


# Columnstore Indexes Worst Practices & Less Known Limitations

by Niko Neugebauer



Niko Neugebauer

Microsoft Data Platform Professional

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

Data Platform MVP

Founder of 3 of Portuguese PASS Chapters

Creator of **CISL – Columnstore Indexes Script Library**

(<https://github.com/NikoNeugebauer/CSIL>)

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

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

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


# Today's Agenda is Simple

- Intro
- Overall Recommendation
- Building Columnstore Indexes
- Loading Data into Columnstore Indexes
- Querying the Columnstore Indexes



# Columnstore Indexes are:

- Verytically separated
- Grouped into Segments
- Extremely compressed
- Tuned for processing large volumes of data
  
- A Row Group CAN contain between 1 and 1048576 rows



# Today's Agenda I

Overall:

- Not using the highest available compatibility level

Building

- Cutting on Memory
- Using Strings
- Partitioning
- Compression Delay



# Today's Agenda II

## Loading:

- Not using Bulk Load API (<102.400 rows)
- Identity
- Merge
- Parallel Loading
- SSIS




# Today's Agenda III

## Querying

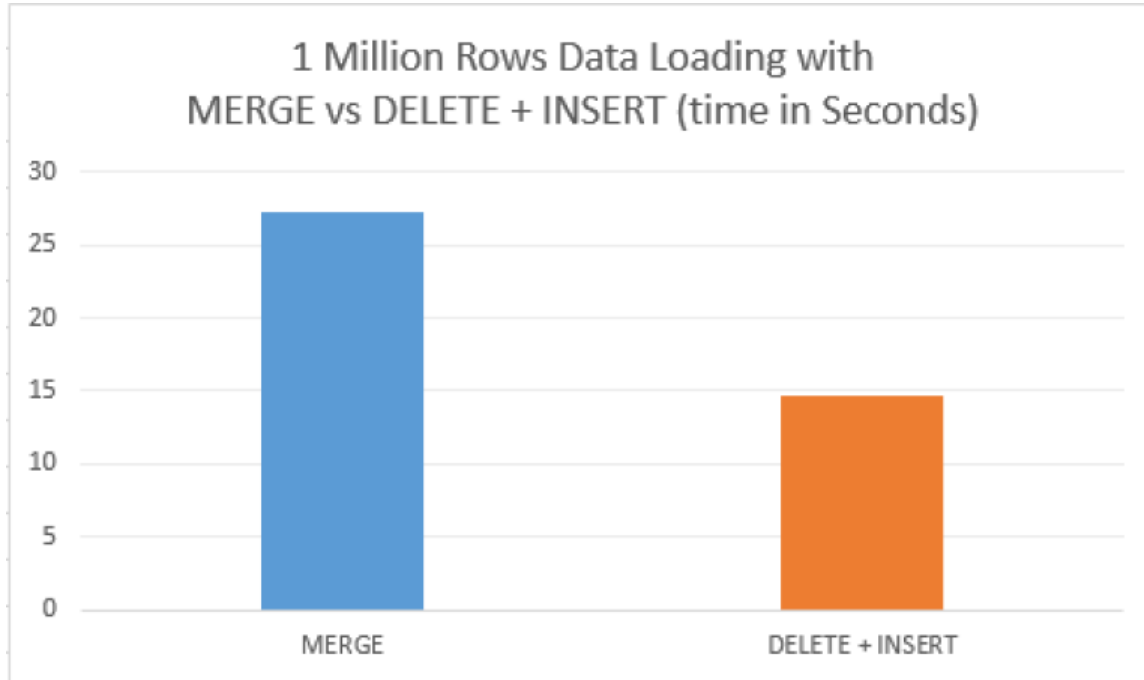
- Casting on the wrong side
- Overusing the NULL Expressions
- Focusing on Cross Apply where it is not needed
- Not using the Advanced Storage Engine Features
- Estimated are similar! Similar How?
- Using Natively Compiled Stored Procedures

# Identity






# Merge




# Scaling Merge ?



# Merge Execution Plans


Query 1: Query cost (relative to the batch): 26%

```
DELETE Target FROM dbo.lineitem_cci as Target WHERE EXISTS ( SELECT 1 FROM dbo.lineitem_cci_stage as Source WHERE Target.L_ORDE
```



Query 2: Query cost (relative to the batch): 74%

```
INSERT INTO dbo.lineitem_cci WITH (TABLOCK) SELECT l_shipdate, l_orderkey, l_discount, l_extendedprice, l_suppkey, l_quantity,
```



Thank you very much!



## Resources:

My Columnstore Blogpost Series (100+):

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

CISL – Open Source Columnstore Library:

<https://github.com/NikoNeugebauer/CISL>