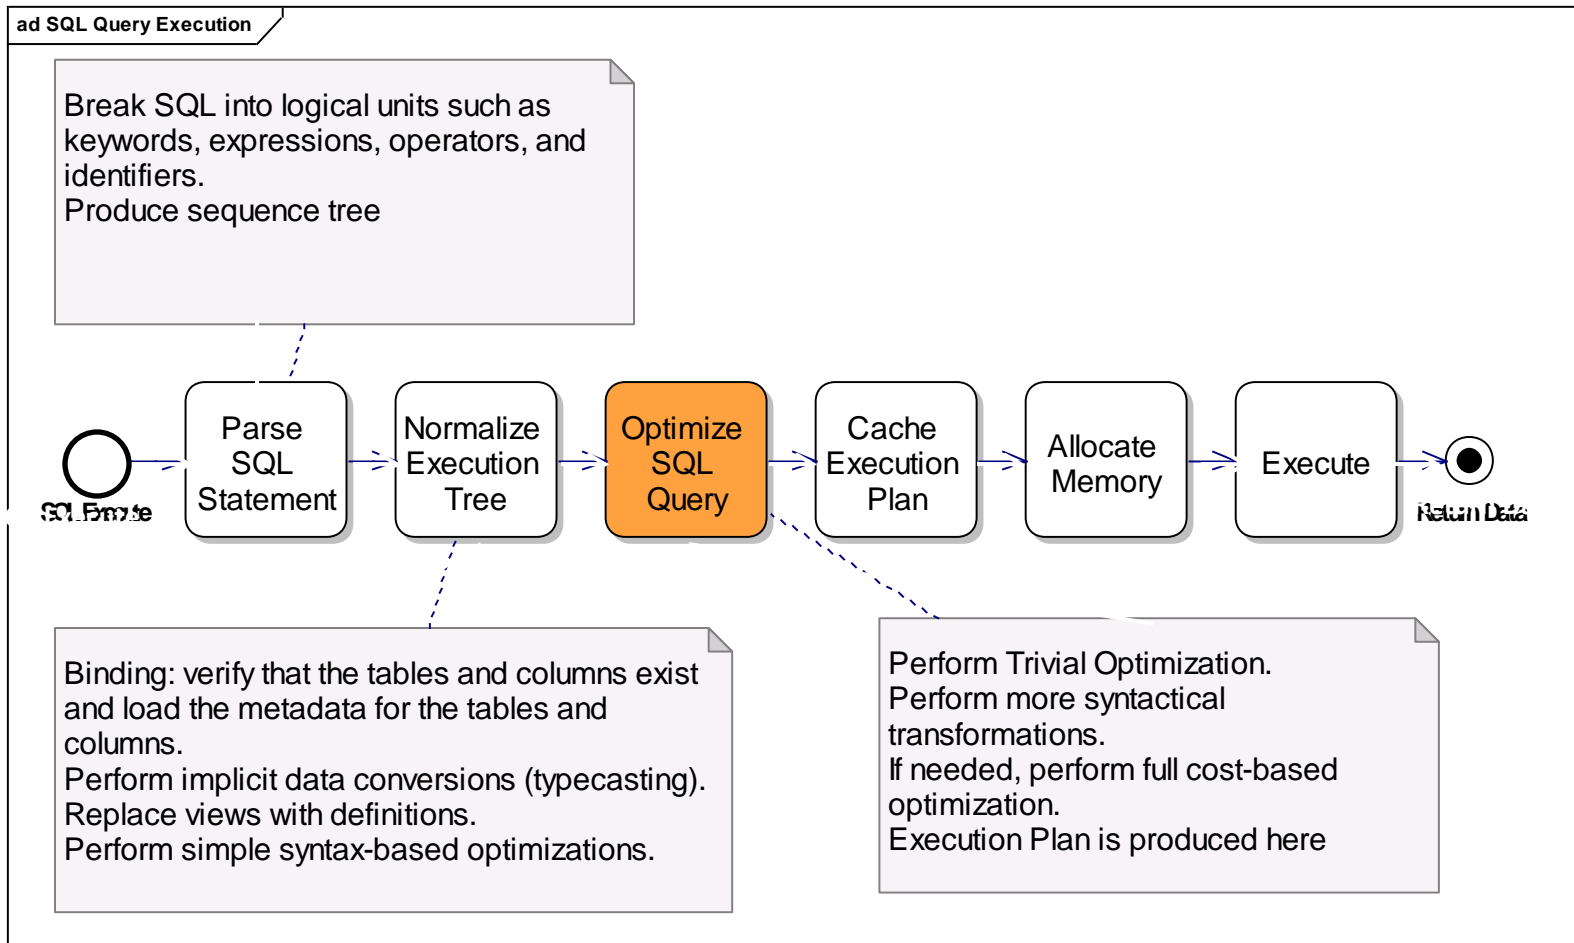
# Agenda

- Basics of Query Execution
- Execution Plan Overview
- Index Tuning

# Agenda

- **Basics of Query Execution**
- Execution Plan Overview
- Index Tuning

# Execution of a query



# Query Optimization

- Execution Plans and cost-based optimizations
- Optimization phases
- Indexes and distribution statistics
- Join Selection

# Execution Plan

- Strategy determined by the Optimizer to access/manipulate data
  - Can be influenced by the developer – query hints
- Key decisions are made
  - Which indexes to use?
  - How to perform JOIN operations?
  - How to order and group data?
  - In what order tables should be processed?
  - Can be cached plans reused?



# Agenda

- Basics of Query Execution
- **Execution Plan Overview**
- Plan Cache
- Plan Caching
- Parameter Sniffing
- Recompilations

# Why understanding Execution Plans?

- Insight into query execution/processing strategy
- Tune performance problems at the source
  - Hardware is not every time the problem
  - High CPU/IO consumption -> poorly tuned Execution Plans
- Understanding Execution Plans is a prerequisite to performance tuning!

# Execution Plan Types 1/2

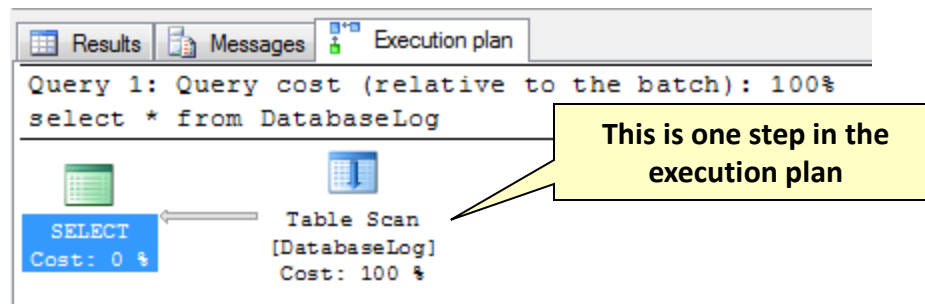
- Estimated Execution Plan
  - Created without ever running the query
  - Uses statistics for estimation
  - Good for long running query tuning
- Actual Execution Plan
  - Created when the actual query runs
  - Uses the real data
- They can be different
  - Statistics out of date
  - Estimated Execution Plan not valid any more

# Execution Plan Types 2/2

- Textual Execution Plan
  - Deprecated in further versions of SQL Server
- XML based Execution Plan
  - Very good for further analysis
  - Can be queried
- Graphic Execution Plan
  - Uses internally the XML based Execution Plan

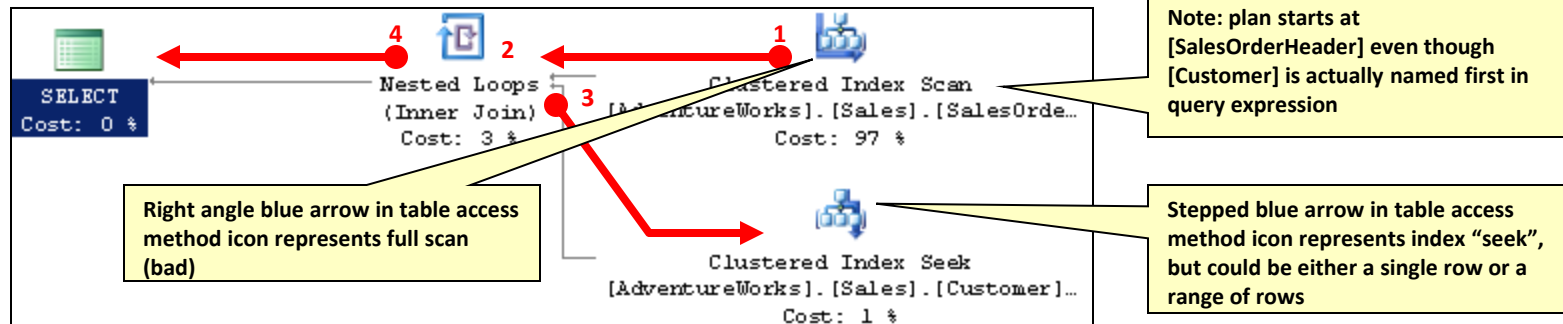
# Reading Execution Plans

- SHOWPLAN permission needed
- Query execution is serial
  - A series of sequential steps/operators
  - Executed one after one
- Execution Plan displays these steps/ operators



# Execution Plan Sample

```
select c.CustomerID
       , soh.SalesOrderID
       , soh.OrderDate
from Sales.Customer c
join Sales.SalesOrderHeader soh on c.CustomerID = soh.CustomerID
where soh.OrderDate > '20050101'
```



# Operator Properties

- Each operator has several properties with additional information

The screenshot displays a query plan in SQL Server Enterprise Manager. The plan consists of three operators: a SELECT operator (Cost: 0%), a Hash Match (Inner Join) operator (Cost: 40%), and an Index Scan (NonClustered) operator (Cost: 5%). The Index Scan operator is highlighted in blue, and its properties are expanded in a yellow window on the right.

**Index Scan (NonClustered)**  
Scan a nonclustered index, entirely or only a range.

Property	Value
Physical Operation	Index Scan
Logical Operation	Index Scan
Actual Number of Rows	19820
Estimated I/O Cost	0,0283102
Estimated CPU Cost	0,021959
Number of Executions	1
Estimated Number of Executions	1
Estimated Operator Cost	0,0502692 (5%)
Estimated Subtree Cost	0,0502692
Estimated Number of Rows	19820
Estimated Row Size	11 B
Actual Rebinds	0
Actual Rewinds	0
Ordered	False
Node ID	1

**Object**  
[AdventureWorks2008].[Sales].[Customer].  
[IX\_Customer\_TerritoryID] [c]

**Output List**  
[AdventureWorks2008].[Sales].[Customer].CustomerID

# Common Properties

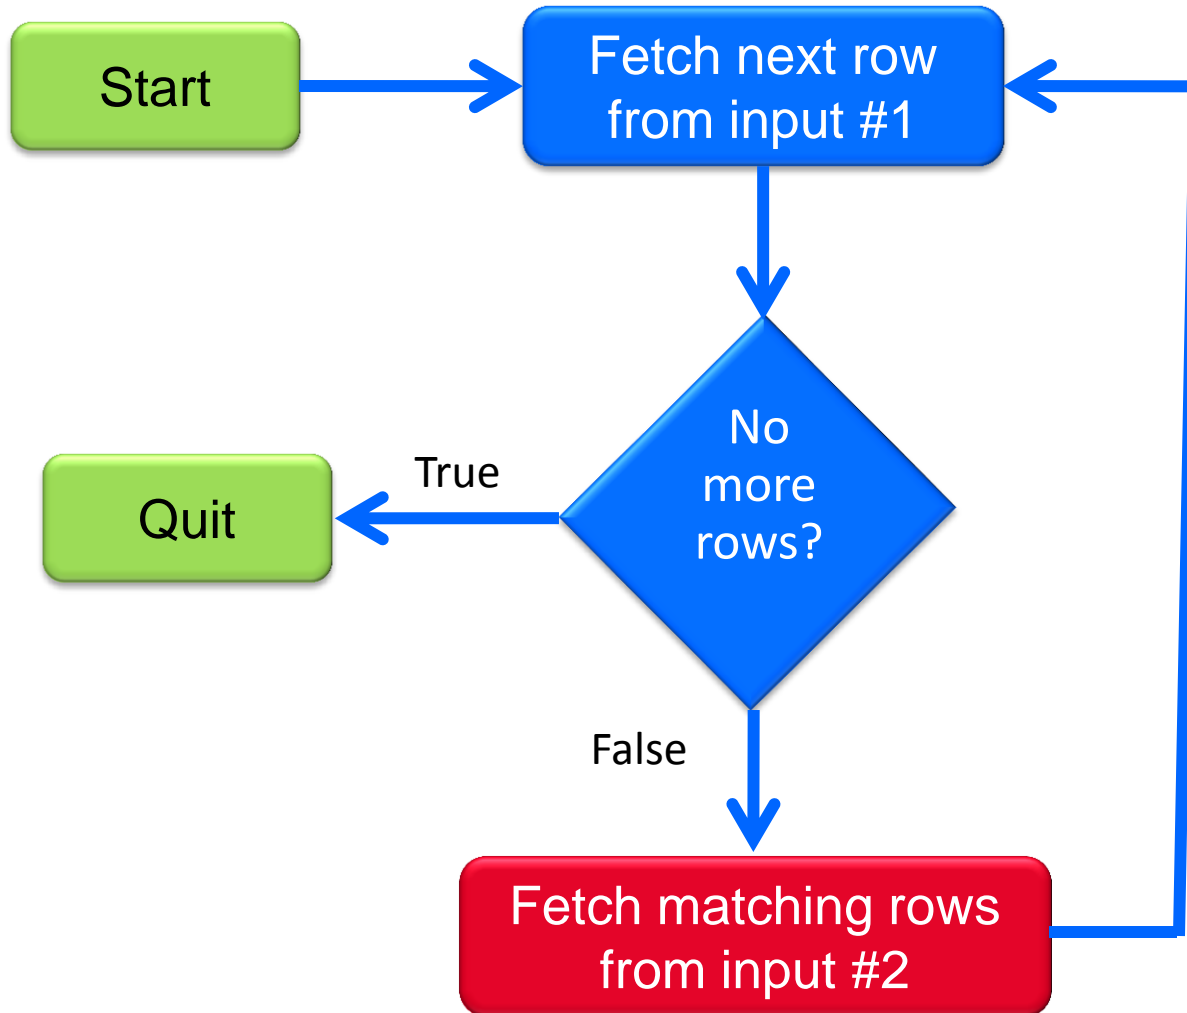
- Physical Operation
- Logical Operation
- Estimated I/O Cost
- Estimated CPU Cost
- Estimated Operator Cost
- Estimated Subtree Cost
- Estimated Number of Rows
- Estimated Row Size



# Common Operators

- Data Retrieval Operators
  - Table Scan (reads a whole table)
  - Index Scan (reads a whole index)
  - Index Seek (seeks into a index)
- Join Operators
  - Nested Loop (outer loop, inner loop)
  - Merge Join (needs sorted input)
  - Hash Match (uses a hash table internally)
- Aggregation Operators
  - Stream Aggregate
  - Hash Aggregate

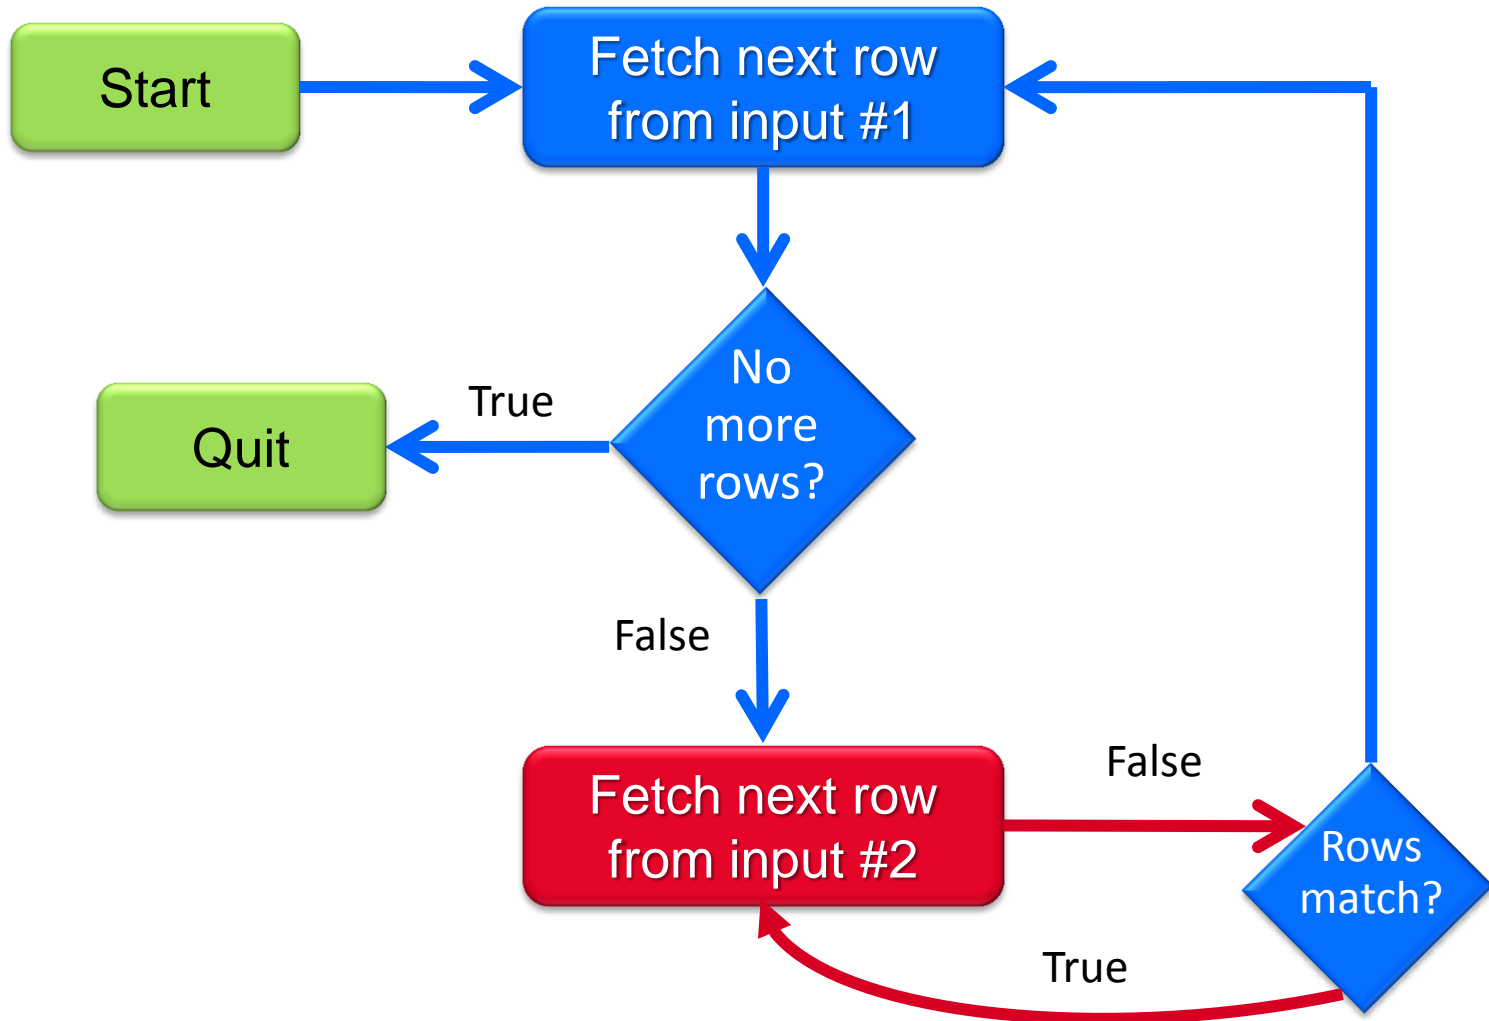
# Nested Loop Join



# Nested Loop

- „For each Row“ operator
- Takes output from one step and executes another operation „for each“ output row
- Outer Loop, Inner Loop
- Only join type that supports inequality predicates

# Merge Join



# Merge Join

- Needs at least one equijoin predicate
- Used when joined columns are indexed (sorted)
- Otherwise (expensive) sorting is needed
  - Plan may include explicit sort

Start

Fetch next row from input #1

Fetch next row from input #2

# Hash Join

No more rows?

No more rows?

Apply Hash Function

Apply Hash Function

Place row in hash bucket

Probe bucketed for matching rows

Quit



# Hash Join

- Needs at least one equijoin predicate
- Hashes values of join columns from one side (smaller table)
  - Based on Statistics
- Probes them with the join columns of the other side (larger table)
- Uses hash buckets
- Stop and Go for the Probe Phase
- Needs memory grants to build the hash table
  - If they exceed, Hash Join is spilled to TempDb
  - Performance decreases!

# Stream Aggreate

- Data must be sorted on the aggregation columns
- Processes groups one at a time
- Does not block or use memory
- Efficient if sort order is provided by index
  - Plan may include explicit sort



# Hash Aggregate

- Data need not be sorted
- Builds a hash table of all groups
- Stop and Go
- Same as Hash Join
  - Needs memory grants to build the hash table
  - If they exceed, Hash Aggregate is spilled to TempDb
  - Performance decreases!
- General better for larger input sets

# Capturing Execution Plans

- SQL Server Management Studio
  - Development
- SQL Profiler
  - Production
  - Event: „Performance/Showplan XML“
  - Dump Execution Plan to file
- `sys.dm_exec_cached_plans`
  - Production

# Demo

## Working with Execution Plans

# Agenda

- Basics of Query Execution
- Execution Plan Overview
- **Index Tuning**
  - Search Arguments
  - Covering Index
  - Tipping Point
  - Index Intersection
  - Index Joins
  - Filtered Indexes
  - Indexed Views

# General

- SET STATISTICS IO ON/OFF
  - Returns the IO page count
- SET STATISTICS TIME ON/OFF
  - Returns the elapsed time
- Index improves read performance
- Index degrades write performance
  - Index maintenance
  - But it can also improve performance, because the record to be updated/deleted must be found in the first step (SELECT)

# Search Arguments

- Can be part of
  - WHERE clause
  - JOIN clause
- Has major impact on Index usage
  - Breaks performance
- Wrong usage can lead to SCANS instead of SEEKS
- Recommendations
  - No functions/expressions on Search Arguments

# Demo

## Search Arguments

# Bookmark Lookup

- Occurs when query is not satisfied by a Covering Index
  - Base table must be touched
  - Index Page + Data Page access occurs
- 2 kinds of Bookmark Lookup
  - Clustered Key Lookup (Clustered Table)
  - RID Lookup (Heap Table)
- Deadlocks are possible
  - Writer: X(CI) => X(NCI)
  - Reader: S(NCI) => S(CI)



# Covering Index

- Non-Clustered Index that satisfies a complete query
  - Without touching the base table
  - No access to the data page containing the data
- The following referred columns must be present in the Non-Clustered index
  - SELECT
  - WHERE
- Otherwise
  - Traditional „Bookmark Lookup“

# INCLUDE Property


- Includes columns in the leaf-pages of the index
- Columns are not included in the navigation hierarchy
- Navigation hierarchy is very small
- Benefits
  - More elegant covering index
  - More than 900 bytes per index possible
  - More than 16 columns per index possible

# Demo

## Covering Index

# Tipping Point

- Defines whether to do a Bookmark Lookup or a Table/Clustered Index Scan
- Bookmark Lookups are only done when the Non-Clustered Index is selective enough
- Covering Index will not have a Tipping Point



# Tipping Point Sample

- 80.000 records per 4.000 bytes
  - 12,5% - 16,7% records
- 80.000 records per 400 bytes
  - 1,25% - 1,67% records
- 80.000 records per 40 bytes
  - 0,125% - 0,167% records
- Doesn't depend on the number of records
- Depends on the size of the records!!!

# Demo

## Tipping Point

# Index Intersection

- SQL Server can use multiple indexes to execute a query
  - Selecting small subsets of data
  - Finally performing an intersection between those subsets
- Why
  - No chance to modify existing indexes
  - Already too much columns/bytes in the Non-Clustered Index
- Create multiple narrow indexes, instead of wide indexes
  - Reuse through index intersection is possible!

# Disadvantages

- Not the fastest option for performance
  - More than 1 index must be accessed
  - Join(s) necessary
  - Hash Join needs (maybe) space in TempDb
- Not recommended for high priority queries



# Demo

## Index Intersection

# Filtered Index

- Specialized version of Non-Clustered Index
  - WHERE clause
  - Creates a highly selective set of keys
  - Reduces the size (page count) of the Non-Clustered Index and therefore the query performance
  - Decreased maintenance cost, because index is smaller
- E.g.
  - Column with large amount of NULL values

# Filtered Indexes - Prerequisites

- ANSI settings ON
  - ANSI\_NULLS
  - ANSI\_PADDING
  - ANSI\_WARNING
  - ARITHABORT
  - CONCAT\_NULL\_YIELDS\_NULL
  - QUOTED\_IDENTIFIER
- ANSI settings OFF
  - NUMERIC\_ROUNDABORT

# Demo

## Filtered Index

# Indexed Views

- Views
  - Stores only the SELECT statement (virtual table)
  - Contains no data
  - Data is retrieved by every call again and again
  - Poor performance if you have to do a lot of calculations (like aggregations)
- Indexed Views
  - Unique Clustered Index on a view
  - The data in the view is materialized into the index
  - Index is stored physically in the database file

# Indexed Views - Restrictions

- First index must be a unique Clustered Index
- Non-Clustered Index after Clustered Index
- View definition must be deterministic
- No reference to other views
- Float columns can't be the Clustered Key
- View must be schema-bound

# Indexed Views - Prerequisites

- ANSI settings ON
  - ARITHABORT
  - CONCAT\_NULL\_YIELDS\_NULL
  - QUOTED\_IDENTIFIER
  - ANSI\_NULLS
  - ANSI\_PADDING
  - ANSI\_WARNING
- ANSI settings OFF
  - NUMERIC\_ROUNDABORT

# Demo

## Indexed Views



# Indexing for OR

- OR returns individual sets
  - Ensures that rows that appears in multiple sets are only returned once
- IN is a simplified version of OR conditions
- Each column must be indexed
  - Condition must be selective enough (regarding Tipping Point)

# Demo

## Indexing for OR

# Indexing for Aggregates

- 2 Types of Aggregates
- Stream Aggregate
  - Needs sorted input on the aggregated columns
- Hash Aggregate
  - Done otherwise, when input is not sorted
  - Can be spilled out to TempDb
  - Can lead to TempDb overhead/problems
- Aggregation on multiple columns
  - All columns must be indexed through ONE index

# Demo

## Indexing for Aggregates

# Summary

- Basics of Query Execution
- Execution Plan Overview
- Index Tuning

# SQL Server Performance & Troubleshooting Workshop

- 19. – 21. Juni in Bern
  - Gurten
- Agenda
  - SQL Server Indexing
  - SQL Server Locking & Blocking
  - SQL Server Performance Monitoring & Troubleshooting
- Weitere Informationen
  - <http://www.SQLpassion.at/events.html>