

SQLskills Immersion Event

IEPTO1: Performance Tuning and Optimization

Module 11: Indexing Strategies

Kimberly L. Tripp
Kimberly@SQLskills.com



Overview

- Indexing for performance
- Indexing for AND
- Indexing for OR
- Indexing for joins
- Indexing for aggregates
- Indexed views (for aggregates)



2

© SQLskills. All rights reserved.
<http://www.SQLskills.com>



Indexing for Performance

- **Extremely challenging**
 - Users lie ☺
 - Workload specific
 - Data modifications are impacted by indexes (indexes add overhead to INSERTs/UPDATEs/DELETEs)
 - The type and frequency of the queries needs to be considered
 - This can change over time
 - This can change over the course of the business cycle
 - Need to have an understanding of how SQL Server works in order to create the “RIGHT” indexes – you CANNOT just guess!
- **To do a good job at tuning you must:**
 - Know your data
 - Know your workload
 - Know how SQL Server works!

Indexing Strategies at Design

- **First and foremost: choose a GOOD clustering key**
- **Create your primary keys and unique keys**
- **Create your foreign keys**
 - Manually index your foreign keys with nonclustered indexes
- **Create any nonclustered indexes needed to help with highly selective queries (lookups are OK for highly selective queries)**
- **STOP: this is your “design” base**
- **Add indexes slowly and iteratively over time while learning and understanding your workload as well as query priorities and always re-evaluate if/when things change!**

But Will YOUR Queries Use Them?

- **Subset of columns = projection**
 - Do not use * (unless against view)
 - Optimizer has more chances for optimizing query when result set is NARROW (only the required columns)
- **Subset of rows = selection**
 - Use positive search arguments
 - Isolate the column to one side of the expression
 - **USE:** MonthlySalary > value/12 (constant, seekable)
 - **DO NOT USE:** MonthlySalary * 12 > value (must scan)
 - Be cautious with LEADING wildcards
 - **USE:** LastName LIKE 'S%'
 - Avoid just appending %val% to every value (from the app)
- **Consider using views, stored procedures and functions to limit the columns/rows**

Using the Tools

- **USE the tools!**
 - STATISTICS IO
 - Showplan/Missing Index DMVs
 - Database [Engine] Tuning Advisor
 - BEWARE of the limitations of the tools!
 - Missing Index DMVs (and therefore showplan) only tune the plan that was executed – they do not “hypothesize” about alternatives (like DTA does)
 - All of the index recommendation from tools tend to go for “the best” choice rather than good enough choices
 - NONE of the tools do index consolidation...
- **Resources:**
 - Search “Bart Duncan Missing”
 - Find Missing Indexes – SQL Server Pedia
 - http://sqlserverpedia.com/wiki/Find_Missing_Indexes
 - A bit of searching – lots of good stuff out there!

SET STATISTICS IO ON (1)

- **Scan count: does not mean table scan**
 - Nothing to do with actual type of access
 - Refers to the number of “accesses” an object
- **Logical reads: number of page accesses in the data cache – specific to this query’s execution**
 - A single page can be accessed many times and EVERY one of these will be counted
 - * NOTE: Profiler vs. STATISTICS IO = Profiler includes I/Os performed during the execution of the query (for example, lookups into the plan cache, accessing metadata, security information, etc.). Profiler should always be greater than or equal to STATISTICS IO.
- **Physical reads: number of page reads that this query had to wait for – from disk**
- **Read-ahead reads: secondary process which accesses pages from disk (“reading ahead” of the query) so that SQL Server/CPU doesn’t have to wait**
- **Lob (logical, physical, and read-ahead) reads: same as the above but for all LOB [(n)text, MAX, XML] as well as limited-LOB data types that have overflowed**

SET STATISTICS IO ON (2)

- **Use logical reads as a general “total”**
 - The cost of getting from A to B in “steps” alone
 - Similar to distance
 - Does not include any traffic [blocking] encountered along the way
 - Does not include any worktables required
 - Doesn’t give you the complete picture
- **Use as a piece in the query execution information/puzzle**
 - Usually set in script
 - Can change it in SSMS
 - Tools, Options, Query Execution, Advanced:
 - SET STATISTICS IO ON
 - SET NOCOUNT ON
 - NOTE: Some of these options can have a profound affect on query performance. Should not change the ANSI options.

Showplan

- **Estimated plan**
 - Gives you the plan that SQL Server came up with through optimization – without actually executing it
- **Actual plan**
 - Gives you the plan that SQL Server came up with through optimization – and, executed it
 - This is EXACTLY the same plan [shape] as estimated but includes actual numbers
 - Extremely beneficial in finding cardinality estimation issues
- **Definitely NOT perfect...**
 - What are you really seeing with cached plans and stored procedures...
 - Plans for COMPILED values not the actual value – these can be the most incorrect

Missing Index Hints

- **The “green hint” in showplan, comes from the missing index DMVs**
- **Helpful, but**
 - Not always listed with the query it benefits (consider using SQL Sentry’s Plan Explorer)
 - Gives you the index that reduces the most I/O for the plan that was executed
 - Doesn’t consider other join types or join orders; doesn’t always give the best plan
- **Good to try**
 - If you’re ready to believe it and implement the suggestion consider checking to see what the Database [Engine] Tuning Advisor (DTA) recommends
 - Don’t just trust it; must consider consolidation
 - Review existing indexes
 - Could you create a slightly-wider but better index? Possibly removing one or more existing indexes?
 - The more you tune – the more you’ll find “similar” recommendations

Indexing for AND


- **AND progressively limits the SET**
- **All conditions MUST be true**
- **Indexing strategies**
 - Evaluate columns in WHERE clause
 - Index any single highly selective set
 - Index a combination of columns to yield a highly selective set
 - Order should be based on the most commonly combined criteria (if all SARGs use equality)
 - Order should be based on the most selective ***predicate*** criteria (if SARGs use varying operators such as >, < or LIKE)
 - If no combination of criteria create a selective set AND it's a high priority query, consider covering the query
 - SQL Server may use index intersection to intersect two relatively small sets (HASH Join), this is likely to be achieved without trying

Index Options

```
SELECT m.Member_No, m.FirstName, m.Region_No
FROM dbo.Member AS m
WHERE m.FirstName LIKE 'K%'
      AND m.Region_No > 6
      AND m.Member_No < 5000
```

- **Table scan (always an option)**
 - Clustered on member_no so a full table scan is unnecessary
 - SQL Server can “seek” with a partial table scan
- **NC index on firstname (K% is not very selective)**
- **NC index on region_no (region_no > 6 is 1/3 of the table)**
- **What does SQL Server do?**


Index Intersection



- Think of each of your nonclustered indexes as sets (as mini tables ordered by the key of the index)
- All nonclustered indexes “include” the clustering key in the leaf level of the index
- If we could “join” (or intersect) these sets on their common element (member_no) then we could find the data that we need...
- And, our query only wants these columns
- The intersection of these two indexes covers our query!

Index Intersection

- **Not the *fastest* option available for performance**
 - Requires more than 1 index to get to the data
 - Potentially requires tempdb space (HASH match)
 - Not something for which I strategize but something that might happen with lower priority queries that aren't covered and for those it's PERFECT
- **If it's a high priority query, you should consider doing with 1 index what you're currently doing with 2!**
 - No temp table
 - Only one index to seek/scan



What's Best Depends On the QUERY!

- An index that has firstname first is better for THIS query because it's the most selective SET (based on the query, not the data itself)
- An index with region_no first is good, possibly better if the firstname might accept leading wildcards such as
 - WHERE firstname LIKE '%e%'
- Not as big of a fan of having member_no first
 - It's the most selective data column (it's unique) but, we already have a clustered index on member_no
 - If a highly selective query were to run then SQL Server could seek into the clustered index...
 - If ALL of the queries supply all three of the parameters then region_no or firstname first would help more queries!
- Remember, ALL 6 are better than 2 or even a partial table scan!

SQLskills
immerse yourself in sql server

16

© SQLskills, All rights reserved.
<http://www.sqlskills.com>

Summary: Key Order – How Do You Decide?

- **First and foremost – it depends on the usage of the columns**
 - If you ALWAYS supply LastName and sometimes supply FirstName
 - LastName, FirstName is better than FirstName, LastName
- **Second – it depends on the types of predicates (equality?)**
 - If EVERY query ALWAYS supplies ALL conditions and those conditions are accessed with equality conditions, it does NOT matter:
`WHERE LastName = 'Tripp' AND FirstName = 'Kimberly'`
 - Then, it doesn't matter (these two indexes are REDUNDANT **in this case**):
 - LastName, FirstName
 - FirstName, LastName
- **Third – what about inequality?**
 - Once you start adding predicates that want inequality (LIKE, <, >, etc.) then you might only benefit (or, be able to seek on) the first condition. So, the 2nd and 3rd condition might be OK just to be in the INCLUDE

Indexing for OR

- **What is OR doing?**
 - Gather individual sets
 - Bring together and ensure that any row that appears in multiple places is only displayed once
 - Sound familiar?
- **IN is just a simplified series of OR conditions**
 - If an index exists to help search on each condition and EVERY specific value is HIGHLY selective, then it will use an index every condition
 - If any condition is not selective enough to use the index then a scan will be performed

Indexing for OR

- For ideal performance tuning, treat each OR as a different query
- Each condition CAN use an index
 - Each condition has to be *selective enough* to use the index
 - If there are 6 conditions and 5 are selective but 1 isn't then why would SQL Server use 5 different indexes and then still do a table scan...
 - If you have an IN then SQL Server can use the *same* index multiple times but if some of the conditions are selective and one are more are not then you hit the same issue as above – why would SQL Server use an index AND do a table scan!
- The final step is that an OR cannot return duplicates – SQL Server MUST determine if any rows are in more than one result set.
 - This often requires a temp table and a sort...

Indexing for OR OR is Similar to UNION

- OR removes duplicate rows based on row's unique identifier (RID or clustering key)
- UNION removes duplicate rows based on the SELECT list
- This is NOT good enough... you must add the row's key to the SELECT list if you choose to use UNION
- If you're joining multiple tables, you should consider adding EACH table's key to the query
- OR always removes duplicates
 - What if you know there are no duplicates
 - What if you don't care if duplicates are returned
- Consider UNION ALL

Be sure to test this thoroughly as your queries are semantically different when you change from OR to UNION

Demo

Indexing for OR

Using the Tools for Join Tuning

- Understand how to break down a join
- Understand how to force your join for performance comparisons
- Understand the pros/cons of the showplan recommendations
- Understand how to best use DTA for a more well-rounded recommendation
 - Use DTA from SSMS to see all of the recommendations
 - Know how to use DTA's recommendations iteratively!

Indexing for Joins

- Multiple possible join strategies: do you need to care?
- Items on which to focus:
 - Most expensive table in the join (you have to start somewhere?!)
 - Most expensive join in the plan (it's probably downstream from the most expensive table and a join on that table)
 - Once you know the problem table AND the problem join, focus on tuning that particular table within that specific join!

Best Options for Joins: Phase I



SARG1
Join Col PK



SARG2
Join Col FK

*Do you already have
individual indexes on
each and all of these
columns?*

Foreign key???

- One join strategy might use Table1's SARG1 index to Table2's join key index (loop join)
- Another could use Table2's SARG1 index to Table1's join key index (loop join)
- Another could use only the join key indexes (merge)
- What's best depends on the data!
- If ALL four indexes exist then the optimizer has the best choices

Cover the Combination: Phase II



SARG1
Join Col PK



SARG2
Join Col FK

Still not working?

- Not using these indexes?
- Performance still stinks?
- Cover the combo
 - Problem table (SARG, join): priority for the SARG
 - Problem table (join, SARG): priority for the join
- Only works when the cardinality of the join is low

Cover the Query: Phase III



SARG1
Join Col PK



SARG2
Join Col FK

- Covering the query/queries
- Cover the combo first, THEN add the additionally requested columns, with INCLUDE
 - Problem table (SARG, join): priority for the SARG
 - Problem table (join, SARG): priority for the join

Bringing It All Together (Long Demo)

- Pulling apart a plan and describing a lot while I do it...
- Hard-coding a query to create a base-line to go against
- Deciding where to start
 - Analyzing where we have a problem(s)
 - Finding the problem table
 - Finding the problem join
- Evaluating whether or not an index is a good idea
 - Reviewing/debating the "green hint"
 - Using DTA from SSMS to see if the hint is different


Demo

Analyzing and optimizing join performance

Table Scans – Are They Necessary?

Query 1: Query cost (relative to the batch): 100%

```
SELECT [c].[statement_no] , [s].[statement_dt] , [c].[charge_amt] , [p].[provider_name] , [m].[lastname] FROM [dbo].[charge] AS [c]
Missing Index (Impact 52.953): CREATE NONCLUSTERED INDEX [<Name of Missing Index, sysname>] ON [dbo].[charge] ([<charge_amt>]) INCLU
```




- Is the table scan because you're returning the entire table/all columns?
- Or, it is because the right indexes don't exist?
- What table has the highest cost?

Joins – Are All of the Joins Hash Joins?

Query 1: Query cost (relative to the batch): 100%

```
SELECT [c].[statement_no] , [s].[statement_dt] , [c].[charge_amt] , [p].[provider_name] , [m].[lastname] FROM [dbo].[charge] AS [c]
Missing Index (Impact 52.953): CREATE NONCLUSTERED INDEX [<Name of Missing Index, sysname>] ON [dbo].[charge] ([<charge_amt>]) INCLU
```



- Are all of these hash joins because you're tables are large?
- Or, are they because the right indexes don't exist?
- What join has the highest cost?

Manual Tuning Process

- **Find the most expensive table in the query (charge)**
 - Are there any SARGs – consider what the key would look like with these SARGs independent of the join conditions
- **Find the most expensive join in the query (charge joining to member)**
 - Figure out which join is the join that your problem table is joining to
- **Phase I should be already done**
- **Phase II should be considered**

```
CREATE NONCLUSTERED INDEX Charge_PriorityForSARG
ON [dbo].[charge] ([charge_amt], [member_no])

CREATE NONCLUSTERED INDEX Charge_PriorityForJoin
ON [dbo].[charge] ([member_no], [charge_amt])
```
- **Phase III adds any columns not already present, to the INCLUDE**

```
INCLUDE ([statement_no], [provider_no])
```

Tuning Goal

- **Find the most expensive table in the query (charge)**
- **Find the most expensive join in the query (charge joining to member)**
- **Try to tune CHARGE for its join to member... How?**
 - Review the green hint:

```
CREATE NONCLUSTERED INDEX MissingIndexDMVRecommendation
ON [dbo].[charge] ([charge_amt])
INCLUDE ([member_no],[provider_no],[statement_no])
```
 - Double-check using DTA (Query, Analyze Query in DTA):

```
CREATE NONCLUSTERED INDEX [DTA_K6_K7_K3_K2]
ON [dbo].[charge]
([member_no], [charge_amt], [statement_no], [provider_no])
```
- **Notice how similar these indexes are?**
- **What do they do, how do they differ?**

Result of Manual and Tool-based Tuning

- The missing index DMVs (via the green hint in showplan) came up with:

```
CREATE NONCLUSTERED INDEX MissingIndexDMVRecommendation
ON [dbo].[charge] ([charge_amt])
INCLUDE ([member_no],[provider_no],[statement_no])
```

- Manually, we came up with:

```
CREATE NONCLUSTERED INDEX Charge_PriorityForSARG
ON [dbo].[charge] ([charge_amt], [member_no])
INCLUDE ([statement_no], [provider_no])
```

```
CREATE NONCLUSTERED INDEX Charge_PriorityForJoin
ON [dbo].[charge] ([member_no], [charge_amt])
INCLUDE ([statement_no], [provider_no])
```

- The Database Tuning Advisor came up with:

```
CREATE NONCLUSTERED INDEX [DTA_K6_K7_K3_K2]
ON [dbo].[charge]
([member_no], [charge_amt], [statement_no], [provider_no])
```



33

© SQLskills. All rights reserved.
<http://www.SQLskills.com>

Benefits of These Indexes?

- The green hint/the Missing Index DMVs recommendation:

- Leads with the column charge_amt
- This removes the table scan and changes to an index seek
- This allows filtering by our search argument (charge_amt > 2500)
- PRO: This helps tune the plan that was chosen/executed
- CON: They did not hypothesize for alternatives

- The DTA's recommendation:

- Leads with the column member_no
- This removes the table scan and changes to an index seek
- This allows the join to change to a loop join
- PRO: This significantly reduces the cost/time for the join

- Of the tools – what's better? What gives better performance?

- DTA, but, our index was even a bit better given that you can't seek beyond charge_amt (as the SARGs against it are range-based)



34

© SQLskills. All rights reserved.
<http://www.SQLskills.com>

Database [Engine] Tuning Advisor

- It's not always perfect
- It sometimes yields the same index recommendation that the missing index DMVs
- It often OVER recommends indexes (this is why you want to use it ITERATIVELY after really analyzing where to begin)
- It doesn't recommend any forms of index consolidation
 - This is one of the reasons that a lot of development environments end up over-indexed
- IF you end up creating the index that was recommended for a particular table, then, go ahead and create the statistics that are recommended for that table
 - DTA can create multi-column statistics
 - They can give the optimizer other (sometimes VERY useful) ways of using the recommended index!

Join Strategies

- **Loop join**
 - Iterative search on the inner table based on the number of rows that match in the driving table
 - Usually best when the driver (outer table [chosen by SQL Server]) is small
- **Merge join**
 - Processing both tables at the same time using suitably sorted indexes
 - Usually best when the RIGHT indexes exist
 - An index on EACH table that LEADS with the same column (the column being joined) is necessary
- **Hash join**
 - Two-phase operation (build, then probe): build table (smaller set) and probe table (larger set) allowing SQL to join extremely large sets – in MEMORY (can spill)
 - Either side *can* use indexes to make the sets smaller
 - When this occurs on reasonably small tables then it sometimes means that good indexes don't exist
- **Key points: the strategy I demonstrated works for ALL join types**
 - You do not *need* to know or care about the specific strategy... just give SQL Server the best information from which to choose!

Indexing for Aggregations

- Two types of aggregates:
stream and hash
- Try to achieve stream to minimize overhead in temp table creation
- Computation of the aggregate still required
- Lots of users, contention and/or minimal cache can aggravate the problem!

Aggregate Query

- Member has 10,000 rows
- Charge has 1,600,000 rows

```
SELECT c.member_no AS MemberNo,  
       sum(c.charge_amt) AS TotalSales  
FROM dbo.charge AS c  
GROUP BY c.member_no
```

Aggregate Query

Table Scan + Hash Aggregate

```
SELECT c.member_no AS MemberNo,  
       sum(c.charge_amt) AS TotalSales  
FROM dbo.charge AS c  
GROUP BY c.member_no
```

- Table scan of charge table
 - Largest structure to evaluate
 - Worst case scenario
- Worktable created to store intermediate aggregated results: OUT OF ORDER (HASH)
- Data returned OUT OF ORDER unless ORDER BY added
- Additional ORDER BY causes another step for SORT, and sorting can be expensive!


Worst Case

Clustered index scan
(table scan)
1,600,000 rows

Hash aggregate
yields 9,114 rows
out of order

Sort
only has to sort
9,114 rows instead
of 1,600,000 rows

Return data



Aggregate Query Index Scan + Hash Aggregate

```
SELECT c.member_no AS MemberNo,
       sum(c.charge_amt) AS TotalSales
FROM dbo.charge AS c
GROUP BY c.member_no
```

- Out of order covering index on charge table
 - Index exists which is narrower than base table
 - Used instead of table to cover the query
- Worktable still created to store intermediate aggregated results: OUT OF ORDER (HASH)
- Data returned OUT OF ORDER unless ORDER BY added
- Additional ORDER BY causes another step for SORT, and sorting can be expensive!

Not as Bad

COVERING

Index scan
1,600,000
narrower rows

Hash aggregate
yields 9,114
rows out of
order

Sort
only has to sort
9,114 rows
instead of
1,600,000 rows
Return data

Microsoft SQL Server Management Studio

Query 1: Query cost (relative to the batch): 100%

```
SELECT member_no AS MemberNo,
       sum(charge_amt) AS TotalSales
FROM dbo.charge
GROUP BY member_no
ORDER BY member_no -- To be fair in the comparison go
```

Query 1: Query cost (relative to the batch): 100%

SELECT member_no AS MemberNo, sum(charge_amt) AS TotalSales FROM dbo.charge GR...

Table 'charge'.
Logical reads 3,770

Index Scan
Scan a nonclustered index, entirely or only a range.

Physical Operation	Index Scan
Logical Operation	Index Scan
Number of Rows	1600000
Estimated Row Size	19 B
Estimated I/O Cost	2.78461
Estimated CPU Cost	1.76016
Estimated Operator Cost	4.54476 (35%)
Estimated Subtree Cost	4.5447602
Estimated Number of Rows	1600000

Object
[credit].[dbo].[charge].Covering1

Aggregate Query Index Scan + Stream Aggregate

```
SELECT c.member_no AS MemberNo,
       sum(c.charge_amt) AS TotalSales
FROM dbo.charge AS c
GROUP BY c.member_no
```

- **Covering index on charge table (in ORDER of GROUP BY clause)**
 - Index exists which is narrower than base table
 - Used instead of table to cover the query
 - Covers the GROUP BY so data is grouped
- **Less work to aggregate results IN ORDER**
- **Data returned IN ORDER unless ORDER BY/ joins added**
- **Adding an ORDER BY identical to the GROUP BY does NOT cause any additional step for sorting!**

Much Better!

COVERING
Index scan
1,600,000
narrower rows

Stream
aggregate
also yields
9,114 rows
IN ORDER

**NO SORT
REQUIRED**

Return data

Microsoft SQL Server Management Studio

Query 1: Query cost (relative to the batch): 100%

```
SELECT member_no AS MemberNo, sum(charge_amt) AS TotalSales FROM dbo.charge GR...
```

Execution plan details:

Physical Operation	Stream Aggregate	Index Scan
Logical Operation	Aggregate	
Number of Rows	9114	9114 rows
Estimated Row Size	19 B	
Estimated I/O Cost	0	
Estimated CPU Cost	0.964557	
Estimated Operator Cost	0.9645596 (18%)	
Estimated Subtree Cost	5.5093198	
Estimated Number of Rows	9114	

Table 'charge'.
Logical reads 3,770

Query executed successfully. ROADRUNNER64/S

See the Difference?

Microsoft SQL Server Management Studio

File Edit View Query Project Tools Window Help

New Query credit Execute

Indexing for A...QLDev01.credit* Indexing for Ag...SQLDev01.credit Indexed V...

Results Messages Execution plan

Query 1: Query cost (relative to the batch): 48%

SELECT member_no AS MemberNo, sum(charge_amt) AS TotalSales FROM dbo.charge WITH (...)

Cost: 0 % Sort Cost: 4 % Hash Match (Aggregate) Cost: 45 % Clustered Index Scan [credit].[dbo].[charge].ChargePK Cost: 51 %

Query 2: Query cost (relative to the batch): 36%

SELECT member_no AS MemberNo, sum(charge_amt) AS TotalSales FROM dbo.charge WITH (...)

Cost: 0 % Sort Cost: 5 % Hash Match (Aggregate) Cost: 59 % Index Scan [credit].[dbo].[charge].Covering1 Cost: 35 %

Query 3: Query cost (relative to the batch): 16%

SELECT member_no AS MemberNo, sum(charge_amt) AS TotalSales FROM dbo.charge WITH (...)

Cost: 0 % Stream Aggregate (Aggregate) Cost: 18 % Index Scan [credit].[dbo].[charge].Covering2 Cost: 82 %

Query executed successfully. ROADRUNNER64\SQLDev01 (9.0 B2) REDMOND\w4tripp (54) credit 00:00:07 27342 rows

Pending Checkins

Ready

© SQLskills, All rights reserved. <http://www.SQLskills.com>

Concerns

- **Hash aggregates**
 - More temp tables
 - More contention in tempdb
 - Larger tempdb required
 - Performance varies on each execution
- **Stream/hash aggregate**
 - Aggregate needs to be computed

Is there a better way?
Indexed views!

Demo

What kinds of gains can you get?
Will it be worth it?



Views/Indexes: Quick Review

- **Views**
 - Named, saved SELECT statement
 - Tabular data set
 - Data definition (no ORDER BY unless TOP is used)
- **Indexes**
 - Clustered (only 1 per table)
 - Defines order and structure of data
 - Leaf level = data (of the table)
 - Nonclustered (249/999 per table)
 - Separate and duplicated data
 - Automatically maintained
 - Order and structure defined per index

Indexed Views Defined

- **A materialized view = a named, saved query whose data is structured and ordered by the definition of the clustered index on the view**
- **Pros = speed (usually exponential gains)**
 - No calculations
 - No joins
 - No aggregations
 - No temp tables
- **Cons**
 - Disk Space (minimal, affected by type of view)
 - INSERT/DELETE may affect the index
 - UPDATES to the data used by the view will also be update in the index
 - Goal: Don't index everything, find the right use

Indexed Views with Aggregations

Query defined uses aggregates for one or more of the columns
Interesting observation: data set often gets smaller

- **Questions to ask**
 - How many people are requesting the aggregation?
 - How complex is the aggregation?
 - What is the rate that columns involved in the computation change?
 - Are you searching on the aggregated value?
 - Is it highly selective?
 - How small does the set become? Hot row?
- **When do you want to use indexed views with aggregations?**
 - When users need ranges of data based on the aggregation
 - When read performance outweighs the disk space requirements
 - When the aggregation does not create significant contention!

Real Power of Indexed Views

- Optimizer is aware of them even when view is NOT referenced in the query (Enterprise/Developer Ed.)
- Calculations, aggregations and complex joins are already computed so only an index SEEK or SCAN is necessary
 - Processing steps might be saved
 - Possibly less disk access
 - Less CPU
 - Less memory (not necessarily with joins)
 - Tempdb space not necessary
- Nonclustered indexes can be created!
- It's a mini clustered table of all the data (including aggregations, joins, computations) that you need!

The Bad News Kind Of...

- Full support in Enterprise Edition
(and Developer Edition) ONLY
 - View does NOT need to be named and optimizer hints are not required
 - ANY query that CAN benefit WILL
- Reduced functionality on other Editions
 - VIEW must be named and an optimizer hint is REQUIRED

`FROM viewName WITH (NOEXPAND)`

 - Base table queries will NOT benefit from Index
 - ONLY queries that directly name and directly access the view will benefit
- Review complete details in the whitepaper

Seems Complex?

- Stay focused on purpose
- Keep indexes on views to a minimum unless READ ONLY database
- Consider impact to INSERT, UPDATE, DELETE performance
- Add indexed views to maintenance scripts

Test, test, test!



Aggregate Query Indexed View

```
SELECT member_no AS MemberNo,  
       sum(charge_amt) AS TotalSales  
FROM dbo.charge  
GROUP BY member_no  
ORDER BY member_no  
go
```

Query 1: Query cost (relative to the batch): 100%

SELECT member_no AS MemberNo, sum(charge_amt) AS TotalSales FROM dbo.charge GROUP...

Clustered Index Scan
Scanning a clustered index, entirely or only a range.

Physical Operation	Clustered Index Scan
Logical Operation	Clustered Index Scan
Number of Rows	9114
Estimated Row Size	19 B
Estimated I/O Cost	0.0268287
Estimated CPU Cost	0.0101824
Estimated Operator Cost	0.0370111 (100%)
Estimated Subtree Cost	0.0370111
Estimated Number of Rows	9114

Object
[credit].[dbo].[SumOfAllChargesByMember].SumOfAllChargesIndex

Table 'SumOfAllCharges'.
Logical reads 35

Query executed successfully.

Ready

Object Explorer

Registered Servers

Properties

SQL Server Enterprise Edition

© SQLskills, All rights reserved.
<http://www.sqlskills.com>

See the Difference?

Query 1: Query cost (relative to the batch): 48%
 SELECT member_no AS MemberNo, sum(charge_amt) AS TotalSales FROM dbo.charge WITH (...)
 Execution plan: SELECT (Cost: 0%), Sort (Cost: 4%), Hash Match (Aggregate) (Cost: 45%), Clustered Index Scan ([credit].[dbo].[charge].ChargePK) (Cost: 51%).

Query 2: Query cost (relative to the batch): 36%
 SELECT member_no AS MemberNo, sum(charge_amt) AS TotalSales FROM dbo.charge WITH (...)
 Execution plan: SELECT (Cost: 0%), Sort (Cost: 5%), Hash Match (Aggregate) (Cost: 59%), Index Scan ([credit].[dbo].[charge].Covering1) (Cost: 35%).

Query 3: Query cost (relative to the batch): 16%
 SELECT member_no AS MemberNo, sum(charge_amt) AS TotalSales FROM dbo.charge WITH (...)
 Execution plan: SELECT (Cost: 0%), Stream Aggregate (Aggregate) (Cost: 19%), Index Scan ([credit].[dbo].[charge].Covering2) (Cost: 82%).

Query 4: Query cost (relative to the batch): 0%
 SELECT member_no AS MemberNo, sum(charge_amt) AS TotalSales FROM dbo.charge GROUP BY member_no
 Execution plan: SELECT (Cost: 0%), Clustered Index Scan ([credit].[dbo].[SumOfAllChargesByMemberNo]) (Cost: 100%).

© SQLskills, All rights reserved.
<http://www.sqlskills.com>

Indexed View with Aggregates

Key points:

- Tempdb access not necessary
- NO worktables are necessary
- Aggregated set should be small but not too small as to create a hot ROW spot of activity (which can create excessive blocking)
 - GROUP BY member_no – probably OK
 - GROUP BY state – too few rows in aggregate
 - GROUP BY country – **AVOID** like the plague!
- Performance of data modification statements should be tested

Database Tuning Advisor

- DTA will include numerous configurable settings:
 - Length of TIME to tune
 - Physical design structures:
 - Indexes and **indexed views**
 - Indexes only
 - **Indexed views only**
 - Nonclustered indexes
 - Partitioned tables and indexes
 - Aligned partitioning
 - Full partitioning
 - Advanced: online vs. offline

Summary

- Ask for **ONLY** the data you need
 - Limit the rows requested with effective WHERE clause criteria
 - Limit the columns requested with effective SELECT lists
- Prioritize OLTP/OLAP and choose the clustering key based on the clustered index debate
- Add “key” indexes
 - Nonclustered for PK (if not the clustered)
 - Nonclustered for the unique keys
 - MANUALLY index your foreign keys
- Add nonclustered for SARGs
- Consider covering for high priority/low selectivity
- **Test, test, test!**

Indexing Strategies Overall

- Good base table indexes and a very small number of indexes to start (*some performance improvements should be handled by good design strategies*)
- General strategies:
 - Narrow indexes have **very few** uses!
 - Be careful that your general strategy is NOT:
 - See a WHERE clause, create a single-column index on it
 - To automatically create an index on every column (horrible!)
 - Guessing... or tuning queries randomly (without workload/index analysis)
 - Wider indexes have MANY, MANY more uses!
 - I'm not saying that you need to create indexes that have all of your columns in them but understanding a lot more about internals and how SQL Server works is VERY important for better performance!

Review

- Indexing for performance
- Indexing for AND
- Indexing for OR
- Indexing for joins
- Indexing for aggregates
- Indexed views (for aggregates)

Questions!

