



Essbase Block Storage Optimization Basics

What Do You Want To Optimize

Marketing Technologies Group | www.mtgny.com

- Loading Data
- Calculations
- Query
- Administration Tasks

Tools

- Outline Design
- Storage Properties
- Block Configuration
- Caches
- Procedure Design

Principles

- Disk I/O is slow
- Smaller is faster
- BSO is prone to data explosion
- BSO databases get big fast

Where to Look For Information

Marketing Technologies Group | www.mtgny.com

- Database Properties
- Application Log
- Set Msg
 - ▣ Summary
 - ▣ Detail

Optimizing Loading Data

Marketing Technologies Group | www.mtgny.com

- Use wide files instead of long files (wide on a dense dimension)
- Sort by sparse dimensions
- Increase Index cache

Outline Design

Marketing Technologies Group | www.mtgny.com

- Fewer Dimensions
- Smaller Dimensions
- Fewer Generations (subtotals)
- Hourglass Shape
- Attribute Dimensions

Storage Properties

Marketing Technologies Group | www.mtgny.com

- Label Only
 - ▣ Non-aggregating parents
- Dynamic Calcs
 - ▣ Dense calcs

Understanding Blocks

Marketing Technologies Group | www.mtgny.com

- ❑ Controlled by sparse dense configuration
- ❑ Smallest Unit of I/O
- ❑ Entire block is in RAM at one time
- ❑ Reading/Writing blocks twice takes twice as long
- ❑ See [Demystifying Sparse and Dense Webcast](#)

Setting Sparse and Dense Dimensions: Rules of Thumb, Part 1

Marketing Technologies Group | www.mtgny.com

- Time and Accounts are dense. Everything else is sparse.
- Target block size is approx 100k.
- If initial block size is:
 - Too Small
 - change a small sparse dimension to dense
 - Too Big
 - change a dense dimension to sparse

Target Block Size: 8k – 80k

Marketing Technologies Group | www.mtgny.com

8k Bytes/Block
/
8 Bytes/Cell
1000 Cells/Block

10 Time Periods
100 Accounts

80k Bytes/Block
8 Bytes/Cell
10,000 Cells/Block

10 Time Periods
1000 Accounts

Setting Sparse and Dense Dimensions: Rules of Thumb, Part 2

Marketing Technologies Group | www.mtgny.com

- ❑ Large dimensions should be sparse.
- ❑ Dimensions with many formulas should be dense.
- ❑ Dimensions where data is loaded sequentially should be sparse.
- ❑ Dimensions with data that is reported together should be dense.

Large Accounts Dimensions

Marketing Technologies Group | www.mtgny.com

- Accounts is calculated first
- Sparse accounts dim require extra passes
- Larger block is usually better than a sparse accounts dimension
- Easily push block size to 160k (perhaps 2000 stored members) or more
- With dynamic calcs maybe 3000 or more

Large Time Dimensions

- Break apart “All Years” and “Year Total”
- Year Total = dense
- All Years = probably sparse

	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Year Total
2003	74	36	72	98	94	34	8	3	75	13	75	60	641
2004	11	121	85	41	41	21	35	44	34	141	95	93	763
2005	105	134	161	193	12	95	98	159	72	132	27	82	1,270
2006	29	30	7	23	187	28	104	242	9	220	234	37	1,150

Sparse Time Dimensions

Marketing Technologies Group | www.mtgny.com

- Probably good for time based incremental loads using intelligent calc
- Probably bad for growth rates, activity based cash flow statements etc.

Caches

- ❑ Index cache
- ❑ Data cache
- ❑ Data File Cache (Direct I/O only)
- ❑ Calculator Cache
- ❑ Dynamic Calc Cache
- ❑ Sort/Retrieval Buffer

Index Cache

- ❑ `ess00001.ind`
- ❑ Stores the index (the list of blocks)
- ❑ Set the Index cache equal to the size of the index
- ❑ Hit rate on Index cache should be near 1

Data Cache

- ❑ `ess0000n.pag`
- ❑ Stores the blocks (the data)
- ❑ Set the data cache equal to remaining available memory (available to application)
- ❑ Hit rate on Index cache should be as high as possible

Procedure Design

Marketing Technologies Group | www.mtgny.com

- Replace out 0's
- Agg missing
- Intelligent calc
- Parallel calc
- Level 0 calc v. upper level calc

Outline Optimize Menu Item

- Optimize menu item “automatically” optimizes outline by:
 - ▣ Dynamic calcs on dense
 - ▣ Label only non-aggregating parents
 - ▣ Dynamic calc on small sparse gen1

