

Sybase vs Oracle: Scalar Subquery & Inline View

Query	Sybase									Oracle								
	Test1	Test2	Test3	Test4	Test5	Avg	Pct	Pct	Pct	Test1	Test2	Test3	Test4	Test5	Avg	Pct	Pct	
1 Subquery COUNT	1.160	1.150	1.170	1.153	1.160	1.159			32.9%	3.600	3.500	3.690	3.390	3.430	3.522			304.0%
3 InlineView COUNT	1.593	1.556	1.583	1.593	1.553	1.576	36.0%		27.7%	5.260	5.300	6.140	5.590	6.100	5.678	61.2%		360.4%
3 InlineView COUNT Improved	1.183	1.173	1.190	1.180	1.183	1.182	2.0%											480.5%
1 Subquery SUM	2.173	2.153	2.163	2.153	2.173	2.163		86.7%		(Abandoned at 120 mins)								
3 InlineView SUM	3.033	2.933	3.033	2.933	3.013	2.989	38.2%	158.0%	3.8%	79.00	79.00				79.00			2,643.0%
3 InlineView SUM Improved	2.183	2.180	2.183	2.143	2.143	2.166	0.2%	87.0%										3,646.6%
Configuration	Sybase ASE 15.0.3 with 680MB DataCache									Oracle 10.2.0.3.0								
“O/S”	On WindowsXP 1.0GB Allocated On Parallels 3.0, on MacOS 10.4.11 (Demo system)									On WindowsServer (details not provided, massive o/s cache)								
Machine	Intel Core Duo 2.16GHz									Unknown								
Tables	REF_Customer			1000 rows; Random Customers						Same as Sybase side								
	CustomerTransaction			2M rows; 1K Credits & 1K Debits per Customer														
	CustomerTransaction_2			2M rows; 1K Credits & 1K Debits per Customer														
	All rows padded to 250B, 8 rows per page																	
Indices	2 + 1 + 1									1 + 1 + 1 (Not used: Tablescan)								
Result Set (Scalar Data Points)	1000 Customers with COUNT/SUM(Credit/Debit)									Same as Sybase side								
	2 Scalar data points per row; 1000 source rows per Scalar data point																	
Sybase Note	We are testing Query Processing, Logical I/O; not Physical I/O																	
	Idle server; 2 x Engines; Parallelism Off; 1 Engine used																	
	Single Client & Server on same system																	
	Improved = Sybase side only; Code Improved: Subquery fat Join & GROUP BY replaced with Subquery.																	
Note	<p>a. Oracle figures <i>recorded</i> in seconds; Sybase figures <i>recorded</i> in millisecs; all figures are seconds</p> <p>b. Row counts now fixed up, they are identical; Oracle Indices are suspect (Tablescan but Tony says it is “normal”)</p> <p>c. Table populations are identical. But the disk space used (488MB per transaction table) does not match the Oracle side.</p> <p>d. Although the physical tables & internal query processing are closer, they are not close enough to be comparable</p> <p>e. Pink/Blue Percentage <i>column</i> uses the yellow <i>cells</i> as base: it compares different flavours of the test, within each product</p> <p>f. Green Percentage <i>column</i> uses the respective <i>cell</i> of the other product as base: it compares the products, per test</p> <p>g. All code used on the Sybase side is located in the same directory as this PDF.</p> <p>h. Although only one transaction table is required, separate tables (the second a copy) were used on the Oracle side. I downgraded my side to suit.</p>																	