

SQLskills Immersion Event

IE1: Internals and Performance

Module 6: Table Design Strategies

Kimberly L. Tripp
Kimberly@SQLskills.com



Overview

- Data type best practices
- Understanding row width (vertical partitioning)
- Application inconsistencies in types



2

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



Use the “Right” Data Type

System supplied data types

- Binary
- Character
- Integers
- Exact numerics
- Monetary
- Date and time types
 - Lots of new date/time data types in 2008/2012
- Legacy LOB (image, (n)text)
- LOB (“max” types, XML)
- Uniqueidentifier (GUID)
- FILESTREAM (vs. LOB)

Find the “right” data type for the job:

- * Use the *smallest (but least restrictive) data type possible*
- * Use *Unicode over ASCII*
- * If the data type varies:
 - < 5 chars should be nchar
 - 5-20 chars – questionable
 - > 20 char – lean towards nvarchar
- * For decimal/numeric data:
 - Find the right range
 - Standardize on decimal or numeric
 - Understand precision and range
 - Consider vardecimal in 2005 or row compression in 2008+
- * Use *uniqueidentifier sparingly*
- * Consider “sparse” attribute for 2008 (for flexible design)


Optimal Row Width

- Consider table usage above all else
- Estimate average row length
 - Overhead
 - Fixed-width columns
 - Estimate average from realistic sample data

```
SELECT avg(data_length(columnname))
FROM tname
```
 - Review min, max and avg. row width of existing and/or sample tables

```
sys.dm_db_index_physical_stats
```
- Calculate page density (rows/page):
 $8096 \text{ bytes/page divided by } ??? \text{ bytes/row} = \text{rows/page}$
- Calculate wasted bytes – on disk and in memory

Consider a Customer Table With 1,600,000 rows



Customer = 12.8GB
Partitioned Tables

= 7.3 GB

- Savings in overall disk space (5.5 GB saved)
- Not reading data into cache when not necessary
- LOB data can be isolated to support online index operations (*prior to 2012 where rebuilds with LOB can be done online*) to more critical data
- Locks are table specific therefore less contention at the row level

* The PRIMARY KEY column(s) must be made redundant for the additional tables.
Above: 47 columns in Customer; 49 columns total between 3 tables.

Vertical Partitioning

- Optimizing row size for:
 - Caching: better page density means less memory required
 - Locking: only locking the columns that are of interest minimizes even row-level conflicts
- Usage defines vertical "partitions" or "sets"
 - Logically group columns to minimize joins
 - Consider read only vs. OLTP columns (LOB separate from OLTP to allow online index maintenance (*prior to 2012*) for the critical/OLTP part of the table)
 - Consider columns often used together
- If every query requires a join, this isn't *as optimal as it could be* but should still be considered

Pushing LOBs “out of row”

- Subtle form of vertical partitioning
- Doesn't affect the application
- May *significantly* improve performance
- When should you do this:
 - You have a lot of “small” LOB values (values under 8K) that actually create large rows
 - LOBs aren't returned on most requests so you're filling cache with LOB values that aren't being used

What about “Place Holder” Rows? Nullability and INSERT Performance


- No default: no specific value required/specified at INSERT
- NULL values DO NOT mean empty space (*NULL bitmap is stored separately from the column data*)
- Working with NULLs
 - Accessing columns which allow NULL values can cause inconsistencies when developers/users are not aware of them
 - Math with NULL values can produce interesting results (value – NULL = NULL)
 - ANSI session settings can affect results sets when accessing columns that allow nulls
- Sometimes it's best to pre-allocate the row if you're using place-holders (so that updates do not cause massive fragmentation)

Inconsistencies in Data Types


- **Query doesn't match the column definition**
 - The case of the `implicit_conversion`
- **Key inconsistencies**
 - "Probe Residual" in showplan for hash join
 - May add a hash value for comparisons
 - May add a converted version of a column
 - Wastes storage space, index size, backups, ...
- **Inconsistencies in any layers can be costly**
 - Tables
 - Stored procedures/functions
 - Ad hoc queries/application interface
- **Consider tools like Visual Studio 2008 Database Edition (VSDE) for refactoring and static code analysis**

Probe Residual in Join

Query 1: Query cost (relative to the batch): 45%
`SELECT m.member_no, m.lastname, c.charge_dt FROM dbo.charge AS c JOIN...`




Query 2: Query cost (relative to the batch): 55%
`SELECT m.member_no, m.lastname, c.charge_dt FROM dbo.charge AS c JOIN...`



Probe Residual in Join


Query 1: Query cost (relative to the batch): 45%

```
SELECT m.member_no, m.lastname, c.charge_dt FROM dbo.charge AS c JOIN...
```



Query 2: Query cost (relative to the batch): 55%

```
SELECT m.member_no, m.lastname, c.charge_dt FROM dbo.charge AS c JOIN...
```



This process adds the residual (and has to compute it). Note: You do NOT see this after 2008 R2. 2012 only shows the PROBE RESIDUAL.




11

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

Probe Residual in Join


Query 1: Query cost (relative to the batch): 45%

```
SELECT m.member_no, m.lastname, c.charge_dt FROM dbo.charge AS c JOIN...
```



Query 2: Query cost (relative to the batch): 55%

```
SELECT m.member_no, m.lastname, c.charge_dt FROM dbo.charge AS c JOIN...
```



12

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

Data Type Best Practices

- **Consistency, consistency, consistency!**
- **Sloppy coding leads to:**
 - Performance problems
 - Difficulty when performance tuning
- **Scalability can only happen with good design**
 - Tables can be created easily but design takes knowledge:
 - Knowing the data
 - Knowing the users
 - Knowing the system
 - Take more time for design/prototyping – the sooner you begin to code, the longer it's going to take!
 - Consider changes over time – if already in place...some tools (Data Dude) might help with refactoring, testing, static code analysis!

Review

- Data type best practices
- Understanding row width (vertical partitioning)
- Application inconsistencies in types

Questions!

