

SQLServerFast.com

Execution Plan Video Training

Block 2: Reading data

Level: Advanced

Chapter 4: Reading data in parallel or batch mode

Serial versus parallel scan

Scan operators

Serial execution plan

Read data

Return data to parent



Clustered Index Scan (Clustered)

Serial versus parallel scan

Scan operators

Serial execution plan

Parallel execution plan

One copy of the operator per thread

Threads run isolated

But we do not want every copy of the scan to read the same data!



Serial versus parallel scan

Scan operators

Serial execution plan

Parallel execution plan

One copy of the operator per thread

Threads run isolated

Parallel Page Supplier (PPS)

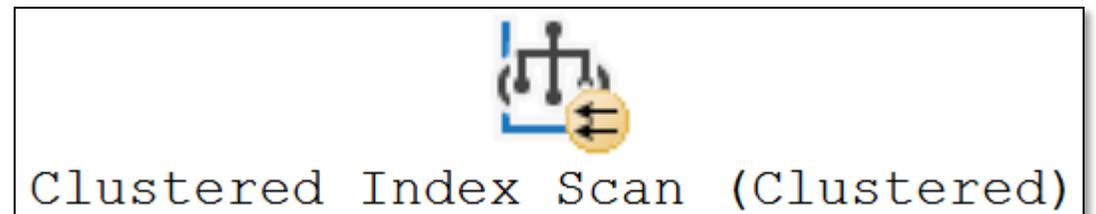
Sits in between scan operators and storage engine

Delivers data to each thread on request

Minimal skew

No control which row goes to which thread

May necessitate Parallelism operator



Serial versus parallel scan

Scan operators

Serial execution plan

Parallel execution plan

One copy of the operator per thread

Threads run isolated

Parallel Page Supplier (PPS)

Sits in between scan operators
and storage engine

Delivers data to each thread on request

One full page at a time

Exception: Large table with selective predicate



Clustered Index Scan (Clustered)



Clustered Index Scan (Clustered)



Clustered Index Scan (Clustered)



Clustered Index Scan (Clustered)

Serial versus parallel scan

Scan operators

Seek operators

~~Singleton lookup~~

Range seek

Also uses Parallel Page Supplier when running on multiple threads

No control which row goes to which thread

May necessitate Parallelism operator



Clustered Index Seek (Clustered)

Batch mode plans

SQL Server 2012

Columnstore indexes

Read only

No deltastore rowgroup

No deleted bitmap

Batch mode plans

SQL Server 2012

Columnstore indexes

Read only

All data in compressed rowgroups

Structure of batch is similar to structure of compressed rowgroup

Batch size is smaller than rowgroup size

Columnstore Index Scan “breaks” rowgroup into batches

Optimized to benefit from similar structure



Columnstore Index Scan (NonClustere...

Batch mode plans

SQL Server 2012

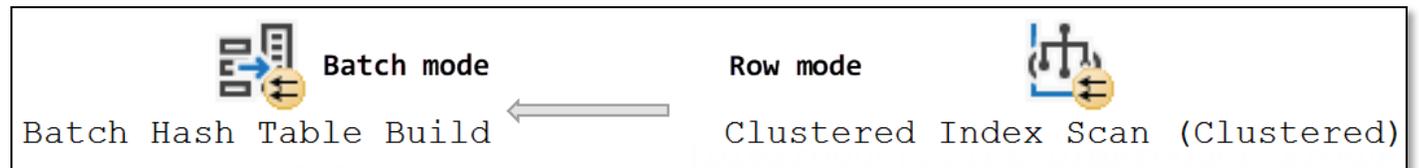
Columnstore indexes

Rowstore indexes

Can only be read in row mode

Execution Mode Adapter

Invisible in execution plan



Batch mode plans

SQL Server 2012

Columnstore indexes

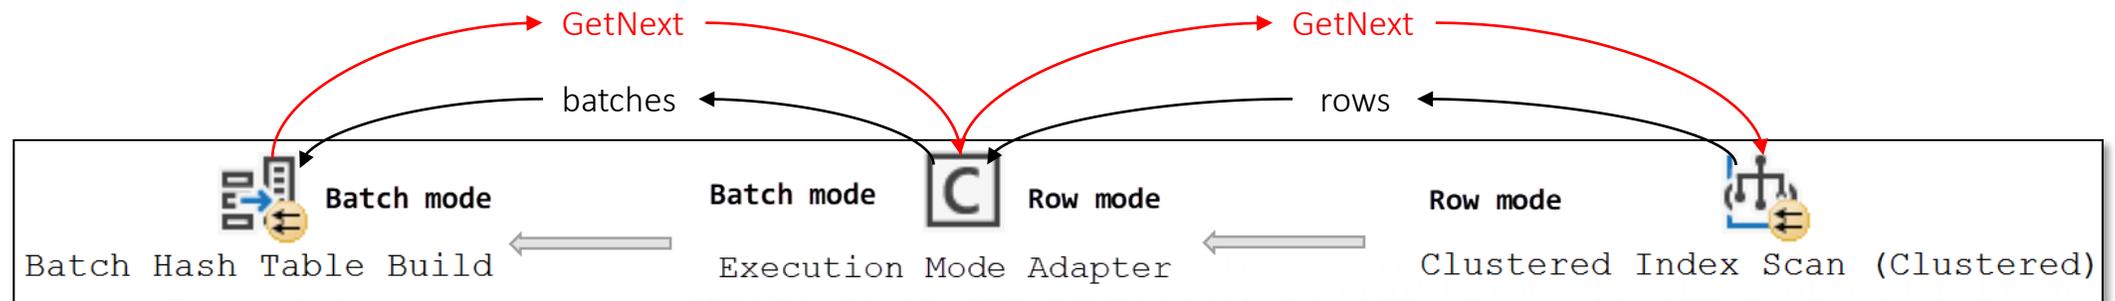
Rowstore indexes

Can only be read in row mode

Execution Mode Adapter

Invisible in execution plan

Converts between row mode and batch mode



Batch mode plans

SQL Server 2012

SQL Server 2014 – 2017

Columnstore indexes

Deleted bitmap

Remove logically deleted data before “chunking” rowgroup into batches



Columnstore Index Scan (NonClustere...

Batch mode plans

SQL Server 2012

SQL Server 2014 – 2017

Columnstore indexes

Deleted bitmap

Deltastore rowgroup

Rowstore B-tree format

Columnstore Index Scan has to convert rows to batches



Columnstore Index Scan (NonClustered...)

Batch mode plans

SQL Server 2012

SQL Server 2014 – 2017

Columnstore indexes

Rowstore indexes

Can only be read in row mode

Execution Mode Adapter

Invisible in execution plan

Converts between row mode and batch mode

Batch mode plans

SQL Server 2012

SQL Server 2014 – 2017

SQL Server 2019

- Columnstore indexes

- Rowstore indexes

 - Can be read in batch mode

 - Batch mode on rowstore

 - More efficient than Execution Mode Adapter

Summary

Reading data in parallel execution plans

Reading data in batch mode execution plans

Summary

Reading data in parallel execution plans

- Scan operators

- Seek operators (range seek only)

- Parallel Page Supplier

Summary

Reading data in parallel execution plans

Reading data in batch mode execution plans

Up to SQL Server 2017

- Reading from columnstore index supported in row or batch mode

- Reading from rowstore index / heap only supported in row mode

Execution Mode Adapter

SQL Server 2019 +

- Reading from rowstore index / heap now also supported in batch mode

Next chapters

Chapter 5: Assorted read optimizations

- Read-ahead reading

- Advanced scan (aka “merry-go-round scan”)

- Dynamic range seeks

- Partitioning