PERFORMANCE TUNING WITH SQL SERVER 2017

MILOŠ RADIVOJEVIĆ,

PRINCIPAL DATABASE CONSULTANT, BWIN GVC, VIENNA, AUSTRIA

CI ENCONTRO DA COMUNIDADE SQLPORT, LISBON, 28TH MAY 2018

About me

Principal Database Consultant, bwin GVC, Vienna, Austria

Data Platform MVP

Co-Founder: SQL Pass Austria

Conference Speaker

Book Co-Author

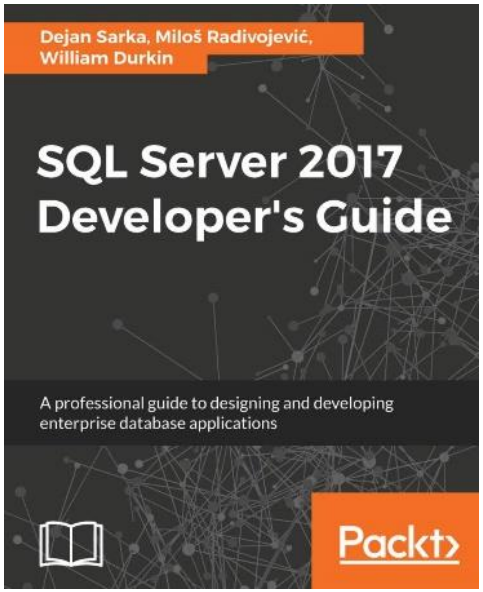

Contact:
MRadivojevic@gvcgroup.com

http://milossql.wordpress.com

Agenda

New SQL Server (Microsoft) release cycles

Adaptive Query Processing

- Interleaved executions
- Batch Mode Adaptive Join
- Memory Grant feedback

Query Store as Game Changer

- Troubleshooting with Query Store
- Automatic tuning

SQLPort

New SQL Server (Microsoft) release cycles

Upgrade challenge (risk, not fully atomated test routines…)
Learning
Too frequent
Quality?
Abandoned services/features

SQLPort

CI ENCONTRO DA COMUNIDADE SQLPORT, LISBON, 28TH MAY 2018

Adaptive Query Processing

Adaptive Query Processing

Query Optimizer

- Chooses physical operators and creates the execution plan
- Estimates memory that is needed for query execution (Memory Grant)
- Based on estimates done by the Cardinality Estimator

Query Execution Issues

- Slow response time
- Intensive resource consuming
- Reduced throughput and concurrency

Adaptive Query Processing

SQL Server 2016 (and prior)

- After the execution plan is created, it is used in consecutive query executions, without changes (with the same operators and memory grants)

SQL Server 2017 Adaptive Query Processing

- Breaking the pipeline between query optimization and execution
- Executing a part of the query during the execution plan creation
- Updating a part of the cached plan during consecutive query executions (Memory Grant)
- Batch mode Adapter Join Operator

Adaptive Query Processing

Interleaved Execution

Batch Mode Memory Grant Feedback

Batch Mode Adaptive Join

Interleaved Execution

Related to queries with multi table valued functions (MTVF)

- Break the optimization process
- Execute the part of the query with function call and get actual cardinality
- Continue with the optimization process

Epilogue

- More appropriate plan (correct cardinality instead of cardinality 100

Costs

- Increased CPU compile time
- Increased costs are acceptable, plan is usually better (sometimes significantly)

SQLPort

MTVF Execution

SQL Server 2016

```
Query 1: Query cost (relative to the batch): 100%
SELECT ol.OrderID, ol.UnitPrice, ol.StockItemID FROM Sales.Orderlines ol INNER JOIN dbo.SignificantOrders() f1 ON f1.Id = ol.OrderID WHERE PackageType
```

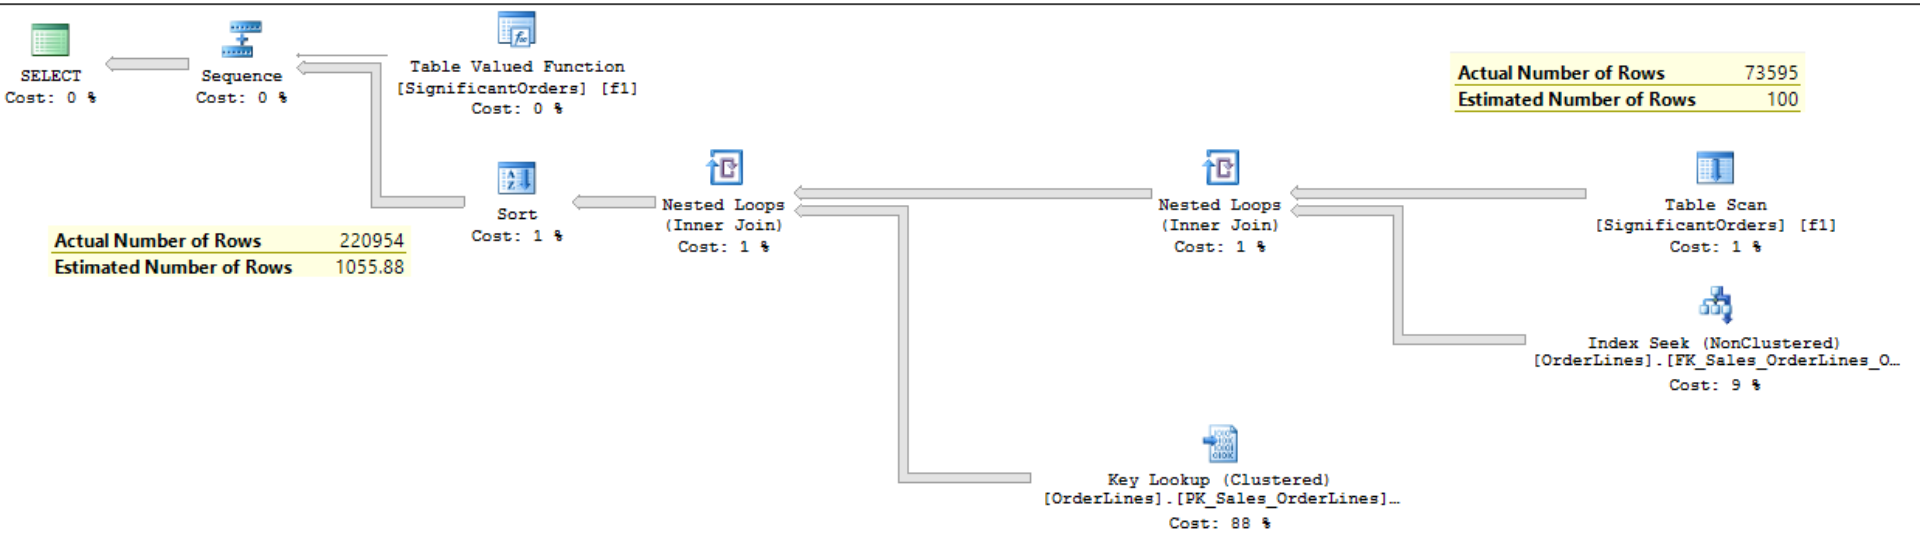


```
Table 'Worktable'. Scan count 0, logical reads 0, physical reads 0, r
Table 'OrderLines'. Scan count 73595, logical reads 865656, physical
Table '#AE4E6FE7'. Scan count 1, logical reads 119, physical reads 0,

CPU time = 937 ms,  elapsed time = 2445 ms.
```

SQLPort

Interleaved Execution

SQL Server 2017

```
Query 1: Query cost (relative to the batch): 100%
SELECT ol.OrderID, ol.UnitPrice, ol.StockItemID FROM Sales.Orderlines ol INNER JOIN dbo.SignificantOrders() f1 ON f1.Id = ol.OrderID WHERE PackageTypeID
```

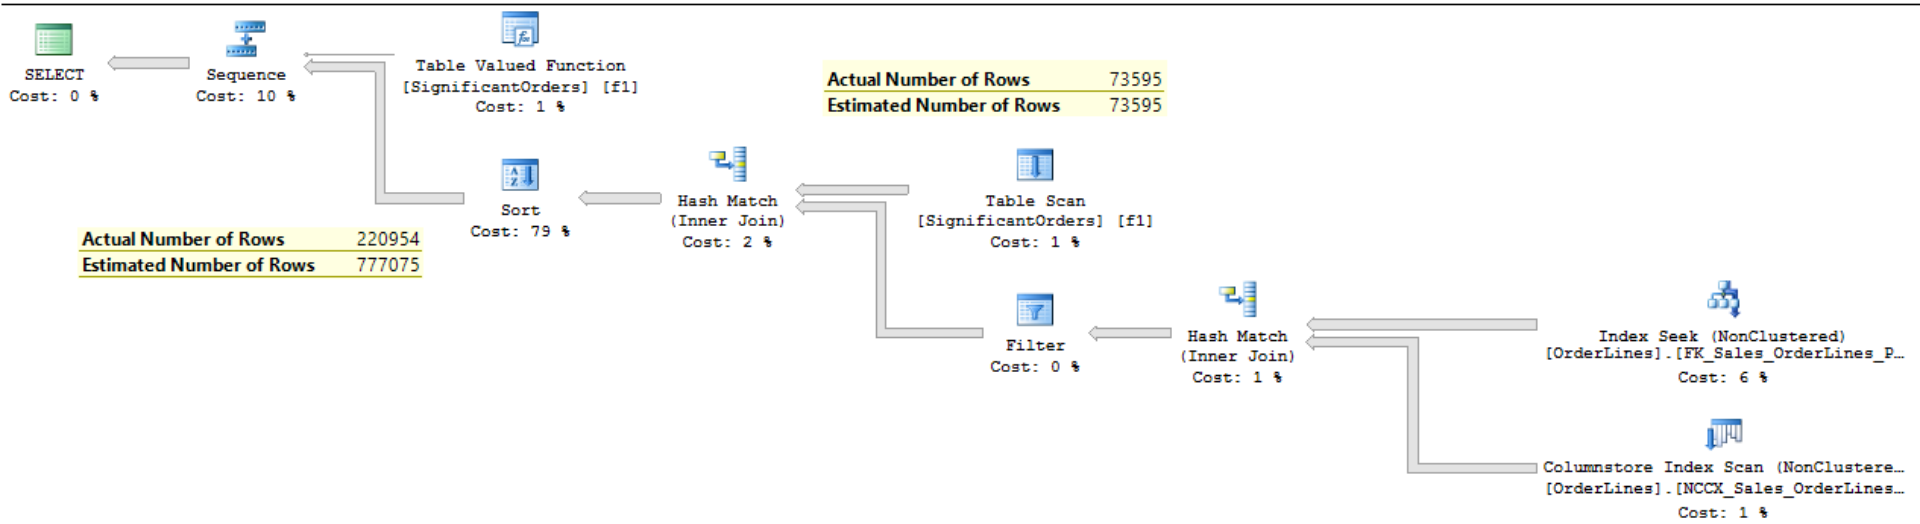


```
Table 'OrderLines'. Scan count 3, logical reads 388, physical reads 0, re
Table 'OrderLines'. Segment reads 1, segment skipped 0.
Table 'Worktable'. Scan count 0, logical reads 0, physical reads 0, read
Table 'Workfile'. Scan count 0, logical reads 0, physical reads 0, read-a
Table '#AF429420'. Scan count 1, logical reads 119, physical reads 0, rea

CPU time = 594 ms,  elapsed time = 1480 ms.
```

SQLPort

Batch Mode Memory Grant Feedback

- Adjust memory grant parameter in the execution plan AFTER the plan is generated
- Monitors the execution of the query and if memory grant is constantly over- or underestimated, it recalculates and adjust it
- Requires a columnstore index on the affected table
- If memory grant memory values oscillate, the feature is disabled

SQLPort

Batch Mode Adaptive Join

- New opeartor- Adaptive Join
  - Allows to choose between Hash Join and Nested Loop Join operators after scanning an input
  - Defines a threshold value that can be use for decision which operator to use
  - It starts as Hash Join and if after input scanning estimated number of rows is less than threshold, it switches to Nested Loop Join

- Generaly, it will better handle some queries with variaty of parameters, but it will not solve all issues caused by wrongly chosen Join operator

- It can be disabled

SQLPort

Adaptive Join Operator

```
EXEC dbo.GetSomeOrderDeatils 112;
```

```
SELECT o.OrderID, o.OrderDate, ol.OrderLineID, ol.Quantity, ol.UnitPrice FROM Sales
Missing Index (Impact 49.4219): CREATE NONCLUSTERED INDEX [<Name of Missing Index,
```

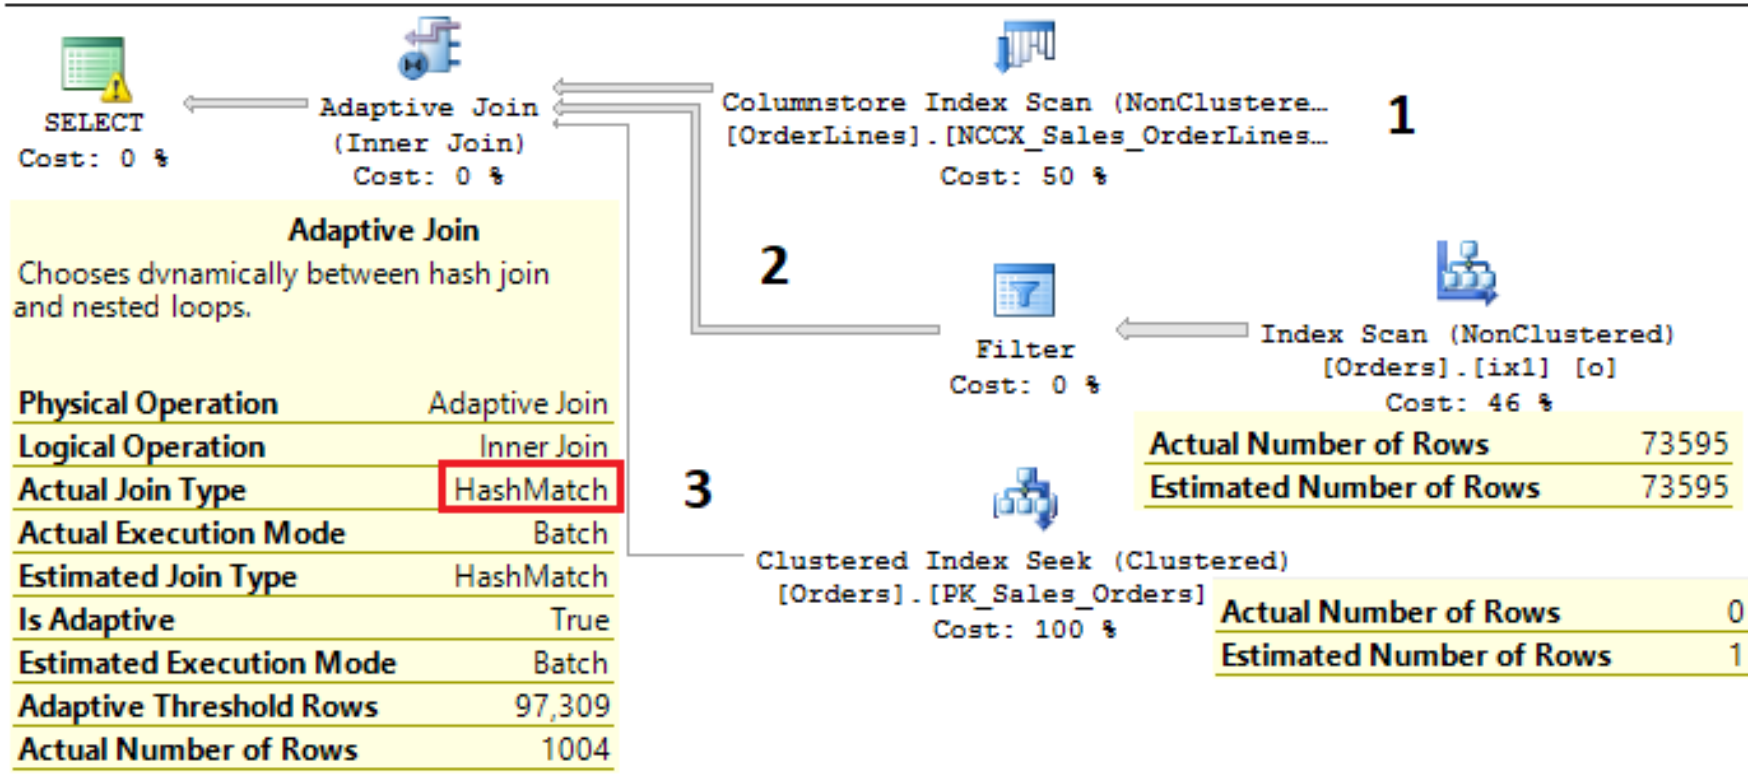

SQLPort

Adaptive Join Operator

```
EXEC dbo.GetSomeOrderDeatils 1;
```

```
SELECT o.OrderID, o.OrderDate, ol.OrderLineID, ol.Quantity, ol.UnitPrice FROM Sales
Missing Index (Impact 49.4219): CREATE NONCLUSTERED INDEX [<Name of Missing Index,
```

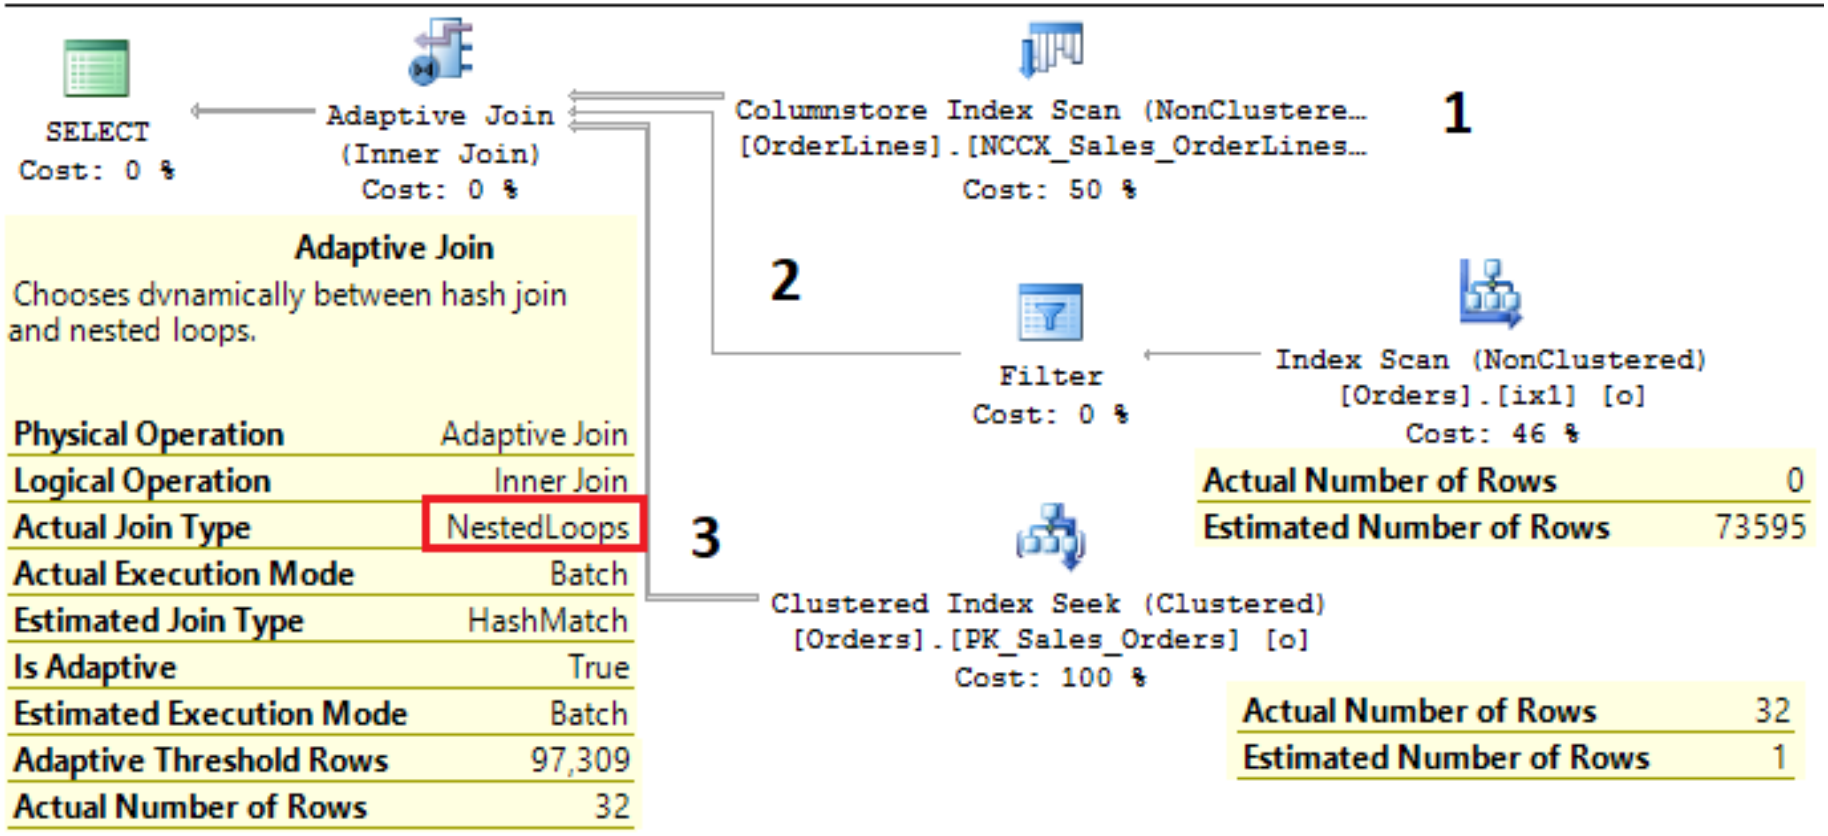

SQLPort

Adaptive Join Operator

Adaptive Join Operator brings overhead

How to disable it:

- OPTION(USE HINT('DISABLE_BATCH_MODE_ADAPTIVE_JOINS'));
- ALTER DATABASE SCOPED CONFIGURATION SET DISABLE_BATCH_MODE_ADAPTIVE_JOINS = ON;

SQLPort

CI ENCONTRO DA COMUNIDADE SQLPORT, LISBON, 28TH MAY 2018

Query Store and Automatic Tuning

Query Store in SQL Server 2016

New troubleshooting tool

- Captures all execution relevant parameters for database queries
- Information are persistent, belongs to the database
- Quick identify performance regressions
- Helps you to learn how your database workload changes over time
- Helps you to identify queries that did not execute successfully
- Allows you to fix some performance issues

SQLPort

Query Store in SQL Server 2017

New features and Enhancements

- Query Store captures wait stats (24 wait stats categories)
- Tuning Recommendations
- Automatic Tuning

SQLPort

Query Store in SQL Server 2017

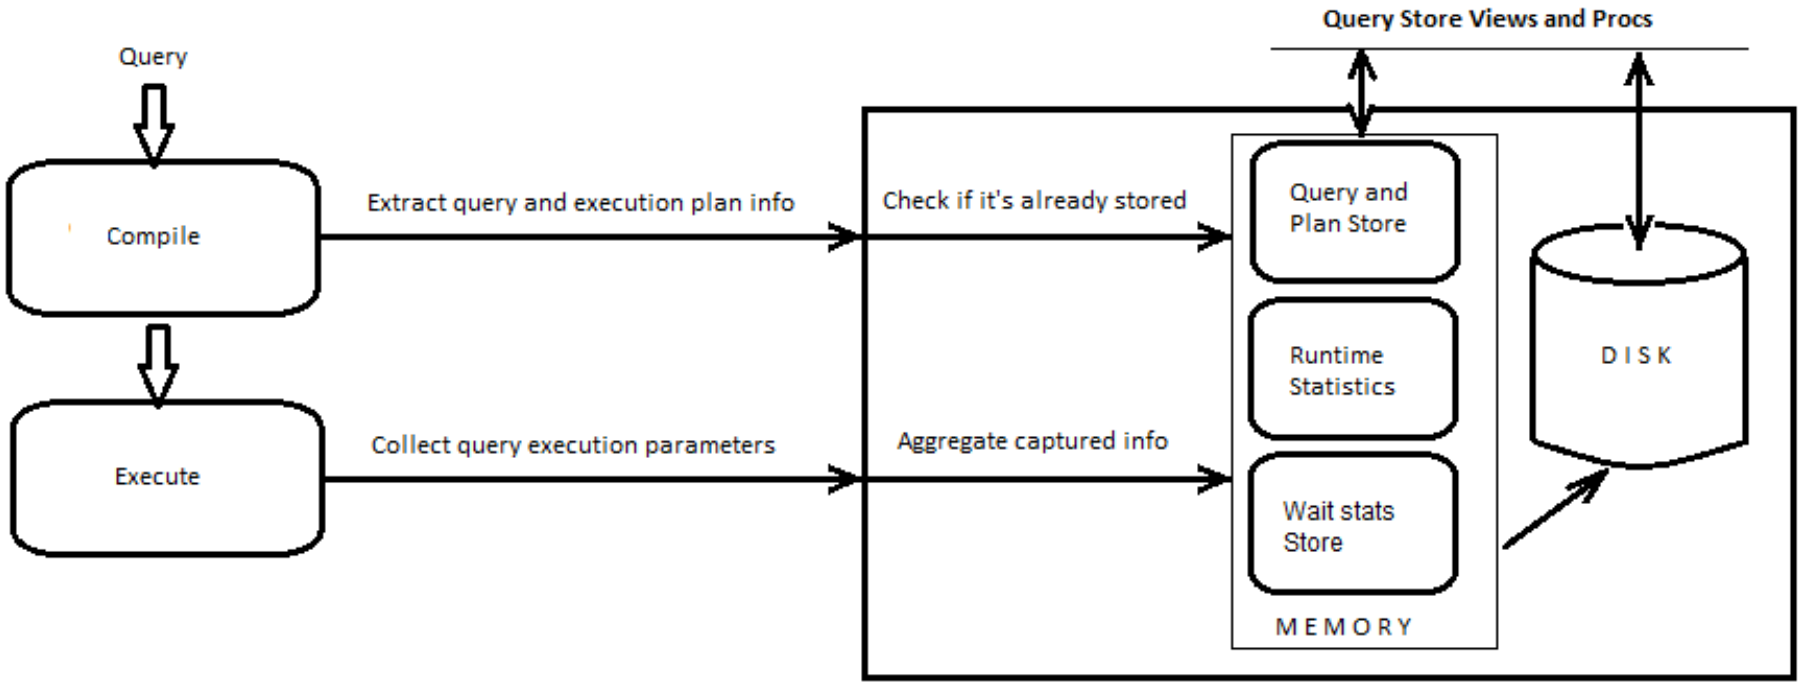

SQLPort

Query Store in SQL Server 2017

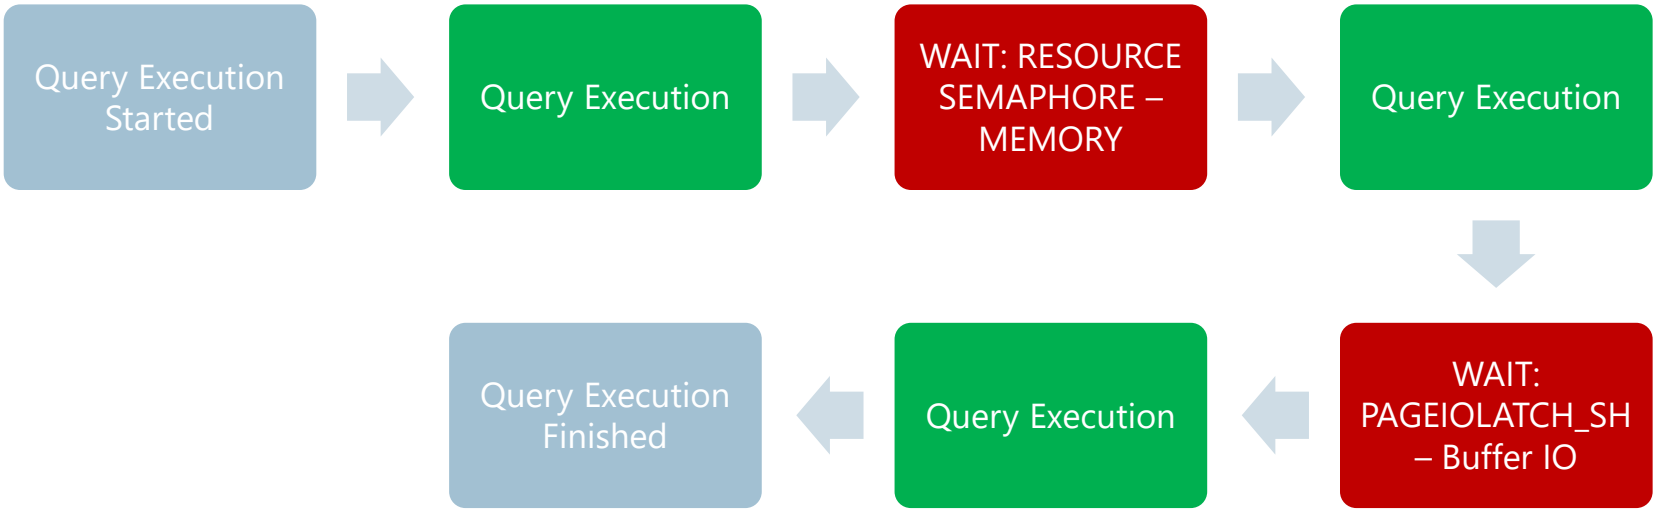

SQLPort

SQL Server 2017 Automatic Tuning

Two options:

- Offline: Recommended actions via DMV **sys.dm_db_tuning_recommendations**
- Online: automatically switch to the last known good plan whenever the regression is detected

```
ALTER DATABASE CURRENT SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN = ON);
```

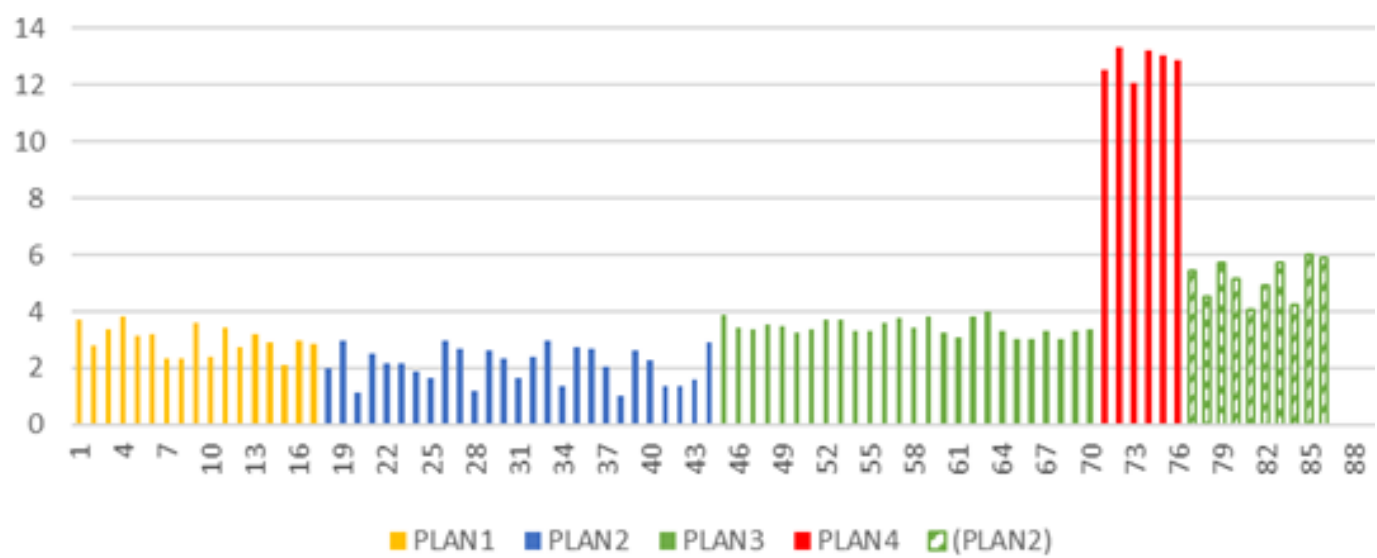

SQLPort

SQL Server 2017 Automatic Tuning

Connect

7\MSSQLSERVER01 (SQL Server 14.0.1000.169 - IC

- Databases
  - System Databases
  - Database Snapshots
  - NewDB
  - WideWorldImporters
    - Database Diagrams
    - Tables
    - Views
    - External Resources
    - Synonyms
    - Programmability
    - Query Store
      - Overall Resource Consumption
      - Top Resource Consuming Queries
      - Queries With Forced Plans
      - Queries With High Variation
      - Tracked Queries
    - Service Broker
    - Storage
    - Security
  - WideWorldImportersDW
- Security
- Server Objects
- Replication
- PolyBase
- Always On High Availability
- Management
- Integration Services Catalogs
- SQL Server Agent (Agent XPs disabled)
- XE Profiler

Queries with forced plans for database WideWorldImporters

Portrait View | Landscape View | Configure

	query id	query sql text	forced plan id	▼force
1	1	SELECT * FROM Sales.Orders o INNER JOIN Sales.Or...	1	0

Plan summary for query 1. Time period: Last month ending at 21.10.2017 23:00

Avg: 80, 60, 40, 20, 0

21.09.2017 24.09.2017 27.09.2017 30.09.2017 03.10.2017 06.10.2017 09.10.2017 12.10.2017 15.10.2017 18.10.2017 21.10.2017

Plan Id: 1, 2, 3

Plan 1 [forced]

Force Plan | Unforce Plan

```
Query 1: Query cost (relative to the batch): 100%
SELECT * FROM Sales.Orders o INNER JOIN Sales.OrderLines ol ON o.OrderID= ol.OrderID INNER JOIN Sales.OrderLines ol2 ON o.…
```

SELECT Cost: 0 % ← Nested Loops (Inner Join) Cost: 0 % ← Nested Loops (Inner Join) Cost: 0 % ← Nested Loops (Inner Join) Cost: 0 % ← Nested Loops (Inner Join) Cost: 0 %

SQLPort

OBRIGADO!