Automatic Data Optimization

Saving Space
and Improving Performance!

Erik Benner, Enterprise Architect

Who am I?

Erik Benner

@erik_benner

TalesFromTheDatacenter.com

Enterprise Architect

Ebenner@mythics.com

- Published Author
- IOUG Solaris SIG Leader
- RAC Attack Ninja
- First Version of Oracle…7
- Linux since 1992
- Solaris since 1996
- OEM 12c since 2011
- Oracle Database Appliance since 2011
- DB 12c in prelaunch 2013

feed the
data
monster

MYTHICS™
ORACLE Platinum Partner

Growth in Data Diversity and Usage

1,800 Exabytes of Data in 2011, 20x Growth by 2020

Today's Drivers

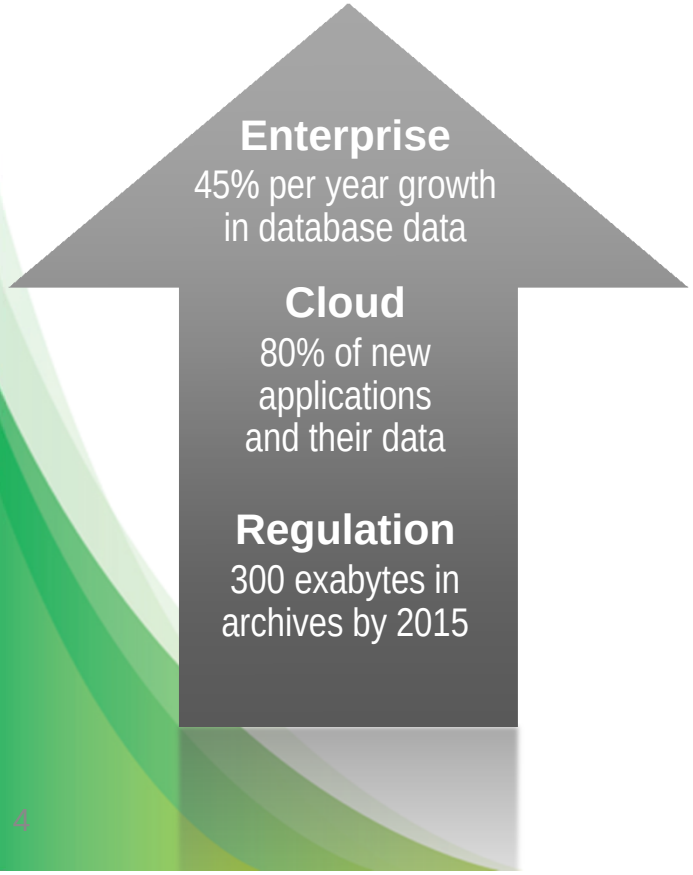

Emerging Growth Factors

Mobile
#1 Internet access device in 2013

Big Data
Large customers top 50PB

Social Business
$30B/year in commerce by 2015

MYTHICS™
ORACLE Platinum Partner

I know how to divide. What I want to learn is how to divide and conquer. That could come in handy one day."

MYTHICS™
ORACLE Platinum Partner

The Concept of Partitioning

Simple Yet Powerful

ORDERS

Large Table

Difficult to Manage

ORDERS

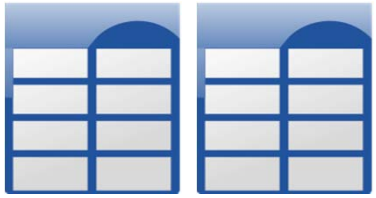

JAN FEB

Partition

Divide and Conquer

Easier to Manage

Improve Performance

ORDERS

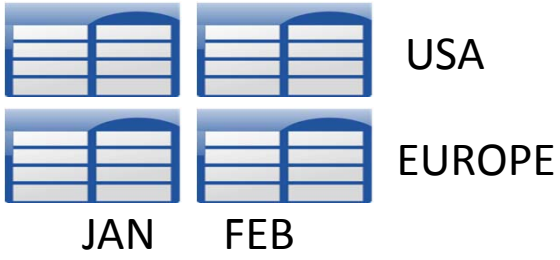

Composite Partition

Better Performance

More flexibility to match business needs

Transparent to applications

MYTHICS™ ORACLE Platinum Partner

What Can Be Partitioned?

- Tables
  - Heap tables
  - Index-organized tables
- Indexes
  - Global Indexes
  - Local Indexes
- Materialized Views

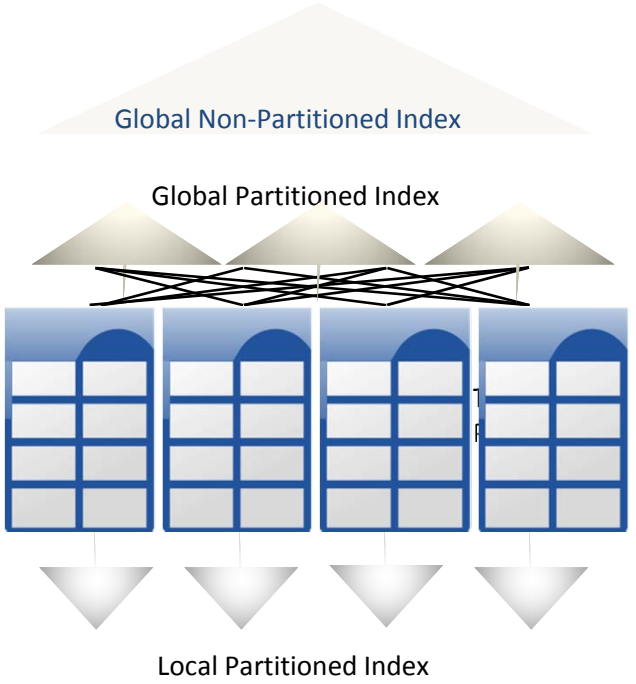

Partition for Tiered Storage

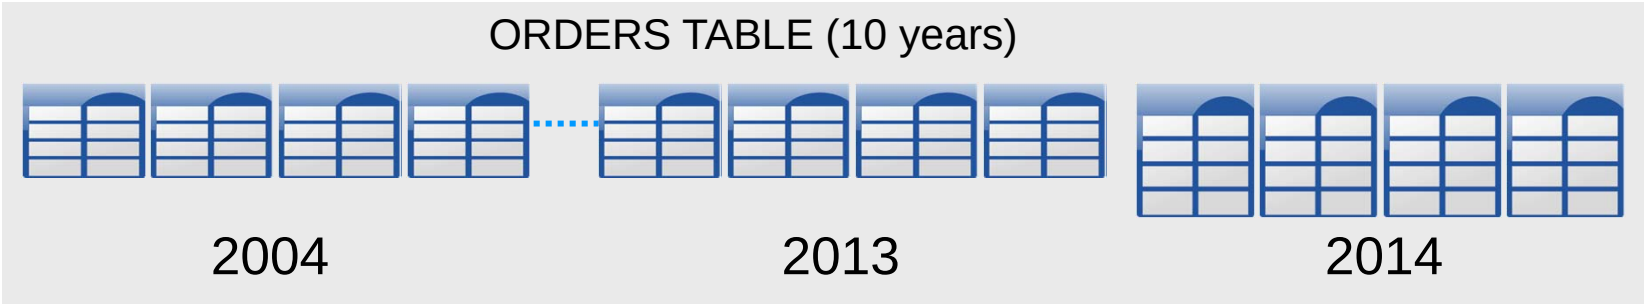

95% Less Active

5% Active

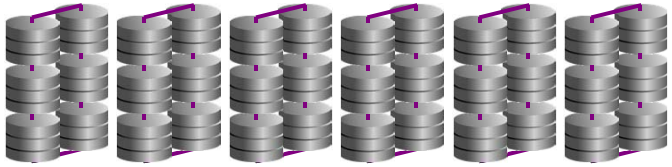

Low End Storage Tier
2-3x less per terabyte

High End Storage Tier

MYTHICS™
ORACLE Platinum Partner

MYTHICS™
ORACLE Platinum Partner

Compression Techniques

COMPRESSION TYPE:	SUITABLE FOR:
Basic Compression	"Mostly read" tables and partitions in Data Warehouse environments or "inactive" data partitions in OLTP environments
Advanced Row Compression	Active tables and partitions in OLTP and Data Warehouse environments
Advanced LOB Compression and Deduplication	Non-relational data in OLTP and Data Warehouse environments
Advanced Network Compression and Data Guard Redo Transport Compression	All environments
RMAN/Data Pump Backup Compression	All environments
Index Key Compression	Indexes on tables for OLTP and Data Warehouse
Hybrid Columnar Compression – Warehouse Level	"Mostly read" tables and partitions in Data Warehouse environments
Hybrid Columnar Compression – Archive Level	"Inactive" data partitions in OLTP and Data Warehousing environments

Advanced Row Compression

- **Partition/table/tablespace data compression**
  - Support for conventional DML Operations (INSERT, UPDATE)
  - *Customers indicate that 2x to 4x compression ratio's typical*
- **Significantly eliminates/reduces write overhead of DML's**
  - Batched compression minimizes impact on transaction performance
- **"Database-Aware" compression**
  - Does not require data to be uncompressed – keeps data compressed in memory
  - Reads often see improved performance due to fewer I/Os and enhanced memory efficiency

Hybrid Columnar Compression

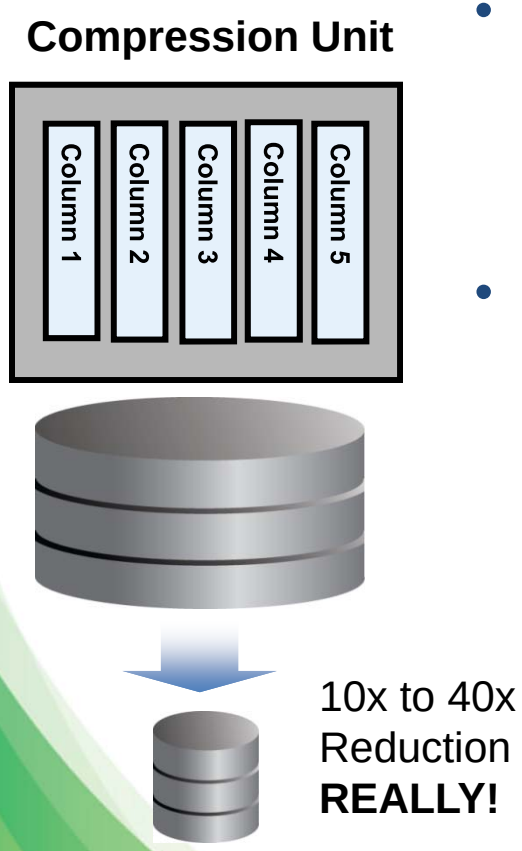

- **Hybrid Columnar Compressed Tables**
  - Compressed tables can be modified using conventional DML operations
  - Useful for data that is bulk loaded and queried
- **How it Works**
  - Tables are organized into Compression Units (CUs)
    - CUs are multiple database blocks
  - Within Compression Unit, data is organized by column instead of by row
    - Column organization brings similar values close together, enhancing compression

Data Compression

- Reduce storage footprint, read compressed data faster

3X
Advanced Row Compression

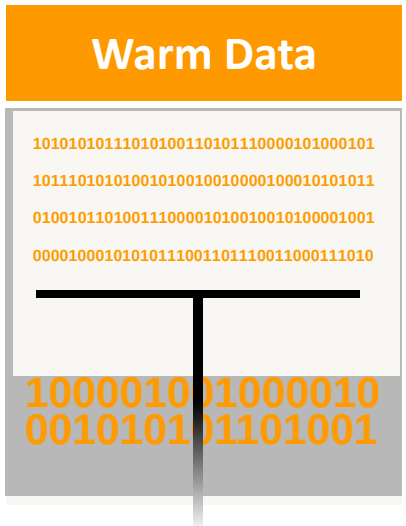

10X
Columnar Query Compression

15X +
Columnar Archive Compression

MYTHICS™
ORACLE Platinum Partner

Compression Benefits

- **Transparent: 100% Application Transparent**
- **Smaller: Reduces Footprint**
  - CapEx: Reduces server & storage costs for primary, standby, backup, test & dev databases …
  - OpEx: Reduces heating, cooling, floor space costs …
    - Additional ongoing savings over life of a database as database grows in size
- **Faster: Transactional, Analytics, DW**
  - Greater speedup from in-memory: 3-10x more data fits in buffer cache & flash cache
  - Faster queries
  - Faster backup & restore speeds
- **End-to-end Cost / Performance Benefits across CPU, DRAM, Flash, Disk & Network**

MYTHICS™
ORACLE Platinum Partner

Automatic Data Optimization (ADO)

- ADO offers capability to move and/or compress data based on observed usage patterns
- Uses heat maps to determine how often data has been access
- Tracks exactly how data has been utilized (DML vs. query, single-block vs. table scan)
- Data usage patterns can be tracked at tablespace, segment, and row level

Automatic Data Optimization

Simple Declarative SQL extension

▪ Advanced Row Compression (2-4x) ▪ Affects ONLY candidate rows ▪ Cached in DRAM & FLASH	row store compress advanced row **after 2 days of no update**
▪ Warehouse Compression(10x) ▪ High Performance Storage	column store compress for query low **after 1 week of no update**
▪ Warehouse Compression(10x) ▪ High Capacity Storage	tier to lower cost tablespace
▪ Archive Compression(15-40X) ▪ Archival Storage	column store compress for archive high **after 6 months of no access**

MYTHICS™ ORACLE Platinum Partner

Heat Maps: “How Hot Is It?”

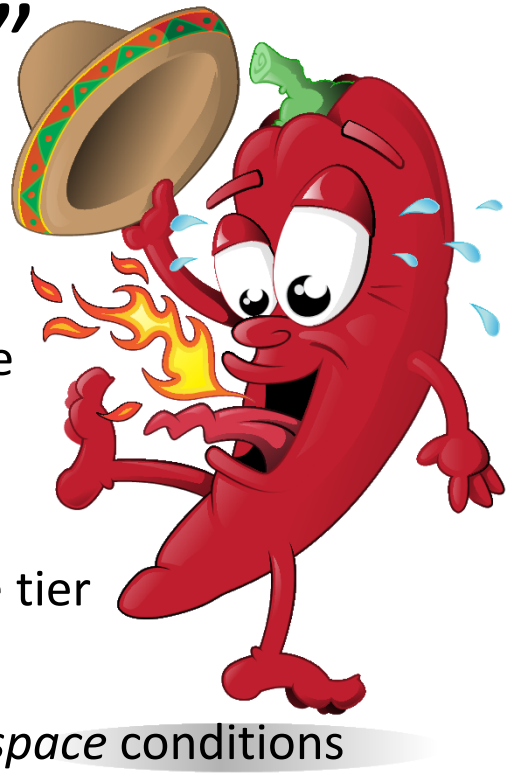

ADO leverages heat maps to:

- Capture data *usage frequencies*
- Determine which *compression level* is most appropriate for *how* data is being used
- Determine which data could be *moved* from a high-performance *storage tier* to a lower-performance tier
- Decide when data should be moved between *different tablespaces* or *partitions* to limit possible *out-of-space* conditions

Heat Maps: An Example

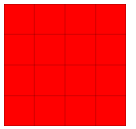

Heat map shows *heavy* DML and queries:
Leave data uncompressed

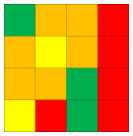

After7 days of more *limited* access:
Enable ADVANCED compression

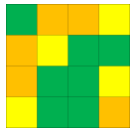

After 30 days of only *limited* access:
Enable HCC QUERY LOW* compression

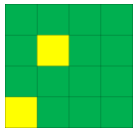

After 90 days of *no* access:
Enable HCC ARCHIVE HIGH* compression

* Requires Exadata Storage Servers, ZFS Appliance, or Pillar Axiom storage

MYTHICS™
ORACLE Platinum Partner

How much space will we save?

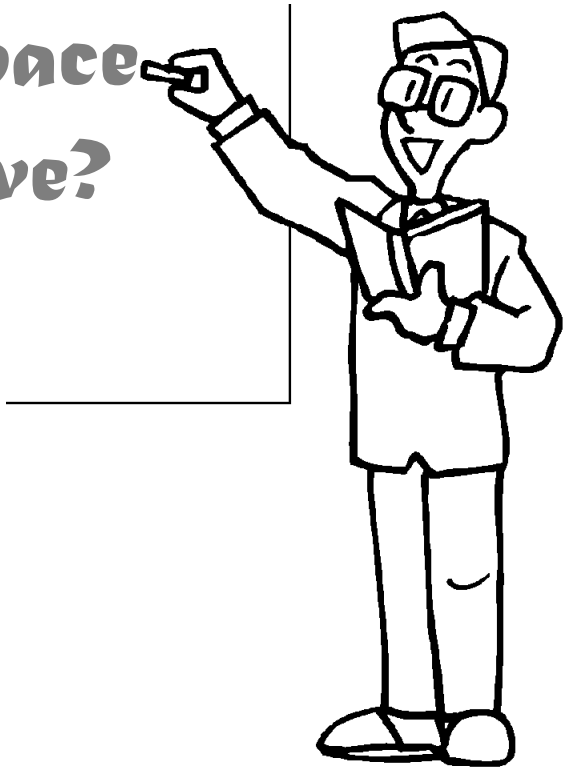

DBMS_COMPRESSION

Compliments of Oracle

The DBMS_COMPRESSION package gathers compression-related information within a database environment. This includes tools for estimating compressibility of a table for both partitioned and non-partitioned tables, and gathering row-level compression information on previously compressed tables. This gives the user with adequate information to make compression-related decision.

Built into 12c, can be downloaded for 11gR2

The Official docs

http://docs.oracle.com/database/121/ARPLS/d_compress.htm#ARPLS236

ADO.SQL

Constant	Value	Description
COMP_NOCOMPRESS	1	No compression
COMP_FOR_OLTP	2	OLTP compression
COMP_FOR_QUERY_HIGH	4	High compression level for query operations
COMP_FOR_QUERY_LOW	8	Low compression level for query operations
COMP_FOR_ARCHIVE_HIGH	16	High compression level for archive operations
COMP_FOR_ARCHIVE_LOW	32	Low compression level for archive operations

```
set serveroutput on

DECLARE
 blkcnt_cmp   PLS_INTEGER;
 blkcnt_uncmp PLS_INTEGER;
 row_comp     PLS_INTEGER;
 row_uncmp    PLS_INTEGER;
 cmp_ratio    NUMBER;
 comptype     VARCHAR2(30);
BEGIN
  dbms_compression.get_compression_ratio('ADO', 'ADO', 'ADO_DEMO', NULL,
32, blk
cnt_cmp, blkcnt_uncmp, row_comp, row_uncmp, cmp_ratio, comptype);

  dbms_output.put_line('Block Count Compressed:   ' || TO_CHAR(blkcnt_cmp));
  dbms_output.put_line('Block Count UnCompressed: ' ||
TO_CHAR(blkcnt_uncmp));
dbms_output.put_line('Compression Type:         ' || comptype);
END;
/
```

Execution

```
SQL> @ado
Compression Advisor self-check validation successful. select count(*)
on both
Uncompressed and EHCC Compressed format = 1000001 rows
Block Count Compressed:    8941
Block Count UnCompressed: 37986
Compression Type:          "Compress Archive Low"

PL/SQL procedure successfully completed.
```

MYTHICS™
ORACLE Platinum Partner