THE SCARY DBA

Query Performance Tuning: Start to Finish

Grant Fritchey

Who?

- Product Evangelist for Red Gate Software
- Microsoft SQL Server MVP
- PASS Chapter President
- Author:
  - SQL Server Execution Plans
  - SQL Server 2008 Query Performance Tuning Distilled

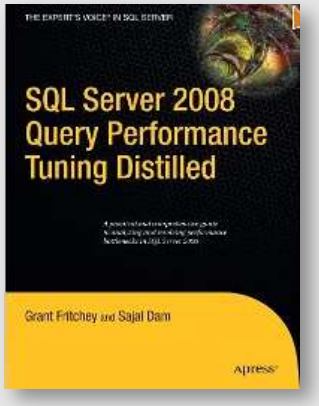

Why Tune Queries?

- Most volatile aspect of a database system
- Subject to changes in data
- Affected by changes in structure
- Impacted by poor coding choices
- Victim of resource contention

Finish Line

- The ability to collect performance metrics on their servers as part of an overall query tuning methodology
- An understanding of how the optimizer works in support of writing better TSQL code as well as troubleshooting poorly performing queries
- The ability to generate, read, and understand execution plans from multiple sources in support of troubleshooting poorly performing queries
- A working knowledge of DMVs that will help them identify and fix performance issues on their servers
- The ability to address common query performance problems

Query Performance Tuning: Start to Finish

OPTIMIZER, STATISTICS, INDEXES & CONSTRAINTS

Finish Line

- The ability to collect performance metrics on their servers as part of an overall query tuning methodology
- An understanding of how the optimizer works in support of writing better TSQL code as well as troubleshooting poorly performing queries
- The ability to generate, read, and understand execution plans from multiple sources in support of troubleshooting poorly performing queries
- A working knowledge of DMVs that will help them identify and fix performance issues on their servers
- The ability to address common query performance problems

Optimizer

- Simply an Amazing piece of Software
- Cost-based
- Not Perfect
- Plan on Helping the Optimizer

Relational Engine

QUERY

Relational Engine

QUERY

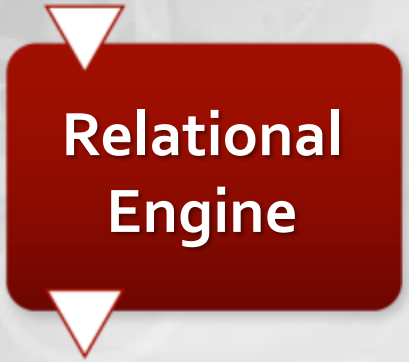

Relational Engine

QUERY

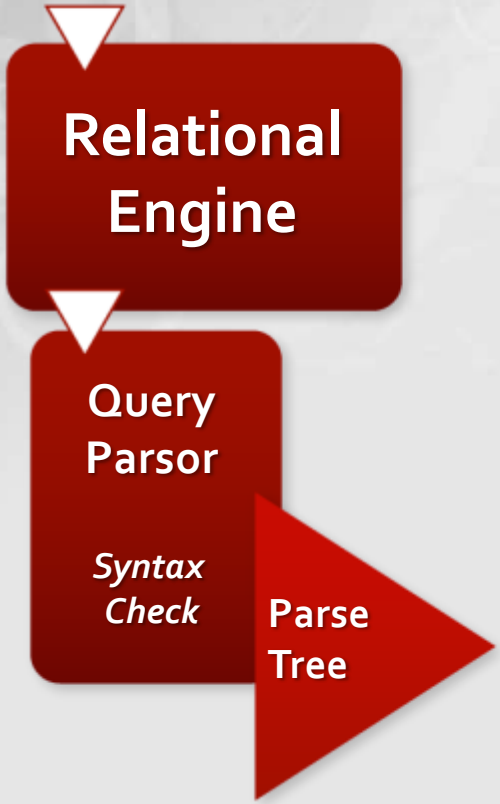

Relational Engine

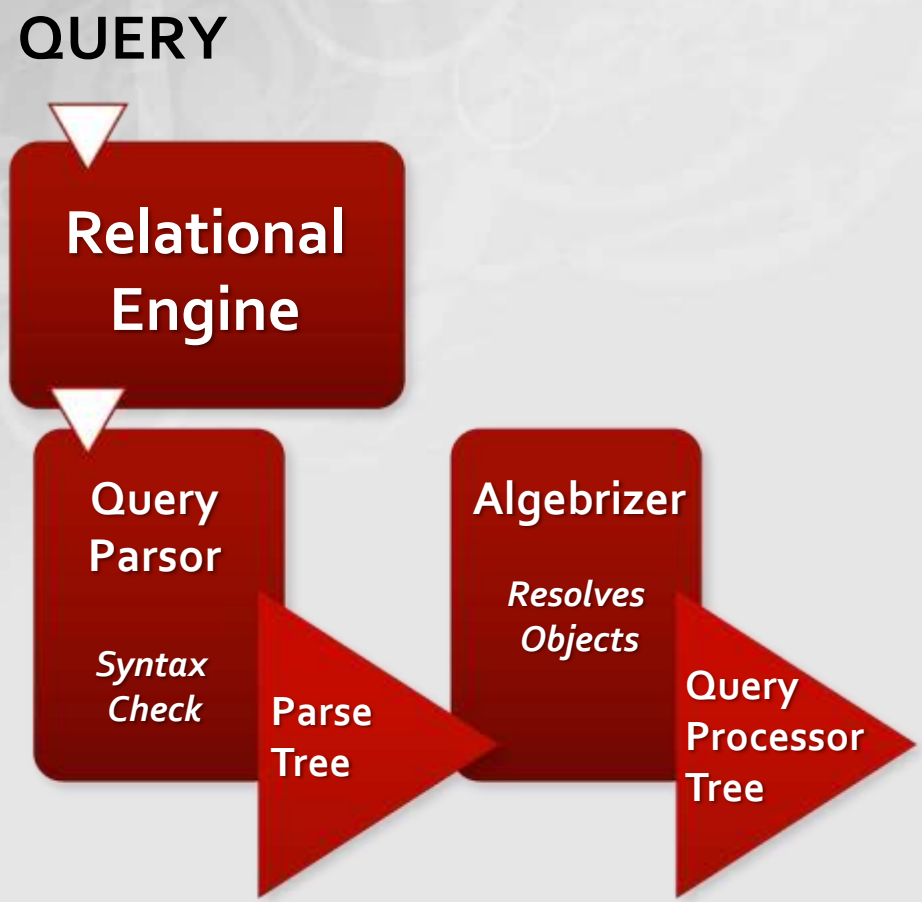

Relational Engine

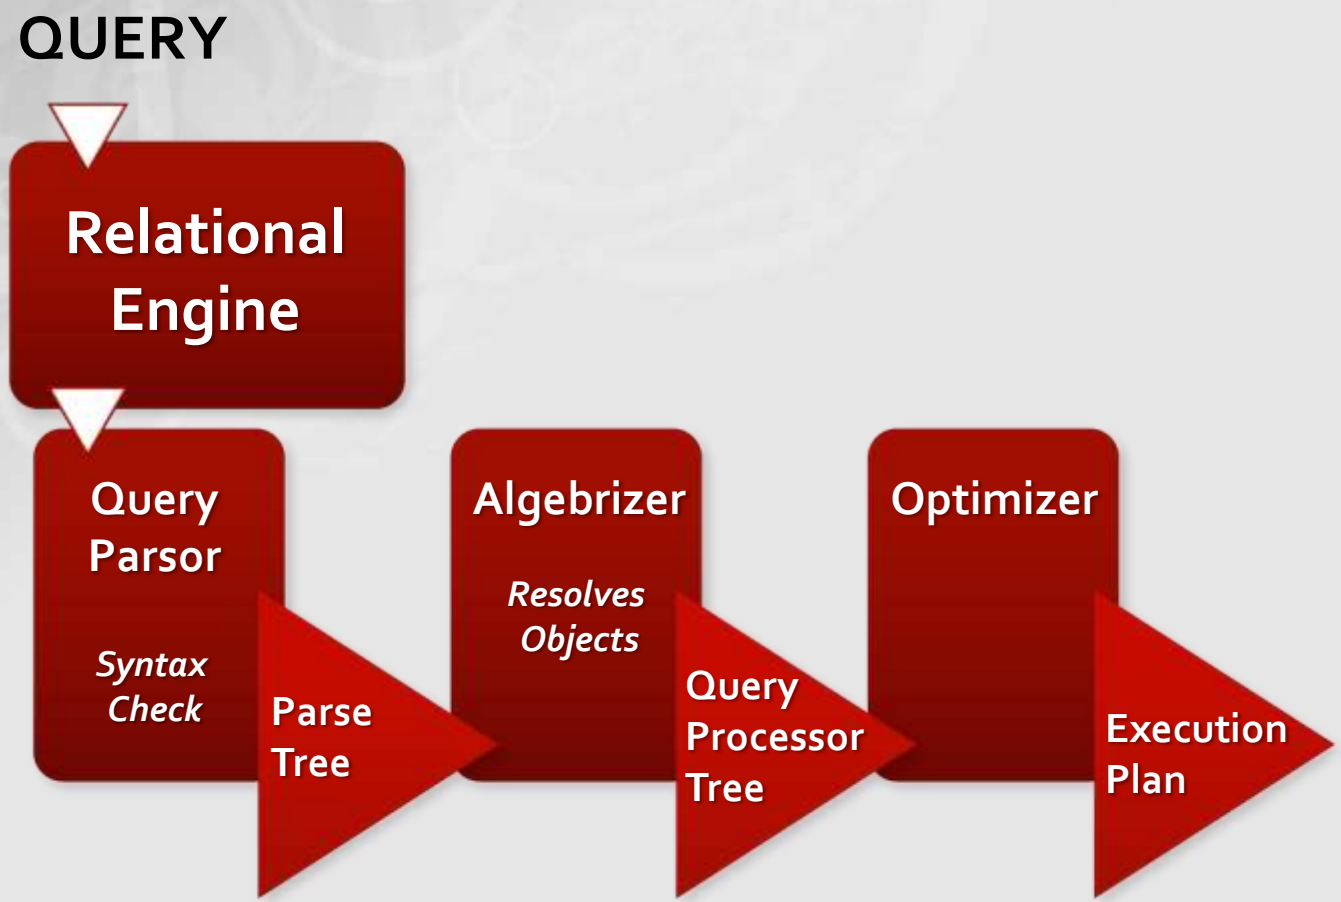

Relational Engine

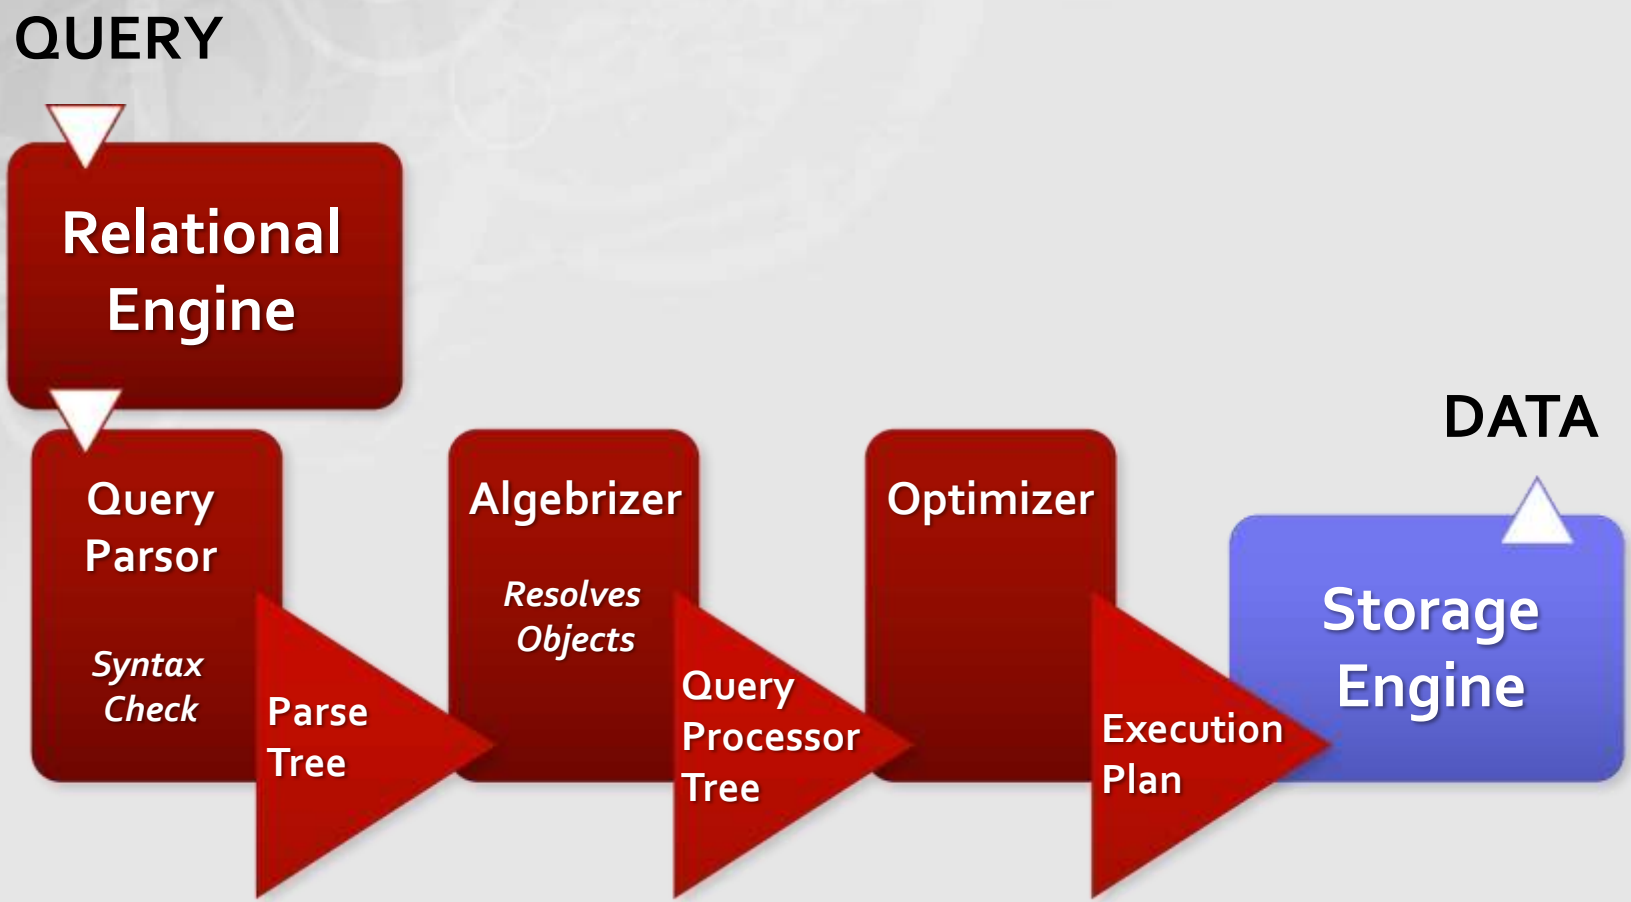

Observing the Optimizer

- Sys.dm_exec_query_optimizer_info
- Execution Plans

Statistics

- Information about the Distribution of the Data
  - Created on Index Keys
  - Created on columns
  - Created manually
- Cardinality
- By Default, Created Automatically
- By Default, Maintained Automatically
- Automatic Maintenance Is Not Enough

Investigating Statistics

- DBCC SHOW_STATISTICS(*table*, *target*)
  - » Header

	Name	Updated	Rows	Rows Sampled	Steps	Density	Average key len...	String Index	Filter Expressi...	Unfiltered Rows
1	IX_TransactionHistoryArchive_ProductID	Jan 19 2011 9:57PM	89253	89253	200	0.04100511	8	NO	NULL	89253

  - » Density

	All density	Average Len...	Columns
1	0.002012072	4	ProductID
2	1.120411E-05	8	ProductID, TransactionID

  - » Histogram

	RANGE_HI_KEY	RANGE_ROWS	EQ_ROWS	DISTINCT_RANGE_ROWS	AVG_RANGE_ROWS
1	1	0	6	0	1
2	3	5	786	1	5
3	316	6	786	1	6
4	324	82	786	7	11.71429
5	327	10	786	2	5
6	328	0	619	0	1
7	329	0	781	0	1
8	331	58	786	1	58
9	350	56	786	10	5.6

Histogram

- 200 steps across the data
- An equal distribution of rows
- Leads to best possible sampling of data
  - » But it's not perfect

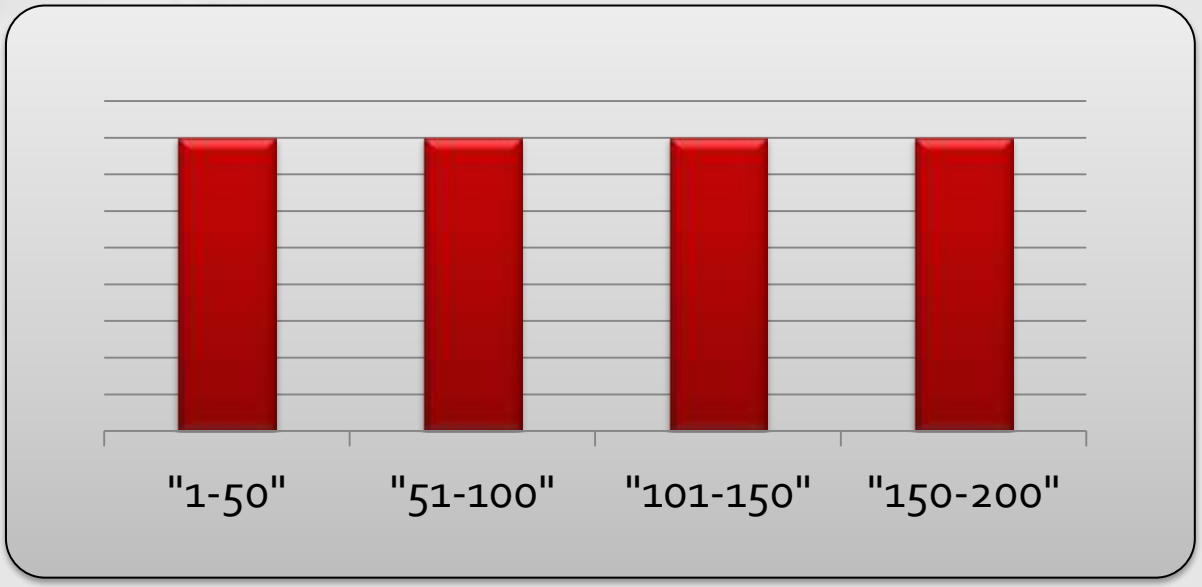

Updating Statistics

- sp_updatestats
  - Can resample
  - Won't run everywhere
- UPDATE STATISTICS X
  - WITH FULLSCAN

Indexes

- Clustered Index
  - Primary Key Default (but not necessary)
  - Data is stored at the leaf level
  - Data is ordered by the key
- Non-clustered Index
  - Uses cluster key or RID of a heap
  - INCLUDE stored at leaf
- And the rest – outside the scope of this session

Constraints

- Primary Key
  - Cluster by default, but doesn't have to be
  - Always an index
- Foreign Key
  - No indexes are created with this constraint
- Unique Constraint
  - This constraint is an index

What's All This Mean?

SELECT ID FROM TableA WHERE ID = 42

Table A			

What's All This Mean?

```
SELECT a.ID, b.Name, c.Value
FROM TableA as a
        JOIN TableB as b
           On a.ID = B.ID
        JOIN TableC as c
           ON b.OtherID = c.OtherID
WHERE a.ID = 42
```

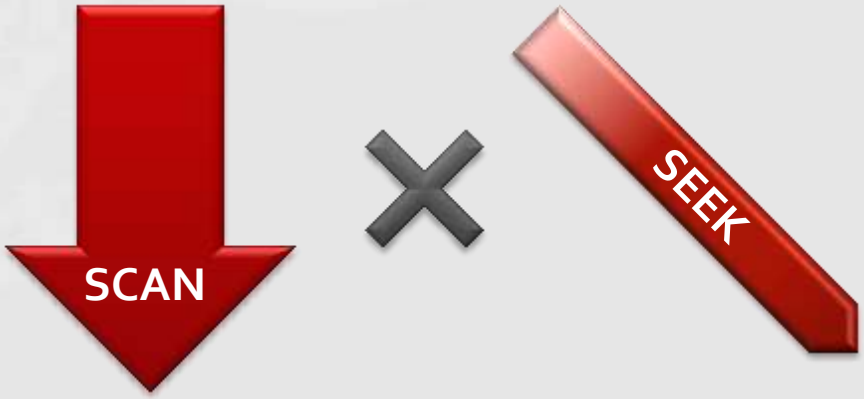

324 Possible Plans

Finish Line

- The ability to collect performance metrics on their servers as part of an overall query tuning methodology
- An understanding of how the optimizer works in support of writing better TSQL code as well as troubleshooting poorly performing queries
- The ability to generate, read, and understand execution plans from multiple sources in support of troubleshooting poorly performing queries
- A working knowledge of DMVs that will help them identify and fix performance issues on their servers
- The ability to address common query performance problems

Optimizer Resources

- Dr. Dewitt's Key Note, PASS Summit 2010 http://www.facebook.com/l.php?u=http%3A%2F%2Fwww.slideshare.net%2FGraySystemsLab%2Fpass-summit-2010-keynote-david-dewitt&h=306f5
- "Inside SQL Server 2008 T-SQL Querying" Itzik Ben-Gan
- "SQL Server 2008 Internals" Kalen Delaney
- "Inside The Query Optimizer" Benjamin Nevarez

Speaker Rate

- http://tinyurl.com/24fgjq9