

Clustered Columnstore - Introduction

Niko Neugebauer



#199 | BULGARIA 2013

Sponsors

In partnership with



Niko Neugebauer

Microsoft Data Platform Professional

OH22 (<http://www.oh22.net>)

15+ years in IT

SQL Server MVP

Founder of **3** Portuguese PASS Chapters

Blog: <http://www.nikoport.com>


Twitter: [@NikoNeugebauer](https://twitter.com/NikoNeugebauer)

LinkedIn: <http://pt.linkedin.com/in/webcaravela>

Email: info@webcaravela.com




Explore Everything PASS Has to Offer




Free SQL Server and BI Web Events



Free 1-day Training Events



Regional Event



This is Community



Business Analytics Training



Local User Groups Around the World



Session Recordings



PASS Newsletter



Free Online Technical Training

Today's program:

- Principles of the Columnstore Indexes & *Batch* Mode
- History
- Columnstore in SQL Server 2012
- Clustered Columnstore in SQL Server 2014
- *Batch* Mode improvements for SQL Server 2014

Principles of the Columnstore Indexes:


- Using *a-kind-of-in-memory*[™] engine
- X-Velocity aka Vertipaq (PowerPivot, SSAS Tabular)
- Vertically separated, extremely compressed & tuned for processing large volumes of data

Principles of the Columnstore Indexes

topics:

- Clustered & Nonclustered Indexes
- RowStore vs ColumnStore
- Columnstore Index creation sequence
- Not-suitable for OLTP, unless you are processing whole table and not altering data

RowStore vs ColumnStore



RowStore vs ColumnStore

ID	First Name	Last Name	Salary
1	Jody	Philips	43.03
2	Mark	Johnson	37.08
3	Matt	Markensen	16.81
4	Gail	Lindberg	24.90

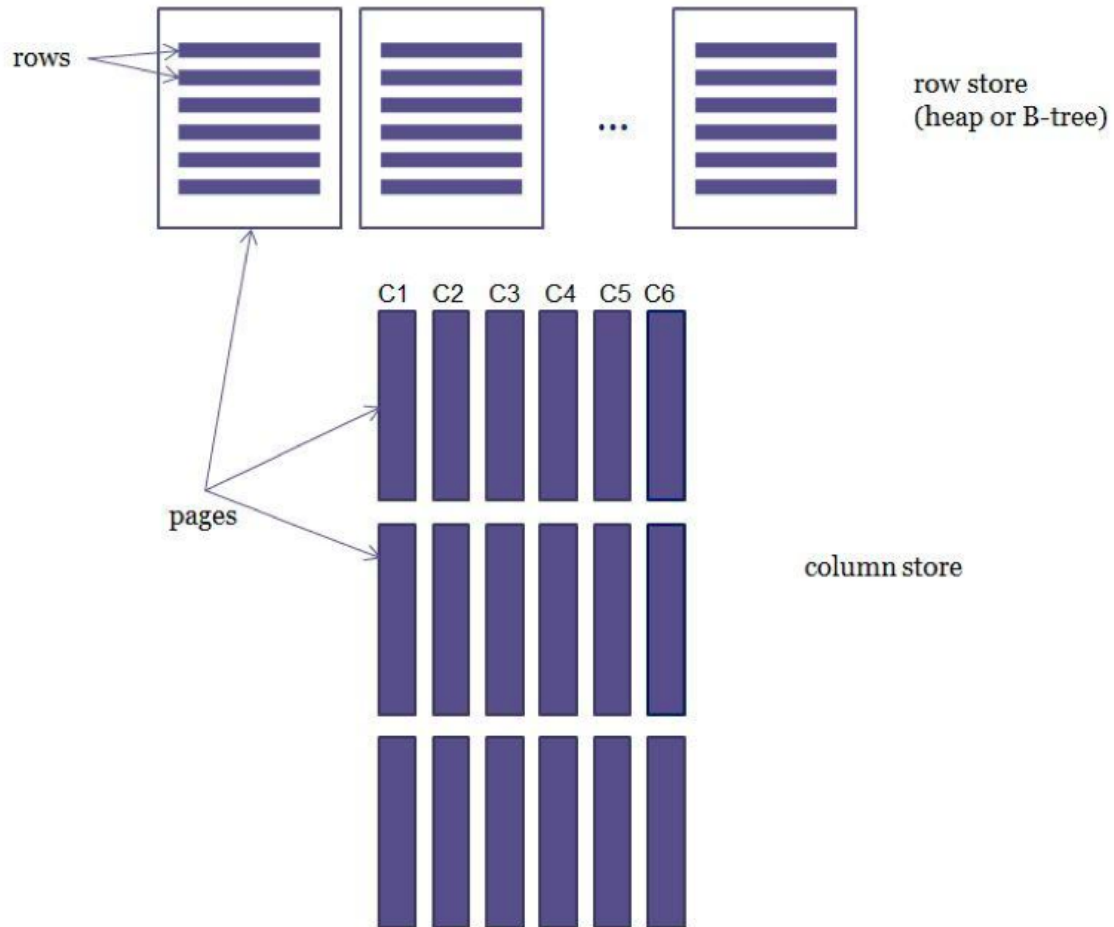
Row Store Index:

1, Jody, Philips, 43.04; 2, Mark, Johnson, 37.08; ...

Column Store Index:

1, 2, 3, 4; Jody, Mark, Matt, Gail; ...

RowStore vs ColumnStore



Columnstore Index creation sequence:


1. Run-length compression
2. Dictionary encoding
3. Binary compression

Run-length compression

ID	ID
1	1:2
1	2
2	3:3
3	4
3	1-5:2
3	
4	
1	
5	
1	
5	



Dictionary encoding



Binary Compression

Super-secret compression turning data into LOB's. 😊

Vertipac aka **xVelocity**

Row Groups


~ 1.000.000 Rows are grouped in Segments,
called Row Groups

Data is Encoded & Compressed

Data is Unsorted

Stored as LOB's

When executed in batch mode the Segments



Principles of the *Batch* mode:

- Vector processing
- **1 row vs ~1000 rows**
- Programmers, its like passing **Array vs 1 param**
- Think about it as if it would be a Factory processing vs Manual Processing (19th vs 18th Century)

History of the Columnstore

- Concept exists since 1969 (or since the beginning of Database era)
- Sybase IQ (since 1990s)
- PDW v1 had the first implementation for Microsoft Databases
- SQL Server 2012 introduced Columnstore
- Besides SQL Server there are Vertica, C-Store, Volt, Terradata, ...
- Michael Stonebraker's vision – In 10 years there will be virtually no Rowstore for Datawarehousing

Columnstore in SQL Server 2012

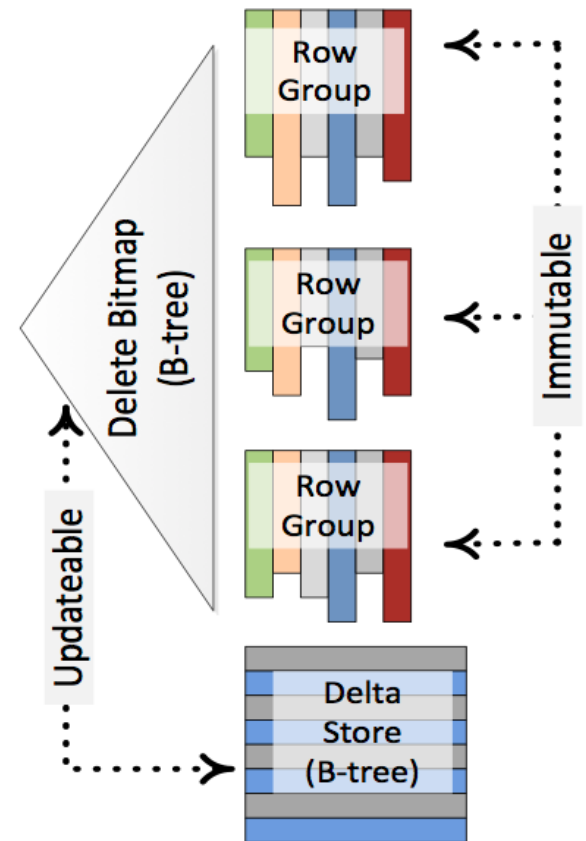
- Only Nonclustered (but supporting multiple indexes)
- Non-Updateable (and no schema-modification)
- Very limited data types (Precision 18 and no LOB's)
- *Batch* mode functions only in certain conditions (no OUTER JOIN, UNION ALL, EXIST, IN, Scalar Aggregates, Distinct Aggregates)

Columnstore in SQL Server 2014

- Clustered & Nonclustered
- Clustered Columnstore are **updateable** & Schema can be modified
- Large range of data types (except LOB)
- Better Memory estimation & consumption
- Batch mode is significantly improved
- NEW! Archival Compression

Columnstore in SQL Server 2014

- Delta-Stores (open & close)
- Deleted Bitmap
- Delete & Update work as a DELETE + INSERT
- Schema modification
- Constraints are supported (Non-Unique)



Important Limitations in SQL Server 2014

- No NonClustered Indexes for Clustered Columnstore
- No ONLINE operations
- Not every operator in Execution plans is supporting *Batch* mode.

Batch Mode Improvements in SQL Server 2014

- Mixed Mode – Row & Batch mode can co-exist
- OUTER JOIN, UNION ALL, EXIST, IN, Scalar Aggregates, Distinct Aggregates
- Some TempDB operations for Columnstore Indexes are running in Batch mode.



Demo Time

Links

- <http://www.nikoport.com/2013/07/05/clustered-columnstore-indexes-part-1-intro/> – My series of blog posts on Clustered Columnstore Indexes
- <http://rusanu.com/2013/06/11/sql-server-clustered-columnstore-indexes-at-teched-2013/> – Remus Rusanu Introduction for Clustered Columnstore
- <http://research.microsoft.com/pubs/193599/Apollo3%20-%20Sigmod%202013%20-%20final.pdf> – White Paper on the Clustered Columnstore



Sponsors

In partnership with

