

SQLServerFast.com

Execution Plan Video Training

Block 2: Reading data

Level: Basic

Chapter 4: Lookup operators

Lookup operators

Key Lookup

Same as Clustered Index Seek



Key Lookup (Clustered)

Lookup operators

Key Lookup

Same as Clustered Index Seek, singleton lookup

Called Clustered Index Seek in SQL Server 2005 RTM

Changed in SQL Server 2005, Service Pack 2

Called “Key Lookup” when ...

... it does a singleton lookup, and

... target value fetched from nonclustered index on same table



Key Lookup (Clustered)

Key Lookup

SELECT
Cost: 0 %

 Nested Loops
(Inner Join)
Cost: 0 %

 Index Seek (NonClus...
[PersonsCluster].[i...
Cost: 35 %

 Key Lookup (Cluster...
[PersonsCluster].[P...
Cost: 65 %

```
SELECT FirstName,
       Other
FROM   dbo.PersonsCluster
WHERE  LastName = 'Visser';
```

Page type: Index		
prev		next
LName	FName	page
		xxxx
Visser	Klaas	xxxx

Page type: Index	
prev	next
ID	page
	xxxx
25	xxxx

Page type: Index		
prev		next
LName	FName	page
		xxxx
Jansen	Jan	xxxx

Page type: Index		
prev		next
LName	FName	page
Visser	Klaas	xxxx

Page type: Index	
prev	next
ID	page
	xxxx
7	xxxx

Page type: Index	
prev	next
ID	page
25	xxxx

Page type: Index			
prev			next
LName	FName	ID	Title
Blok	Wim	4	Mr.
Boer, de	Wilma	7	Mrs.
Jansen	Klaas	9	Mr.

Page type: Index			
prev			next
LName	FName	ID	Title
Jansen	Jan	1	Mr.
Rekels	Els	16	Ms.
Rekels	Els	25	Mrs.
Visser	Annet	81	Dr.

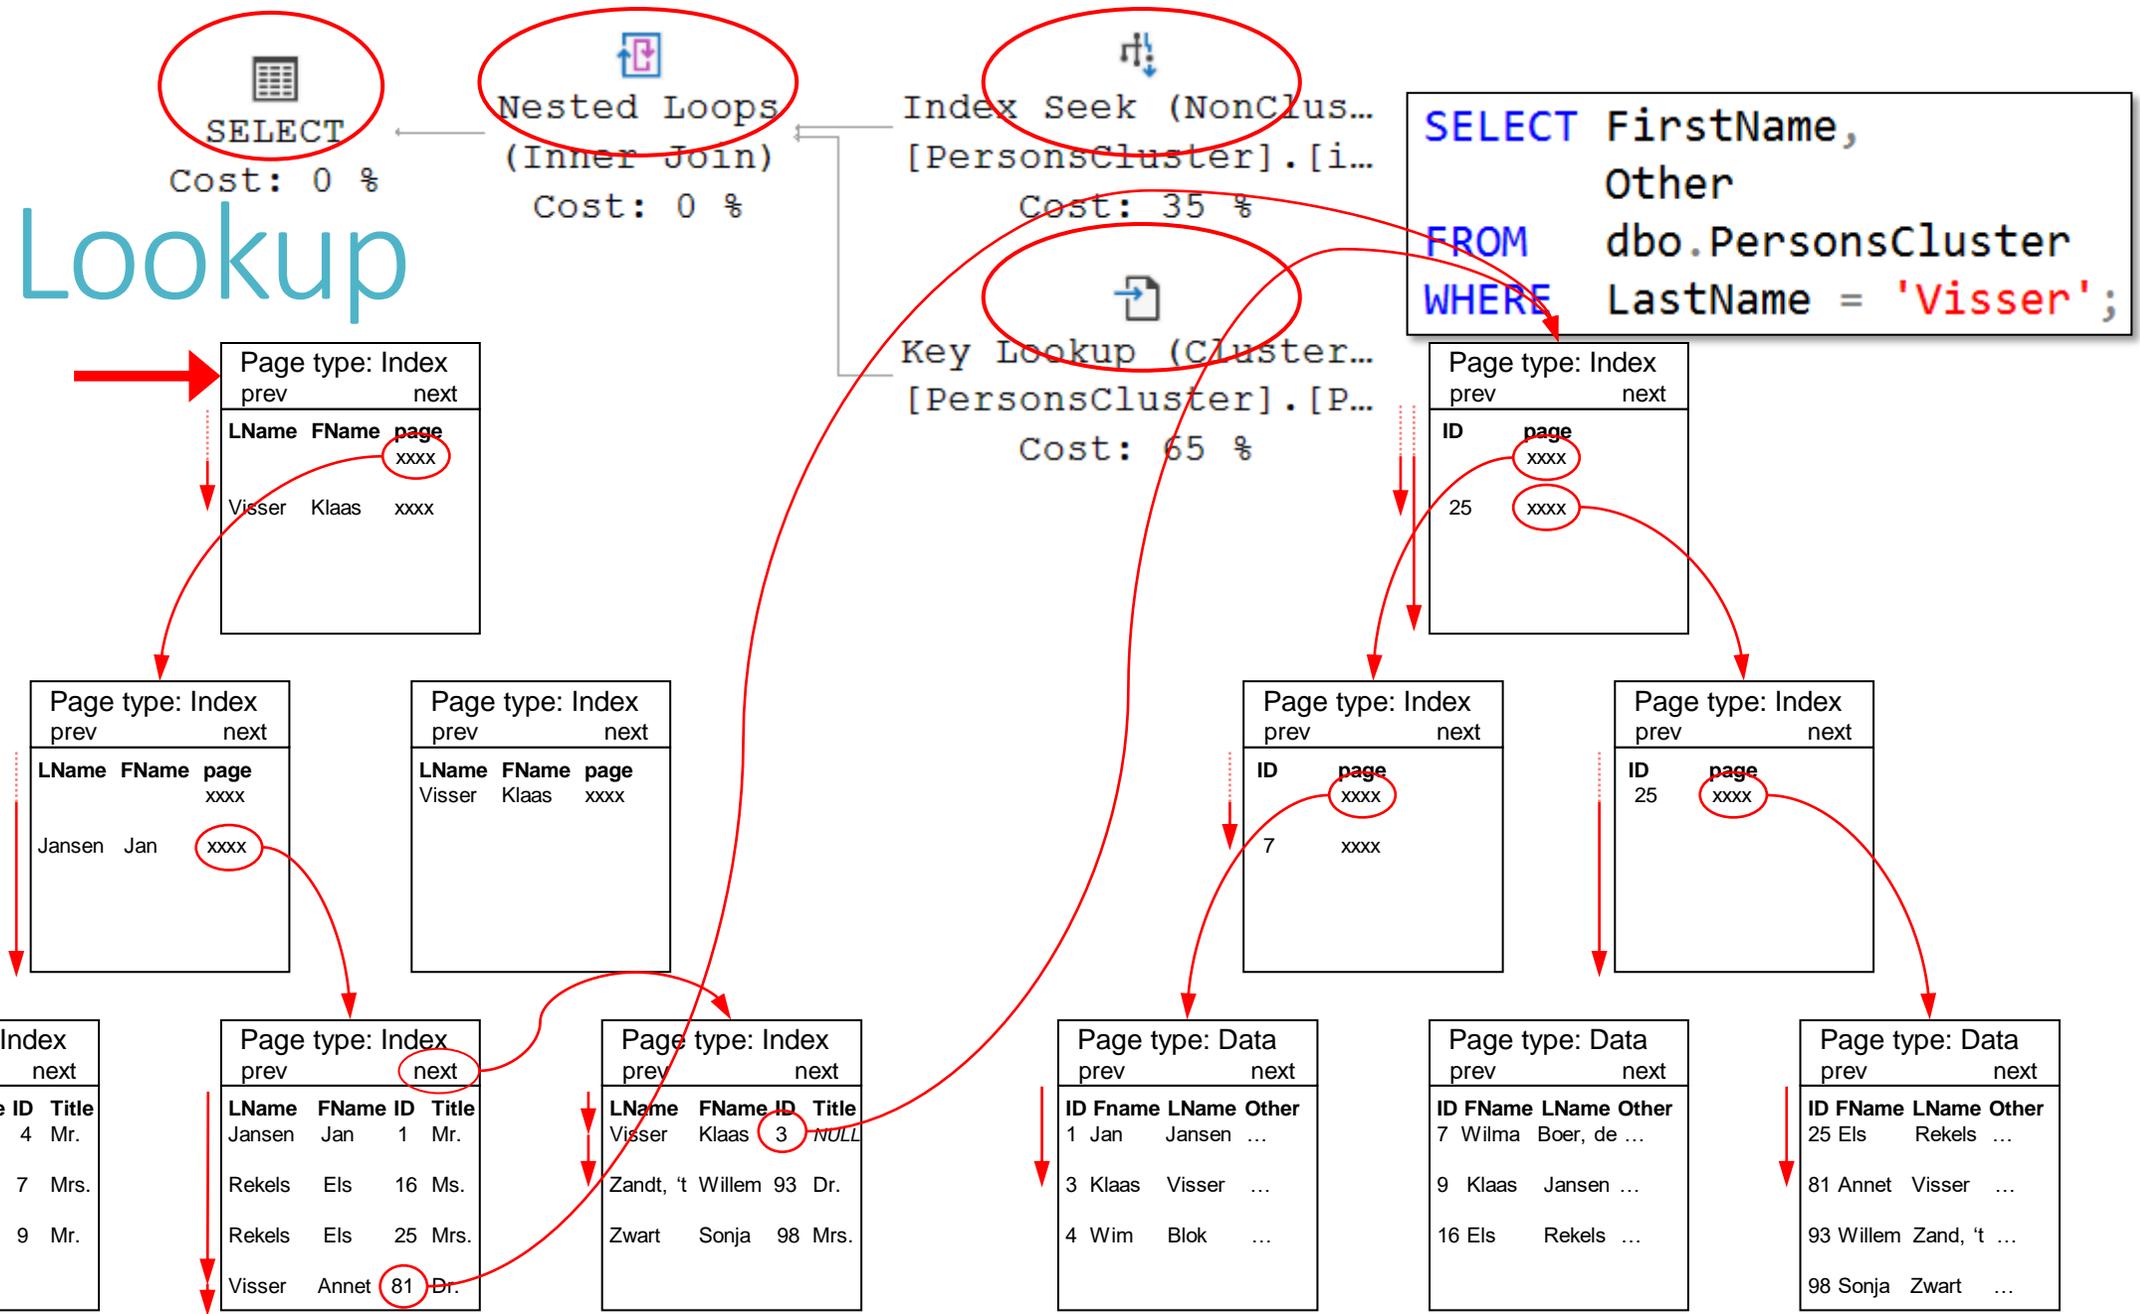
Page type: Index			
prev			next
LName	FName	ID	Title
Visser	Klaas	3	NULL
Zandt, 't	Willem	93	Dr.
Zwart	Sonja	98	Mrs.

Page type: Data			
prev			next
ID	Fname	LName	Other
1	Jan	Jansen	...
3	Klaas	Visser	...
4	Wim	Blok	...

Page type: Data			
prev			next
ID	FName	LName	Other
7	Wilma	Boer, de	...
9	Klaas	Jansen	...
16	Els	Rekels	...

Page type: Data			
prev			next
ID	FName	LName	Other
25	Els	Rekels	...
81	Annet	Visser	...
93	Willem	Zandt, 't	...
98	Sonja	Zwart	...

Key Lookup



Lookup operators

Key Lookup Properties (popup)

Actual row count
(total for all executions)

Estimated row count
(total for all executions)

Target clustered index
(database, schema, table, index,
plus optionally alias)

Key Lookup (Clustered)	
Uses a supplied clustering key to lookup on a table that has a clustered index.	
Physical Operation	Key Lookup
Logical Operation	Key Lookup
Actual Execution Mode	Row
Estimated Execution Mode	Row
Storage	RowStore
Number of Rows Read	14
Actual Number of Rows for All Executions	14
Actual Number of Batches	0
Estimated Operator Cost	0,0280383 (89%)
Estimated I/O Cost	0,003125
Estimated CPU Cost	0,0001581
Estimated Subtree Cost	0,0280383
Number of Executions	14
Estimated Number of Executions	9,5
Estimated Number of Rows for All Executions	9,5
Estimated Number of Rows Per Execution	1
Estimated Row Size	22 B
Actual Rebinds	0
Actual Rewinds	0
Ordered	True
Node ID	3
Object	[AdventureWorks2012].[Person].[Person]. [PK_Person_BusinessEntityID] [p]
Output List	[AdventureWorks2012].[Person].[Person].NameStyle; [AdventureWorks2012].[Person].[Person].Suffix
Seek Predicates	Seek Keys[1]: Prefix: [AdventureWorks2012].[Person]. [Person].BusinessEntityID = Scalar Operator([AdventureWorks2012]. [Person].[Person].[BusinessEntityID] as [p].[BusinessEntityID])

(Clustered)

Lookup operators

Key Lookup Properties (popup)

Actual row count
(total for all executions)

Estimated execution count

Estimated row count
(single execution)

Key Lookup (Clustered)	
Uses a supplied clustering key to lookup on a table that has a clustered index.	
Physical Operation	Key Lookup
Logical Operation	Key Lookup
Actual Execution Mode	Row
Estimated Execution Mode	Row
Storage	RowStore
Number of Rows Read	14
Actual Number of Rows for All Executions	14
Actual Number of Batches	0
Estimated Operator Cost	0,0280383 (89%)
Estimated I/O Cost	0,003125
Estimated CPU Cost	0,0001581
Estimated Subtree Cost	0,0280383
Number of Executions	14
Estimated Number of Executions	9,5
Estimated Number of Rows for All Executions	9,5
Estimated Number of Rows Per Execution	1
Estimated Row Size	22 B
Actual Rebinds	0
Actual Rewinds	0
Ordered	True
Node ID	3
Object [AdventureWorks2012].[Person].[Person]. [PK_Person_BusinessEntityID] [p]	
Output List [AdventureWorks2012].[Person].[Person].NameStyle; [AdventureWorks2012].[Person].[Person].Suffix	
Seek Predicates Seek Keys[1]: Prefix: [AdventureWorks2012].[Person]. [Person].BusinessEntityID = Scalar Operator([AdventureWorks2012]. [Person].[Person].[BusinessEntityID] as [p].[BusinessEntityID])	

9.5 * 1 vs 14

(Clustered)

Lookup operators

Key Lookup Properties (popup)

Key Lookup (Clustered)	
Uses a supplied clustering key to lookup on a table that has a clustered index.	
Physical Operation	Key Lookup
Logical Operation	Key Lookup
Actual Execution Mode	Row
Estimated Execution Mode	Row
Storage	RowStore
Number of Rows Read	14
Actual Number of Rows for All Executions	14
Actual Number of Batches	0
Estimated Operator Cost	0,0280383 (89%)
Estimated I/O Cost	0,003125
Estimated CPU Cost	0,0001581
Estimated Subtree Cost	0,0280383
Number of Executions	14
Estimated Number of Executions	9,5
Estimated Number of Rows for All Executions	9,5
Estimated Number of Rows Per Execution	1
Estimated Row Size	22 B
Actual Rebinds	0
Actual Rewinds	0
Ordered	True
Node ID	3
Object	
[AdventureWorks2012].[Person].[Person]. [PK_Person_BusinessEntityID] [p]	
Output List	
[AdventureWorks2012].[Person].[Person].NameStyle; [AdventureWorks2012].[Person].[Person].Suffix	
Seek Predicates	
Seek Keys[1]: Prefix: [AdventureWorks2012].[Person]. [Person].BusinessEntityID = Scalar Operator([AdventureWorks2012]. [Person].[Person].[BusinessEntityID] as [p].[BusinessEntityID])	

Specification of rows to be read
(always equality on all key columns)

(Clustered)

Lookup operators

Key Lookup Properties (popup)

Key Lookup (Clustered)	
Uses a supplied clustering key to lookup on a table that has a clustered index.	
Physical Operation	Key Lookup
Logical Operation	Key Lookup
Actual Execution Mode	Row
Estimated Execution Mode	Row
Storage	RowStore
Number of Rows Read	14
Actual Number of Rows for All Executions	14
Actual Number of Batches	0
Estimated Operator Cost	0,0280383 (89%)
Estimated I/O Cost	0,003125
Estimated CPU Cost	0,0001581
Estimated Subtree Cost	0,0280383
Number of Executions	14
Estimated Number of Executions	9,5
Estimated Number of Rows for All Executions	9,5
Estimated Number of Rows Per Execution	1
Estimated Row Size	22 B
Actual Rebinds	0
Actual Rewinds	0
Ordered	True
Node ID	3
Object	
[AdventureWorks2012].[Person].[Person]. [PK_Person_BusinessEntityID] [p]	
Output List	
[AdventureWorks2012].[Person].[Person].NameStyle; [AdventureWorks2012].[Person].[Person].Suffix	
Seek Predicates	
Seek Keys[1]: Prefix: [AdventureWorks2012].[Person]. [Person].BusinessEntityID = Scalar Operator([AdventureWorks2012]. [Person].[Person].[BusinessEntityID] as [p].[BusinessEntityID])	

What columns are returned?

Are these needed in the query?
Can they be INCLUDED in an index?
Should a different index be clustered?



Lookup operators

Key Lookup Properties (popup)

Key Lookup (Clustered)	
Uses a supplied clustering key to lookup on a table that has a clustered index.	
Physical Operation	Key Lookup
Logical Operation	Key Lookup
Actual Execution Mode	Row
Estimated Execution Mode	Row
Storage	RowStore
Number of Rows Read	46
Actual Number of Rows for All Executions	8
Actual Number of Batches	0
Estimated Operator Cost	0,11109 (50%)
Estimated I/O Cost	0,003125
Estimated CPU Cost	0,0001581
Estimated Subtree Cost	0,11109
Number of Executions	46
Estimated Number of Executions	34,9278
Estimated Number of Rows for All Executions	617,3383867
Estimated Number of Rows Per Execution	17,6747
Estimated Row Size	11 B
Actual Rebinds	0
Actual Rewinds	0
Ordered	True
Node ID	4
Predicate	[AdventureWorks2012].[Person].[Person].[EmailPromotion] as [p]. [EmailPromotion]=(1)
Object	[AdventureWorks2012].[Person].[Person]. [PK_Person_BusinessEntityID] [p]
Seek Predicates	Seek Keys[1]: Prefix: [AdventureWorks2012].[Person]. [Person].BusinessEntityID = Scalar Operator ([AdventureWorks2012].[Person].[Person].[BusinessEntityID] as [p]. [BusinessEntityID])

Estimated number of rows is wrong when a predicate is pushed into a lookup!!!

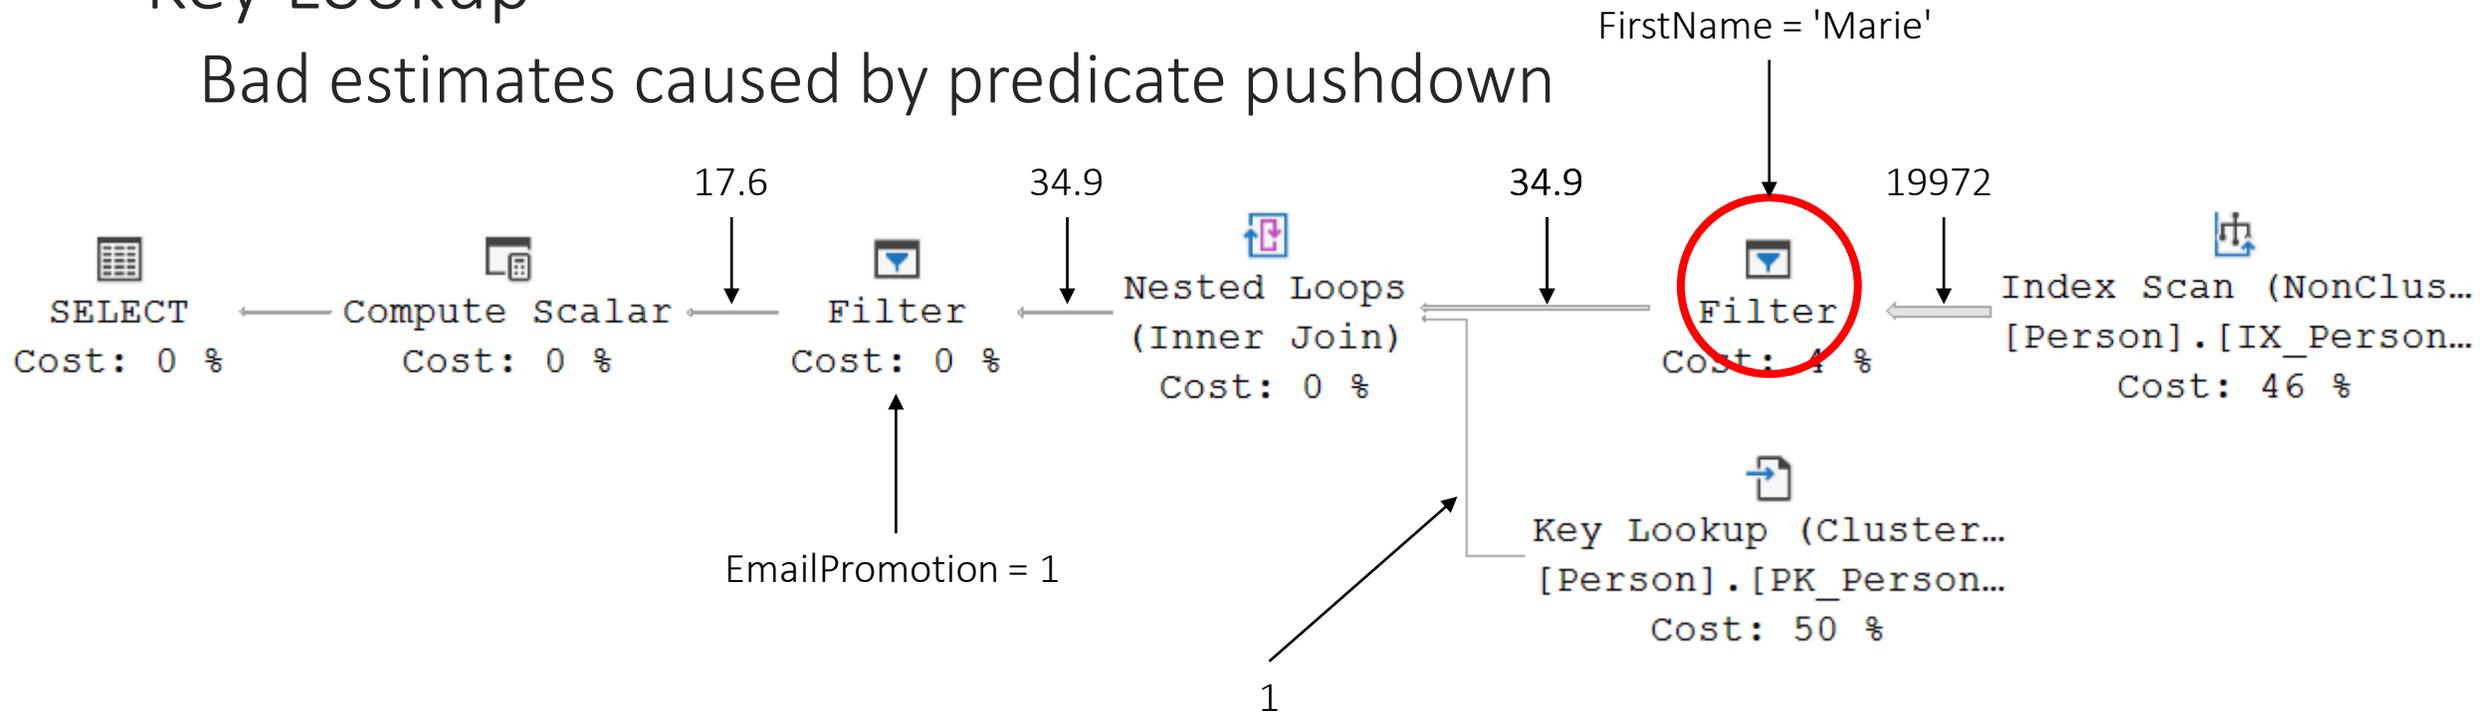
Additional filter pushed into lookup

(Clustered)

Lookup operators

Key Lookup

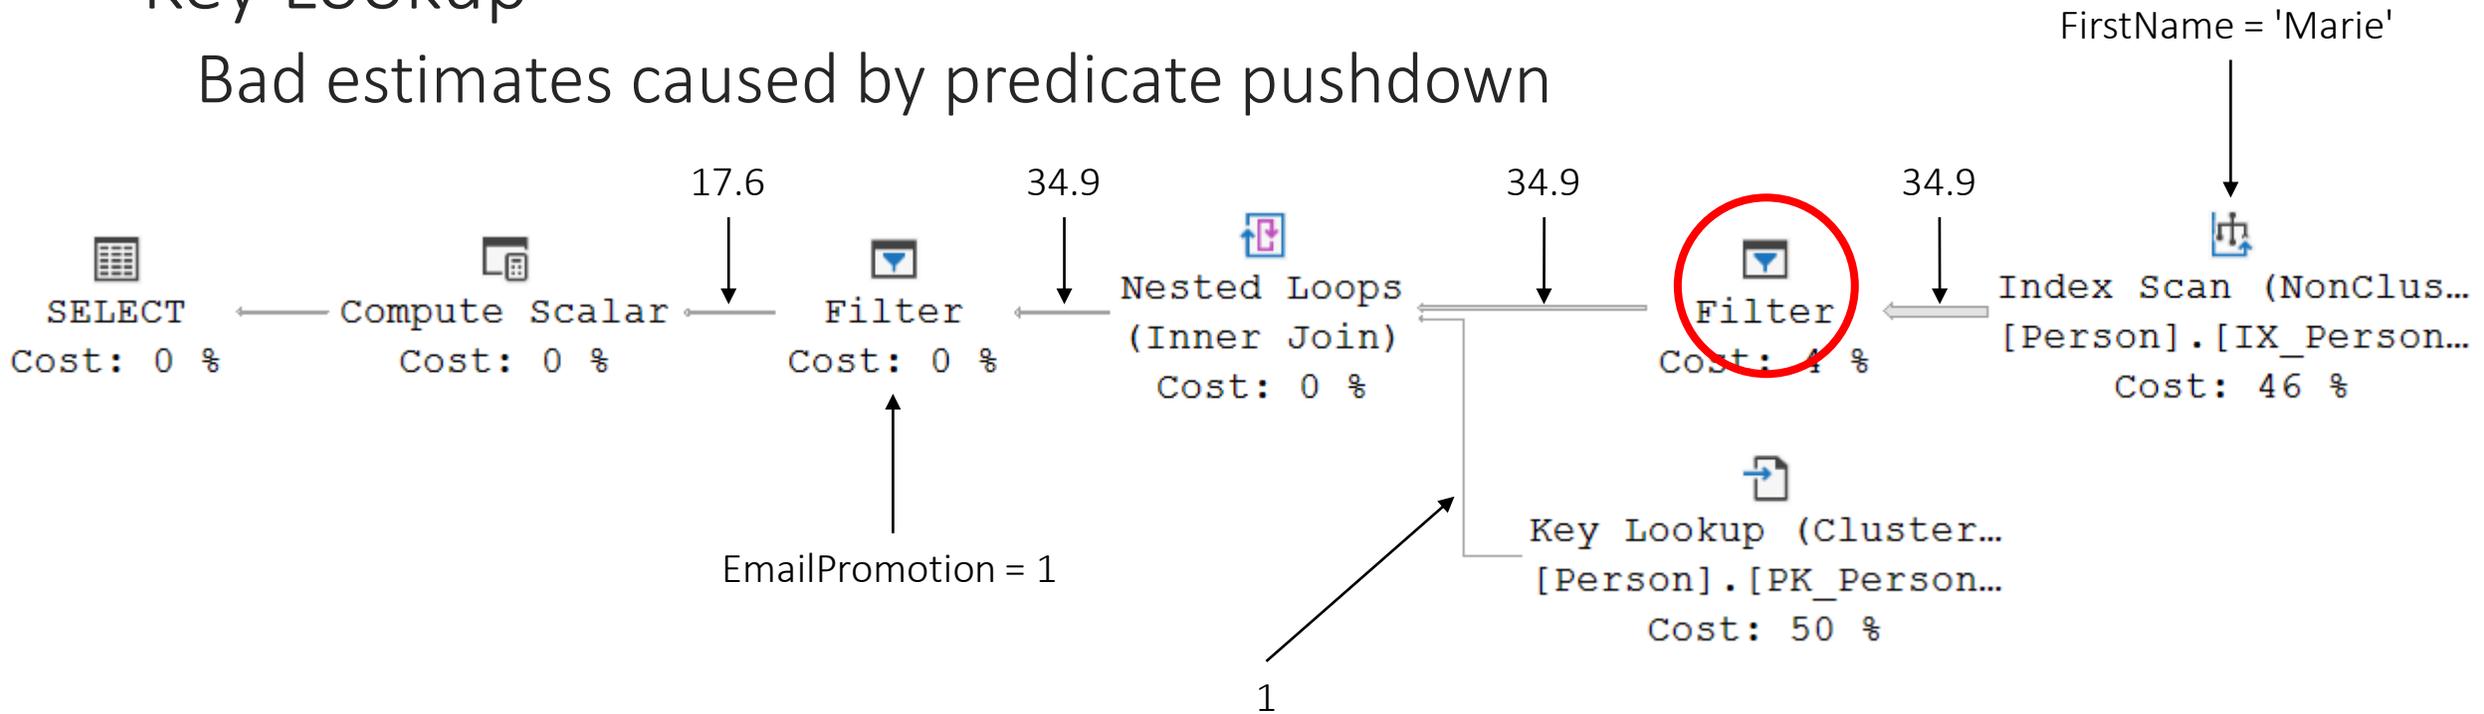
Bad estimates caused by predicate pushdown



Lookup operators

Key Lookup

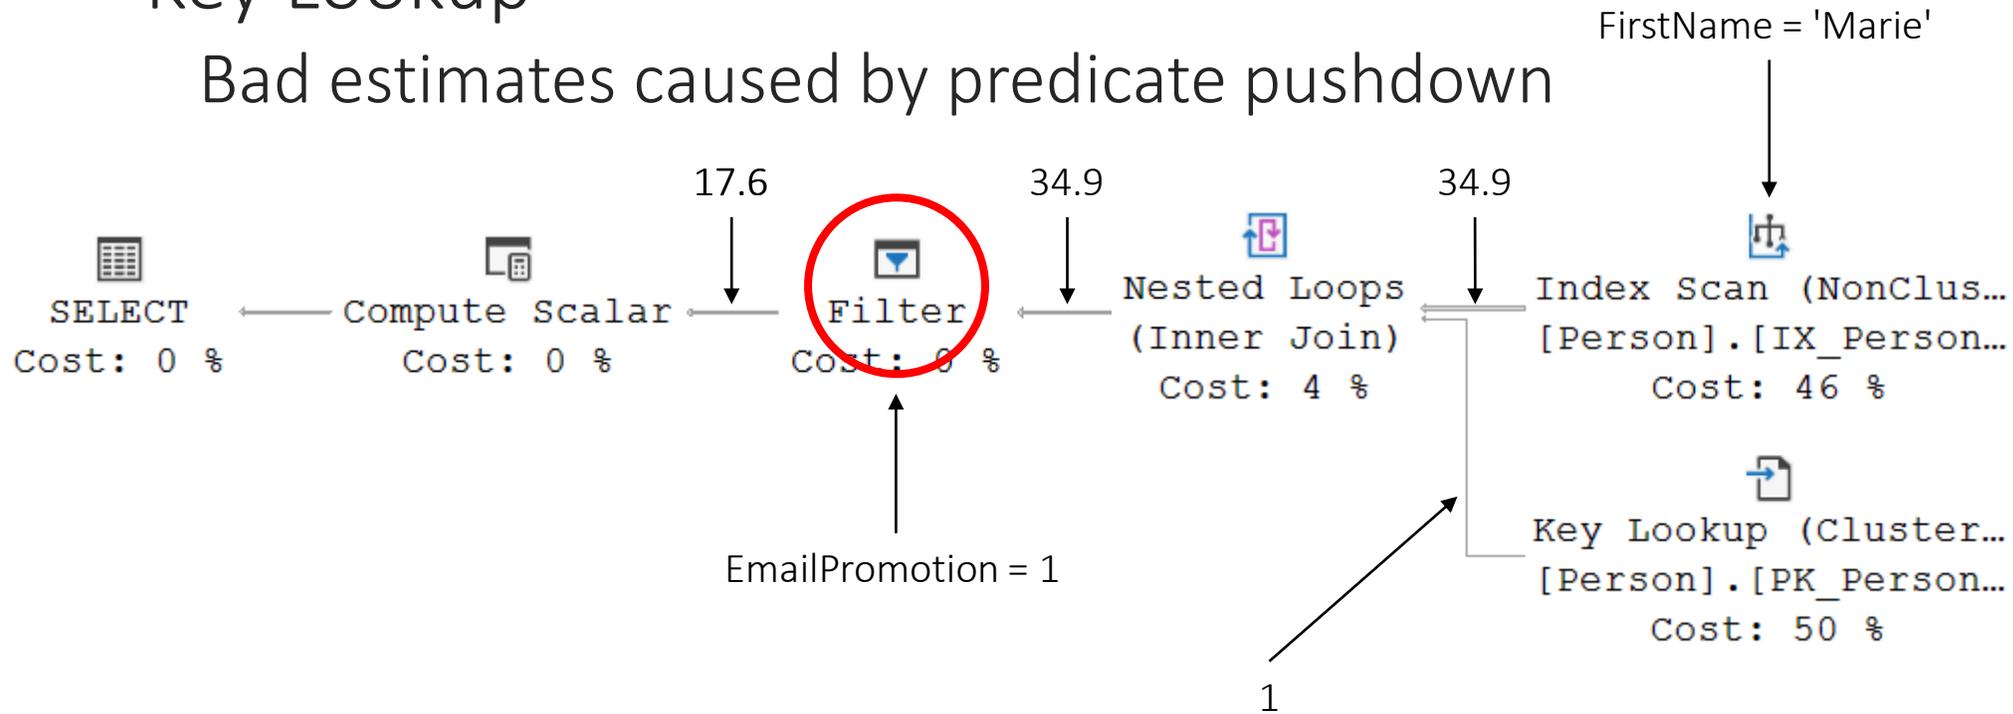
Bad estimates caused by predicate pushdown



Lookup operators

Key Lookup

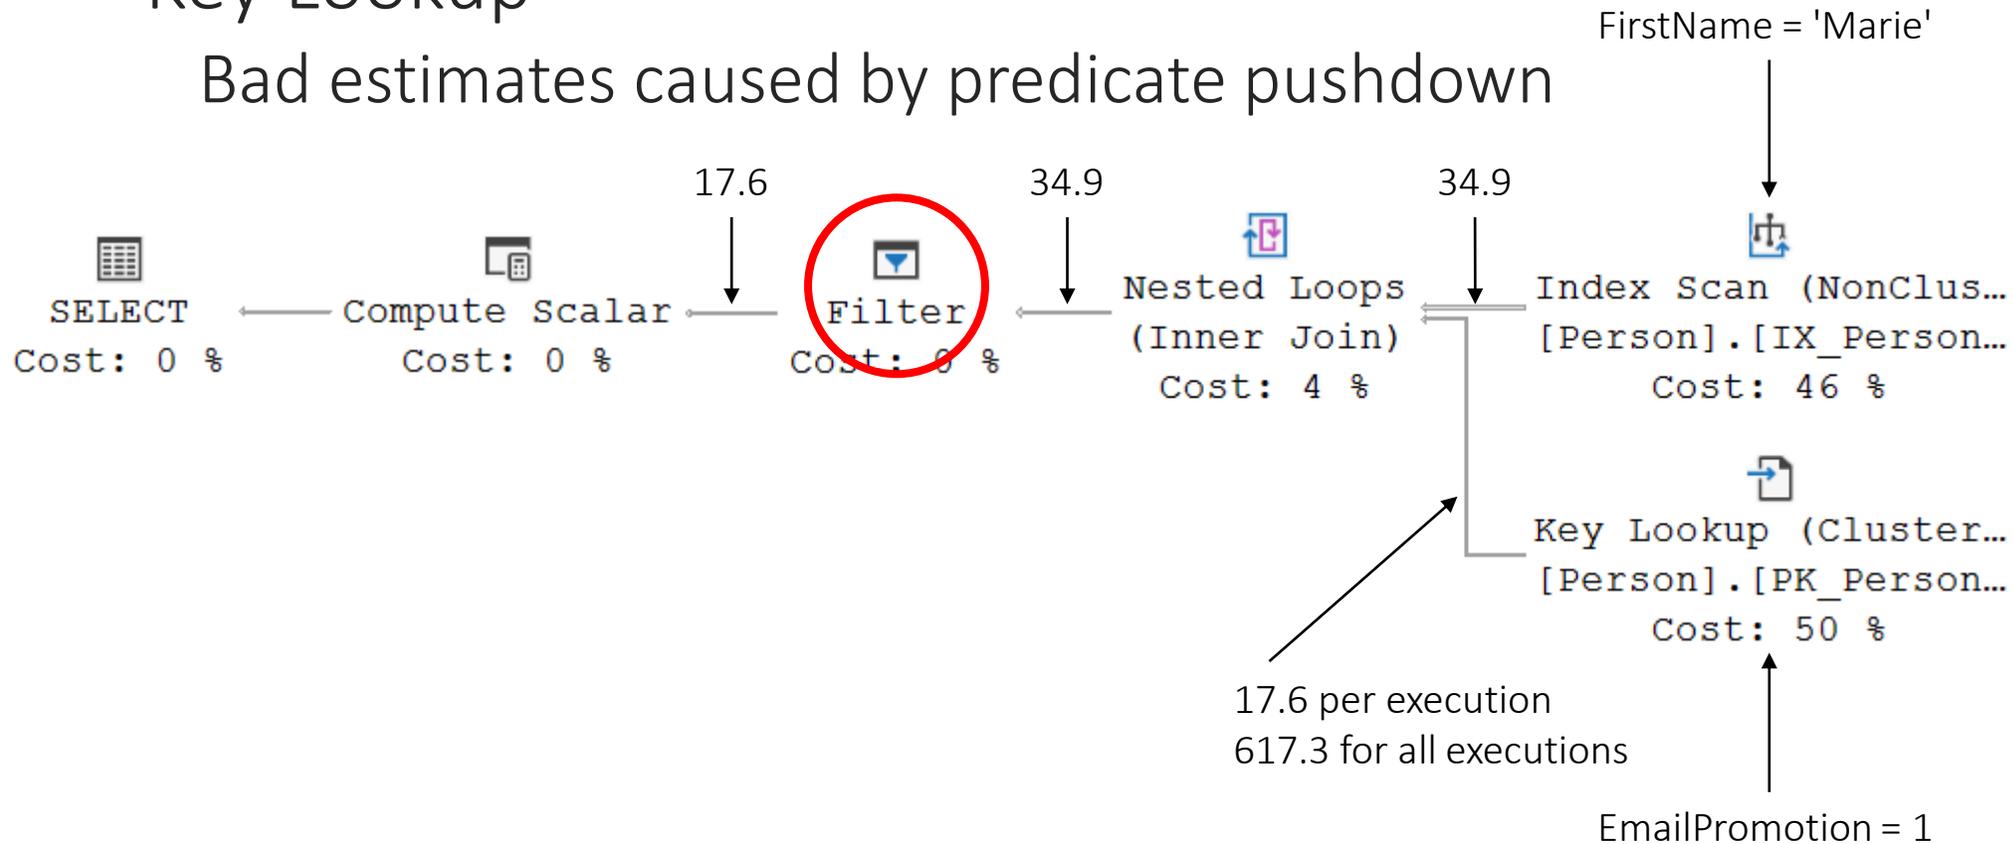
Bad estimates caused by predicate pushdown



Lookup operators

Key Lookup

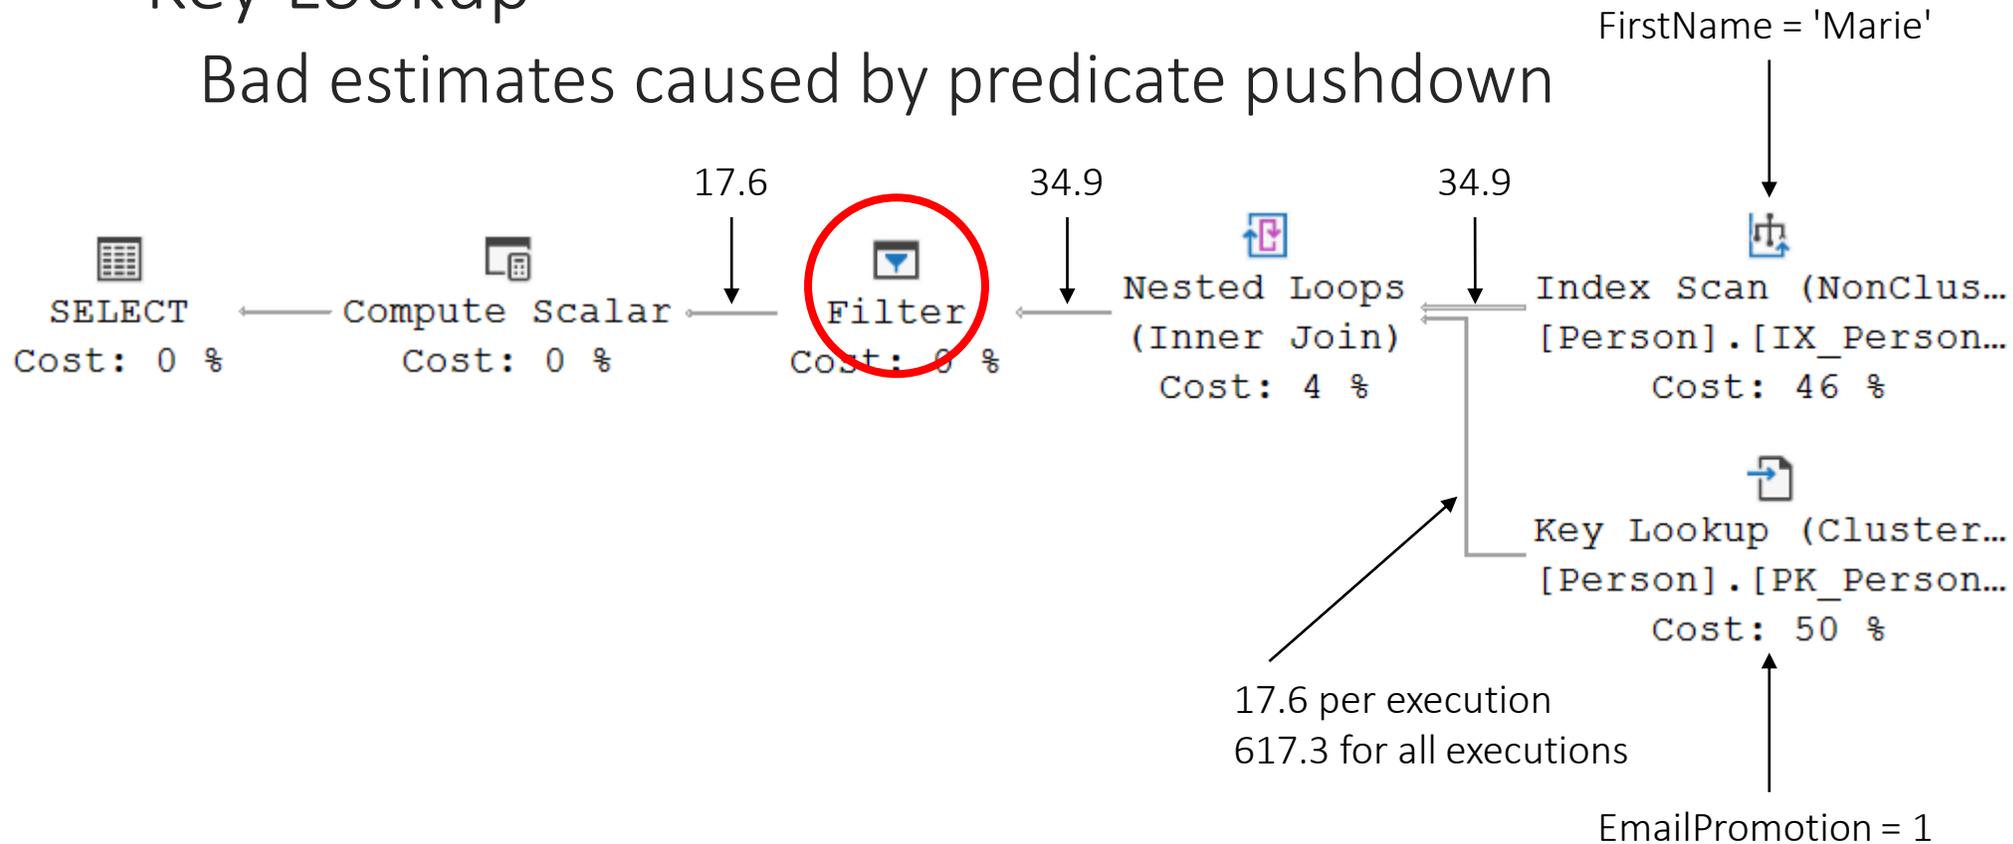
Bad estimates caused by predicate pushdown



Lookup operators

Key Lookup

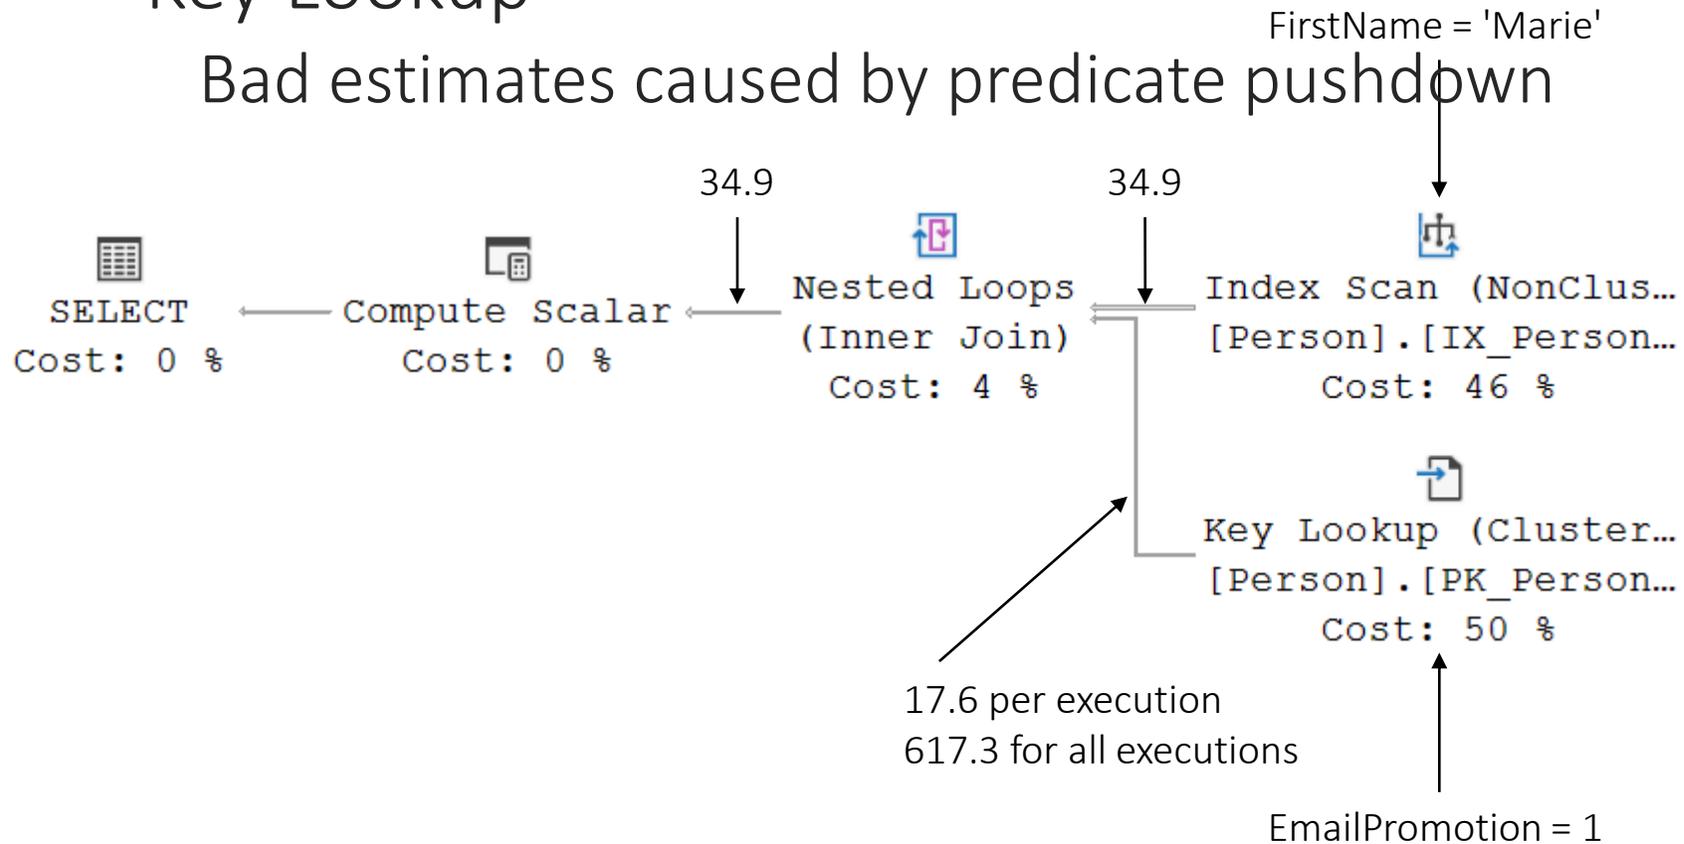
Bad estimates caused by predicate pushdown



Lookup operators

Key Lookup

Bad estimates caused by predicate pushdown



Lookup operators

Key Lookup Properties (popup)

Key Lookup (Clustered)	
Uses a supplied clustering key to lookup on a table that has a clustered index.	
Physical Operation	Key Lookup
Logical Operation	Key Lookup
Actual Execution Mode	Row
Estimated Execution Mode	Row
Storage	RowStore
Number of Rows Read	46
Actual Number of Rows for All Executions	8
Actual Number of Batches	0
Estimated Operator Cost	0,11109 (50%)
Estimated I/O Cost	0,003125
Estimated CPU Cost	0,0001581
Estimated Subtree Cost	0,11109
Number of Executions	46
Estimated Number of Executions	34,9278
Estimated Number of Rows for All Executions	617,3383867
Estimated Number of Rows Per Execution	17,6747
Estimated Row Size	11 B
Actual Rebinds	0
Actual Rewinds	0
Ordered	True
Node ID	4
Predicate	[AdventureWorks2012].[Person].[Person].[EmailPromotion] as [p]. [EmailPromotion]=(1)
Object	[AdventureWorks2012].[Person].[Person]. [PK_Person_BusinessEntityID] [p]
Seek Predicates	Seek Keys[1]: Prefix: [AdventureWorks2012].[Person]. [Person].BusinessEntityID = Scalar Operator ([AdventureWorks2012].[Person].[Person].[BusinessEntityID] as [p]. [BusinessEntityID])

Estimated number of rows is wrong when a predicate is pushed into a lookup!!!

Additional filter pushed into lookup

Don't even try!!
???

(Clustered)

Lookup operators

RID Lookup

Same *function* as Key Lookup

Targets heap, not clustered index



RID Lookup (Heap)

RID Lookup

SELECT
Cost: 0 %

Nested Loops
(Inner Join)
Cost: 0 %

Index Seek (NonClus...
[PersonsHeap].[ix_P...]
Cost: 49 %

RID Lookup (Heap)
[PersonsHeap]
Cost: 51 %

```
SELECT FirstName,
       Other
FROM   dbo.PersonsHeap
WHERE  LastName = 'Visser';
```

Page type: Index
prev next

LName	FName	page
Visser	Klaas	xxxx

Page type: Index
prev next

LName	FName	page
Jansen	Jan	xxxx

Page type: Index
prev next

LName	FName	page
Visser	Klaas	xxxx

Page type: Index
prev next

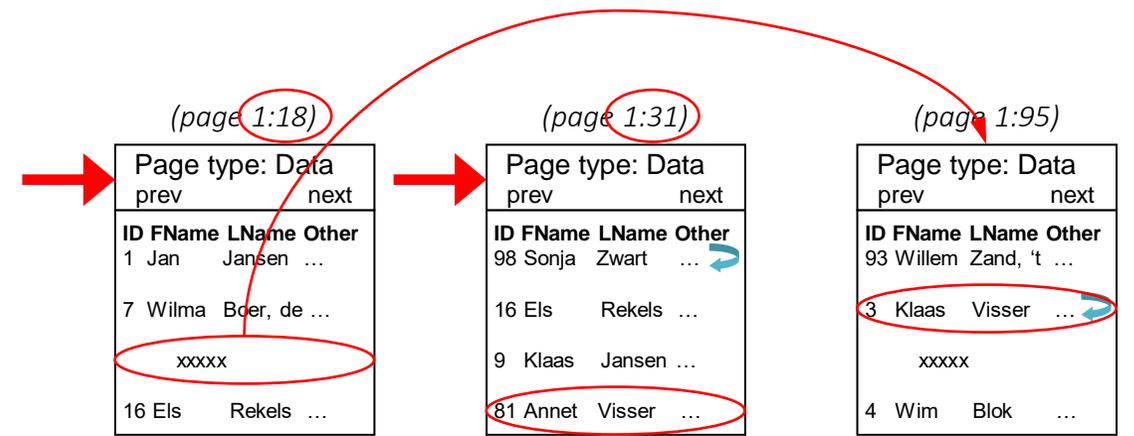
LName	FName	RID
Blok	Wim	1,95,4
Boer, de Wilma		1,18,2
Jansen	Klaas	1,31,3

Page type: Index
prev next

LName	FName	RID
Jansen	Jan	1,18,1
Rekels	Els	1,31,2
Rekels	Els	1,18,4
Visser	Annet	1,31,4

Page type: Index
prev next

LName	FName	RID
Visser	Klaas	1,18,3
Zandt, 't Willem		1,95,1
Zwart	Sonja	1,95,4



Lookup operators

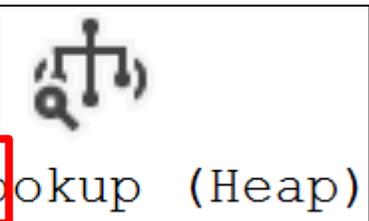
RID Lookup Properties (popup)

Actual and estimated row count

Pushed down filter
(causes wrong values in estimates)

Seek Predicates:
always on Bmknⁿnⁿ
(which is the RID)

RID Lookup (Heap)	
RID Lookup	
Physical Operation	RID Lookup
Logical Operation	RID Lookup
Actual Execution Mode	Row
Estimated Execution Mode	Row
Storage	RowStore
Number of Rows Read	5
Actual Number of Rows for All Executions	5
Actual Number of Batches	0
Estimated Operator Cost	0,0132665 (80%)
Estimated I/O Cost	0,003125
Estimated Subtree Cost	0,0132665
Estimated CPU Cost	0,0001581
Estimated Number of Executions	5
Number of Executions	5
Estimated Number of Rows for All Executions	25
Estimated Number of Rows Per Execution	5
Estimated Row Size	277 B
Actual Rebinds	0
Actual Rewinds	0
Ordered	True
Node ID	3
Predicate	[AdventureWorks2017].[dbo].[DatabaseLog].[DatabaseUser]=N'dbo'
Object	[AdventureWorks2017].[dbo].[DatabaseLog]
Output List	[AdventureWorks2017].[dbo].[DatabaseLog].PostTime; [AdventureWorks2017].[dbo].[DatabaseLog].DatabaseUser; [AdventureWorks2017].[dbo].[DatabaseLog].Event
Seek Predicates	Seek Keys[1]: Prefix: Bmk1000 = Scalar Operator([Bmk1000])



Index tuning

Key Lookup and RID Lookup

Indicates chosen nonclustered index does not cover the query

Some columns are needed, but not available in that index

Missing columns can be found in Column List

Index tuning

Key Lookup and RID Lookup

Indicates chosen nonclustered index does not cover the query

Some columns are needed, but not available in that index

Missing columns can be found in Column List and Predicate

Index tuning

Key Lookup and RID Lookup

Indicates chosen nonclustered index does not cover the query

Can be “fixed” by adding those columns

E.g. through the INCLUDE keyword

No need to do this for each and every lookup operator

Try to optimize as many queries as possible with as little overhead as possible

Index tuning

Key Lookup and RID Lookup

Indicates chosen nonclustered index does not cover the query

Can be “fixed” by adding those columns

Clustered index choice plays a large role

- Narrow

- Unique

- Static

- Ever increasing

- Used for filtering and sorting

- Frequently returned

Index tuning

Key Lookup and RID Lookup

Indicates chosen nonclustered index does not cover the query

Can be “fixed” by adding those columns

Clustered index choice plays a large role

Narrow

Unique

Static

Ever increasing

Used for filtering and sorting

Frequently returned

Few columns

Data types that use few bytes

Storage reduction for all nonclustered indexes

More rows in the buffer pool

Index tuning

Key Lookup and RID Lookup

Indicates chosen nonclustered index does not cover the query

Can be “fixed” by adding those columns

Clustered index choice plays a large role

Narrow

Unique

Static

Ever increasing

Used for filtering and sorting

Frequently returned

SQL Server adds 8-byte uniquifier to nonunique indexes
Makes the index less narrow

Index tuning

Key Lookup and RID Lookup

Indicates chosen nonclustered index does not cover the query

Can be “fixed” by adding those columns

Clustered index choice plays a large role

Narrow

Unique

Static

Ever increasing

Used for filtering and sorting

Frequently returned

Update affects all nonclustered indexes

Move data in clustered index

Often used for foreign key

Index tuning

Key Lookup and RID Lookup

Indicates chosen nonclustered index does not cover the query

Can be “fixed” by adding those columns

Clustered index choice plays a large role

Narrow

Unique

Static

Ever increasing

Used for filtering and sorting

Frequently returned

Each insert in “last” page

Reduces page splits

May cause insert hotspot

Index tuning

Key Lookup and RID Lookup

Indicates chosen nonclustered index does not cover the query

Can be “fixed” by adding those columns

Clustered index choice plays a large role

Narrow

Unique

Static

Ever increasing

Used for filtering and sorting

Frequently returned

Filter benefits from seek

Ordering benefits from scan (*Ordered* = True)

When done on clustered index, no lookup needed

For nonclustered index, benefit offset by cost of lookup

Index tuning

Key Lookup and RID Lookup

Indicates chosen nonclustered index does not cover the query

Can be “fixed” by adding those columns

Clustered index choice plays a large role

Narrow

Unique

Static

Ever increasing

Used for filtering and sorting

Frequently returned

Less lookups needed

Clustered key columns are in all nonclustered indexes

Index tuning

Key Lookup and RID Lookup

Indicates chosen nonclustered index does not cover the query

Can be “fixed” by adding those columns

Clustered index choice plays a large role

Changing clustered index on existing database is painful

All nonclustered indexes need to be updated too

Try to pick correct clustered index in initial design

No simple “one size fits all” solution

IDENTITY as (clustered) primary key for all tables?

Can be VERY bad for performance!!!

Summary

Lookup operators

Key Lookup

Read one row, based on clustered key passed in

RID Lookup

Read one row, based on RID value passed in

Pushed down filter affects *Estimated Number of Rows* properties

Not always bad

When needed, make indexes covering or revise clustered index choice

Next chapters

Chapter 5: Special scans

Constant Scan

Inserted Scan

Deleted Scan