SQLServerFast.com Execution Plan Video Training

Block 1: Understanding execution plans

Level: Basic

Chapter 6: Cardinality in the execution plan

Cardinality in the execution plan

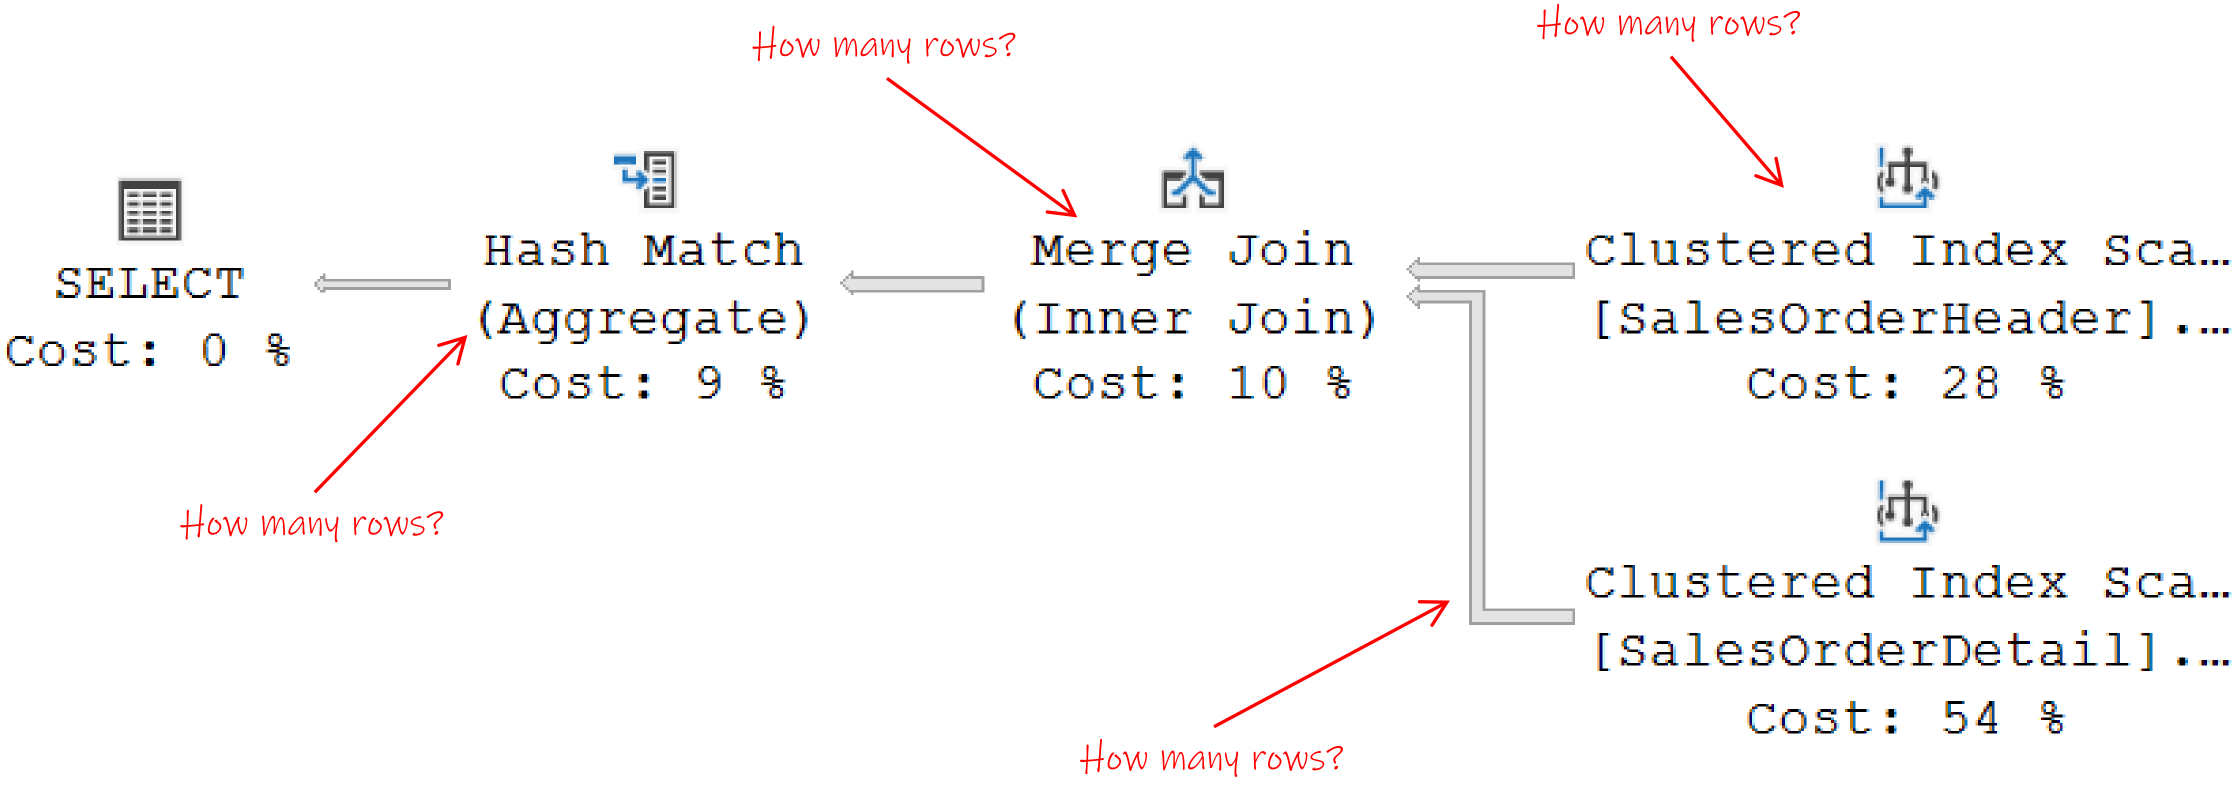

How many rows?

How many rows?

How many rows?

How many rows?

Cardinality in the execution plan

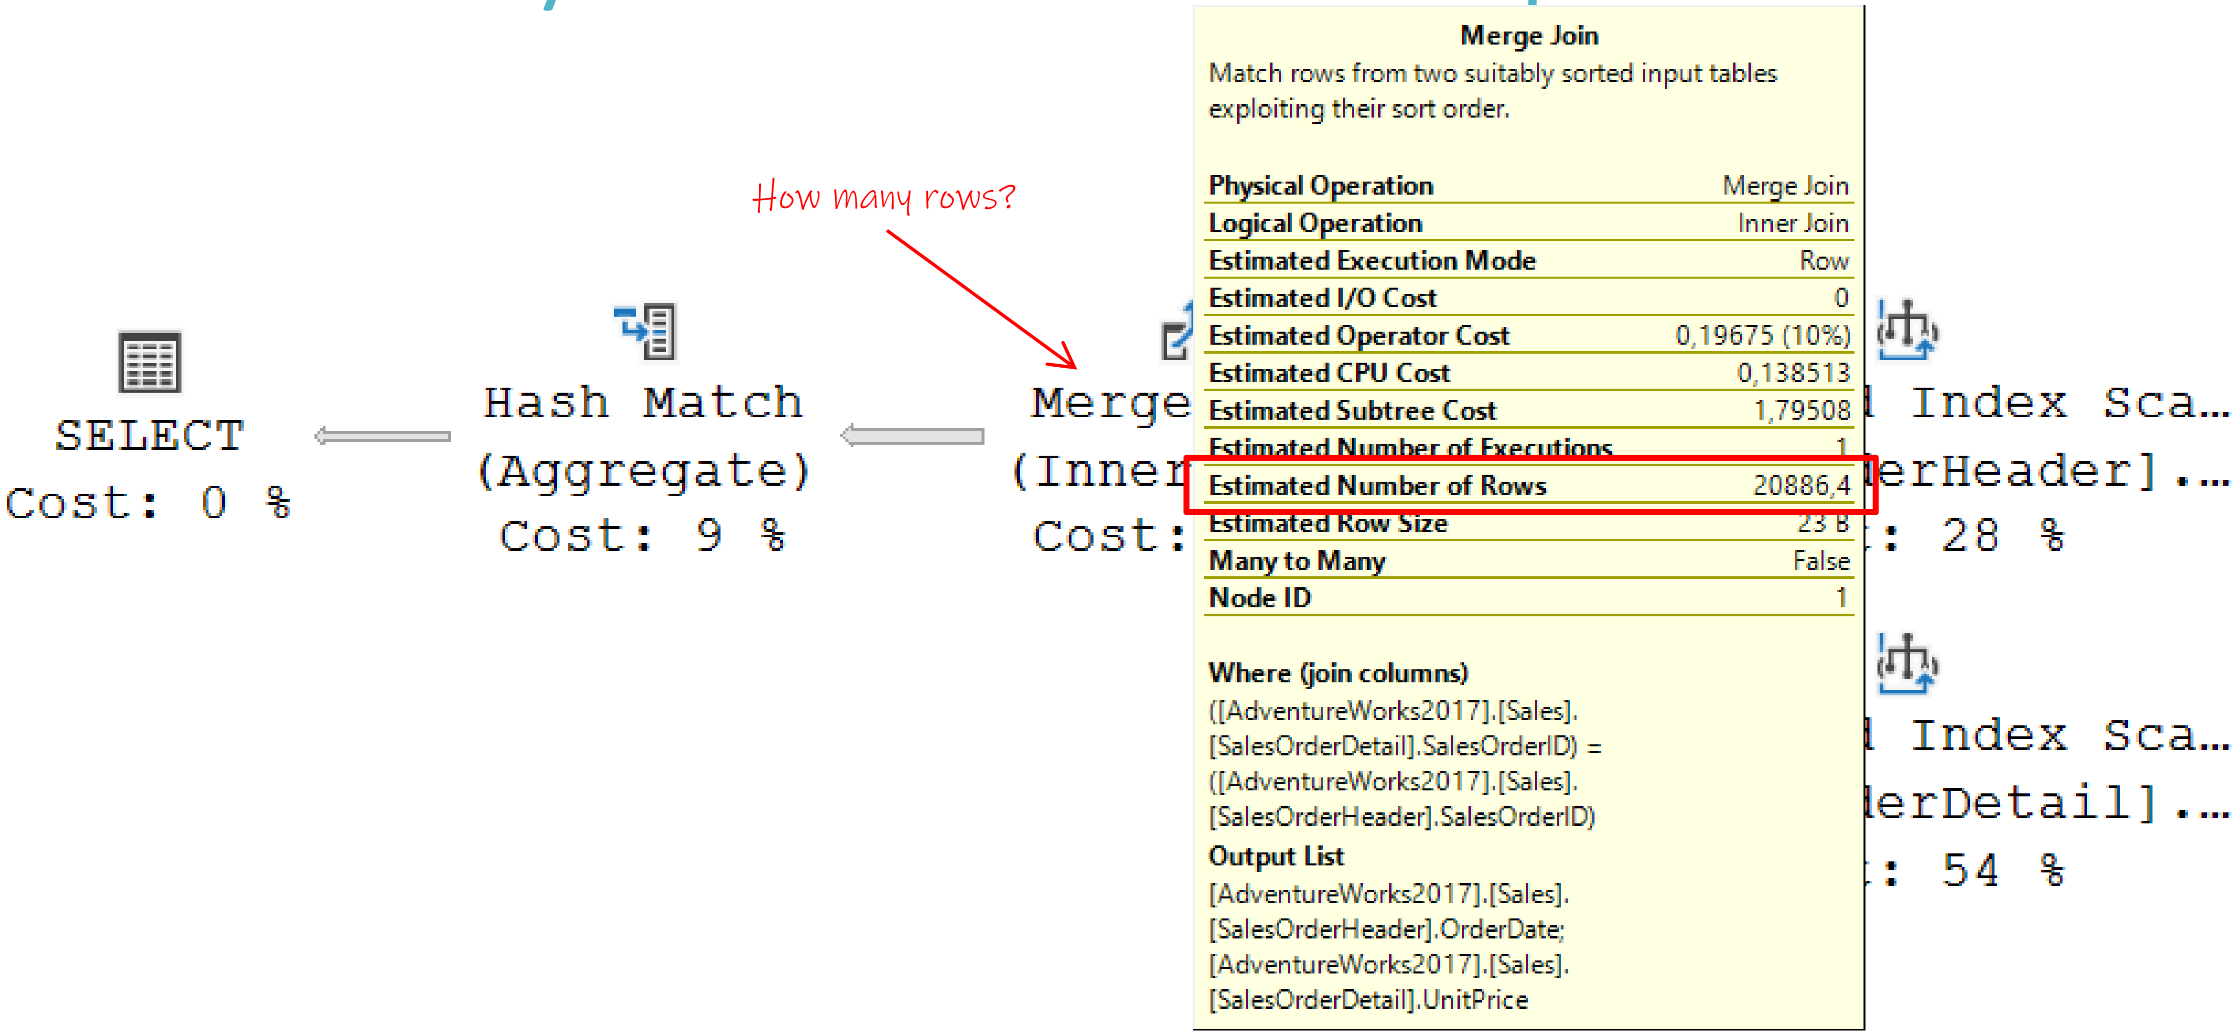

How many rows?

Merge Join

Match rows from two suitably sorted input tables exploiting their sort order.

Physical Operation	Merge Join
Logical Operation	Inner Join
Estimated Execution Mode	Row
Estimated I/O Cost	0
Estimated Operator Cost	0,19675 (10%)
Estimated CPU Cost	0,138513
Estimated Subtree Cost	1,79508
Estimated Number of Executions	1
Estimated Number of Rows	20886,4
Estimated Row Size	23 B
Many to Many	False
Node ID	1

Where (join columns)

([AdventureWorks2017].[Sales].[SalesOrderDetail].SalesOrderID) = ([AdventureWorks2017].[Sales].[SalesOrderHeader].SalesOrderID)

Output List

[AdventureWorks2017].[Sales].[SalesOrderHeader].OrderDate; [AdventureWorks2017].[Sales].[SalesOrderDetail].UnitPrice

Cardinality in the exec

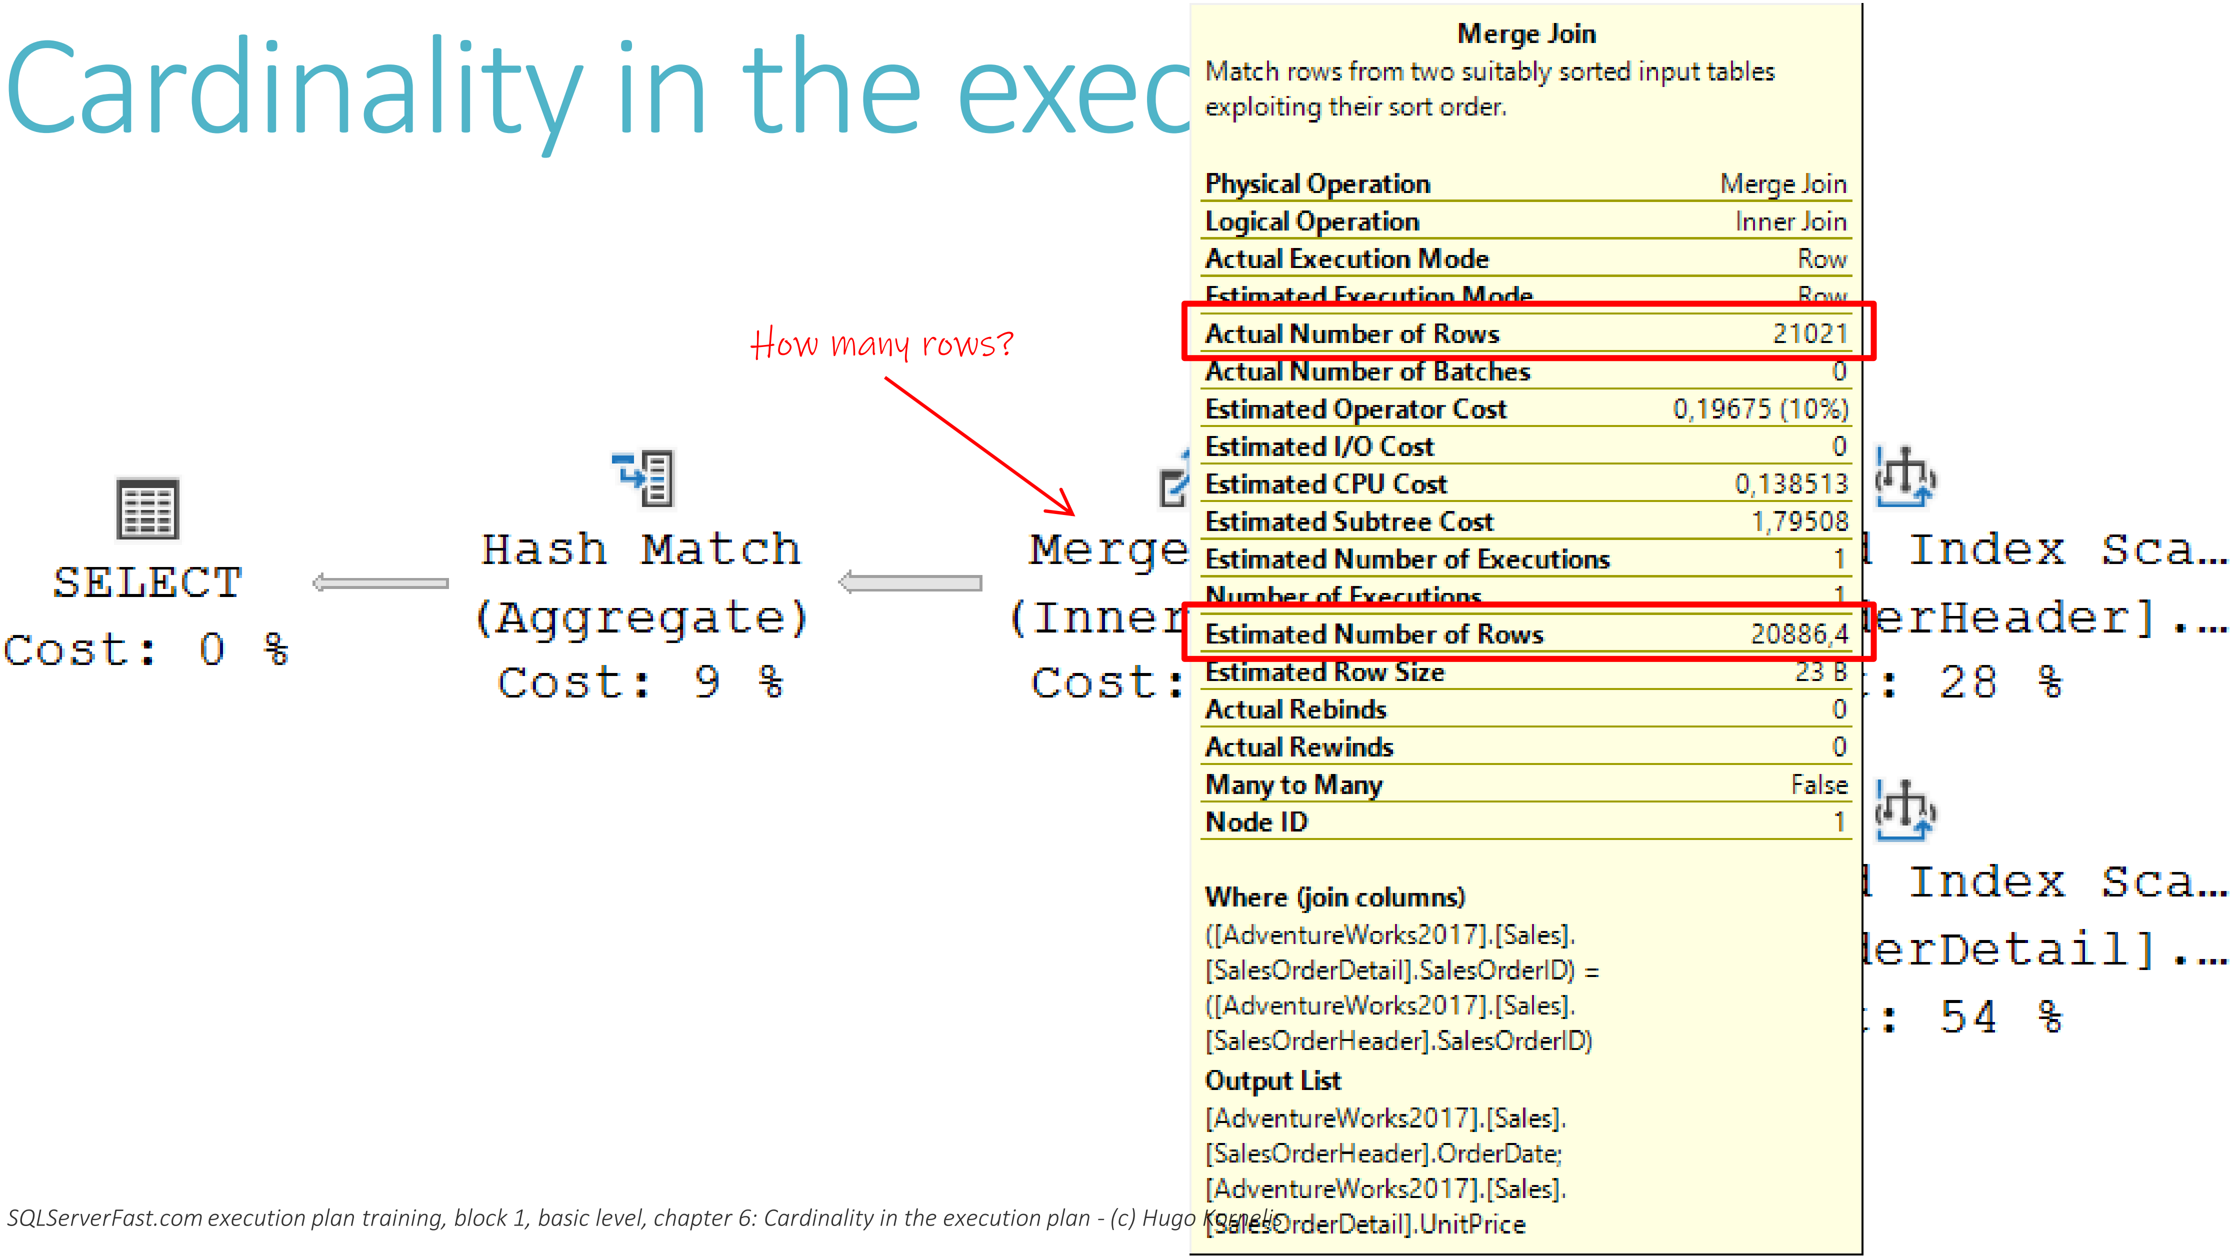

How many rows?

Merge Join

Match rows from two suitably sorted input tables exploiting their sort order.

Physical Operation	Merge Join
Logical Operation	Inner Join
Actual Execution Mode	Row
Estimated Execution Mode	Row
Actual Number of Rows	21021
Actual Number of Batches	0
Estimated Operator Cost	0,19675 (10%)
Estimated I/O Cost	0
Estimated CPU Cost	0,138513
Estimated Subtree Cost	1,79508
Estimated Number of Executions	1
Number of Executions	1
Estimated Number of Rows	20886,4
Estimated Row Size	23 B
Actual Rebinds	0
Actual Rewinds	0
Many to Many	False
Node ID	1

Where (join columns)

([AdventureWorks2017].[Sales].[SalesOrderDetail].SalesOrderID) = ([AdventureWorks2017].[Sales].[SalesOrderHeader].SalesOrderID)

Output List

[AdventureWorks2017].[Sales].[SalesOrderHeader].OrderDate; [AdventureWorks2017].[Sales].[SalesOrderDetail].UnitPrice

Cardinality in the execution plan

For more information about parallelism,
see block 6: Technical operators

Property: *Estimated Number of Rows*

- How many rows will the operator return?
  - Always an estimate
  - Based on “cardinality estimation”

Property: *Actual Number of Rows*

- How many rows did the operator really return?
  - Always the actual, real number
- Includes breakdown per thread in parallel plan
  - Can be used to detect parallel skew

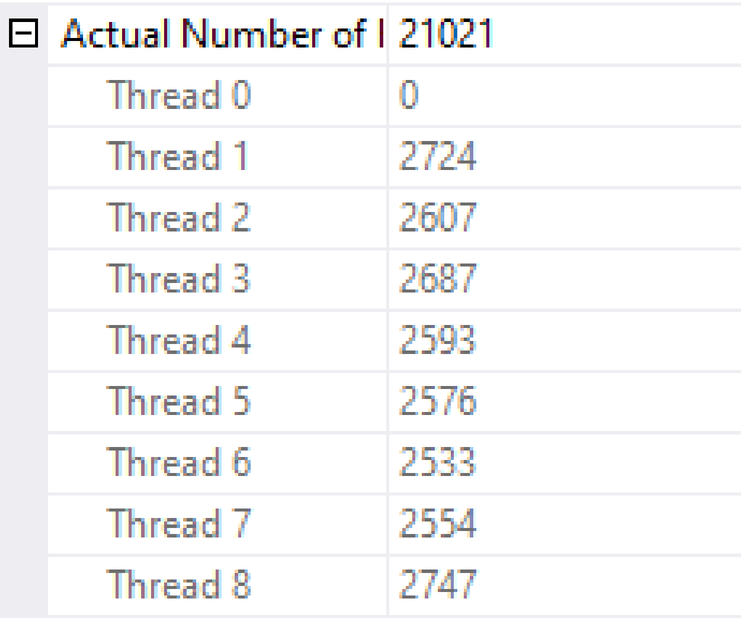

⊟ Actual Number of I	21021
Thread 0	0
Thread 1	2724
Thread 2	2607
Thread 3	2687
Thread 4	2593
Thread 5	2576
Thread 6	2533
Thread 7	2554
Thread 8	2747

Cardinality in the execution plan

- Compare *Actual Number of Rows* to *Estimated Number of Rows*
  - Same or small difference?
    - Cardinality estimate was good or good enough
  - Huge difference?
    - Cardinality estimate was wrong
    - This *can* be the reason the query runs slow
  - Beware: Comparison may not be straightforward!

For more information about how cardinality estimates influence plan choice, see advanced chapter 2: Common operator properties

Cardinality in the execution plan

Property: *Estimated Number of Executions*

- Also based on cardinality estimation
- How often will the operator be initialized?
- Usually equal to 1
- Exception: Nested Loops operators
  - Operators on lower input of Nested Loops get initialized multiple times
  - Every Initialize() call is counted as an “execution”

See also chapter 3: How to read an execution plan

For more information about Nested Loops, see block 3: Combining data

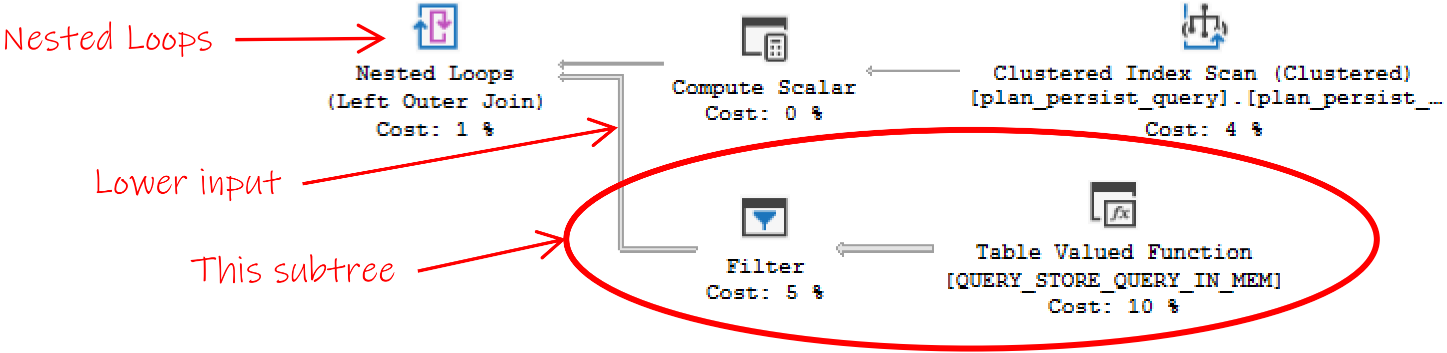

Cardinality in the execution plan

For more information about parallelism,
see block 6: Technical operators

Property: *Estimated Number of Executions*

Property: *Number of Executions*

- Only in execution plan plus run-time statistics
- How often was the operator actually initialized?
- Think of it as *"actual"* number of executions
  - On lower input of Nested Loops, compare to *Estimated Number of Executions*
  - In parallel sections of an execution plan, *Number of Executions* is number of threads
    - In this case, comparison with *Estimated Number of Executions* is **NOT** useful

Cardinality in the execution plan

Compare *Actual Number of Rows* to *Estimated Number of Rows*

- Not always computed the same!
  - Operators on lower input of Nested Loops

For more information about Nested Loops, see block 3: Combining data

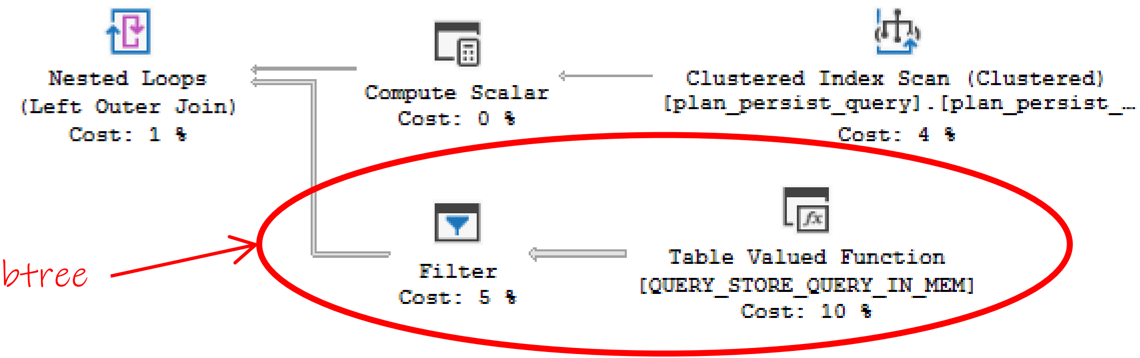

This subtree

Cardinality in the execution plan

- Compare *Actual Number of Rows* to *Estimated Number of Rows*
  - Not always computed the same!
    - Operators on lower input of Nested Loops
      - *Estimated Number of Rows*: Expected average per execution

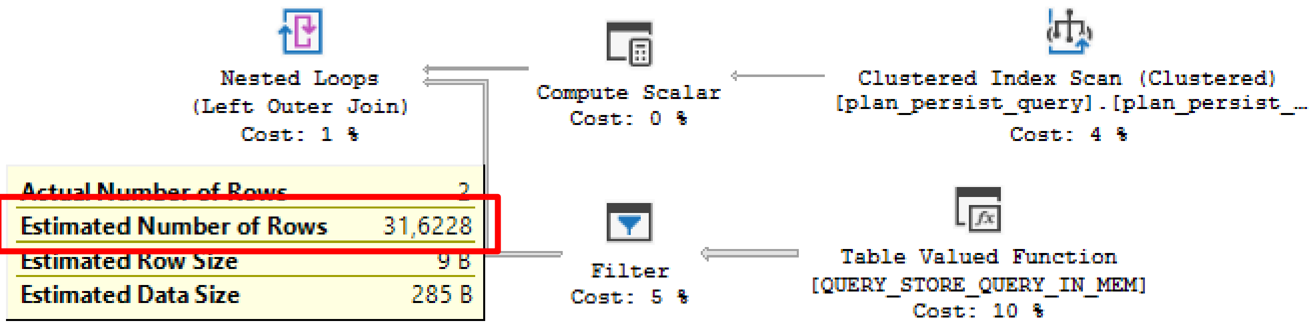

Actual Number of Rows	2
Estimated Number of Rows	31,6228
Estimated Row Size	9 B
Estimated Data Size	285 B

Per execution

Cardinality in the execution plan

- Compare *Actual Number of Rows* to *Estimated Number of Rows*
  - Not always computed the same!
    - Operators on lower input of Nested Loops
      - *Estimated Number of Rows*: Expected average per execution
      - *Actual Number of Rows*: Total across all executions

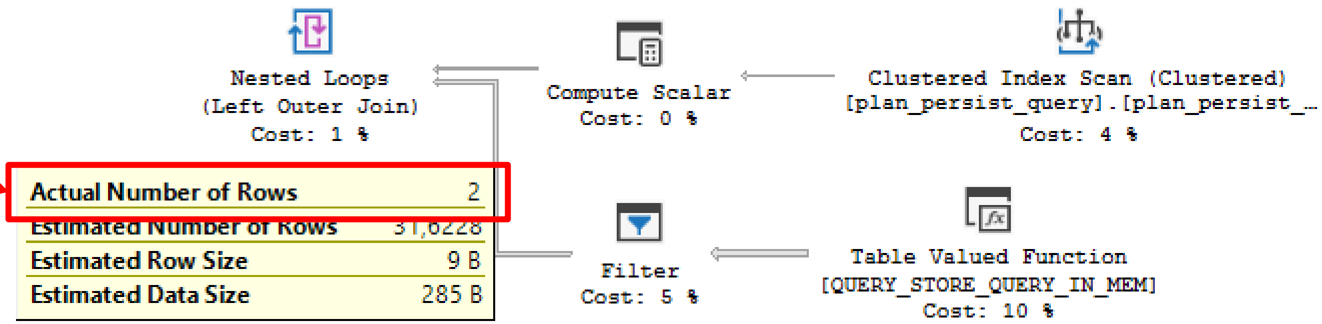

All executions

Actual Number of Rows	2
Estimated Number of Rows	31,6228
Estimated Row Size	9 B
Estimated Data Size	285 B

Cardinality in the execution plan

- Compare *Actual Number of Rows* to *Estimated Number of Rows*
  - Not always computed the same!
    - Operators on lower input of Nested Loops
      - *Estimated Number of Rows*: Expected average per execution
      - *Actual Number of Rows*: Total across all executions
      - Can't compare *Actual Number of Rows* to *Estimated Number of Rows* directly!
        - Should compare *Actual Number of Rows* to *Estimated Number of Rows* * *Estimated Number of Executions*

~~2 vs 31.6228~~

2 vs (31.6228 * 8)

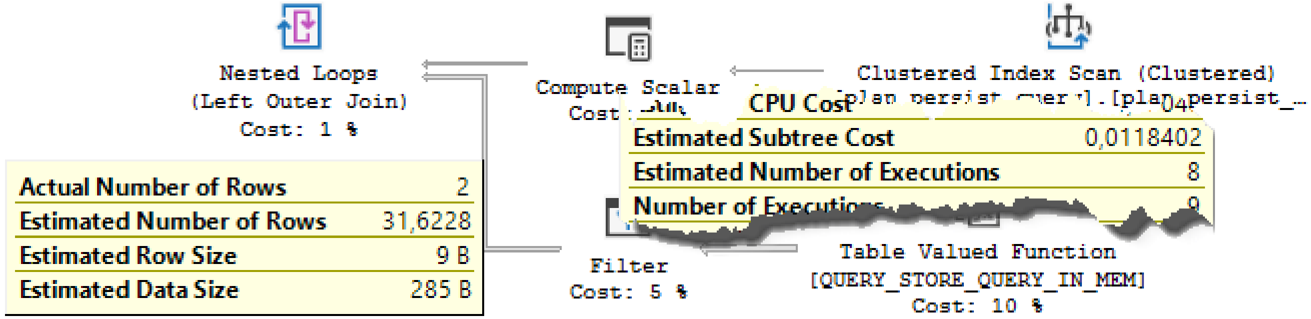

Actual Number of Rows	2
Estimated Number of Rows	31,6228
Estimated Row Size	9 B
Estimated Data Size	285 B

CPU Cost	...04
Estimated Subtree Cost	0,0118402
Estimated Number of Executions	8
Number of Executions	9

Cardinality in the execution plan

- Compare *Actual Number of Rows* to *Estimated Number of Rows*
  - Not always computed the same!
    - Operators on lower input of Nested Loops
      - *Estimated Number of Rows*: Expected average per execution
      - *Actual Number of Rows*: Total across all executions
      - Can’t compare *Actual Number of Rows* to *Estimated Number of Rows* directly!
        - Should compare *Actual Number of Rows* to *Estimated Number of Rows* * *Estimated Number of Executions*
      - Improved labels in Management Studio 18.5
        - *Actual Number of Rows* ***for All Executions***, *Estimated Number of Rows* ***Per Execution***
        - Added (computed) *Estimated Number of Rows for All Executions*
        - Does not help for older versions, other tools, or your own queries with XQuery logic on execution plans

Cardinality in the execution plan

- Look at *Actual Number of Rows* or *Estimated Number of Rows*
  - Indication of amount of work done
    - Each row returned *by* an operator represents work done to produce the row
    - Each row returned *to* an operator represents work done to process the row
  - *Usually* (not always!), high row count means lot of time/resources used
    - Can be used to focus on plan areas with high *Actual Number of Rows*
      - Or, for execution plan only, high *Estimated Number of Rows* * *Estimated Number of Executions*
  - Indicated visually in Management Studio

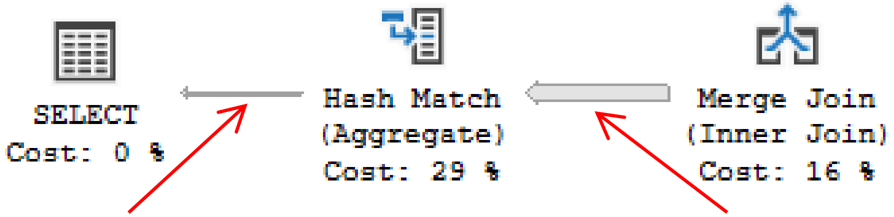

"Thin pipe" = few rows

"Fat pipe" = many rows

Cardinality in the execution plan

- Issues with *Actual* and *Estimated Number of Rows*
  - Total for all executions versus average per execution
    - Improved in SQL Server Management Studio 18.5
  - Arrow width uses *Estimated Number of Rows per Execution*
    - Even in SQL Server Management Studio 18.5
  - Arrow width uses *Number of Rows Read*
  - Compute Scalar often doesn't actually execute

For more information about Compute Scalar, see block 5: Various data manipulations

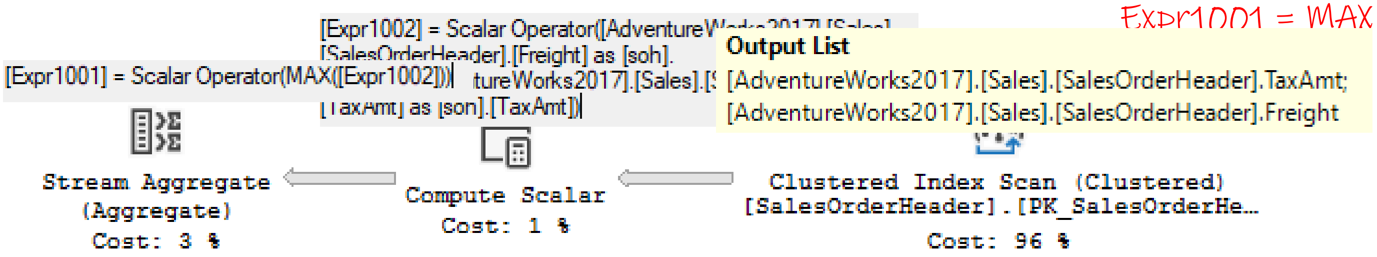

[Expr1002] = Scalar Operator([AdventureWorks2017].[Sales].
[SalesOrderHeader].[Freight] as [soh].
[Expr1001] = Scalar Operator(MAX([Expr1002])) tureWorks2017].[Sales].[S
[TaxAmt] as [soh].[TaxAmt])

Output List
[AdventureWorks2017].[Sales].[SalesOrderHeader].TaxAmt;
[AdventureWorks2017].[Sales].[SalesOrderHeader].Freight

Stream Aggregate
(Aggregate)
Cost: 3 %
Compute Scalar
Cost: 1 %
Clustered Index Scan (Clustered)
[SalesOrderHeader].[PK_SalesOrderHe...
Cost: 96 %

Expr1001 = MAX(Freight + TaxAmt)

Output List: TaxAmt ; Freight

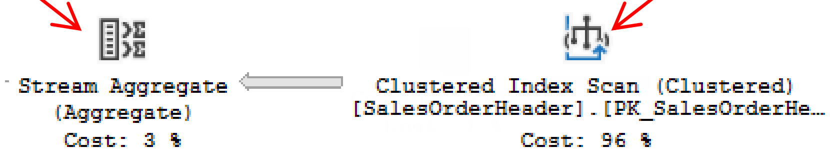

Stream Aggregate
(Aggregate)
Cost: 3 %
Clustered Index Scan (Clustered)
[SalesOrderHeader].[PK_SalesOrderHe...
Cost: 96 %

Cardinality in the execution plan

Issues with *Actual* and *Estimated Number of Rows*

Total for all executions versus average per execution

Improved in SQL Server Manageme

Arrow width uses *Estimated Nu* *xecution*

E nagement St

Ar *mber of Rc*

Co

F an plus!

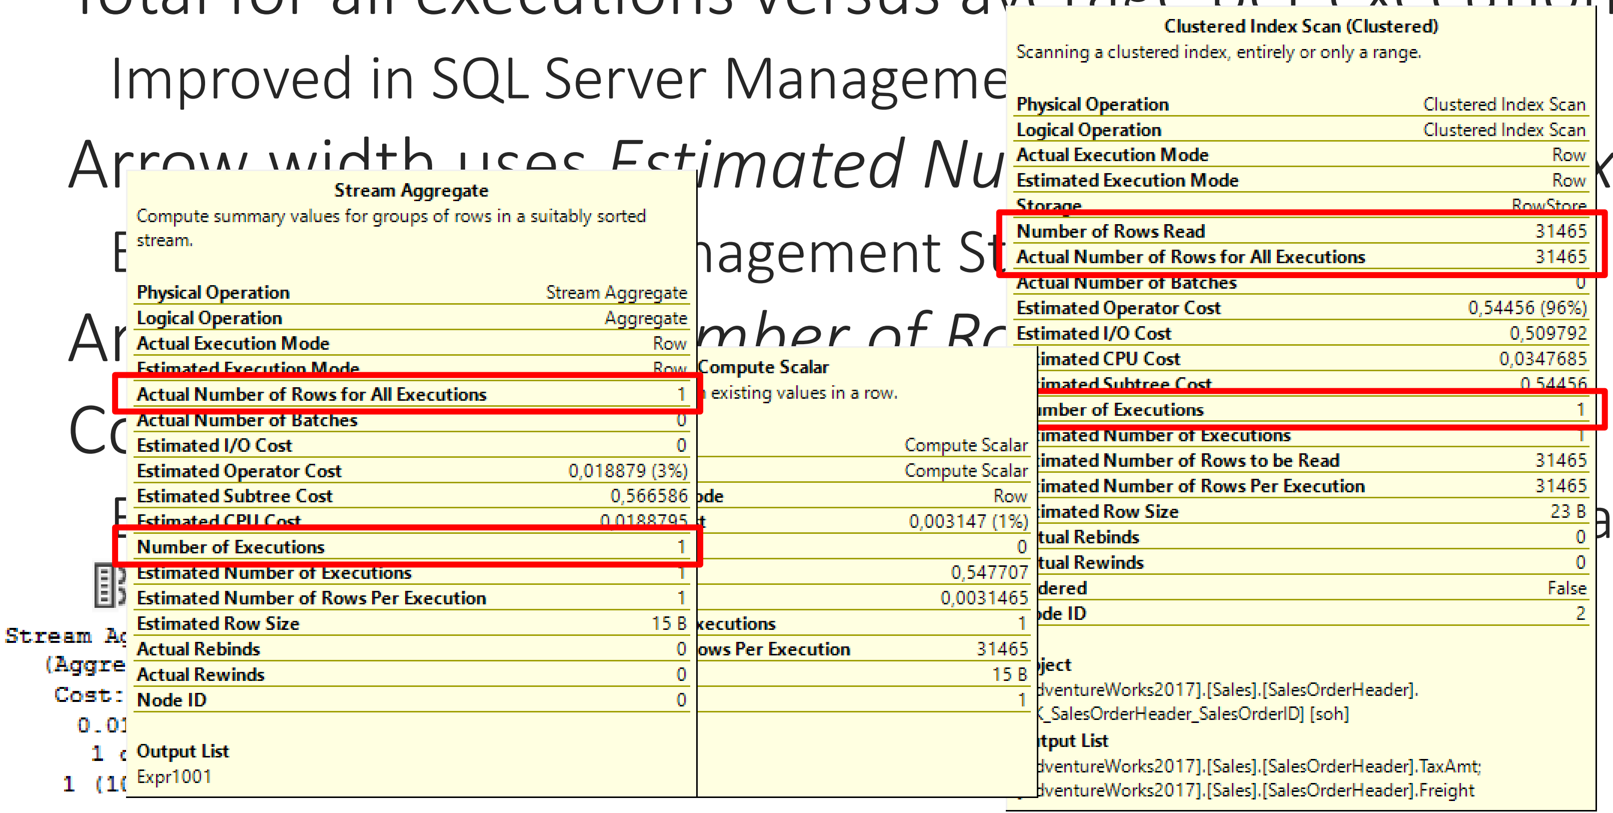

Stream Ag
(Aggre
Cost:
0.01
1 c
1 (1(

Stream Aggregate

Compute summary values for groups of rows in a suitably sorted stream.

Physical Operation	Stream Aggregate
Logical Operation	Aggregate
Actual Execution Mode	Row
Estimated Execution Mode	Row
Actual Number of Rows for All Executions	1
Actual Number of Batches	0
Estimated I/O Cost	0
Estimated Operator Cost	0,018879 (3%)
Estimated Subtree Cost	0,566586
Estimated CPU Cost	0,0188795
Number of Executions	1
Estimated Number of Executions	1
Estimated Number of Rows Per Execution	1
Estimated Row Size	15 B
Actual Rebinds	0
Actual Rewinds	0
Node ID	0

Output List
Expr1001

Compute Scalar

existing values in a row.

	Compute Scalar
	Compute Scalar
ode	Row
t	0,003147 (1%)
	0
	0,547707
	0,0031465
xecutions	1
ows Per Execution	31465
	15 B
	1

Clustered Index Scan (Clustered)

Scanning a clustered index, entirely or only a range.

Physical Operation	Clustered Index Scan
Logical Operation	Clustered Index Scan
Actual Execution Mode	Row
Estimated Execution Mode	Row
Storage	RowStore
Number of Rows Read	31465
Actual Number of Rows for All Executions	31465
Actual Number of Batches	0
Estimated Operator Cost	0,54456 (96%)
Estimated I/O Cost	0,509792
imated CPU Cost	0,0347685
imated Subtree Cost	0,54456
mber of Executions	1
imated Number of Executions	1
imated Number of Rows to be Read	31465
imated Number of Rows Per Execution	31465
imated Row Size	23 B
tual Rebinds	0
tual Rewinds	0
dered	False
de ID	2

ject
dventureWorks2017].[Sales].[SalesOrderHeader].
(_SalesOrderHeader_SalesOrderID] [soh]
tput List
dventureWorks2017].[Sales].[SalesOrderHeader].TaxAmt;
dventureWorks2017].[Sales].[SalesOrderHeader].Freight

Cardinality in the execution plan

- Issues with *Actual* and *Estimated Number of Rows*
  - Total for all executions versus average per execution
    - Improved in SQL Server Management Studio 18.5
  - Arrow width uses *Estimated Number of Rows per Execution*
    - Even in SQL Server Management Studio 18.5
  - Arrow width uses *Number of Rows Read*
  - Compute Scalar often doesn't actually execute
    - But …. no *Actual Number of Rows*, even in an execution plan plus!
    - Arrow width always based on *Estimated Number of Rows*
    - Change in width? Check for bad estimates!

Cardinality in the execution plan

Issues with *Actual* and *Estimated Number of Rows*

- Total for all executions versus average per execution
- Arrow width uses *Estimated Number of Rows per Execution*
- Arrow width uses *Number of Rows Read*
- Compute Scalar often doesn’t actually execute
- (Rare): related to missing nodes

For more information on missing nodes
and how this affects properties,
see advanced chapter 4: Missing nodes

Cardinality in the execution plan

- Issues with *Actual* and *Estimated Number of Rows*
  - Total for all executions versus average per execution
  - Arrow width uses *Estimated Number of Rows per Execution*
  - Arrow width uses *Number of Rows Read*
  - Compute Scalar often doesn’t actually execute
  - (Rare): related to missing nodes
    - Operator with *Predicate* property **may** report *Estimated Number of Rows* before applying the predicate
    - Key Lookup / RID Lookup with *Predicate* **may** report total estimate as *Estimated Number of Rows* ***Per Execution***
  - Specific cases only; usually *Actual* and *Estimated Number of Rows* are reported fine

Summary

- *Actual Number of Rows* and *Estimated Number of Rows*
  - Use to find execution plan sections that do a lot of work
  - Compare to find incorrect cardinality estimates
    - These mistakes might have caused bad plan choices
  - Use to find skew in parallel processing
- Be aware of the caveats and special cases

Next chapters

- Chapter 7: Percentages in the execution plan
  - Compare plan and operator costs
  - Compare cardinality estimates