

ORACLE®

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

ORACLE®

Automatic Data Optimization with Oracle Database 12c

Gregg Christman
Senior Product Manager, ACO/HCC

ORACLE®
DATABASE 12^c




Plug into the **Cloud**.

Oracle Database 12c

Automatic Data Optimization

- Data Growth Trends
- Oracle Compression
- Heat Map
- Automatic Data Optimization
- Summary

Data Growth Challenges



- **IT must support exponentially growing amounts of data**
 - Explosion in online access and content
 - Government data retention regulations
- **Performance often declines as data balloons**
- **IT budgets are flat or decreasing**
- **Need to grow data**
 - Without hurting performance
 - Without growing storage/administrative cost
 - Storage tiering and compression tiering are key

Growth in Data Diversity and Usage

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

Today's Drivers

Enterprise

45% per year growth
in database data

Cloud

80% of new
applications
and their data

Regulation

300 exabytes in
archives by 2015

Emerging Growth Factors

Mobile

#1 Internet access device in 2013

Big Data

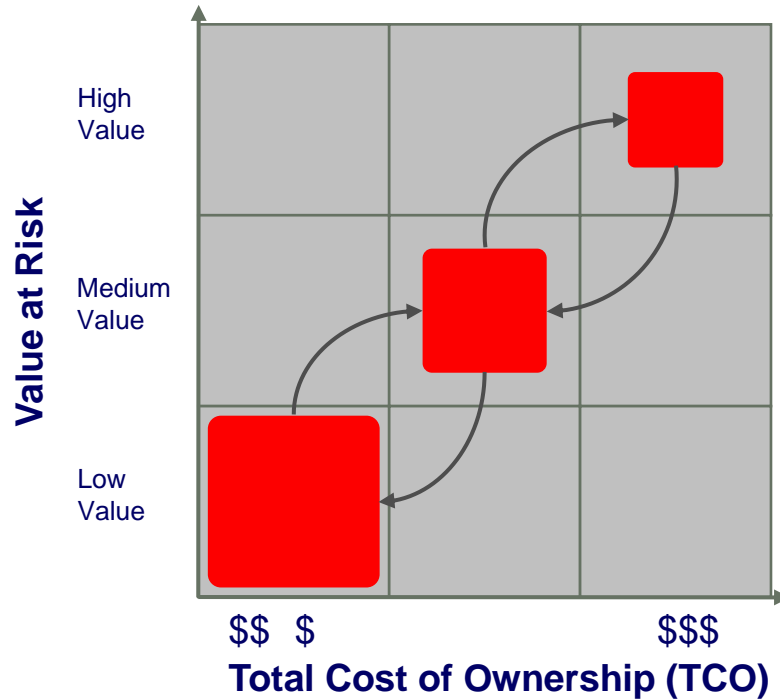
Large customers top 50PB

Social Business

\$30B/year in commerce by 2015

Information Lifecycle Management

Managing Data Over its Lifetime



“The policies, processes, practices, and tools used to align the business value of information with the most appropriate and cost effective IT infrastructure from the time information is conceived through its final disposition.”

*Storage Networking Industry Association
(SNIA) Data Management Forum*

Compression Techniques

COMPRESSION TYPE:	SUITABLE FOR:
Basic Compression	Read only 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	Read only tables and partitions in Data Warehouse environments
Hybrid Columnar Compression – Archive Level	“Inactive” data partitions in OLTP and Data Warehousing environments

Oracle Advanced Compression

Transparent, Smaller, *Faster*

- 100% Application Transparent
- End-to-end Cost/Performance Benefits across CPU, DRAM, Flash, Disk & Network
- Runs Faster: OLTP Apps (Transactional & Analytics) & DW
- Reduces Database Footprint
 - CapEx & OpEx savings
 - Increases Cloud ROI through Database Footprint reduction in DRAM Memory

Managing Storage Challenges

Manage more data
without incurring
additional cost

Compress data,
without impacting
performance

Tier and
compress data
based on usage

Oracle Advanced Compression

New Features, New Feature Names

	Oracle Database 11g	Oracle Database 12c
Oracle Advanced Compression	OLTP Compression	Advanced Row Compression
	Secure Files Compression	Advanced LOB Compression
	Secure Files De-duplication	Advanced LOB Deduplication
	Hybrid Columnar Compression	Hybrid Columnar Compression
	NEW	Heat Map (Object and Row Level)
	NEW	Automatic Data Optimization
	NEW	Temporal Optimization

Understanding Data Usage Patterns

Database 'heat map'

0100110101001101 1001011101010101 1001011101010101
0011010111010101 0101010101010101 0101010101010101
0111010101010110 0111010101010110 0111010101010110
1001000101010101 1001000101010101 1001000101010101
1001011101010101 1101000101011101 1001011101010101
0101010101010101 0101110101011101 0101010101010101
0111110100011110 0111110100011110 0111110100011110
1001000101010101 1001000101010101 1001000101010101
1011011100010101 1011011100010101 1011011100010101
0101010101010101 0101010101010101 0101010101010101
0101110000011111 0101110000011111 1001000101010101
1001000101010101 1001000101010101 1001000101010101
1001000101010101 1001011101010101 0111110101011111
0101010101010101 1101010101010101 1101010101010101
1001000101010101 0111110101011111 0111110101011111
1001100101010101 1001100101010101 1001100101010101

Understanding Data Usage Patterns

Database 'heat map'


0100110101001101 1001011101010101 1001011101010101
0011010111010101 0101010101010101 0101010101010101
0111010101010110 0111010101010110 0111010101010110
1001000101010101 1001000101010101 1001000101010101
1001011101010101 1101000101011101 1001011101010101
0101010101010101 0101110101011101 0101010101010101
0111110100011110 0111110100011110 0111110100011110
1001000101010101 1001000101010101 1001000101010101
1011011100010101 1011011100010101 1011011100010101
0101010101010101 0101010101010101 0101010101010101
0101110000011111 0101110000011111 1001000101010101
1001000101010101 1001000101010101 1001000101010101
1001000101010101 1001011101010101 0111110101011111
0101010101010101 1101010101010101 1101010101010101
1001000101010101 0111110101011111 0111110101011111
1001100101010101 1001100101010101 1001100101010101

12c Heat Map – Visualize Data Usage

Insight identifies opportunities & drives automation



Table Level Heat Map



Partition Level Heat Map



Row Level Heat Map

Heat Map

Enterprise Manager

ORACLE Enterprise Manager Cloud Control 12c

Setup Help SYSMAN Log Out

Enterprise Targets Favorites History Search Target Name

SALES Oracle Database Performance Availability Schema Administration

Logged in as SYS adc2101088.us.oracle.com

Information Lifecycle Management

Object Activity Policy

Top 100 Tablespace Activity Heat Map: SALES > Top 100 Object Activity Heat Map by Tablespace: SALES_USERS1

Top 100 Object Activity Heat Map by Tablespace: SALES_USERS1

View by: Last Modified Time

- Last 1 Week
- Last 1 Week - Last 1 Year
- > 1 Year

Search for Object Activity

Tablespace: SALES_USERS1

Schema: ALL

Name: ALL

Type: ALL

Partition:

Access Type: Last Modified Time

From: To:

Search

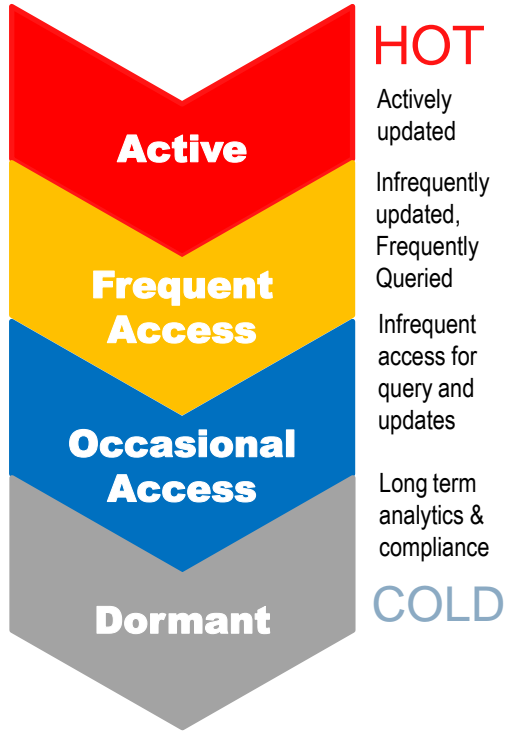
Search Results

Object Policy Tablespace Policy View Policy

Schema	Name	Partition	Type	Size (MB)	Tablespace	Policies
No search conducted.						

Heat Map

What it tracks



- **“Heat Map” tracking**

- Database level Heat Map shows which tables and partitions are being used
- Block level Heat Map shows last modification at the block level

- **Comprehensive**

- Segment level shows both reads and writes
- Distinguishes index lookups from full scans
- Automatically excludes stats gathering, DDLs or table redefinitions

- **High Performance**

- Object level at no cost
- Block level < 5% cost

Automatic Data Optimization

Usage based data compression

01110101010010
10000100010101

Hot Data

10101010111010100110101
11000010100010110111010
10100101001001000010001
01010110100101101001110
00010100100101000010010
00010001010101110011010
10100101001001000010001

3X

Advanced Row Compression

Warm Data

1010101011101010011010111000010100
01011011101010100101001001000001000
1010101101001011010011100001010010
0101000010010000100010101011010010
1000010100100101001010111000010
111001010010010100101011101101010

10X

Columnar Query Compression

Archive Data

10101010111010100110101110000101000101011
101010100101001001000010001010101101001011
010011100001010010010100001001000010001010
101010101110101001101011100001010001011011
101010101110101001101011100001011101011001

15X

Columnar Archive Compression

Automatic Data Optimization

Policy-based management

- **Declarative Policy Specification: Condition → Action**
 - alter table sales ilm add **row store compress advanced row after 3 days of no modification**;
 - Conditions are time period after creation, access, modification of data, or **customized**
 - Actions can be Compression Tiering or Tablespace Tiering
- **Policies are inherited from the tablespace or table**
 - New tables inherit from tablespace; can also be applied to existing tables
 - New partitions (including interval partitions) inherit from table

Simple Declarative SQL extension

Automatic Data Optimization

ALTER TABLE sales ILM add

Active	<ul style="list-style-type: none">▪ Advanced Row Compressed (2-4x)▪ Affects ONLY Candidate Rows▪ Cached in DRAM & FLASH	compress for advanced row after 2 days of no update
Frequent Access	<ul style="list-style-type: none">▪ HCC Query Compressed (10x)▪ High Performance Storage	compress for query low after 1 week of no update
Occasional Access	<ul style="list-style-type: none">▪ HCC Query Compressed (10x)▪ Low Cost Storage	tier to SATA Tablespace
Dormant	<ul style="list-style-type: none">▪ HCC Archive Compressed (15-50X)▪ Archival Storage	compress for archive high after 6 months no access

Up to 15x Smaller Footprint & Faster Queries

Automatic Data Optimization for OLTP

- Hybrid Columnar Compression now complements Advanced Row Compression
- Best Practice:
 - Step 1: **Use Advanced Row Compression for entire DB** and then
 - Step 2: **ADO automatically converts into HCC Query compressed** once the updates cool down, and is used mainly for reporting
 - => Query speed of columnar & 10x smaller footprint
 - Step 3: **ADO automatically converts into HCC Archive compressed** once data cools down further and is no longer frequently queried
 - => 15-50x smaller footprint

Optimizes Data Based on Heat Map

Automatic Data Optimization for DW

- Data generally comes in via Bulk Loading
- Workload dominated by queries, even during loading

Step 1: Bulk Load directly into HCC Query Compressed

- 10x smaller footprint, Query speed of Columnar

Step 2: ADO automatically converts to HCC Archive Compressed and moves to Lower Cost Storage once its queried infrequently

- Data remains online, with 15-50x smaller footprint, & lower storage cost

Fast, Flexible Loads & Queries on Columnar

Automatic Data Optimization – Mixed Use

- Fastest Load with uncompressed & Fastest Queries with columnar
 - Mixed workloads often use Java app or 3rd party tools to insert and update data that does not use Bulk Loads, so cannot use Columnar
- Step 1: **Load into uncompressed**, conventional inserts & updates
 - Fast loading, & flexibility of using a regular OLTP app for loading
- Step 2: **ADO moves to Row Compressed or Columnar Compressed or Low Cost Storage** once updates cool down
 - Faster Queries, 3-10x smaller footprint

Scheduled Policy Execution

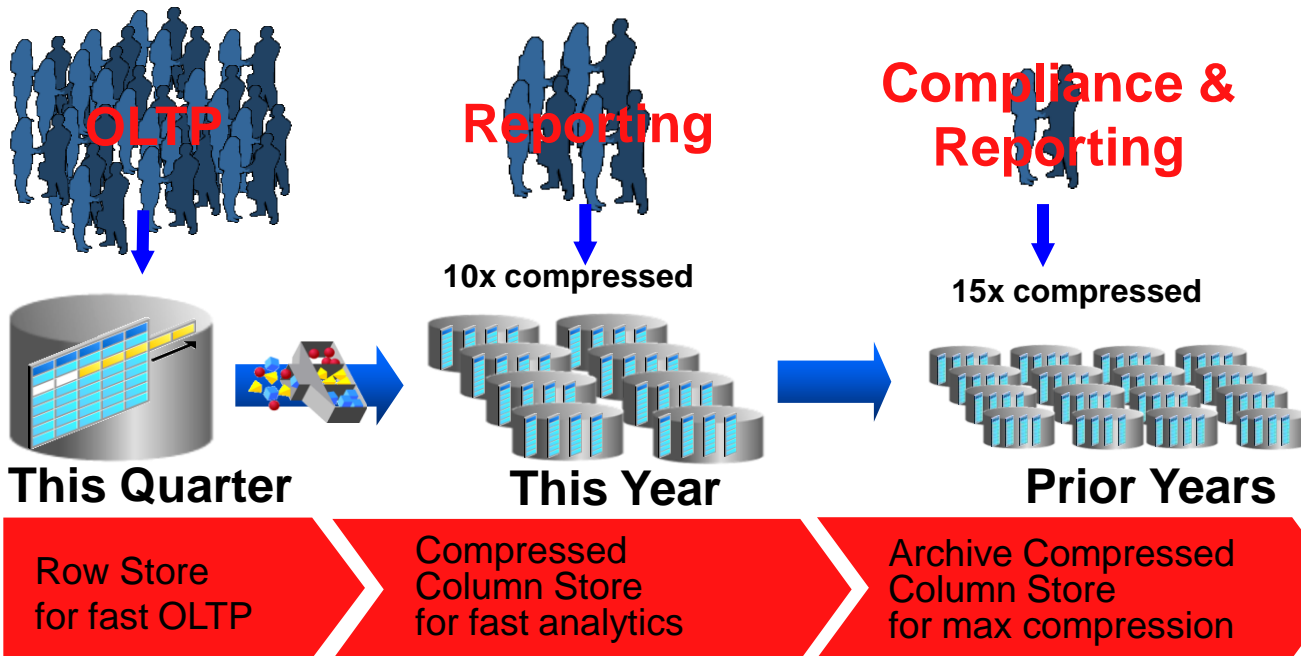
Automatic Data Optimization

- **Immediate and background policy execution**
 - Policies are executed in maintenance windows
 - Can be manually executed as needed

- **Policies can be extended to incorporate Business Rules**
 - Users can add custom conditions to control placement
(e.g. 3 months after the ship date of an order)

Automatic Data Optimization

As data cools, online conversion to columnar compression is automatic



As data ages:

- Activity declines
- Volume grows
- Older data primarily for reporting

alter table ... add policy

... compress for query after 3 months of no modification

... compress for archive after 1 year ...

Summary

Heat Map & Automatic Data Optimization

- **Heat Map**

- Automatically tracks access
- Database-aware: maintenance jobs, backups, etc don't affect Heat Map

- **Automatic Data Optimization**

- Declarative easy-to-use syntax to define data compression & movement policies
- Extensible with business-specific logic

ORACLE[®] 12^c
DATABASE



Plug into the **Cloud.**

ORACLE®