

SQLskills Immersion Event

IEPTO2: Performance Tuning and Optimization

Whiteboard Drawings plus Other Interesting Slides/Resources

Kimberly L. Tripp

Kimberly@SQLskills.com



**NOTE: Drawings from class PLUS a BUNCH of stuff that I discussed in
other IE courses – IEPTO2 Training**

Why is this Such an Issue?


EVIL ORMS!

ORM = Object Relational Mapper

Pro – Rapid Application Development

Con – Fast to production, SLOW for the life of the application

How not to structure your database-backed web applications: a study of performance bugs in the wild
<https://bit.ly/2Ixpki4>



I totally agree!

Session Settings and Client Connectivity

SET Options	Required for Perf Features	Default Server Value	SSMS	SQLCMD	SQL Server Agent	Default OLE DB and ODBC Value	Default DB-Library Value	.NET / Your App
ANSI_DEFAULTS ²	NO	OFF	OFF	OFF	OFF	OFF	OFF	?
ANSI_NULL_DFLT_ON ²	NO	OFF	ON	ON	ON	ON	OFF	?
ANSI_NULLS ^{1,3}	YES = ON	OFF	ON	ON	ON	ON	OFF	?
ANSI_PADDING ^{2,3}	YES = ON	ON	ON	ON	ON	ON	OFF	?
ANSI_WARNINGS ²	YES = ON	OFF	ON	ON	ON	ON	OFF	?
ARITHABORT ²	YES = ON	ON	ON	OFF	OFF	OFF	OFF	?
CONCAT_NULL_YIELDS_NULL ^{2,3}	YES = ON	OFF	ON	ON	ON	ON	OFF	?
NUMERIC_ROUNDABORT ²	YES = OFF	OFF	OFF	OFF	OFF	OFF	OFF	?
QUOTED_IDENTIFIER ¹	YES = ON	OFF	ON	OFF	OFF	ON	OFF	?

- (1) The only state that's important is how it's set when the stored procedure is CREATED; the runtime setting is irrelevant.
- (2) If different than existing plan in cache, will be recompiled and added to the plan cache; performance may vary at execution.
- (3) In a future version of SQL Server, these options will always be ON and any applications that explicitly set the option to OFF will produce an error. Avoid using this feature in new development work and plan to modify applications that currently use this feature.

These two slides come from the Pluralsight course, Optimizing Procedural Code, Part 2: <https://www.pluralsight.com/courses/sqlserver-optimizing-stored-procedure-performance-part2>

Troubleshooting Session Settings

- Currently connected users (and currently set session settings) can be viewed
 - At the user level:

```
SELECT *  
FROM [sys].[dm_exec_sessions]  
WHERE [session_id] = @@SPID
```
 - Across sessions:

```
SELECT *  
FROM [sys].[dm_exec_sessions]  
WHERE [is_user_session] = 1
```
 - This does not guarantee that the session settings are not being changed elsewhere in the application (or, by the user [e.g. ODBC DSN settings])

These two slides come from the Pluralsight course, Optimizing Procedural Code, Part 2: <https://www.pluralsight.com/courses/sqlserver-optimizing-stored-procedure-performance-part2>

I included these next 5 slides as we discussed a few details about what specifically causes the cache to be flushed vs. plan invalidation. These come from the Pluralsight course listed below, specifically: Module 4: Statement Caching. [Here's a link to the course.](#)

SQL Server: Optimizing Ad Hoc Statement Performance

Module 4: Statement Caching

Kimberly L. Tripp

Kimberly@SQLskills.com

<http://www.SQLskills.com/blogs/Kimberly>



pluralsight
hardcore developer training

Stored Procedure Plans and Caching

- A plan is generated when no plan already exists in cache
- Plans are never saved on disk and may not persist within the cache
 - Some operations completely remove / evict the plan from the cache
 - **Server-level:** Server restart | DBCC FREEPROCCACHE | some RECONFIGURE changes
 - See recordings starting with **Side Effect: Plan Cache Flush** for version-specific details and demo
 - **Database-level:** DBCC FLUSHPROCINDB (undocumented) | sp_dbcmptlevel
 - **Procedure-level:** sp_recompile or they're aged out through non-use
 - Some operations cause the plan to be invalidated (but it still remains an object in the cache); these can be tracked
 - Schema of base object changes (ALTER TABLE / ALTER <object>)
 - Including when indexes are added to the base object
 - Statistics of base objects change
 - See recordings starting with **Plan Invalidation** for version-specific details and demos
- When a plan exists in cache, all subsequent executions use that plan

Side Effect: Plan Cache Flush (1)

- In SQL Server 2005 (only), numerous operations cause the entire plan cache to be cleared / flushed (all plans evicted from cache)
 - Database-level operations that cause an entire plan cache flush include:
 - When a database auto-closes (from AUTO_CLOSE option)
 - When a database is detached
 - When a database snapshot is dropped
 - When a database state is changed (OFFLINE, ONLINE, READ_ONLY, READ_WRITE)
 - When a database backup is restored
 - ...
 - **There are quite a few others;** see KB article 917828 for complete list
- In all versions, numerous server-level operations that cause the entire plan cache to be cleared include:
 - max server memory (MB)
 - min server memory (MB)
 - **... There are quite a few others;** more information coming up

Side Effect: Plan Cache Flush (2)

- **From SQL Server 2005 SP2 onward, messages written to the error log show you that the cache was flushed:**
 - SQL Server has encountered 4 occurrence(s) of cachestore flush for the 'Object Plans' cachestore (part of plan cache) due to some database maintenance or reconfigure operations.
 - SQL Server has encountered 4 occurrence(s) of cachestore flush for the 'SQL Plans' cachestore (part of plan cache) due to some database maintenance or reconfigure operations.
 - SQL Server has encountered 4 occurrence(s) of cachestore flush for the 'Bound Trees' cachestore (part of plan cache) due to some database maintenance or reconfigure operations.

Side Effect: Plan Cache Flush (3)

- From SQL Server 2008 onward, many database-level operations that caused the entire plan cache to be flushed no longer do so
 - Instead, they only cause the plans for that database to be flushed
- Database-level operations that cause the database plan cache to be flushed include:
 - DBCC FLUSHPROCINDB (dbid)
 - When a database is detached / dropped / restored
 - When a database (or filegroup with the database) changes between READ_ONLY and READ_WRITE
 - When a database changes between OFFLINE and ONLINE, or auto closes
 - There are others... (for example, changing recovery model ...)

NOTE: There are many places where a few of these are documented incorrectly. Some aren't detailed at all and others say that many of these operations **STILL** cause the entire plan cache to be cleared.

Side Effect: Plan Cache Flush (4)

- Many server-level configuration changes cause the entire plan cache to be cleared (at the time of RECONFIGURE, not sp_configure)
- Configuration options that could cause the entire cache to be cleared:

<i>Some Versions</i>	<i>Most Versions</i>	<i>SQL Server 2014</i>
cost threshold for parallelism	cross db ownership chaining	access check cache bucket count
max text repl size (B)	index create memory (KB)	access check cache quota
query governor cost limit	max degree of parallelism	clr enabled
remote query timeout (s)	min or max server memory (MB)	max worker threads
	min memory per query (KB)	
	query wait (s)	
	user options	

- Key point: be careful making database-level or server-level configuration changes during production hours
 - Test them in development first
 - sp_configure, then RECONFIGURE – check the cache messages
 - Database changes – check [sys].[dm_exec_procedure_cache]

SQLskills Immersion Event

IEPTO2: Performance Tuning and Optimization

Module 8: Statement Execution, Stored Procedures, and the Plan Cache

Kimberly L. Tripp

Kimberly@SQLskills.com



OK/good


med

Horrible

Parameter sensitive plans can wildly vary...

This was a case where a procedure had 15+ parameters and while there were certain common patterns and the plan was usually ok/good... there were some cases where a “mediocre” plan would end up in the cache. The mediocre plans wouldn’t be quite as problematic but the system would be slow. But, if an uncommon and particularly bad plan ended up in the cache it would effectively kill server performance. No one could connect, nothing would “finish” and the server was pretty much useless. The administrators didn’t know about the DAC (dedicated admin connection) so their only choice (that they could figure out) was to shutdown and restart. To add fuel to the fire, it was a failover cluster so the shutdown/restart was 30-90 seconds.

What was it – ONE – really nasty stored procedure that effectively brought their whole server down!



We talked a little bit about the benefits of isolating your tables from direct access by creating a layer of stored procedures, functions, and views – instead of directly accessing them (as some ORMs do). With a “layer” you can isolate your table changes and design changes by just making sure the “layer” continues to act as it did before. A second problem is that ORMs might not fully understand more complicated designs so that adds another restriction to your future changes.



Sch_M on table

Sp_recompile
tname

NOLOCK longrunning


Sch_S


20 min



We had a side discussion on the blocking that you could have with a Sch_M lock. The key point is that ANY lock blocks a SCH_M lock but a SCH_S held from a NOLOCK query that's extremely long running – will cause nasty secondary blocking while the SCH_M waits for ALL locks to drain off (where the SCH_S lock is probably the longest running)

The next slide is from IE1.






We had a side discussion on the blocking that you could have with a Sch_M lock. The key point is that ANY lock blocks a SCH_M lock but a SCH_S held from a NOLOCK query that's extremely long running – will cause nasty secondary blocking while the SCH_M waits for ALL locks to drain off (where the SCH_S lock is probably the longest running) The next slide talks about my script (PreventLongBlockingChains.sql). The following slide is a graphic from IEPTO1.

sp_recompile tname
WANTS Sch-M

Every other request WANTS

Be sure to check out the script:
PreventLongBlockingChains.sql



WANT sp_recompile frame

any long running statements?
(configurable)
30 sec

if not

then get in queue

if I can't get lock in 5 sec

Get out

Nasty Blocking CHAINS

Relaxed FIFO isn't Always Enough

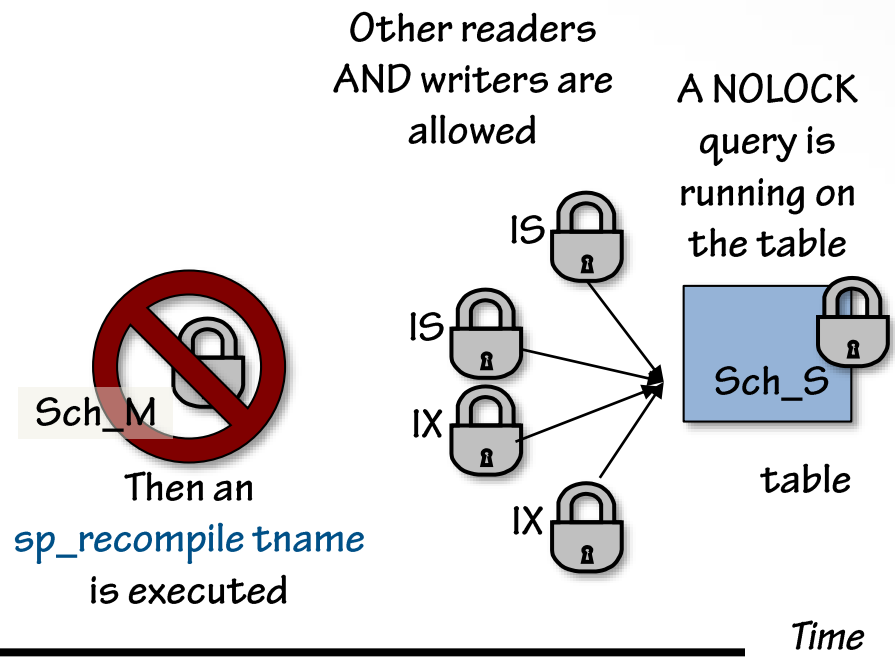
We had a side discussion on the blocking that you could have with a Sch_M lock. This slide is from IEPTO1.


- Imagine someone's running a NOLOCK query against a very large table (the query takes 4 minutes to run [ok, not that nasty])
- This query's running off-hours but the standard is to use NOLOCK
- An **sp_recompile tname** is requested

Unfortunately now...
everybody waits.....

Relaxed FIFO only let's through the folks that are compatible with ALL pending locks (not just those that are currently held).

This can create terribly long blocking chains.





Using EXEC (string) [DSE] and/or sp_ExecuteSql [ES] inside of stored procedures

DSE and ES are not the same.

DSE is unknown until runtime. At that point the statement is treated as though it's adhoc. Because most statements are **not** going to be safe, the statement will end up being recompiled (and optimized) for each execution.

ES is forced parameterization. When a statement is stable you can use this to reduce compilation / CPU costs.

OSFA Dialog (An attempt at *One Size Fits All* Programming)

fast
reason
Oh no...
OMG

where (c1 = @p1
or @p1 is null)
AND (c2 = @p2
or @p2 is null)
AND

CustID
LName FName
SSN
City State

Multipurpose screens/dialog boxes that lead to multipurpose procedures

They looked/sounded good at the time?

But, they're often VERY hard to optimize with where clauses that look like:

```
WHERE (Column1 = @Value1 OR @Value1 IS NULL)  
AND (Column2 = @Value2 OR @Value2 IS NULL)
```

...

Check out the demo scripts in order to fully see how to better optimize this – using `sp_executesql` (for combinations that are stable) and `sp_executesql WITH RECOMPILE` for combinations that could generate different plans (e.g. unstable plans).

Proc 3 params

mno → CL index seek STABLE
In
fn IF leading wc \ 3 char by wc
mno In → stable
mno fn → stable
fn In
mno In fn → stable

Not all parameters NEED to be recompiled!


The demo scripts, blog posts, and videos (from PASS UGs + my PASS Session)

To access the PASS sessions, go here: <https://www.pass.org/Learn/Recordings/Listing.aspx>

Then, search for the session by title or author (e.g. just search for TRIPP)

PASS Summit 2014: **Dealing with Multipurpose Procs and PSP the RIGHT Way!**

The beauty of building the statement dynamically and only supplying the non-NULL values means that – at least – you'll have more versions of the statements in cache (for 3 parameters is 7 combinations and therefore 7 statements). Then, add the logic: **IF @member_no IS NOT NULL** then concatenate **OPTION (RECOMPILE)** to that statement!



limited

at 1 1/2 GB

Clear


If you turn on "optimize for ad hoc workloads" then the number of "stubs" that can fit in cache is much higher.


At that point, searching to find whether or not the statement has executed is harder... this is why you want to make sure that you don't let this cache get really large (clearing it at 1 or 2GB).

Options for different executions

If we largely execute with “standard” users and we absolutely want THAT plan to be in cache and to be re-used then we could:

- 1) Use option (optimize FOR...) and use a literal that will generate a plan that’s great for the “typical” scenario. However, the performance won’t be great for the atypical (super user) scenario.
- 2) We could use EXEC for the typical scenario and EXEC proc WITH RECOMPILE for the super user scenario. The benefit of this is:
 - 1) We get a cached plan for the typical user (no CPU/plan stability)
 - 2) We get a specific plan for the atypical user (the super user) as opposed to a possibly inefficient plan (so, we have to compile but we’re only compiling this one type of plan)





Proc

_____ ✓

_____ ✓

Conditional logic and modularization

This is another way of looking at the same thing from the prior diagram. Really, don't think of the "logic" in your procedure; think only of the procedure as a list of statements (regardless of conditional logic). SQL Server tries to optimize ANY statement that's "*optimizable*" with the parameters that were supplied (regardless of those specific parameters apply for *that* execution). This is another case where you can end up with an inefficient plan because of parameter sniffing.

IF

WILDCARD

Select... Like TS

INDEX

ELSE

EQUALITY

Select... = %e%

INDEX

1st exec

%e%

1st exec

TRIPP

Conditional Logic example from the demo

When we ran our procedure to get members tied to the input of the name supplied – the conditional logic would execute only ONE branch... but, SQL Server would have already had to optimize BOTH branches

If the first execution was EXEC GetMembers '%e%' then the wildcard statement can be optimized using a table scan. The equality case will know that '%e%' (as a LITERAL name) is HIGHLY selective so there, they'll use an index...

However, if the first execution happens to be 'Tripp' then BOTH statements will be using an index and this will make the first statement inefficient.

The point, you can't control who gets there first and you never want statements to be optimized for parameters that weren't intended for them. **SOLUTION? Modularize your code!**

IF

B



E

Else

B



E

Conditional logic and modularization

SQL Server will never step into a sub-procedure unless it's executed and in this case it will only be executed with exactly the right parameters.

Proc
QS
Dec. Qv1
literal
where col = 12

Input @p1, @p2... - Known
literal sniffed
↳ histogram

param
where col = @p1

param sniffed
↳ hist

Variable
where col = @v1

UNKNOWN
AVG Density Vector

Parameters are “sniffed”

SQL Server uses parameters to optimize. Sniffing implies that SQL Server uses the histogram to evaluate the data selectivity. For the FIRST execution, parameter sniffing is great. It's the subsequent executions that can suffer from parameter sniffing (and therefore end up with PSP = parameter sniffing problems).

Variables are “unknown”

Variables are only known at runtime as the statements execute and the variables are assigned. As a result, they are unknown during optimization. Without an actual value, SQL Server cannot use the histogram. Instead SQL Server uses the density vector for the “average” numbers of rows that meet that criteria.

PROC

param list

@p1 value

@p2 value

SQL → uses a variable unknown
D.V.

SQL → uses a literal Known → special features

SQL → uses a parameter — histogram sniffing
(no special features)

If

[SQL → uses param — hist. sniffing

Else

[SQL → uses param — hist sniffing

SQL Server optimizes everything that's optimizable (sp?) ☺

Again, parameters are known. Parameters can be sniffed. SQL Server will optimize every statement that it can — regardless of what actually ends up executing. At the time of optimization they do not know what will or will not execute and they do not know the value of a variable.

DEFAULT
Proc

DSE — protects against SQLinj
execute under context
CALLER

Dynamic String Execution and SQL Injection

To prevent SQL Injection, SQL Server requires that strings execute under the context of the caller.

Dynamic String Execution and Execute As

I have a love/hate relationship with EXECUTE AS. On the one hand, you can give someone rights to something that you wouldn't otherwise be able to give...

However, be careful giving elevated privs around DSE. Generally, my recommendation around DSE is to only create a login-less user with very low privs.

tname

sp_ruc

EXEC(~)
Delete ~ tname

Executes As?

CALLER


USER

(Higher)

low


owner

self stupid



Without Dynamic String Execution permissions flow through the ownership chain

The idea is that when the ownership chain is unbroken then direct permissions (for example, to do a delete) are not necessary. As an owner you are giving rights to a process; this procedure can be protected with more business rules/logic and even reduced/isolated to specific rows. The end user does NOT need direct delete permissions to be able to execute the procedure as long as they've been given execute rights on the spoc.



Dynamic String Execution, EXECUTE AS and SQL Injection

People complained that granting permissions directly to the user was a problem... so, they wanted an alternative. Enter: EXECUTE AS.

While it solves the problem of permissions, it adds more potential for SQLInjection.

So..... Check out these posts:

[Little Bobby Tables, SQL Injection and EXECUTE AS](#)

["EXECUTE AS" and an important update your DDL Triggers \(for auditing or prevention\)](#)

✓ fn
✓ (fn + ln) ?

Stable plans vs. unstable plans – how do you know?

- You could test different combinations using execute with recompile.
- You might know because of selectivity?

* ~~fn + ln + mno~~ ← stable/cons
✓ ln

* ~~ln + mno~~ ←
* ~~mno~~ ←
* ~~mno + fn~~ ←
} ← stable/cons

Options that better balancing recompiling too much vs. not recompiling enough

This is from the multipurpose procedure demo. There are 3 parameters and 7 possible combinations that can be submitted. Instead of adding `OPTION (RECOMPILE)` to the statement (which doesn't always work [remember, it has a very checkered past]) you can build the statement dynamically and then programmatically (by setting a `RECOMPILE` flag to 0 or 1) decide whether or not `OPTION(RECOMPILE)` should be added to the dynamically constructed statement. Then, you can use `sp_executesql` to place ALL seven statements in cache. Those without the `OPTION(RECOMPILE)` clause will be cached. Those with `OPTION (RECOMPILE)` will get a plan on each execution.

15 executions?

15 plans?

2 plans

n 6/5

no good /
Single plan

OPTION
(RECOMPILE)

DSE ?

8/7

IF Subproc
Else Subproc

14

(1) Sep proc
→ OPTION
recompile only work if you did not
opt for

How many plans COULD be generated?

When you're testing your procedures you might have 15 test cases. Executing each of the procedures (with the 15 different combinations of parameters) WITH RECOMPILE. Then, review the number of different plans

If there are 15 different plans for 15 different executions – Use OPTION (RECOMPILE)

If there are 2 primary plans that are fairly split – see if you can create subprocedures and then control the behavior that way (especially if they either don't need to recompile OR only one needs to recompile).

If there are only 2 plans where one is 14 of the executions (and, is the norm) then you might use OPTION (OPTIMIZE FOR ...) but it depends on how bad the performance is for 15th. Can you control that through [other] parameters or programmatically?

① recompile always

- more CPU

+ Plan spec. to those parameters

Default

PSP

+ less CPU

- possibly horrible plans

- not consistent (who gets there first)

UNKNOWN (Avg)

- Bad plan? for atypical values

How many plans COULD be generated?

This is another way to describe what was on the prior page...

These are the pros/cons of the different procedure (and statements within) strategies:

The default

- Prone to parameter sniffing problems
- Possibly HORRIBLE plans for subsequent executions
- No guarantee of what plan actually gets into cache
- + Less CPU (not always though – you do save compilation time but if the execution is really expensive then you might lose all of the gains of having compiled/saved a plan)

Forcing a recompile for EVERY execution

- More CPU
- + Plan specific to those parameters

Optimize for UNKNOWN

- Guarantees an AVERAGE plan
- + Works well when the data is evenly distributed (but, the problems are less likely as well – but, if plans were getting into cache from atypical values then this might solve the problem)

SQLskills Immersion Event


IEPTO2: Performance Tuning and Optimization

Module 9: Index Analysis

Kimberly L. Tripp
Kimberly@SQLskills.com




(these are just “refresh/reminder” drawings about internals, etc.)



Should you drop redundant indexes? It depends. 😊

It's TRUE that query that uses a left-based subset of one index (let's say that a query uses index 2) that is COULD use a wider version of the same (left-based) index. It could use Index 1 instead of 2.

The problem is that there could be specific uses that warrant the smaller version. Questions to ask: Is the smaller version used by scans? Are they executed VERY frequently? How important are they?



Index consolidation

No production environment can support THE BEST index for every query, you'd have too many indexes. To be really successful with indexing – you need to find a balance. This includes consolidating some of your index recommendations to have one index handle multiple purposes!

Key Inc

C1 = C2 ~~C3~~

C1, C2, C3 no inc

Key = navigation


Key must be pres.

Key Inc

<u>LN, FN</u>	
<u>LN</u>	
<u>LN, FN, MI</u>	
LN	SSN
<u>LN, FN, MI (Inc SSN)</u>	

Consolidation?


You can CONSIDER removing indexes that are left-based subsets of others (in terms of the key). However, be sure to preserve the key as that's used for navigation/seeking. The columns that are included can be combined in any order.



Index Health – Checking for fragmentation

To get the value: **avg_fragmentation_in_percent** for **sys.dm_db_index_physical_stats**, SQL Server uses a scan of the first intermediate level (index level 1). All of the modes (limited, sampled, and detailed), use this same method. Then, SAMPLED and DETAILED do additional work (on the leaf level of the index) to give you things like average page density, etc. If your analysis/rebuild scripts only use **avg_fragmentation_in_percent** to determine fragmentation then you should only do a LIMITED mode scan. SAMPLED and DETAILED offer more page-specific details (page density, forwarding pointers, etc.) and if you're NOT using those during your automated maintenance process (and most people aren't) don't waste time!

Sys. dm-db-index-physical-staff



Good for avg-fragmentation-in-percent fragmentation?

Sampled } Avg page density, page fullness
detailed } row size

→ 1/100 of leaf / detailed reads all pages
all levels

FN like K%

rno > 6

mno < 5000

fn, mno, rno

Katie 6 4


Katie 7892 6

Kiera

Kiera


Kimberly

Kimberly




FN = Kimberly

LN = TREPP



LN, FN



FN, LN

Adding more indexes

Here we were discussing the order of columns. If there are range-based predicates then it doesn't usually matter which column is second. You tend to want to put the most selective first (in terms of the PREDICATES supplied). If all of the predicates are equality-based then it doesn't really matter. You're best ordering them by the LEFT-based combination that's used most commonly.

KEY

Density_Vector

c1

c2

c3

c4

c2

c4

c2

c3

c3

c4

c2

c3

c4

Multi-column, column-level statistics

Only DTA recommends multi-column, column-level statistics...

If you are about to trust (and create) a “green hint” recommendation (which comes from the Missing Index DMVs) then you might want to first run the query through DTA. If you end up creating the index that DTA recommends then you should also create the recommended stats (for that same table). These statistics can be helpful to the optimizer to better understand non-left-based subsets of the index for other possible uses.

Step

		DIST.	AVG
50			
100		57	
200	250K	5	50K

This was around our discussion about the limitation of SQL Server's knowledge of the data with gaps in the step values.

If the step describes the range from 101 to 200 (not including those values)

- There are 250K rows over that range
- There are 5 distinct values

The average will show 50K and EVERY value supplied between 101 and 199 will get an estimate of 50K. There's no way for SQL Server to know WHICH values actually exist.

SQLskills Immersion Event

IEPTO1: Performance Tuning and Optimization

**Discussions / drawings around VLDB
and “partitioning” large data sets (VLTs)**

Kimberly L. Tripp

Kimberly@SQLskills.com



Partitioning: PVs vs. PTs

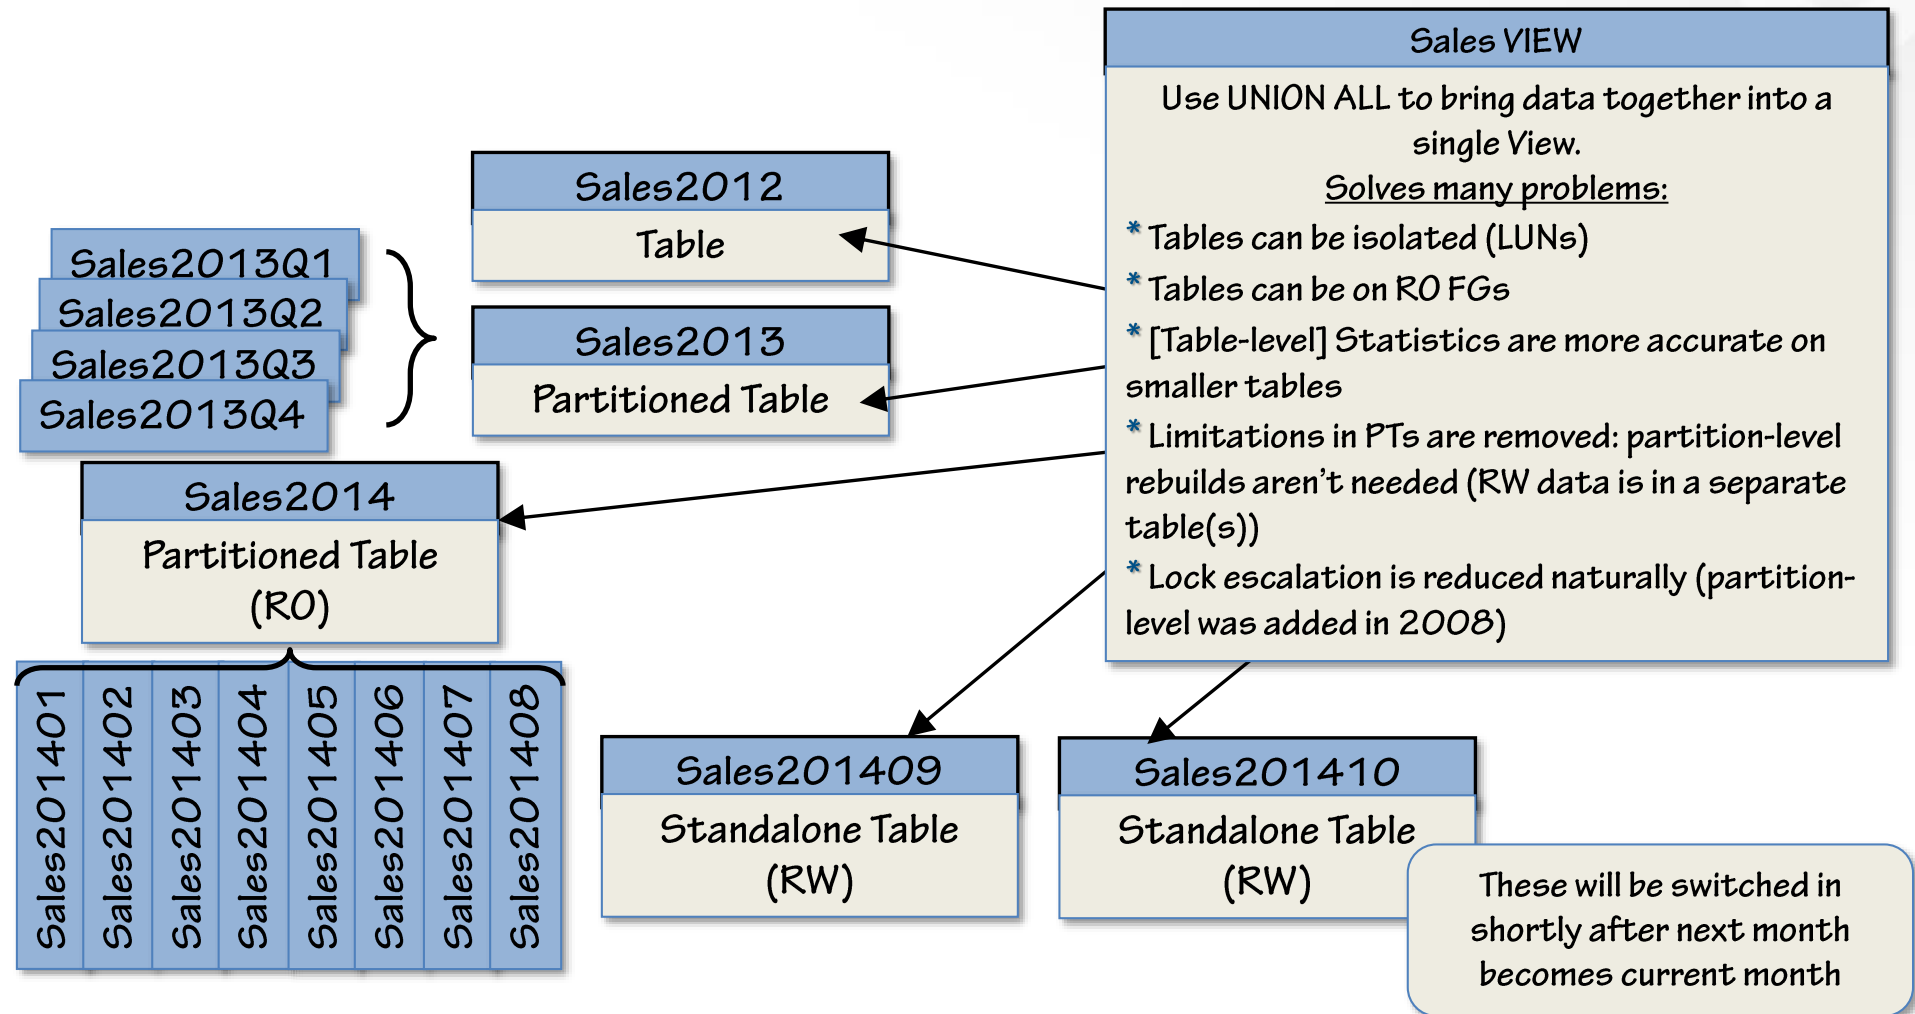
Partitioned views

- Any edition
- Lots of tables to administer
 - Must create/drop indexes on all base tables
 - Can have different indexes
 - Harder for the optimizer to optimize with so many indexes
 - Must verify business logic so that there are no gaps or overlapping values
 - Each table has [potentially] better statistics as the tables are smaller
- Can rebuild any of the tables ONLINE (if using EE)
- Can support multiple constraints on one or more columns

Partitioned tables

- Enterprise Edition only
- Only 1 table to administer
 - Only 1 table to create/drop indexes
 - All partitions have same indexes (which is easier for the optimizer to optimize)
 - Can create different indexes with filtered indexes
 - No possibility of errors (or gaps or overlapping values)
 - Table-level statistics can be less accurate for very large tables when there's a lot of skew to the data
- Partition-level rebuilds are offline but can rebuild the ENTIRE table online (not desirable)
- Can only support partitioning over a single column


Functionally Partitioning Data



Sharding/Scale-out – also, often called Service-oriented Database Architectures (SODA)

Partitioning in our workshop was all within a single database. However, we had a side/note about sharding and scale-out design. The most important thing that I can highlight is that scaling out is most ideal through middle-tier DDR (data-driven routing) where the applications are directed to the appropriate server. If every user randomly goes to any of these instances and all requests go through [distributed partitioned] views then performance will likely be worse!

Check out the whitepapers on SODA (service-oriented database architectures)



From <https://www.sqlskills.com/sql-server-resources/sql-server-whitepapers/>

- [SODA: Service Oriented Database Architecture: App Server-Lite?](#)
- [SODA: How SQL Server 2005 Enables Service Oriented Database Architectures](#)
- [Scalability: Planning, Implementing, and Administering Scaleout Solutions with SQL Server 2005](#)
- [Scalability: Solutions for Highly Scalable Database Applications: An analysis of architectures and technologies](#)

UPV

For this to be
"Updateable"
the partitioning
column must be
leading col of
PK

"PRIMARY"

Check constr.
for filtering

1M 1m 2m 2m-3m 3m-4m

Secondary
constraints
for filtering

June July Aug Sept

Indexed

Views are
on base tables
not PVs

	June	July	Aug
sum	2	sum	sum
sum	2	sum	2 sum
sum	3	sum	

Optimizing VLTs (Very Large Table)

A very large table has many "problems"... to reduce those – don't have just ONE VLT – have smaller tables that are unioned (using UNION ALL) into a view. If you also have restrictive constraints (CHECK constraints) across all of the base tables this combination is called Partitioned Views. You get numerous benefits with this architecture including: better control / manageability, better statistics on each table, online operations, and the reduction of some operations all together (on the historical data). Indexed views are created on the base table (not the partitioned view). And, with partitioned views, you can have secondary filtering criteria (with additionally created constraints).