
ORACLE®

ORACLE

ORACLE DATABASE 12c

Compression and Performance in Oracle Database 12*c*

Kevin Jernigan
Product Manager – Advanced Compression++

Vineet Marwah
Product Development – Advanced Compression++

Plug into the Cloud.

Safe Harbor Statement

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.

Program Agenda

- Data Growth Challenges
- Advanced Compression Option
- Index Compression
- Customer Experience
- Questions

Data is Growing Exponentially

While IT Professionals Grow at 1.5X

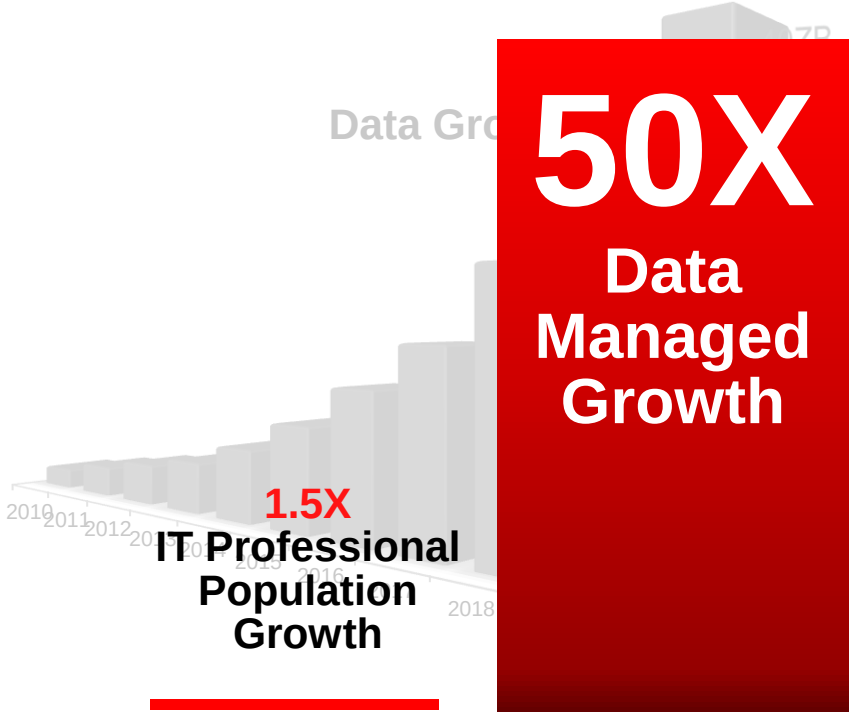

Source: IDC's Digital Universe Study, Dec 2012

Data Growth Challenges

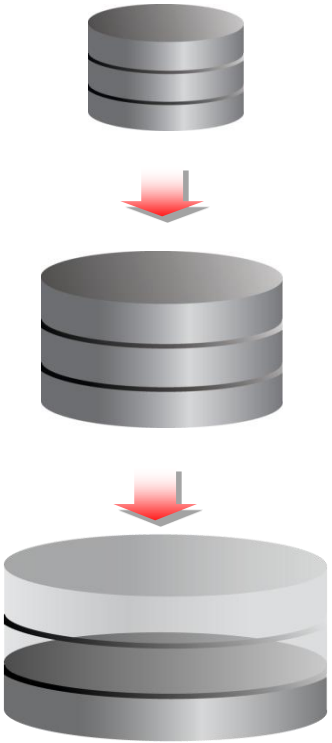

- **IT must support fast-growing amounts of data**
  - Explosion in online access and content
  - Government data retention regulations
- **Performance often declines as data grows**
- **Automate the compression and movement of data**
  - Reduce storage costs
  - Improve performance
- **Need to manage data growth**
  - Without hurting performance
  - Without growing costs
  - With minimal administrative intervention

Oracle Database Storage Optimization Features

Oracle Advanced Compression Option

Oracle Storage

FEATURE:	SUITABLE FOR:
Basic Table 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 Index Compression	Suitable for all supported indexes, not key dependent.
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 – columnar queries for analytics
Hybrid Columnar Compression – Archive Level	"Inactive" data partitions in OLTP and Data Warehousing environments – maximum compression

7

Oracle Database Storage Optimization Features

	FEATURE:	SUITABLE FOR:
Advanced Compression Option	Heat Map	Automatically track access and usage patterns at the block level (for modifications), and at the table or partition level (for reads and modifications)
	Automatic Data Optimization	Declarative policy language to automate compression and movement of data based on Heat Map data, or on custom conditions
	Storage Snapshot Optimization	Better support for storage snapshot-based backups (new in 12.1)
	Online Move Partition to Any Compressed Format	ONLINE Partition Move is new in 12.1
	Flashback Data Archive Optimization	Compression of Flashback Data Archive history tables (11.2.0.4+)
	Direct NFS Client - dNFS	Optimize configuration and performance of NAS storage
Oracle Storage	HCC Row Level Locking (RLL) – Oracle storage (requires Advanced Compression)	Row level locking for HCC-compressed data
	Oracle Intelligent Storage Protocol (OISP)	Optimize and automate configuration of ZFSSA for Oracle Database workloads, when using Direct NFS Client (dNFS)

8

Advanced Compression option

"Database Aware" Compression

- **Advanced Row Compression**
  - Compress data partitions/tables
- **Advanced LOB Compression**
  Advanced LOB Deduplication
  - Compress and deduplicate unstructured data
- **RMAN/DataPump Compression**
- **Data Guard Redo Transport Compression**
- **Advanced Index Compression**
- **Advanced Network Compression**
- **Flashback Data Archive History Table Optimization**

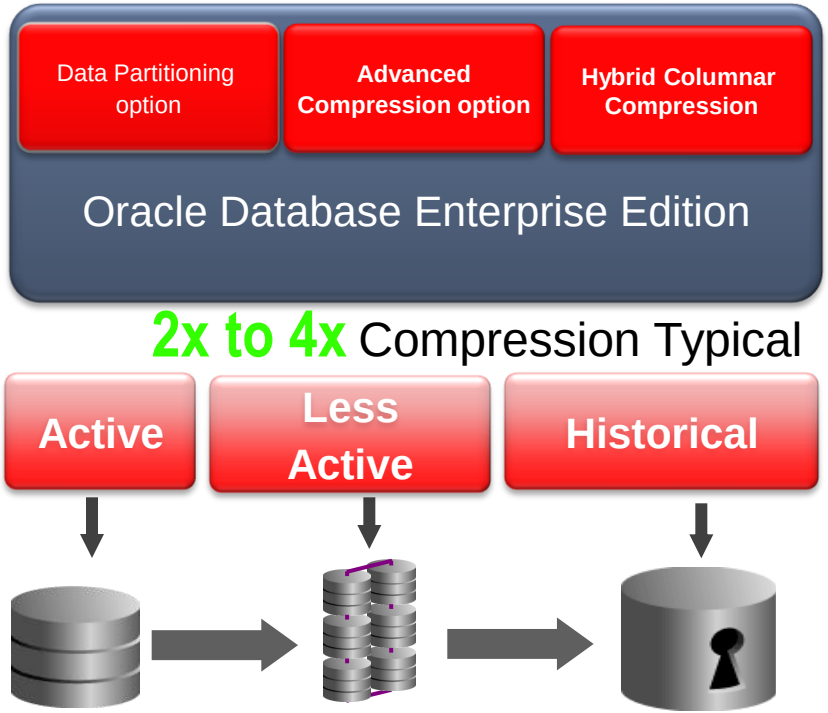

Advanced Row Compression

- **Partition/table/tablespace data compression**
  - Support for conventional DML Operations (INSERT, UPDATE)
  - *Customers indicate that 2x to 4x compression ratio 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

Advanced Row Compression

Employee Table

ID	FIRST_NAME	LAST_NAME
1	John	Doe
2	Jane	Doe
3	John	Smith
4	Jane	Doe

Initially Uncompressed Block

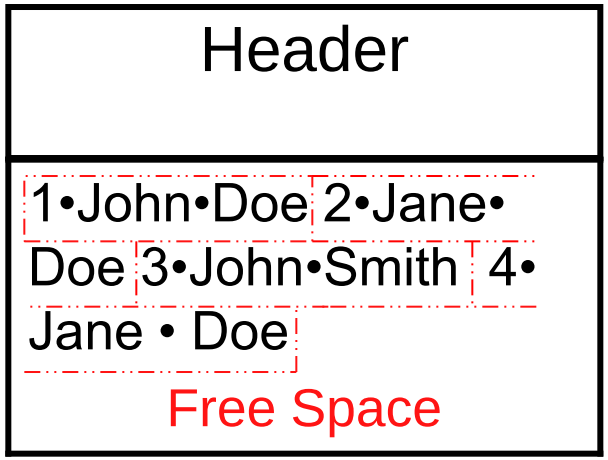


```
INSERT INTO EMPLOYEE
    VALUES (5, 'Jack', 'Smith');
COMMIT;
```

Advanced Row Compression

Employee Table

ID	FIRST_NAME	LAST_NAME
1	John	Doe
2	Jane	Doe
3	John	Smith
4	Jane	Doe
5	Jack	Smith

Compressed Block Header

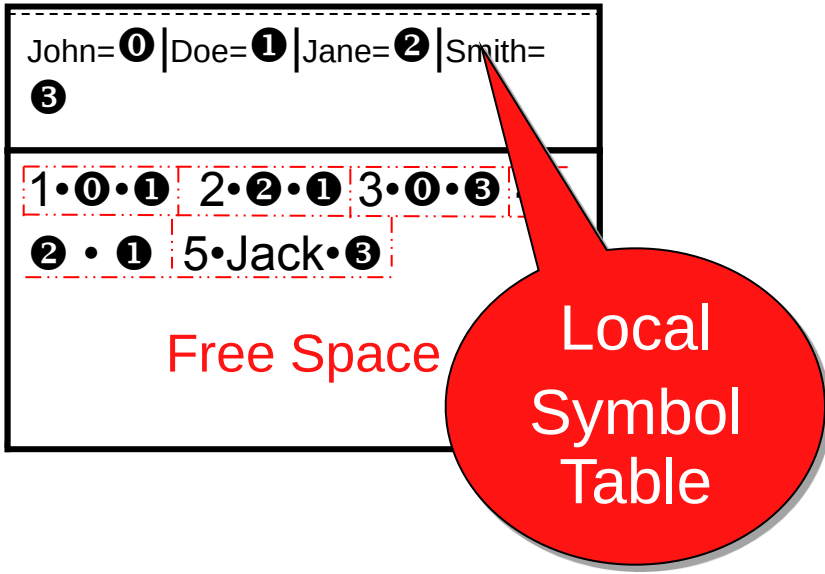

Block-Level *Batch* Compression

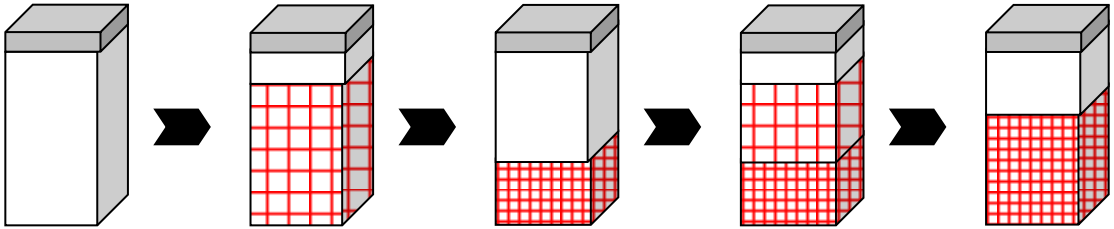

- Minimizes performance overhead and maximizes compression
- Individual INSERTs and UPDATEs do not cause recompression
- Compression cost is amortized over several DML operations
- Block-level (local) compression keeps up with frequent data changes in OLTP environments

Real World Compression Results - ERP Database 10 Largest Tables

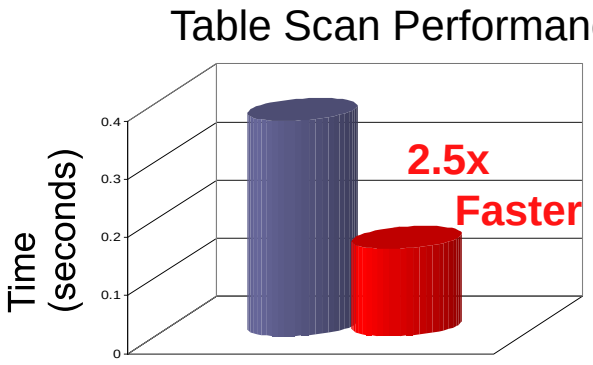

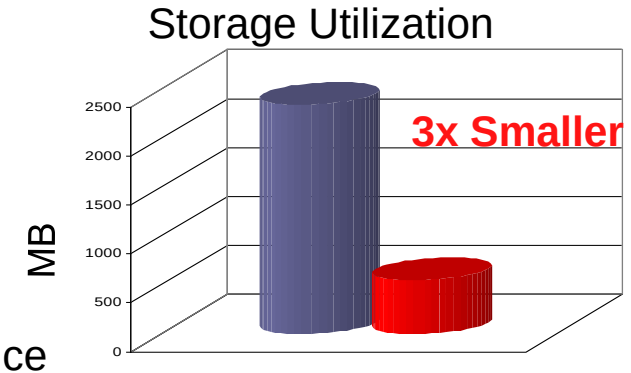

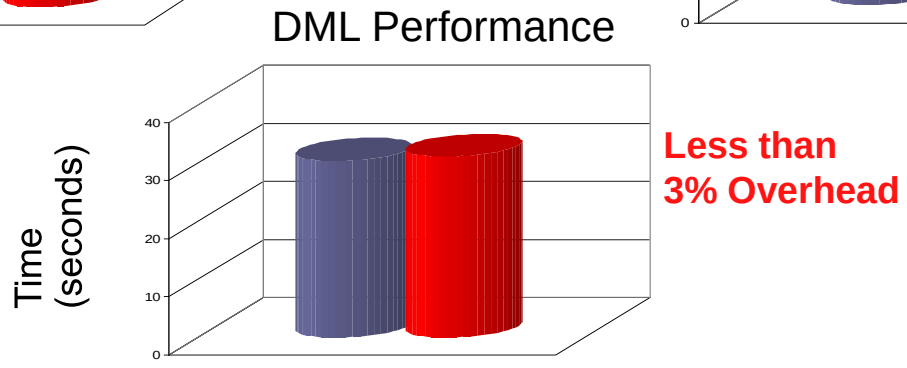

Advanced LOB Compression/Deduplication

- **LOBS typically experience a reduction of 2x to 3x times in size**
  - Automatically avoids compressing data that would not benefit from compression
  - Useful for content management, email applications and data archival applications
- **No adverse impact on read operations**
  - Often improves read performance for cache data
- **Enables storage of a single physical image for duplicate data**
  - Significantly reduces space consumption
  - Dramatically improves writes and copy operations

DataPump Compression

- **Metadata compression since Oracle Database 10*g***
- **Compression for table data during exports**
  - No need to decompress before import
- **Single step compression of both data and metadata**
  - Compressed data directly hits disk resulting in reduced disk space requirements
  - Internal tests reduced dump file size up to 75%
- **Application transparent**
  - Complete DataPump functionality available on compressed files

RMAN Backup Compression

- **Dramatically reduce the storage requirements for backup data**
  - Backup data is compressed before it is written to disk or tape and doesn't need to be uncompressed before recovery
- **Three levels of RMAN Compression: LOW, MEDIUM, and HIGH**
  - **Compression Level LOW**

    Best suited when backup is constrained by CPU
  - **Compression Level MEDIUM**

    Balance between CPU usage and compression ratio
  - **Compression LEVEL HIGH**

    Best compression ratio and highest CPU utilization

    Best suited when backup is constrained by network or I/O

Data Guard Redo Transport Compression

- **Compress network traffic between primary and standby databases**
- **Lower bandwidth networks (<100Mbps)**
  - 15-35% less time required to transmit 1 GB of data
  - Bandwidth consumption reduced up to 35%
- **High bandwidth networks (>100 Mbps)**
  - Compression will not reduce transmission time
  - But will reduce bandwidth consumption up to 35%

Advanced Network Compression

- **Compresses network data to be transmitted at the sending side and then uncompressed at the receiver**
  - Reduces the size of the session data unit (SDU) transmitted over a data connection
    - Reducing the size of data reduces the time required to transmit the SDU
    - On narrow bandwidth connections, with faster CPU, it can significantly improve performance
      - SQL query response becomes faster due to the reduced transmission time

Flashback Data Archive History Table Optimization

- **Flashback Data Archive - FDA**
  - Transparently tracks changes when they are made
  - Tracked in separate history tables
  - History tables are secure, tamper-proof
  - FDA History Table Optimization
    - Advanced Row Compression
    - Advanced LOB Compression and Deduplication
    - Compression Tiering
- **Use Cases**
  - Governance / regulatory / compliance (GRC)
  - Substitute for app development
  - Error correction
  - Replacement for CDC and DW

Index Key Compression

- **Included with Oracle Database Enterprise Edition**
- ***Customers report 2x compression is typical***
  - Compression ratio depends on how many columns are selected/compressibility of those columns
- **ANALYZE INDEX will give advice on whether / how many columns to choose**
  - Index data is NOT decompressed when read from disk into memory
- **Does not require ACO or an ACO license – can compress just indexes and not data**

Advanced Index Compression

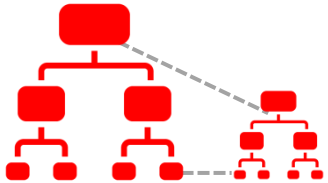

- **Advanced Index Compression automatically chooses the right compression per block**
- **Compresses indexes to reduce their overall storage requirement**
  - Less space required on disk
  - Better use of the database cache
- **Average compression ratio for indexes is 2x to 3x**
- **Little or no discernible overhead**
- **Compression Advisor extended to provide estimated compression ratio**

Customer Experience -- CSX

Improves Data Compression by 7x, Reduces Database Storage Requirements by 21%

Business Objective:

Reduce database footprint, while ensuring high performance and business continuity.

CSX maintains a large Oracle Database footprint—with more than 400 databases supporting critical, commercially-available, off-the-shelf and proprietary business applications, including payroll, dispatching, and a customer-facing order entry systems.

Results:

- Lowered its database storage requirements by 21% and stemmed its storage growth by 19%
- Achieved a compression ratio range of 45% to 70%, freeing up approximately 1.4 TBs of disk space while reducing data storage demands.
- Improved batch job completion time by 30%, freeing available resources to accommodate additional jobs.
- Improved query performance by 20% and enabled CSX to complete database backups 33% faster.
- In the data warehouse environment ACO was applied selectively to older partitions of the largest tables. This was done online with zero downtime.

"With Oracle Database's Advanced Compression option, we have achieved sevenfold data compression, reduced our database storage requirements by 21%, and can complete batch jobs 30% faster."

– Maritza Gonzalez, Technical Director, Data Management, CSX Corporation

http://www.oracle.com/us/corporate/customers/customersearch/csx-1-oac-sl-2005803.html

Other Customers….

- AmerisourceBergen Corporation
  - Reduced the storage footprint for the company's SAP ECC application and SAP business warehouse environment by up to 60% with no negative impact on performance while significantly lowering storage-management costs to achieve a return on investment in six to nine months
- Suguna Foods (EBS)
  - "We chose Oracle E-Business Suite Release 12.1.3 because it offered a robust, easy-to-use, secure, scalable, and integrated enterprise resource planning platform that could modernize our poultry business and support future growth. Together with Oracle Database's Advanced Compression option, we've reduced our cost-accounting time by three days, increased storage capacity by 60%, and cut storage costs by US $32,700 in the first year." – S. Karunanithi, CIO, Suguna Foods Ltd.

http://www.oracle.com/us/products/database/options/advanced-compression/overview/index.html

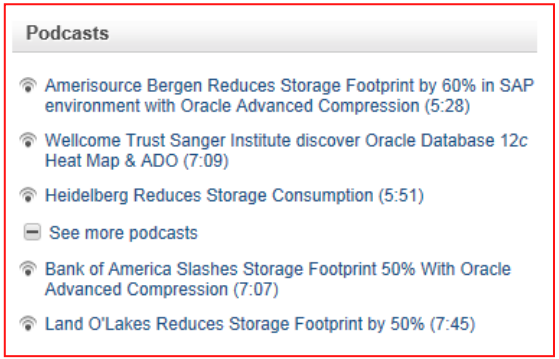

Oracle Database Storage Optimization Features

	FEATURE:	SUITABLE FOR:
Advanced Compression Option	Heat Map	Automatically track access and usage patterns at the block level (for modifications), and at the table or partition level (for reads and modifications)
	Automatic Data Optimization	Declarative policy language to automate compression and movement of data based on Heat Map data, or on custom conditions
	Storage Snapshot Optimization	Better support for storage snapshot-based backups (new in 12.1)
	Online Move Partition to Any Compressed Format	ONLINE Partition Move is new in 12.1
	Flashback Data Archive Optimization	Compression of Flashback Data Archive history tables (11.2.0.4+)
	Direct NFS Client - dNFS	Optimize configuration and performance of NAS storage
Oracle Storage	HCC Row Level Locking (RLL) – Oracle storage (requires Advanced Compression)	Row level locking for HCC-compressed data
	Oracle Intelligent Storage Protocol (OISP)	Optimize and automate configuration of ZFSSA for Oracle Database workloads, when using Direct NFS Client (dNFS)

Automatic Data Optimization

Usage Based Data Compression

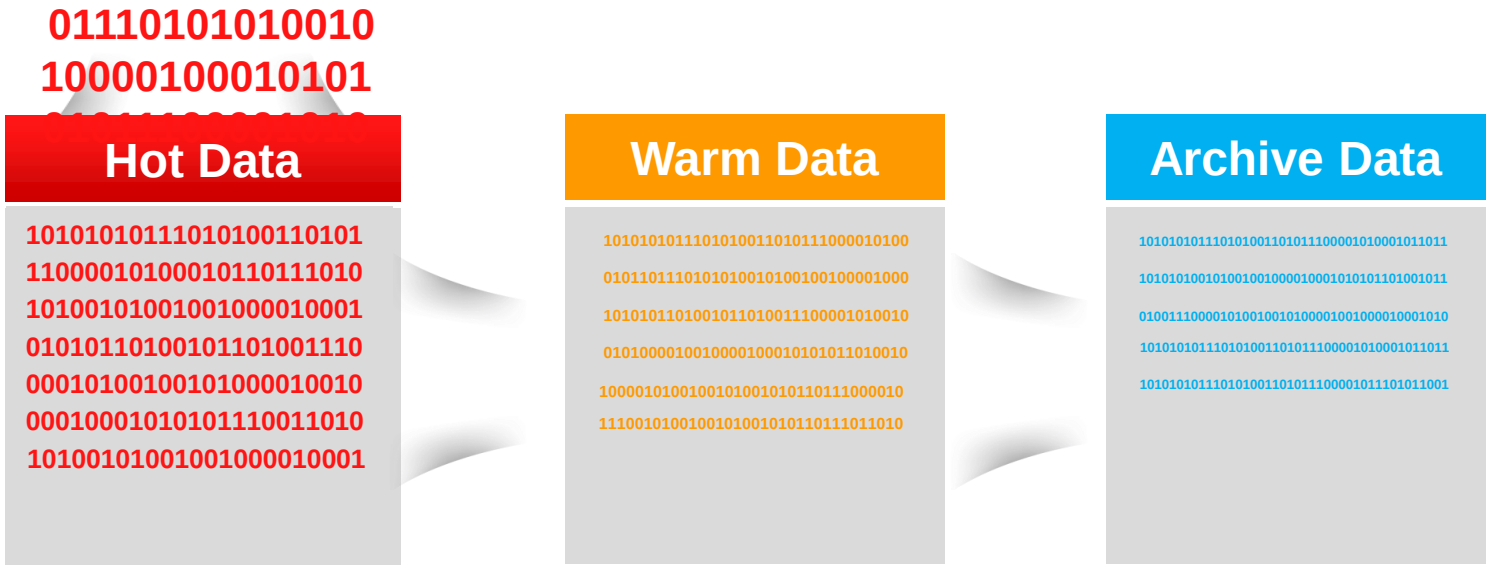

3X
Advanced Row Compression

10X
Columnar Query Compression

15X
Columnar Archive Compression

Heat Map

- **Heat Map gives you a detailed view of how your data is being accessed, and how access patterns are changing over time**

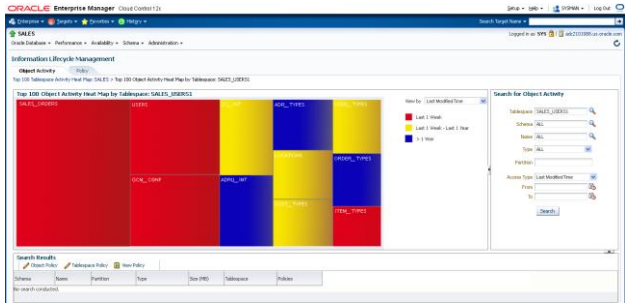

- **Programmatic access to Heat Map data is available through a set of PL/SQL table functions, as well as through data dictionary views**
  - Data modification times are tracked at the row level and aggregated to the block level
  - Modification times, full table scan times, and index lookup times are tracked at the segment level

Automatic Data Optimization (new/existing tables)

	ALTER TABLE sales ILM add
▪ Advanced Row Compressed (2-4x) ▪ Affects ONLY Candidate Rows ▪ Cached in DRAM & FLASH	**row store compress advanced row after 2 days of no modification**
▪ HCC Query Compressed (10x) ▪ High Performance Storage	**compress for query low after 1 week of no modification**
▪ Advanced/HCC Compressed ▪ Low Cost Storage	tier to low_cost_storage Tablespace (Must be Exadata, ZFSSA, or Axiom for HCC)
▪ HCC Archive Compressed (15-50X) ▪ Archival Storage	**compress for archive high after 6 months no access**

Automatic Data Optimization

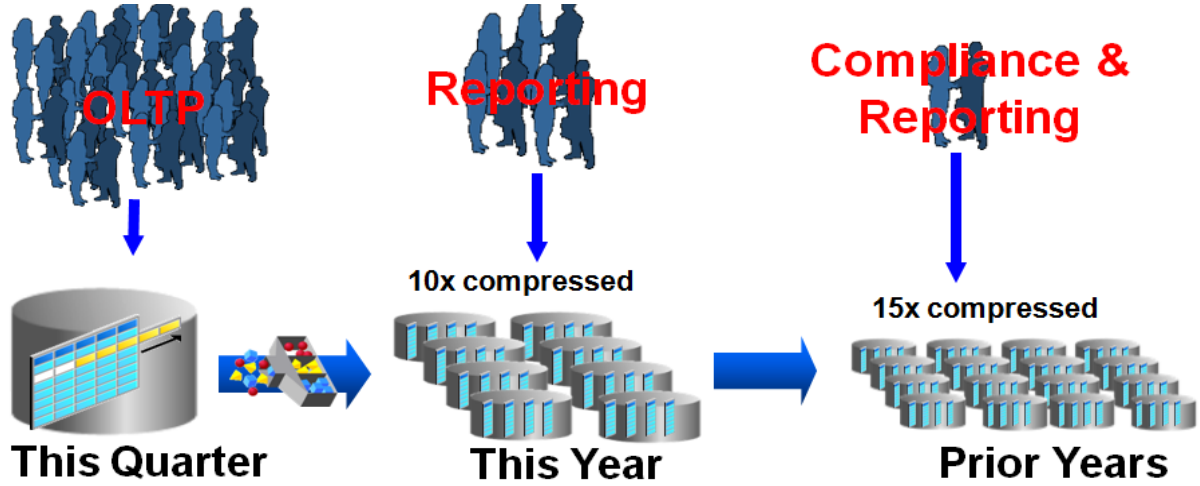

As data ages:

- 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 of no modification**

As data cools down, Automatic Data Optimization automatically converts data to columnar compressed

Advanced Compression Benefits

- **Transparent: 100% Application Transparent**
- **Smaller: Reduces Footprint**
  - CapEx: Lowers server & storage costs for primary, standby, backup, test & dev databases …
  - OpEx: Lowers 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**

Questions?

Hardware and Software

ORACLE

Engineered to Work Together


ORACLE®