

techcast: ASE Performance and Tuning



Eric Miner,
Senior Engineer, Sybase, Inc.
DST Engineering Optimizer Group



Caveats

- Optimizer analysis is not something you'll have to do every day. But, knowing the basics of how to do it can be helpful when the need arises
- These are general steps and suggestions - your mileage may vary
- Should you need to call Tech Support having done an analysis yourself will help them focus the issue

Make Sure “The Same” is actually “The Same”

When comparing the same query on different datasets/servers “the same” may not be

Make sure that datasets and configurations are as identical as is possible – if not the differences may result in differing plans

- Use dump and load, if possible, to insure datasets are the same
- Row & page counts
- Data distribution – are the values the same?
- Server configs – total memory, cache pools and cache sizes, parallel

Make sure to use the same version of ASE

Tools And Techniques Of Optimizer Analysis

There are several tools at your disposal to perform an optimizer analysis

ASE contains most of what you'll need, such as:

- showplan
- statistics io and statistics time
- traceon 302 and 310

There are external utilities that are important too

- optdiag (11.9.2 and above), ISQL
- sp__optdiag – a non-Sybase proc that will output optdiag output without going to the command line. Not supported by Sybase, but very useful. Available at - <http://www.sypron.nl/optdiag.html>

Using A Force or An Abstract Plan As A Resolution or a Tool

A force option or an AP may be the best short term resolution – force index, forceplan, Abstract Plan

Analysis can take time. Consider using a force or and AP while you are analyzing the situation or working with Tech Support

- A force or an AP may be the best resolution for a handful of inefficient queries
- You won't be able to use a force option if you can't get to the query, however an AP may be usable
- A force or an AP can be used to verify that an access is efficient
- See the P&T Guide for syntax

Tools And Techniques Of Optimizer Analysis – noexec and fmtonly

When you want to get optimizer outputs without executing anything

noexec – when on anything following will not execute. Batch SQL will be compiled (optimized)

- Useful when you want to analyze optimization but don't want the query to run
- **Syntax** set noexec {on | off}

fmtonly – when on stored procedures will not execute but will be compiled. Only procs that return a result set

- Useful when you don't want the proc to run but need to get optimizer outputs
- Execute proc with recompile
- **Syntax** set fmtonly {on | off}

Tools And Techniques Of Optimizer Analysis - showplan

Showplan prints the query plan used to execute the query

Easy to read description of the query plan

- Well known and documented extension of the set command - See P&TG for good write up on the output
- **Syntax** set showplan {on | off}
- sp_showplan returns showplan information for spids

Tools And Techniques Of Optimizer Analysis – set statistics io and time

set statistics io prints the type and quantity of I/O performed by a query

- **Syntax** set statistics io {on | off}
- Used as a fundamental measurement of performance - well documented
- Statistics io will show how much physical and logical I/O is done for the query
- If the query is long running don't bother with statistics io - use noexec

Tools And Techniques Of Optimizer Analysis - stats time

- Statistics time measures parse and compile time, server time and elapsed time - well documented
- **Syntax** set statistics time {on | off}
 - Times reported in milliseconds and CPU ticks
- Useful when you need to see how long optimization is taking - parse and compile time
- Can be used when noexec is on to measure parse and compile time

dbcc Traceon 302 & 310 Are at The Heart Of An In-Depth Analysis

dbcc traceon 302 and 310 print the optimizer's cost estimates

Prints the optimizer's index selection and 'search engine' phases

- Traceon 302 – index costing and selection information. Compares the cost of a table scan to the cost of accessing an index
- Traceon 310 – search engine phase. Physical I/O costing, Cache costing, Join order costing, can demote and index to a table scan
- Both are documented in the “ASE P&T Guide - Query Tuning”
- It's best to use them together
 - **Syntax** dbcc traceon (3604,302,310) use 3604 to send output to client, 3605 to send to errorlog (not recommended)

Tools And Techniques Of Optimizer Analysis - traceon 302

Traceon on 302 output is divided into five information blocks:

- **Table information block** - row & page counts, datapage cluster ratio, number of partitions, size of largest partition, partition skew (used in costing a parallel query) and correlation name if one is used
 - Check that row and page counts are reasonable

Tools And Techniques Of Optimizer Analysis - traceon 302

Beginning selection of qualifying indexes for table
'li_2', varno = 0, objectid 304004114.

**The table (Allpages) has 600572 rows, 40646
pages,**

Data Page Cluster Ratio 0.999990

The table has 4 partitions.

The largest partition has 40639 pages.

The partition skew is 3.999114.

Tools And Techniques Of Optimizer Analysis - traceon 302

Base cost block - contains the cost of a table scan in rows and pages. Chosen I/O size, cache name and id and the buffer replacement strategy.

- Used to compare to all other access that are costed

**Table scan cost is 600572 rows, 40646 pages,
using no data prefetch (size 2K I/O),
in data cache 'default data cache' (cacheid 0)
with MRU replacement**

Tools And Techniques Of Optimizer Analysis - traceon 302

The Clause Block - Contains info about the clause(s) being costed for the current table. SARGs are costed first followed by joins. For joins current table is always the outer (left) in the clause.

- Are the clauses of the query listed? SARGs and joins. If you don't see it, the optimizer isn't costing it.
- Are there any unknown SARG values?
 - Will use a default selectivity value or total density based on the SARG operator
 - Usually caused by local variable in a query or proc
- Are there any mismatched datatypes?
 - See the P&T Guide for more information on datatype mismatches

Tools And Techniques Of Optimizer Analysis - traceon 302

Clause Block cont. -

Selecting best index for the **SEARCH CLAUSE:**

```
t1.1_partkey <= 2000
```

```
t1.1_orderkey > unknown-value
```

- If SARGs/joins not printed may be datatype mismatch

Selecting best index for the **JOIN CLAUSE:**

```
l.1_orderkey = o.o_orderkey
```

```
l.1_orderkey < 10000
```

Tools And Techniques Of Optimizer Analysis - traceon 302

Column Block - Contains selectivity of all columns in the clause(s). Selectivity is based on the statistics and the clause(s).

`Estimated selectivity for B, selectivity = 0.001283,
upper limit = 0.052642.`

‘selectivity’ - Range cell density

‘upper limit’ – weight of qualifying cell

No statistics available for colA, using the default range selectivity to estimate selectivity.

`Estimated selectivity for colA, selectivity = 0.330000.`

- No statistics - optimizer will have no info about this column
- Consider adding statistics to the column – Highly recommended
- Default selectivity value used will depend on the operator

Tools And Techniques Of Optimizer Analysis - traceon 302

Index Selection Block - prints info on all qualifying indexes for the clause(s). Scan and filter selectivity values, estimated number of rows and pages, index height and cluster ratios.

- **Scan selectivity** used to estimate the number of index rows and pages to be read
- **Filter selectivity** used to estimate the number of data pages to be accessed
 - Used to estimate the cost of accessing composite indexes when SARG includes more than one column of the index
 - Adds another layer of filtering to index costing

Tools And Techniques Of Optimizer Analysis - traceon 302

- The cluster ratios for the index are printed -

Data Row Cluster Ratio 0.030877,

Index Page Cluster Ratio 0.998118,

Data Page Cluster Ratio 0.023310

- **Data row cluster ratio** - how well clustered index rows are in relation to data rows. Used to cost a non-clustered or DOL clustered index
- **Data page cluster ratio** - used to estimate the effectiveness of large I/O when using this index to access data pages
- **Index page cluster ratio** - used to estimate the cost of large I/O for queries that need to read many leaf pages from a non-clustered or DOL clustered index
- The lower the cluster ratio the less clustering there is, and the less likely the index or a large I/O will be used

Tools And Techniques Of Optimizer Analysis - traceon 302

Estimating selectivity of index 'lineitem_sdate',
indid 5

```
scan selectivity 0.002300,  
filter selectivity 0.000000  
1 rows, 4 pages, index height 2,  
Data Row Cluster Ratio    0.030877,  
Index Page Cluster Ratio 0.998118,  
Data Page Cluster Ratio   0.023310
```

Tools And Techniques Of Optimizer Analysis - traceon 302

Best Access Block - prints the best access for the clauses examined in this section of 302 for the current table. Repeats much of the index selection info also contains the chosen I/O size, the cache to be used and the replacement strategy

The best qualifying index is ind_name

Best qualifying access is a table scan

- This is not the final word on which access will be used, that happens later in the search engine phase. This access will however be used in that phase

Tools And Techniques Of Optimizer Analysis - traceon 310

The **FINAL PLAN** at the end of 310 is the plan that will be executed – this is the optimizer's final word in the plan

- Prints the first plan costed and each subsequently cheaper plan
- Traceon 310 is not currently as well organized as 302
- This is the end of the optimization phase of query processing

Tools And Techniques Of Optimizer Analysis - traceon 302

QUERY IS CONNECTED

Number of tables in join: 6

Number of tables considered at a time: 4

Table count setting: 4

- Indicates that the query is not a Cart product. In this case there are 6 tables in the join and 4 at a time will be costed

IGNORING THIS PERMUTATION

0 - 1 - 2 - 5 -

- Indicates that this join order has no plans cheaper than the cheapest one found to this point

Tools And Techniques Of Optimizer Analysis - traceon 310

NEW PLAN (total cost = 1080) :

- A plan that is cheapest to this point has been costed, the cost is printed

TOTAL # PERMUTATIONS: 864

- This is the total number of join orders examined by the optimizer

TOTAL # PLANS CONSIDERED: 1997

- This is the number of plans that were costed by the optimizer

CACHE USED BY THIS PLAN:

CacheID = 0: (2K) 50920 (4K) 0 (8K) 0
(16K) 43520

- This is the amount and sizes of the caches to be used

Tools And Techniques Of Optimizer Analysis - traceon 310

PARALLEL:

`number of worker processes = 20`

`max parallel degree = 5`

`min(configured,set) parallel degree = 5`

`min(configured,set) hash scan parallel
degree = 5`

- This is parallelism that will be used for this query plan

Tools And Techniques Of Optimizer Analysis - traceon 310

This is query plan that will be executed (only a portion in this case)

```
FINAL PLAN (total cost = 641):
```

```
varno=5 (nrn_cis_cust) indexid=0 ()
```

```
path=0xf7740800 pathtype=sclause
```

```
method=NESTED ITERATION
```

```
scanthreads=1 outerrows=1 rows=105
```

```
joinrel=1.000000 scanpgs=63978 data_prefetch=YES
```

```
data_iosize=16 data_bufreplace=LRU
```

```
scanlio=142 scanpio=469
```

Tools And Techniques Of Optimizer Analysis - Optdiag

Command line utility designed to allow you to read write and simulate statistics

Most statistics are stored as varbinary and are hard, at best, to read and write in the system tables.

- Use to check the state of the statistics
 - Last time statistics were modified
 - Check table/index fragmentation - empty pages, forwarded rows, deleted rows, space utilization and large I/O efficiency
 - Check for highly duplicated values

Check P&TG for syntax

Tools And Techniques Of Optimizer Analysis - Optdiag

Statistics for table: "t2"

Data page count:	14249
Empty data page count:	16
Data row count:	200000.00000
Forwarded row count:	527.00000
Deleted row count:	42.000000
Data page CR count:	5811.000000
Index page CR count:	311.000000
Data row CR count:	44579.000000

Derived statistics:

Data page cluster ratio:	0.9930327868852459
Index page cluster ratio:	0.9986085343228200
Data row cluster ratio:	0.9995110244218134

Tools And Techniques Of Optimizer Analysis - Optdiag

Statistics for column: "col_A"

Last update of column statistics: Mar 22 2000
11:44:16:053AM

Range cell density:	0.0000050000000000
Total density:	0.0000050000000000
Range selectivity:	default used (0.33)
In between selectivity:	default used (0.25)

When changing the default selectivity values -

Range selectivity:	0.00033
--------------------	---------

Tools And Techniques Of Optimizer Analysis - Optdiag

The histogram - describes the distribution of values in the column.

- Two types of cells - **Range Cells** contain more than one value. **Frequency Count Cells** contain only one value, will appear for highly duplicated values. FC are the most accurate form of cell
- **Weight** - the percentage of the column occupied by the value(s) in the cell. The higher the weight the more the value(s) is duplicated
- A very high weight will indicate highly duplicated values - not bad but be aware of it

Tools And Techniques Of Optimizer Analysis - Optdiag

- **Operator** - helps describe the type of cell. The weights are what you want to look at.
- **Value** - this is the boundary value of the cell. Cell includes the upper boundary (value) but not the lower boundary (value of previous cell)

Check the number of cells. If you have duplicated values or range SARGs consider increasing the requested step count.

Tools And Techniques Of Optimizer Analysis - Optdiag

Step	Weight		Value
1	0.00000000	<=	0
2	0.05263982	<=	31423
3	0.05263316	<=	63045
4	0.01974451	<=	74982
5	0.00000000	<	170016
6	0.15908001	=	170016
7	0.05263815	<=	201861
8	0.10263316	<=	201862
9	0.05263982	<=	265249

Frequency count cells can be identified by higher than average weights

Conclusion

Optimizer analysis is not something you need to do all the time

- But, if you need to do it having information on how to do it may help you resolve issues quickly and avoid them in the future

More Optimizer Resources and Help

- The latest Performance and Tuning Guide
 - Don't be put off by the ASE 12.0 in the title, it covers the 11.9.2 features/functionality too
 - <http://sybooks.sybase.com/onlinebooks/group-as/asg1200e>
- Any "What's New" docs for a new ASE release
- Tech Docs at Sybase Support
 - <http://techinfo.sybase.com/css/techinfo.nsf/Home>

More Optimizer Resources and Help

- The Sybase Customer newsgroups
 - <http://support.sybase.com/newsgroups>
- The Sybase list server
 - **SYBASE-L@LISTSERV.UCSB.EDU**
- The external Sybase FAQ
 - http://www.isug.com/Sybase_FAQ/
- Join the ISUG, ISUG Technical Journal, feature requests
 - <http://www.isug.com>

Sybase Developer Network (SDN)

Additional Resources for Developers/DBAs

- Single point of access to developer software, services, and up-to-date technical information:
 - White papers and documentation
 - Collaboration with other developers and Sybase engineers
 - Code samples and beta programs
 - Technical recordings
 - Free software
 - Join today: www.sybase.com/developer