
ORACLE®

ORACLE DATABASE 11g

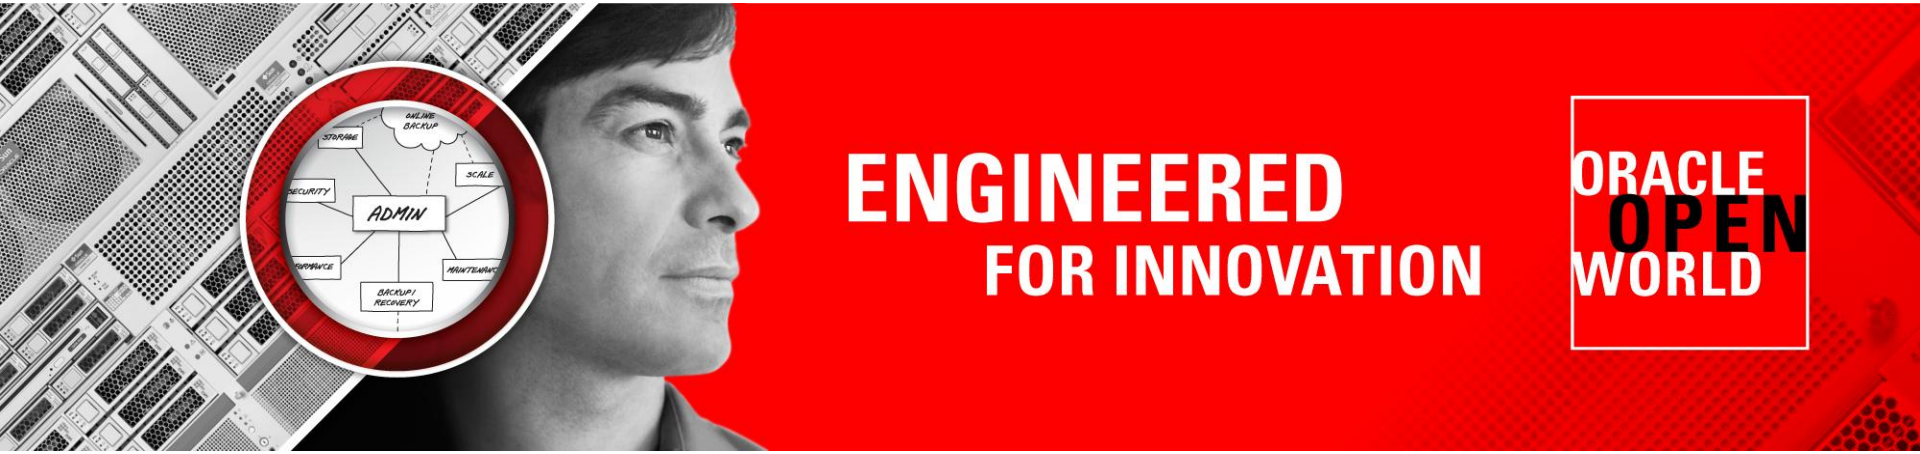

ORACLE®

Oracle Advanced Compression: Deep Dive into OLTP Table Compression

Gregg Christman -- Senior Product Manager

Presenting with

THOMSON REUTERS

TVS

Bank of America

ORACLE OPEN WORLD

ENGINEERED FOR INNOVATION

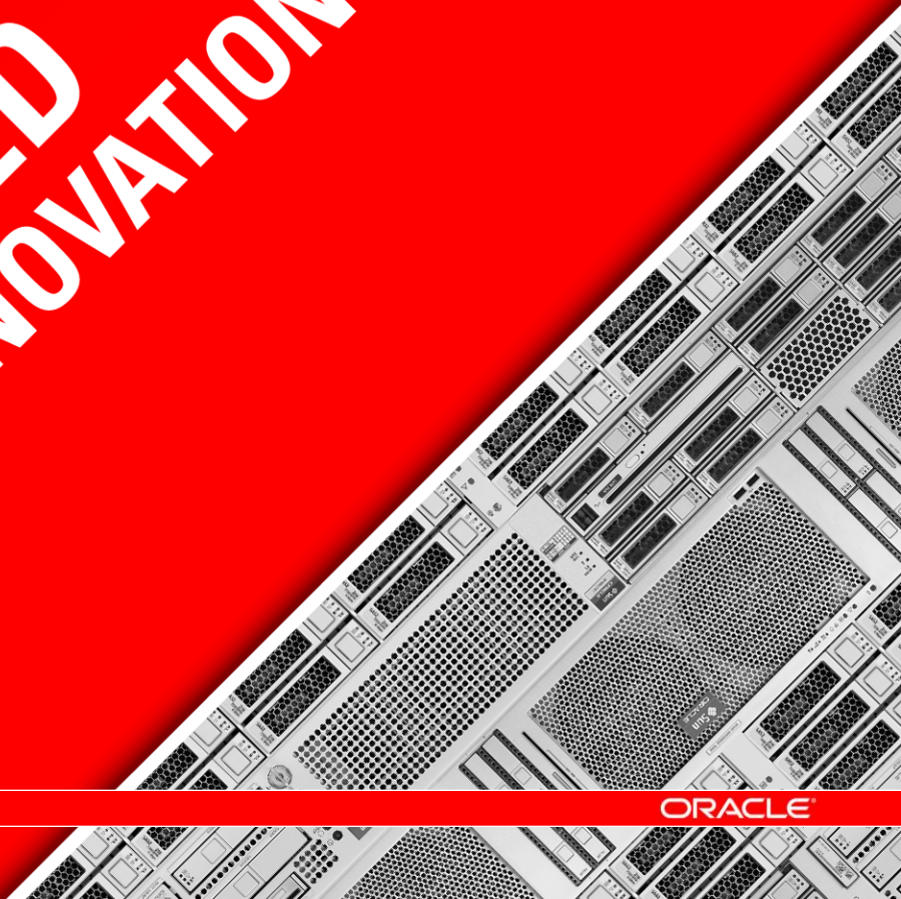

ORACLE

Latin America 2011

December 6–8, 2011

Tokyo 2012

April 4–6, 2012

Oracle OpenWorld Bookstore

- Visit the Oracle OpenWorld Bookstore for a fabulous selection of books on many of the conference topics and more!
- Bookstore located at Moscone West, Level 2
- All Books at 20% Discount

DigitalGuru
Technical Bookshop

Program Agenda

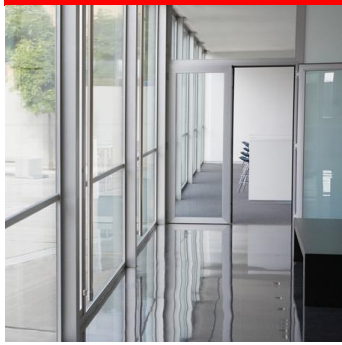

- Data Growth Challenges
- OLTP Table Compression
- Customer Experience's
- Questions and Answers

Data Growth Challenges

- **Explosion in Data Volumes**
  - Government Regulations (Sarbanes-Oxley, etc)
  - User Generated Content (Web 2.0)
  - Application Consolidation
- **IT Managers Must Support Larger Volumes of Data with Limited Technology Budgets**
  - Need to optimize storage consumption
  - Also maintain acceptable application performance
- **Intelligent and Efficient Compression Technology can Help Address These Challenges**

Oracle Advanced Compression Option

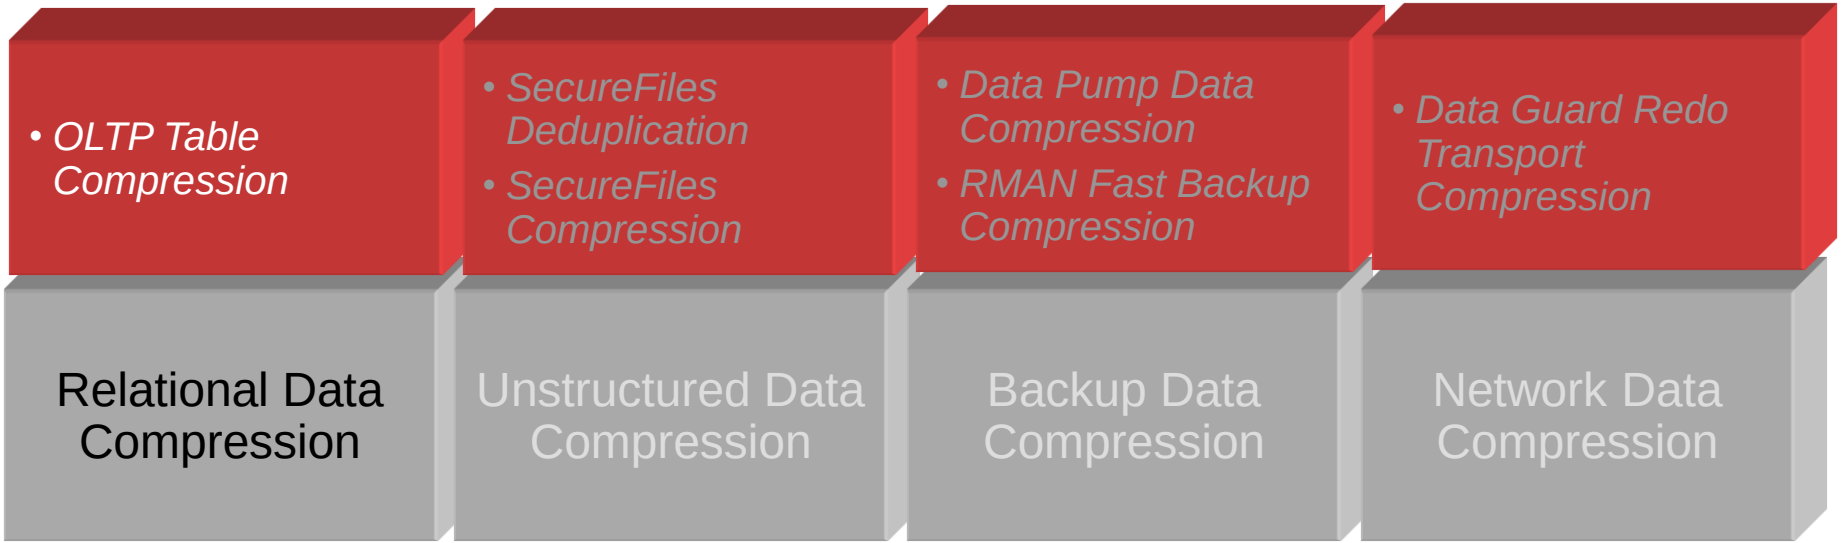

Basic Table Compression

- **Introduced in Oracle Database 9*i* Release 2**
  - Compression during bulk load operations (Direct Load, CTAS)
  - Data modified using conventional DML not compressed
- **Improved performance for queries accessing large amounts of data**
  - Fewer I/Os
  - Buffer Cache efficiency
- **Data is compressed at the database block level**
- **Compression enabled at either the table or partition level**
- **Completely transparent to applications**

OLTP Table Compression

- **Oracle Database 11*g* extends table compression for OLTP data**
  - Support for conventional DML Operations (INSERT, UPDATE)
  - 2x to 4x compression ratio's typical
- **New algorithm significantly reduces write overhead**
  - Batched compression minimizes impact for OLTP transactions
- **No impact on reads**
  - Oracle Database does not require data to be uncompressed – it keeps data compressed in memory
  - Reads may actually see improved performance due to fewer I/Os and enhanced memory efficiency

OLTP Table Compression

Employee Table

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

Initially Uncompressed Block

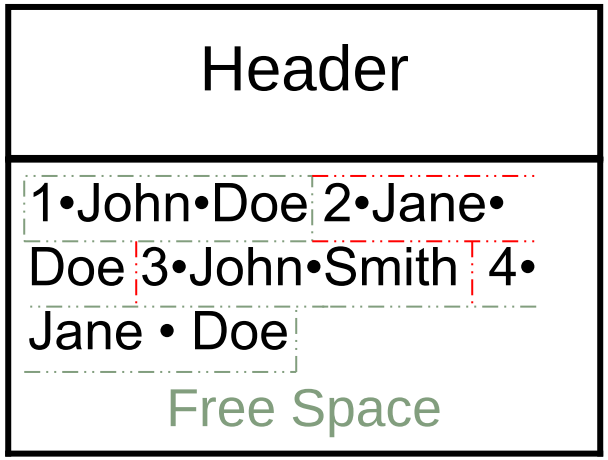


```
INSERT INTO EMPLOYEE
    VALUES (5, 'Jack', 'Smith');
COMMIT;
```

OLTP Table 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

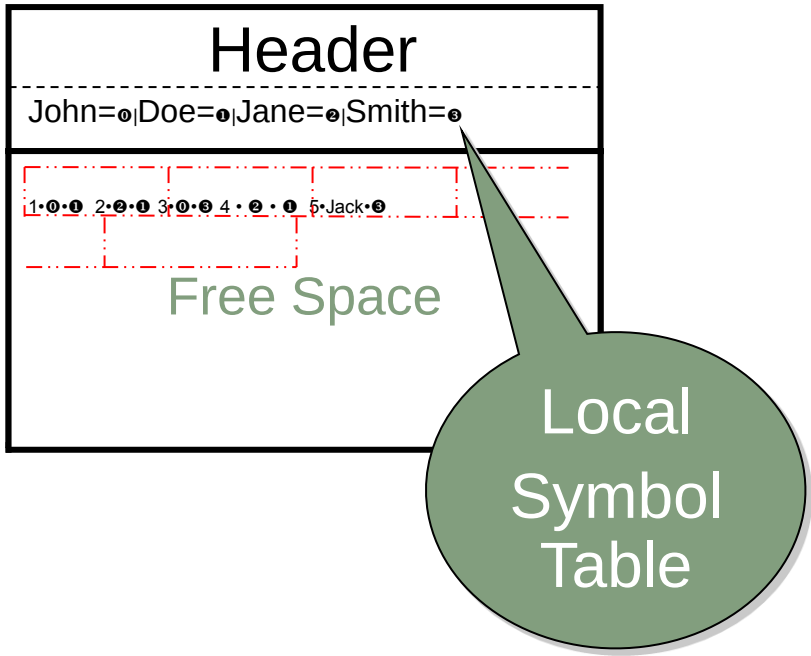

OLTP Table Compression Process

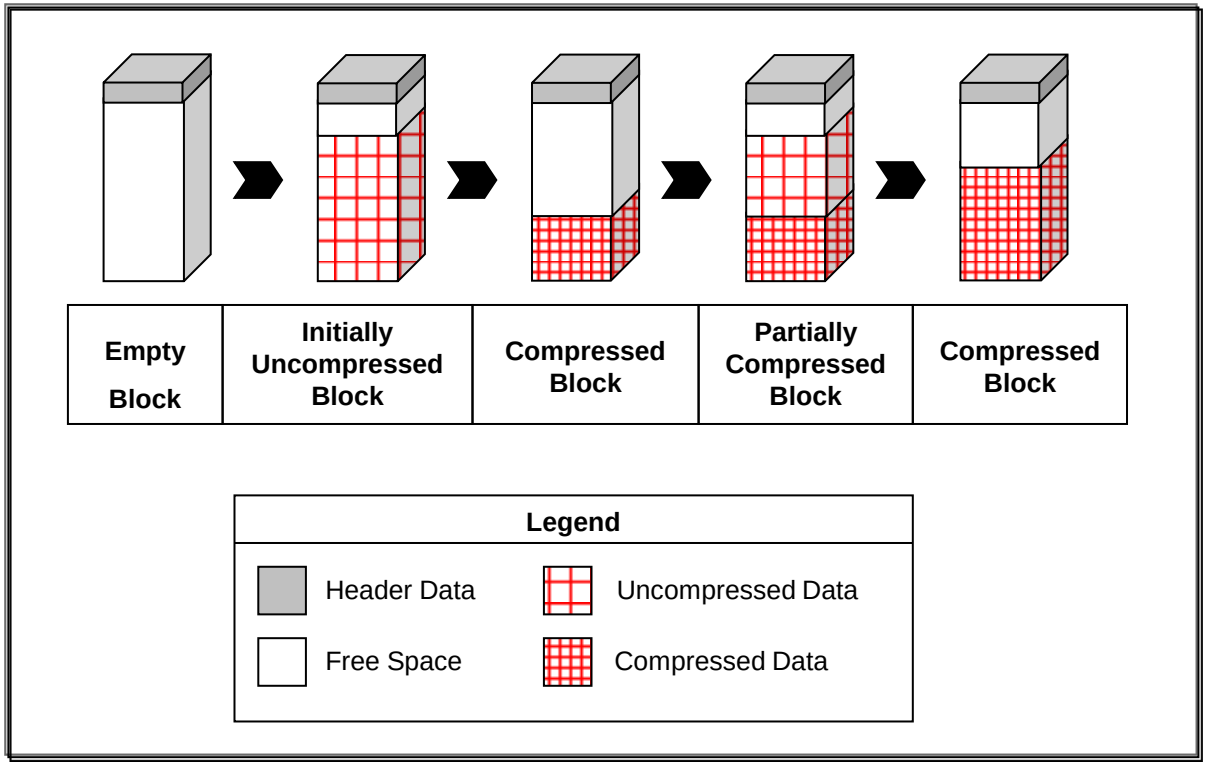

Block-Level *Batch* Compression

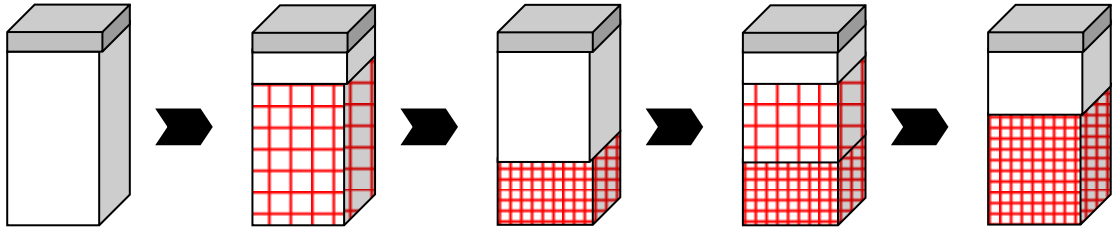

- Patent pending algorithm minimizes performance overhead and maximizes compression
- Individual INSERTs and UPDATEs do not cause recompression
- Compression cost is amortized over several DML operations
- Block-level (Local) compression keeps up with frequent data changes in OLTP environments

Large Objects and Compression

- Use SecureFiles for LOB’s
- A LOB locator always exists for any LOB instance regardless of the LOB storage properties or LOB value - NULL, empty, or otherwise
  - This means that each LOB value is unique because there's a unique LOB ID stored in the beginning of the LOB
  - So OLTP compression won't find any duplicate LOB values
  - SecureFiles can provide both compression and deduplication benefits for LOBs stored out-of-line.

Monitoring Table Compression

- View: DBA_TABLES, columns:
  - COMPRESSION (**ENABLED** / **DISABLED**)
  - COMPRESS_FOR (**OLTP** / **BASIC**)

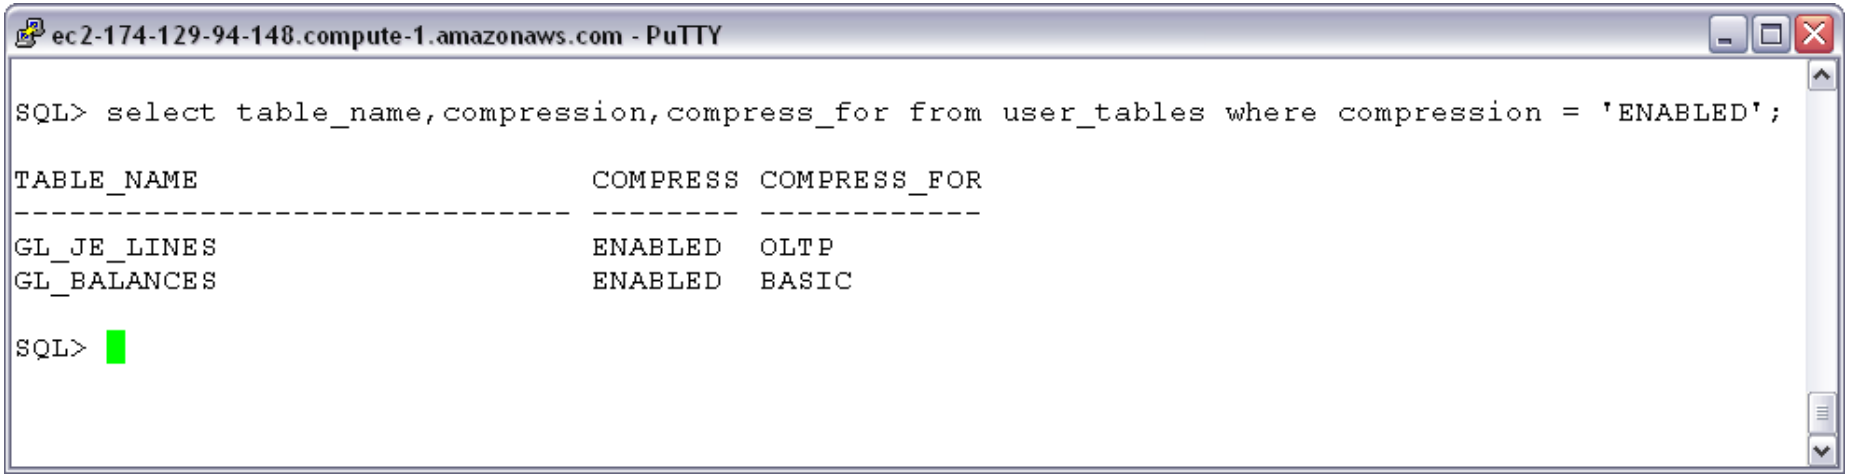

Table Compression Advisor

- Available in 11g Release 2
- Available on OTN *
  - Supports Oracle Database 9*i* Release 2 through 11g Release 1
  - Shows projected compression ratio for uncompressed tables
  - Reports actual compression ratio for compressed tables (11*g* Only)

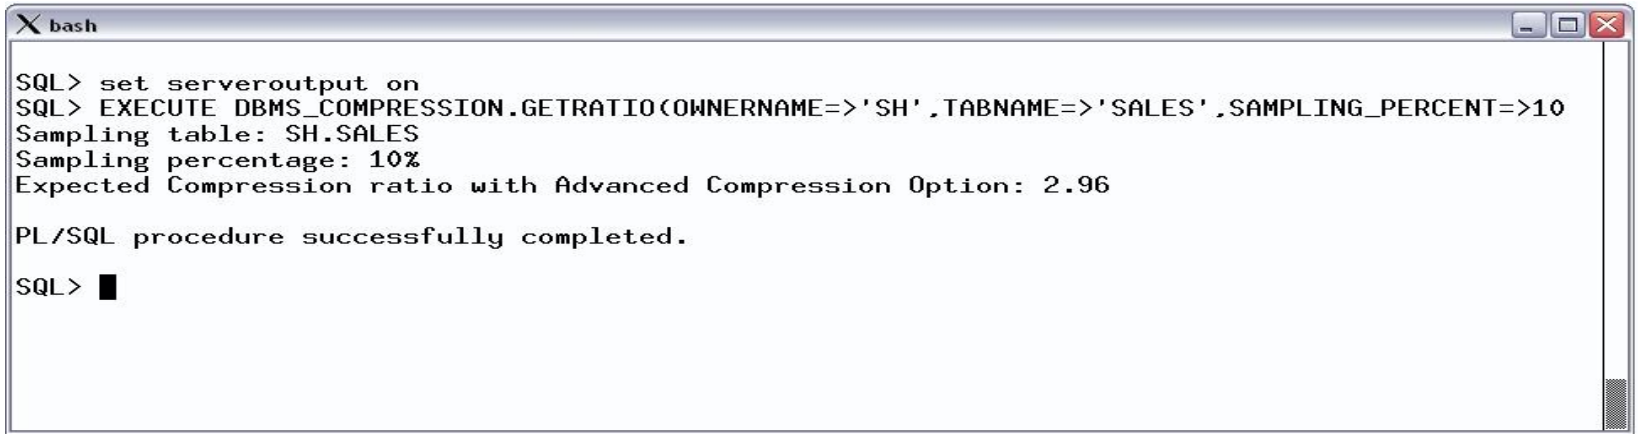

Customer Experience's

Presenting with

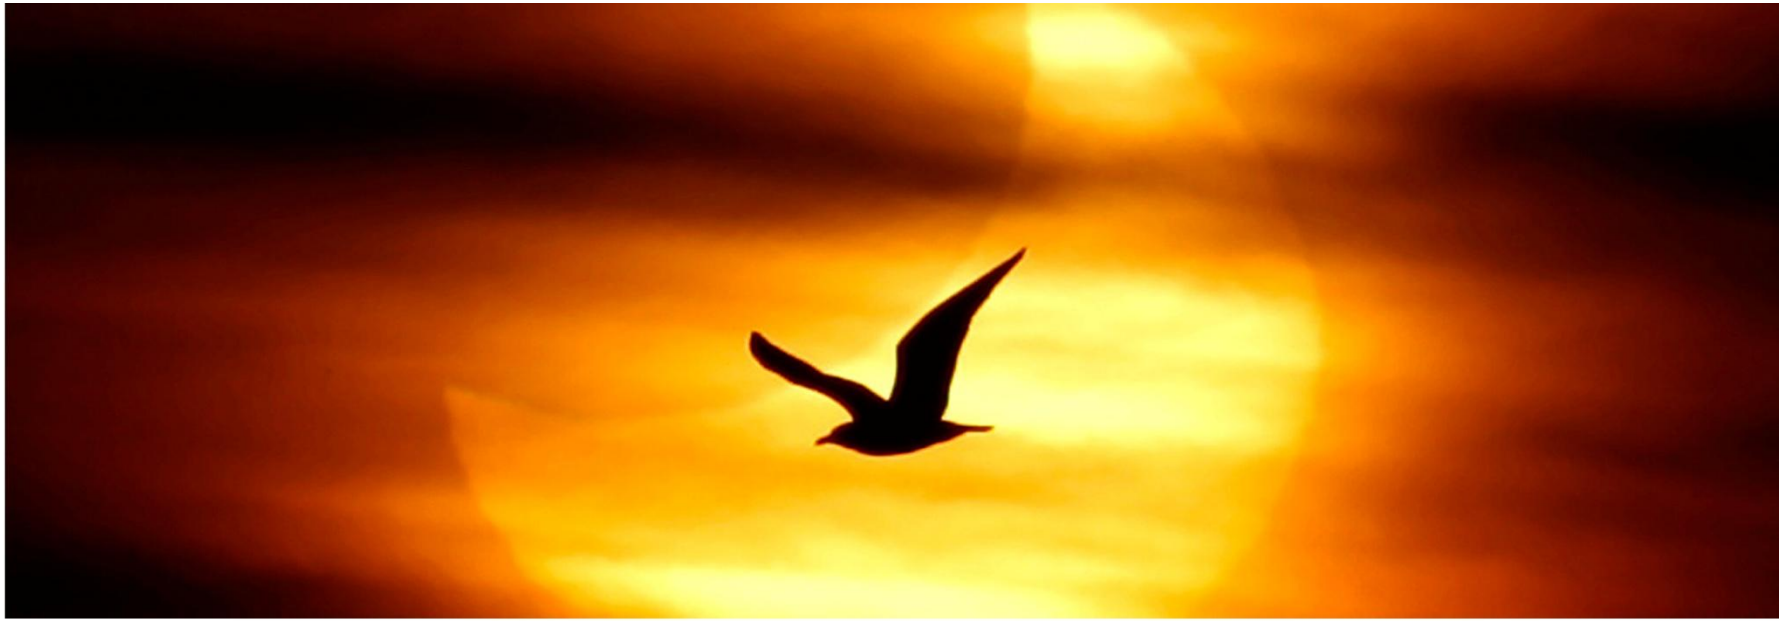

REUTERS/Jon Nazca

Oracle Advanced Compression

Compression Implementations

Dan Dressel, Architect Thomson Reuters

THOMSON REUTERS

Business Challenge

The Knowledge Effect

- New search algorithms that emulate the best practices of legal researchers to deliver in minutes what used to take hours.

THOMSON REUTERS

Processing Flow

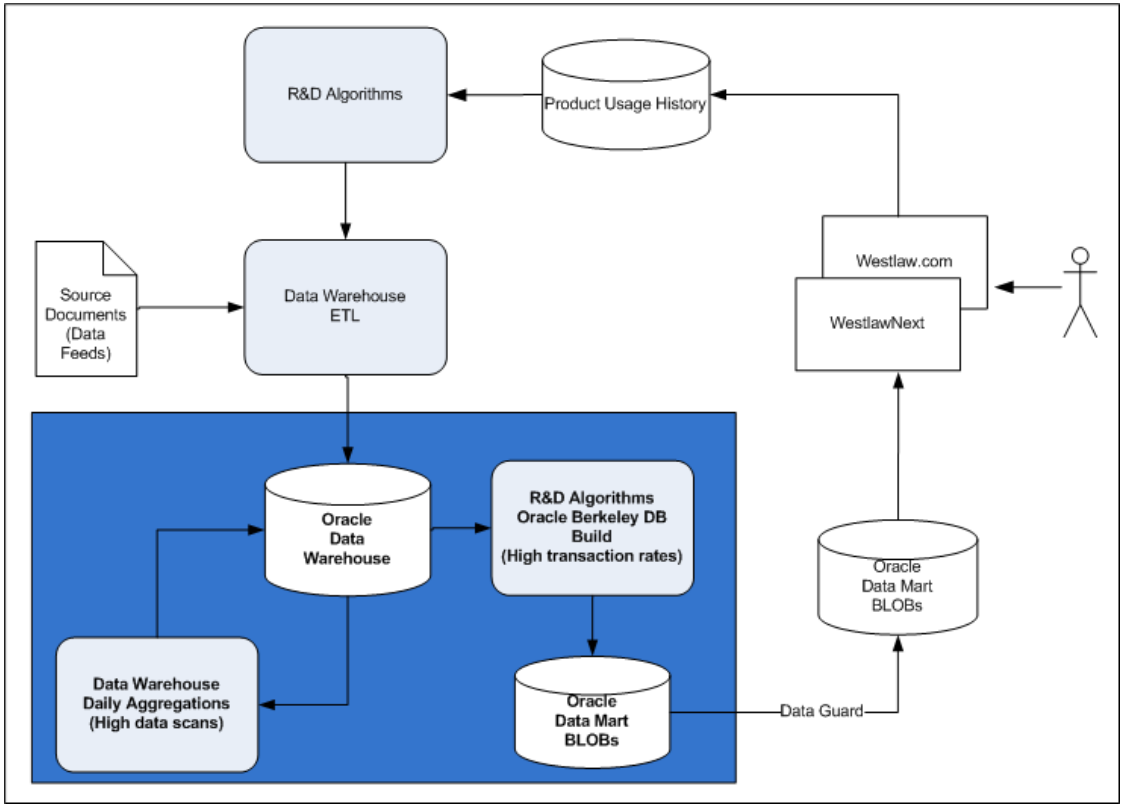

THOMSON REUTERS

Technology Deployed

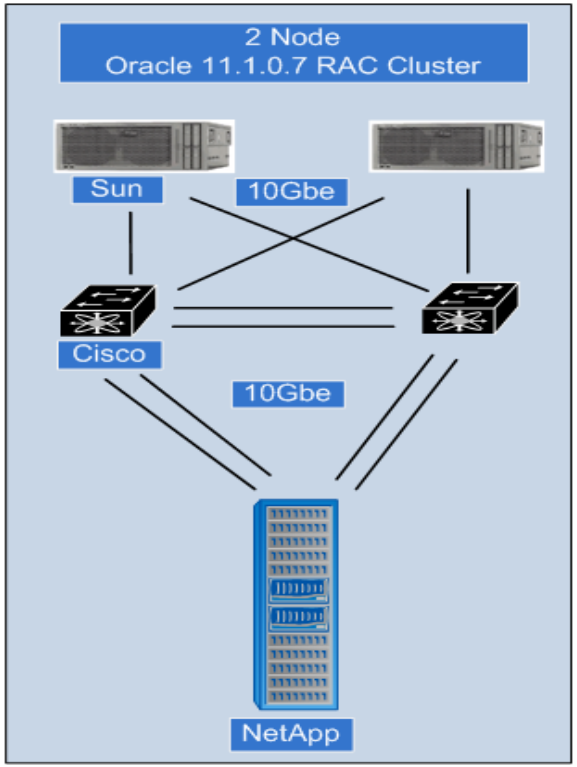

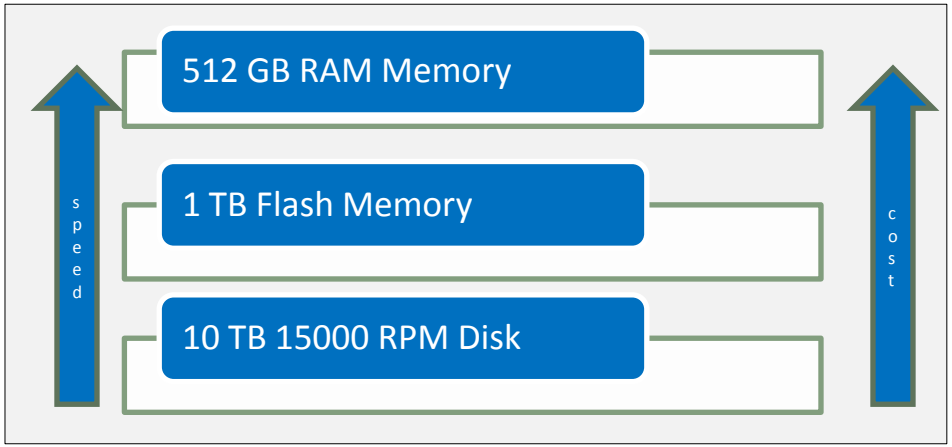

THOMSON REUTERS

Compression Savings

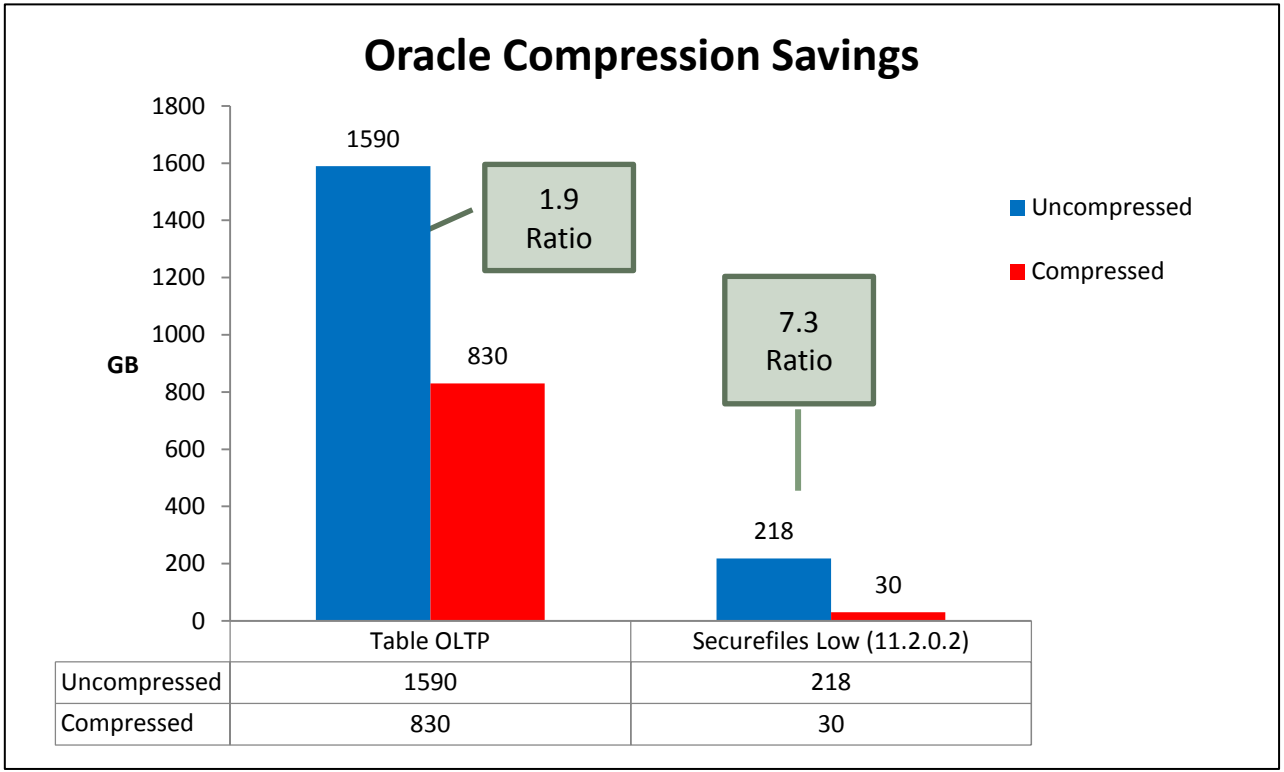

	Table OLTP	Securefiles Low (11.2.0.2)
Uncompressed	1590	218
Compressed	830	30

THOMSON REUTERS

Processing Time

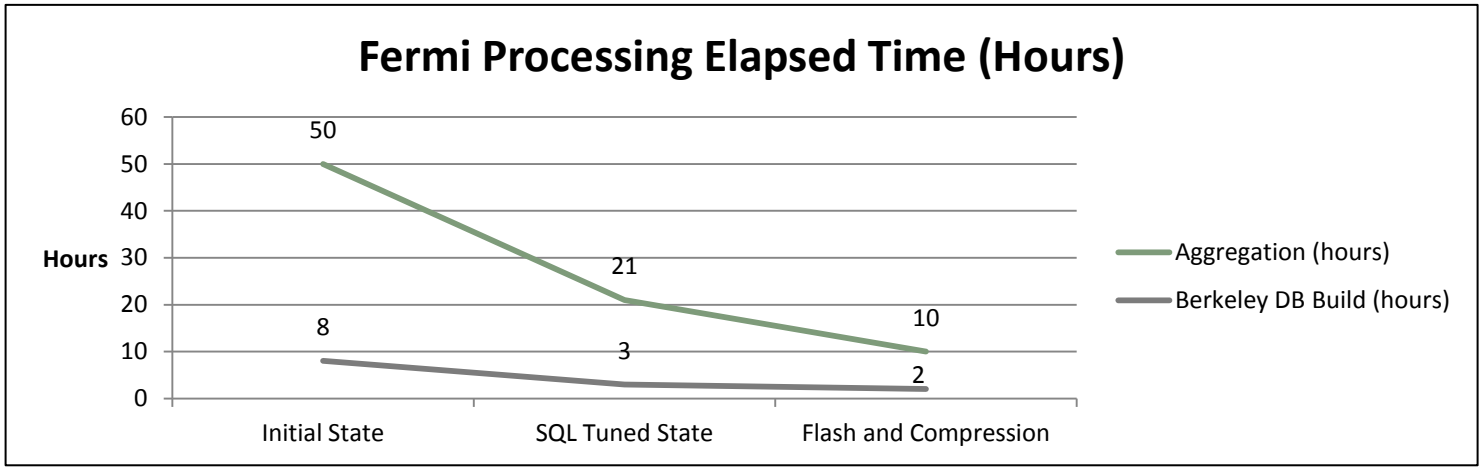

Current State: 12 hours to process on a daily basis which met business requirements for WestlawNext product release. American Association of Law Libraries names WestlawNext new product of the year for 2011.

Future State: 2 hours to process with potential for more daily incremental processing, future state will have more Data, more RAM, more Flash, and Oracle Advanced Compression

THOMSON REUTERS

FlashCache Compressed Data Benefits

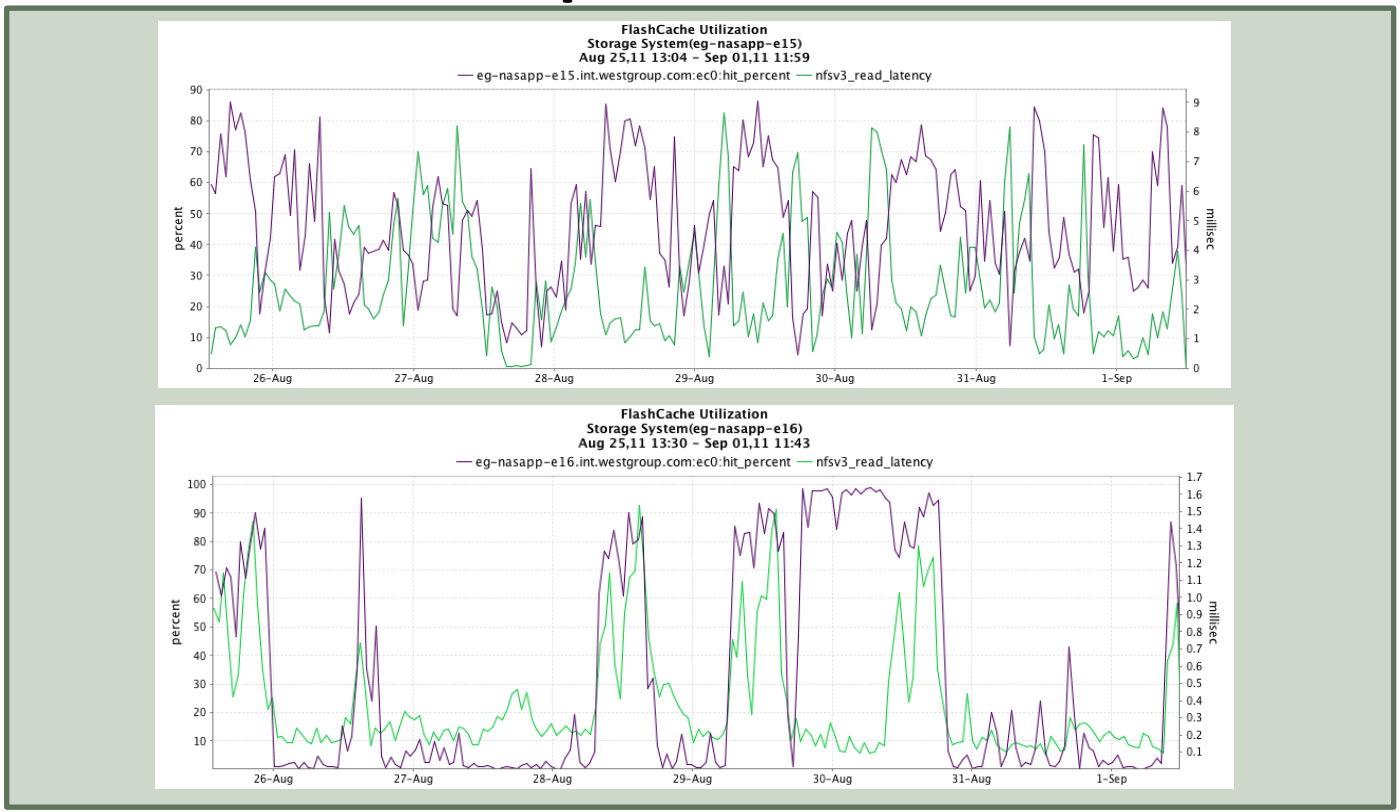

THOMSON REUTERS

SecureFiles Compression

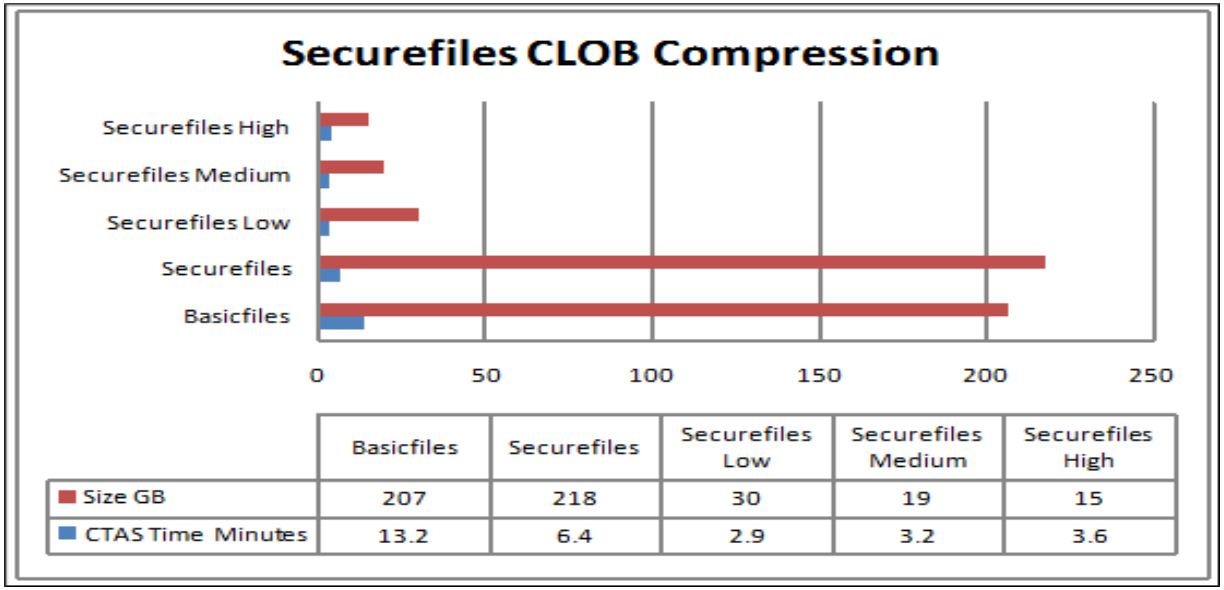

	Basicfiles	Securefiles	Securefiles Low	Securefiles Medium	Securefiles High
Size GB	207	218	30	19	15
CTAS Time Minutes	13.2	6.4	2.9	3.2	3.6

Compression Type	Compression Ratio	Space Saving Rate
Securefiles High	14.5	93%
Securefiles Medium	11.5	91%
Securefiles Low	7.3	86%

THOMSON REUTERS

SecureFiles SQL Monitoring

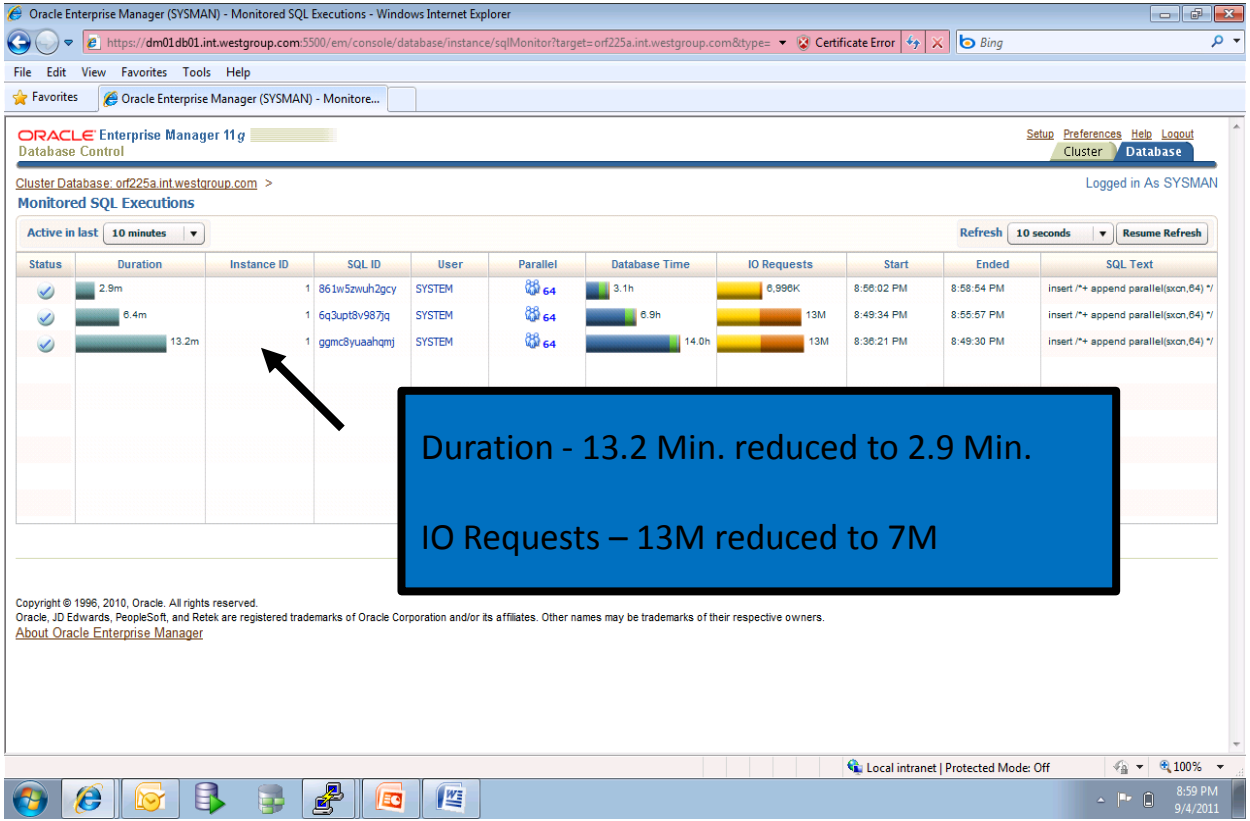

THOMSON REUTERS

Current Use Cases

- **Exadata (HCC and OLTP)**
- **Fusion IO Databases OLTP Compression**
- **Export Dump Files For Data Marts Loads**
  - External table dump with type oracle_datapump access parameters (compression enabled)
- **Data Guard Log Traffic**
- **Securefiles Implementations**
- **RMAN Backups**

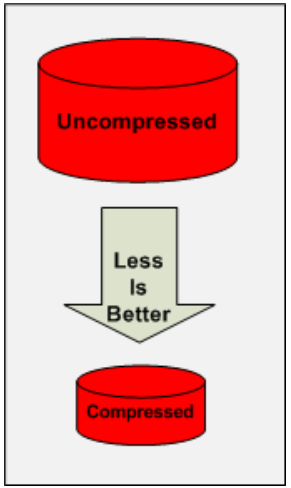

Reduced Physical DB Size = (Lower Cost and Better Performance)

THOMSON REUTERS

Oracle Advanced Compression (ACO)

TVS Motor Company Limited

T G Dhandapani, Chief Information Officer

Introduction

Third largest producer of two wheelers in India

- 7800 employees - 3 Locations in India and 1 location in Indonesia
- $2 billion in revenue (2011-12)
- 2.5 million vehicles (2011-12)
- Deming Prize Winner for Quality - First and only two-wheeler manufacturer in the world
- Exports to over 60 countries

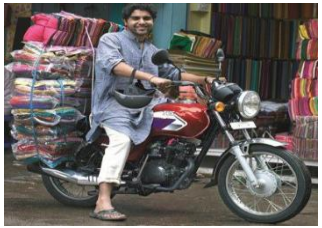

IT Architecture

1100 Dealers

200 Field Staffs

400 Suppliers

CD TVS

INTERNET

30 Sales Office

MPLS Network

Hosur Plant
2000 users

Mysore Plant
200 users

Indonesia Plant
150 users

Nalagarh Plant
80 users

SAP ITS SERVER

SAP

ORACLE

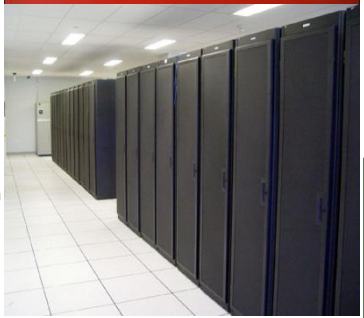

DATA CENTER

- **10000 transactions per day**
- **2GB DB growth per day**

Hardware, OS & DB

HP rx6600 servers &
EVA 4400 storage
HPUX 11.31 IA
ORACLE 11.2.0.2
ERP DB size – 2.3TB
BI, PI, DEV&QA DB size – 6.3TB

Applications

SAP
SAP BI
SAP NetWeaver Process Integration 7.1
DMS
Microsoft Exchange
Altair
PTC
ANSYS
Delcam

Challenges in DB Administration

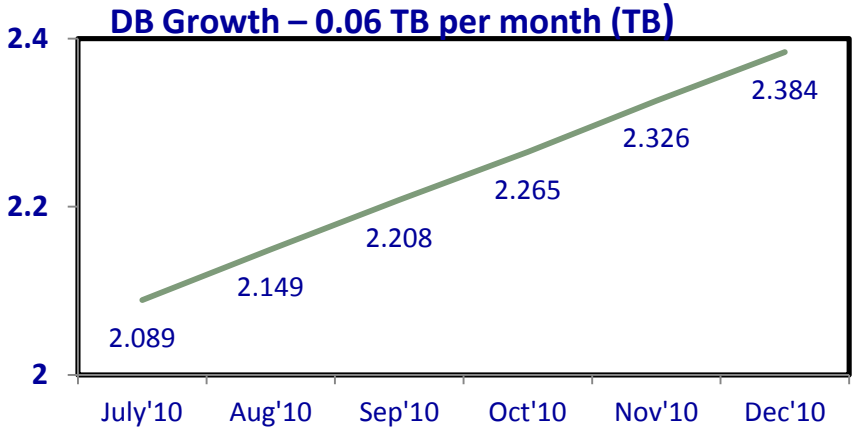

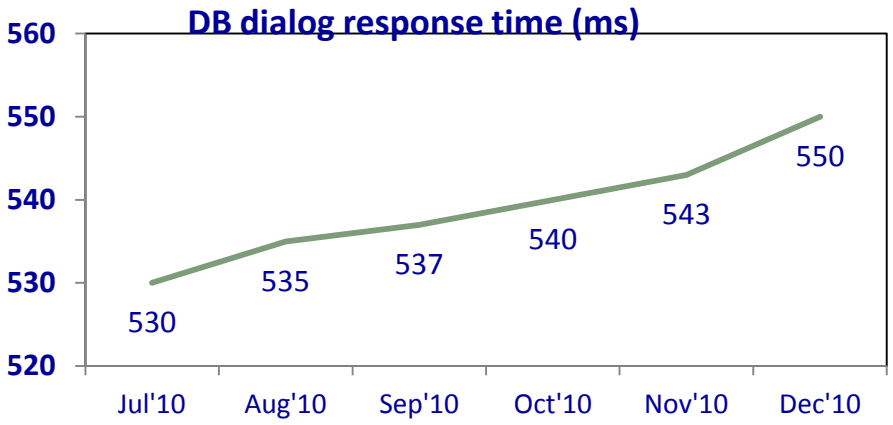

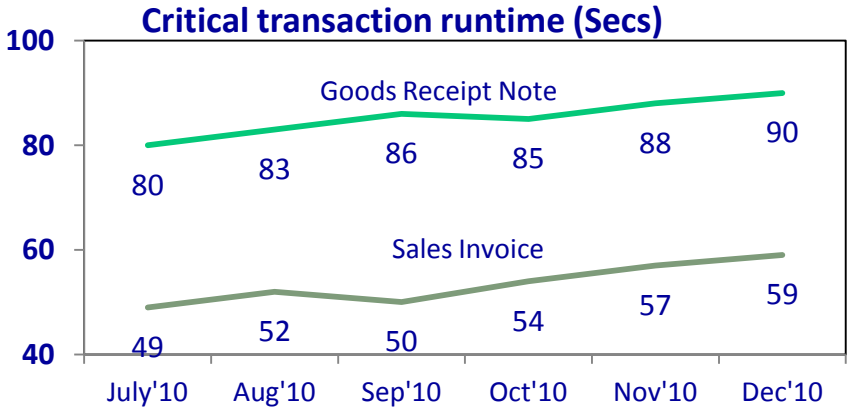

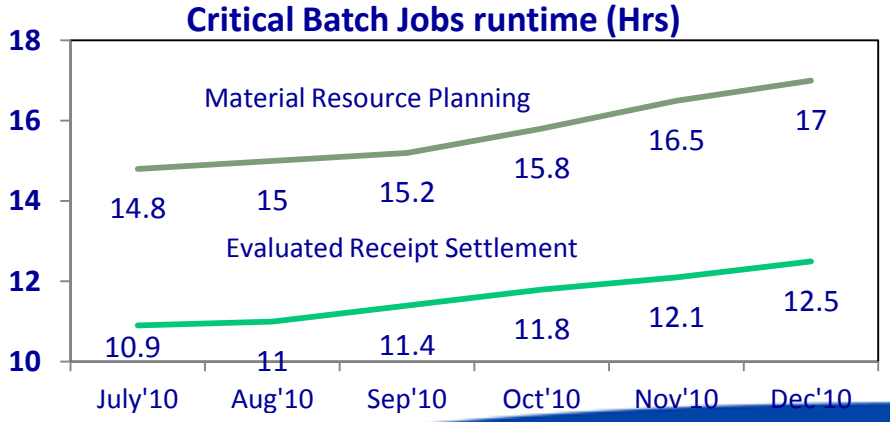

Development of Solution

Upgrade storage	• Needs investment of Rs.6 million (USD 150K) • Long delivery and implementation lead time
Archive data	• Space saving marginal • Long implementation time • Involves investment
Compress database	• ZERO investment • Short implementation time

Hence, we chose Database Compression Option using Oracle's ACO

ACO - Process

Compression in SAND BOX

- Error in online reorg of LOB segments due to bug
- High reorg time of table SOFFCONT1 – 2 hrs

Compression in Development System

- Hyper threading introduced to improve performance of re-org

Compression in Quality System

- Duplicate entries in 2 of the customer tables. Identified and corrected.

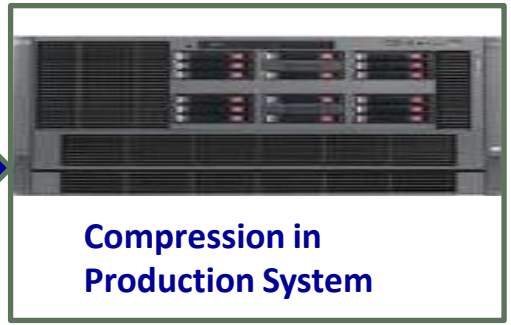

Compression in Production System

- High CPU load (98%) during DML operations in compressed tables on 11.2.0.1. This was solved in 11.2.0.2
- Tables with more than 255 columns should not be compressed – Doing so impacted the transaction times

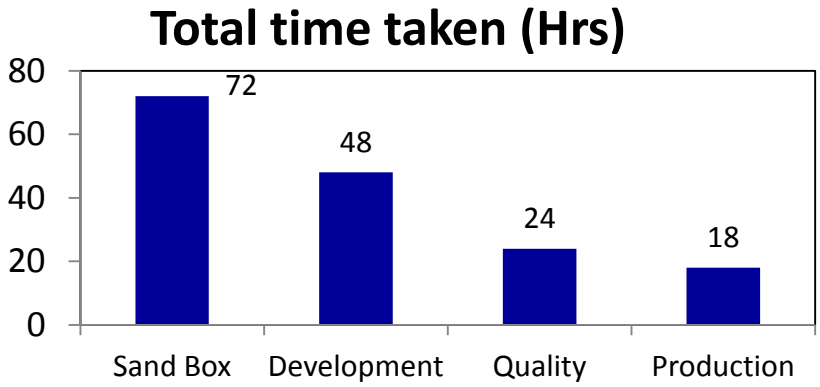

Compress Benefits

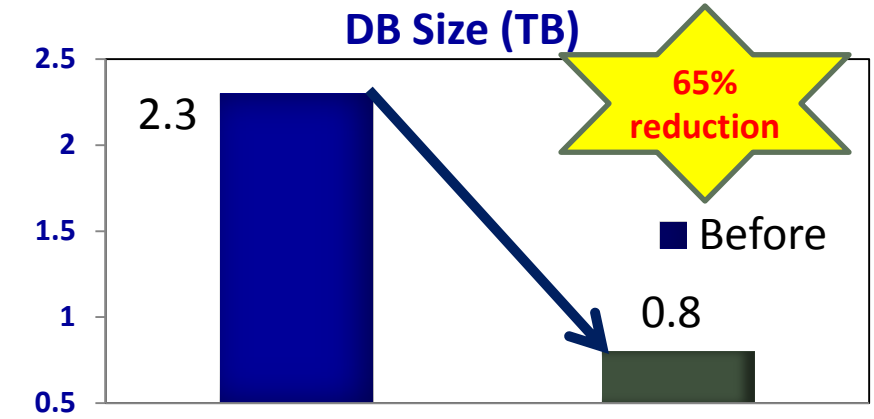

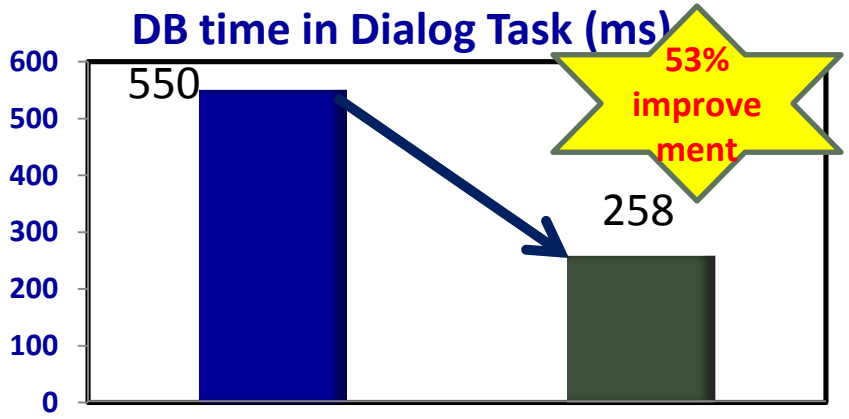

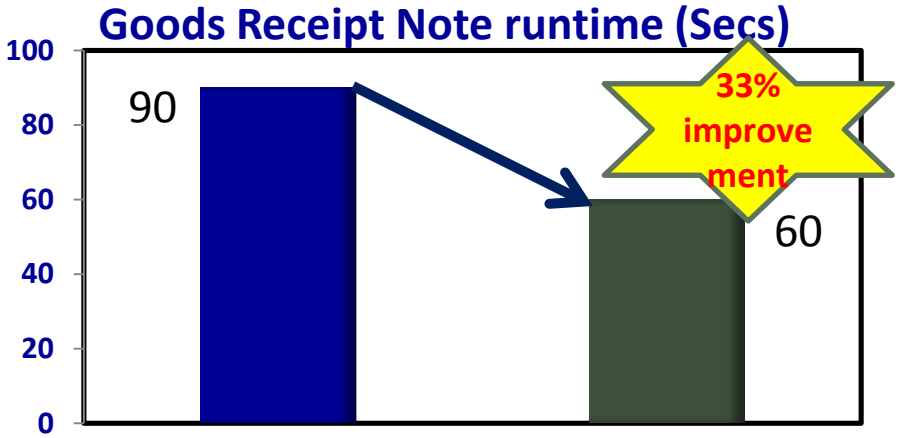

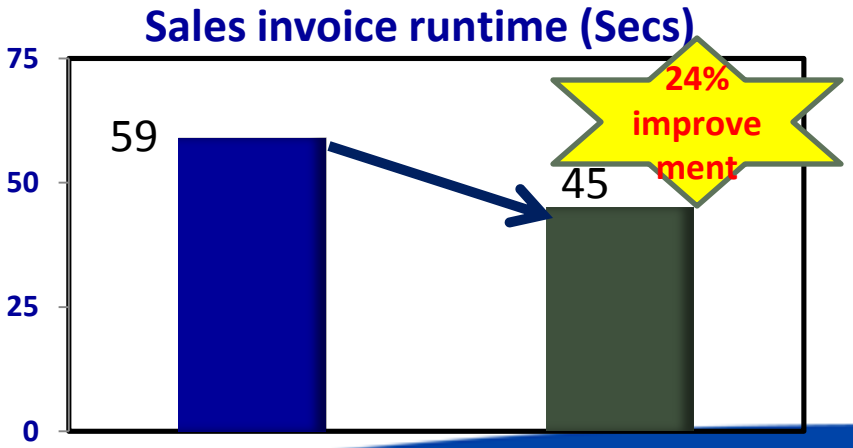

Compress Benefits

Material Resource Planning Batch job runtime (Hrs)

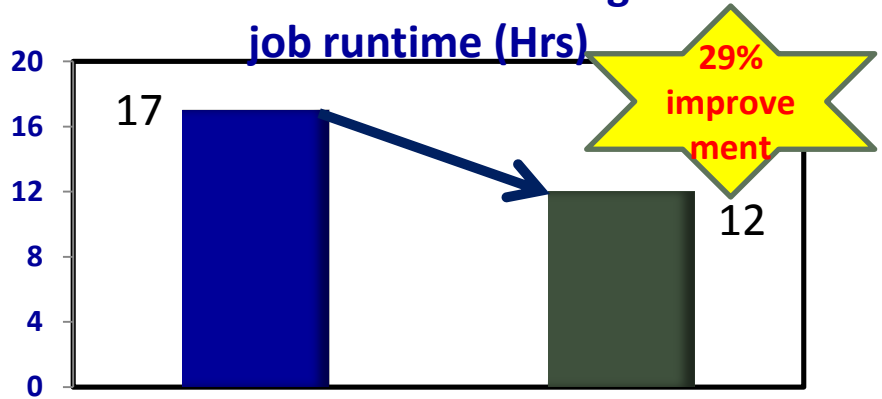

Evaluated Receipt Settlement Batch job runtime (Hrs)

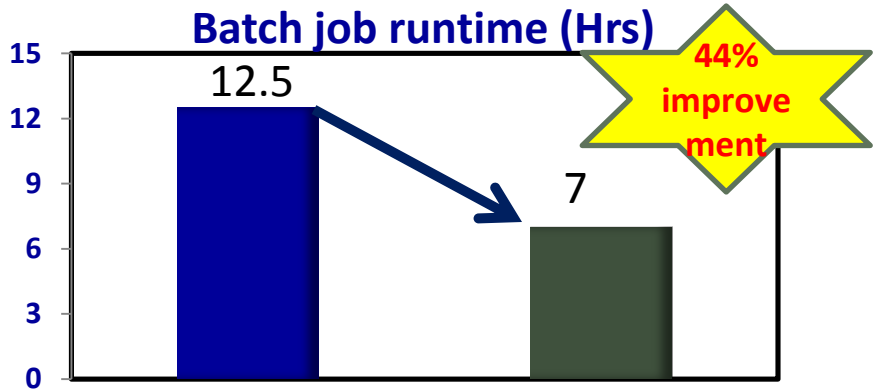

Backup Time (Hrs)

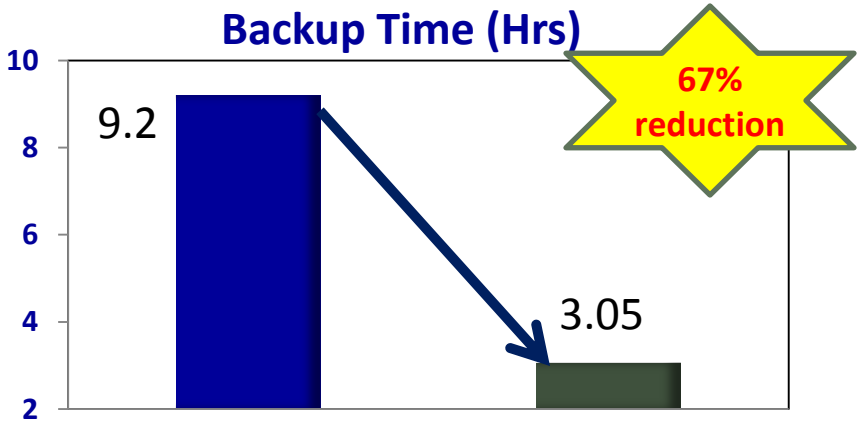

Buffer Quality (%)

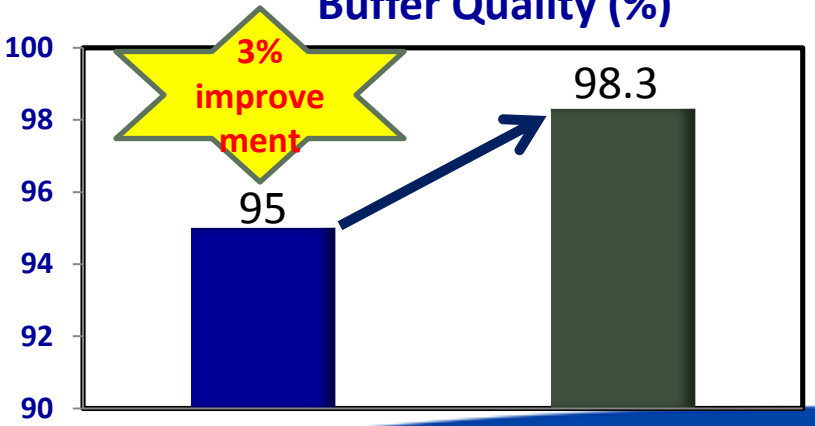

Horizontal Deployment & Space Savings (TB)

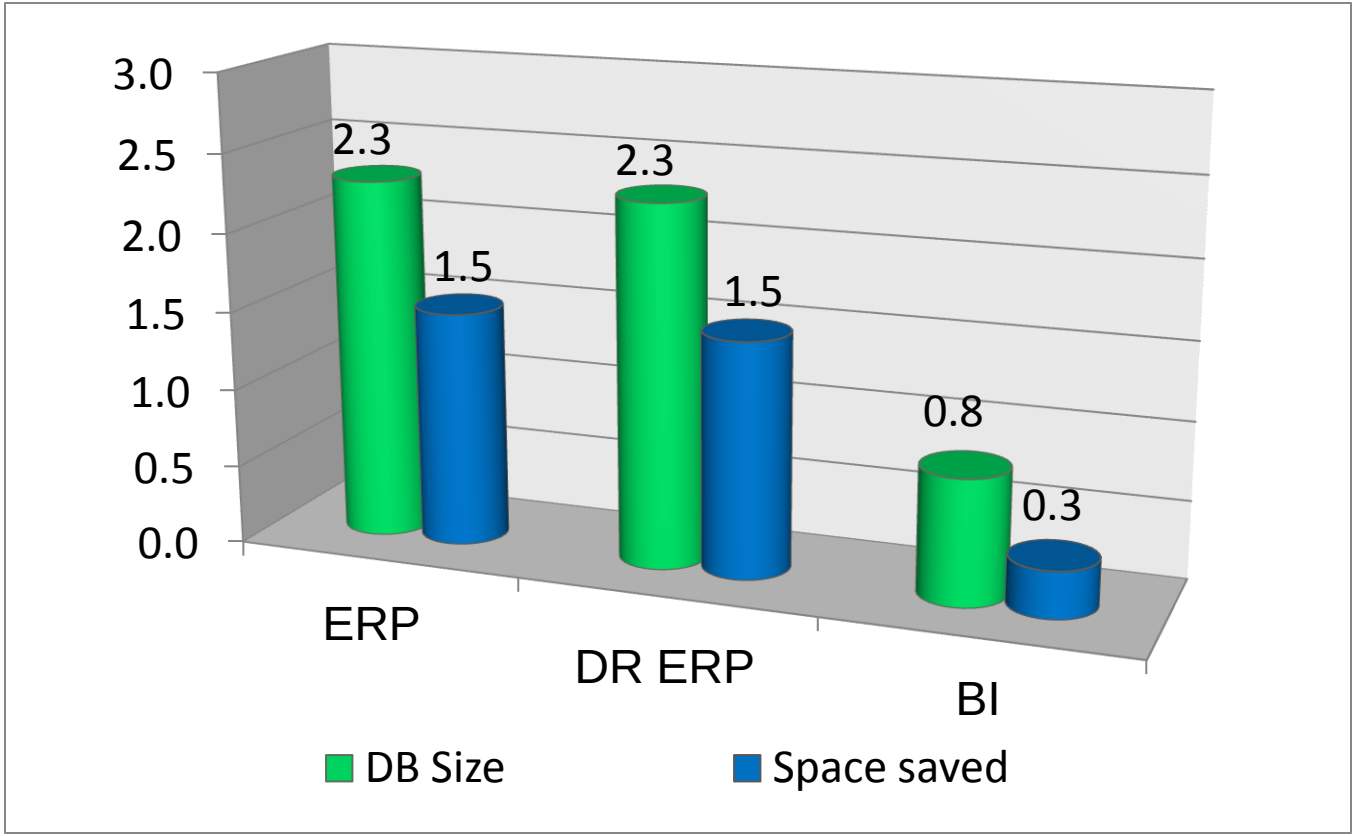

Learning

* Piloting the project in test system to gain experience.

* Tables with more than 255 columns not to be compressed
  * Bug in SAP BRTools used for compression subsequently corrected.

* Upgrade the DB to 11.2.0.2 for best performance.

* Upgrade the DB with latest patches & apply latest parameters.

Way Forward

* Implement Database Vault

* Network and backup compression

Oracle Advanced Compression

Optimize Storage Utilization and lower cost

Omar Abbasi

Business Process Modeling Platform

- Tibco iProcess BPM suite for automated workflow apps and related services
- Single E2E app deployment platform providing Business Process and simulation and execution
- Process Integration and business rules
- Activity Monitoring
- Analysis and Optimization

Environment

- 3 Node RAC Cluster RHES 5.5 Oracle 11gR2
- IBM 3650M2 2x6 core, 48 GB RAM per node
- 4 Instances on 3 nodes
- ASM and OCFS
- 13 Different Env 10 RAC and 3 non RAC total 33 servers using 384 cores
- 17 Applications
- Total Disk Space 50TB

Advanced Compression Uses

- Table Level Compression
- Secure Files
- RMAN
- Data Pump

Advance Compression Stats

Table	No Compression	Medium Compression	High Compression
Table A	456 GB	74 GB	59 GB
Table B	1.7 GB	752 MB	734 MB
Table C	11 GB	4 GB	3 GB
Tablespace	**No Compression**	**Medium Compression**	**High Compression**
ABC_LOB_TBS	456 GB	43 GB	29 GB
ABC_DATA_TBS	26 GB	37 GB	36 GB
ABC_IDX_TBS	3.5 GB	1.8 GB	1.8 GB
RMAN	**No Compression**	**Medium Compression**	**High Compression**
BACKUP SET SIZE	80 GB	47 GB	42 GB
DATAPUMP	**No Compression**	**Medium Compression**	**High Compression**
expdp	447	447	447
PERFORMANCE (in Sec)	**No Compression**	**Medium Compression**	**High Compression**
Query 1	14.31	12.21	14.34
Query 2	4.43	2.05	2.32

Advanced Compression

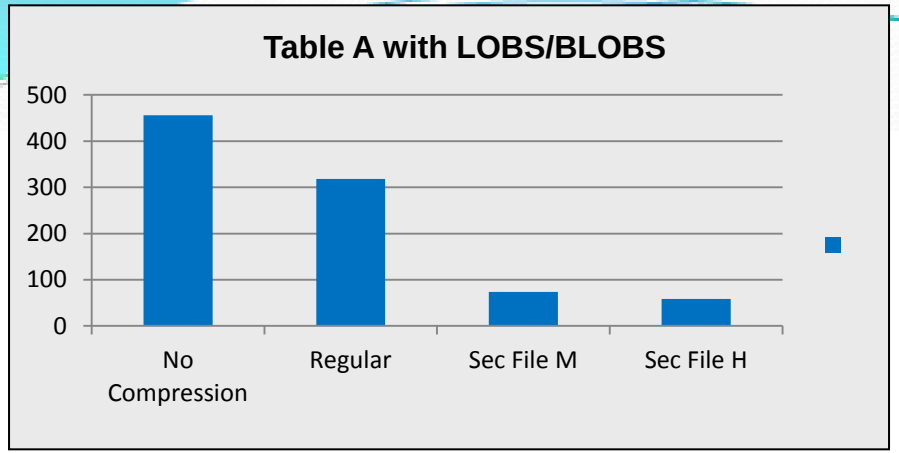

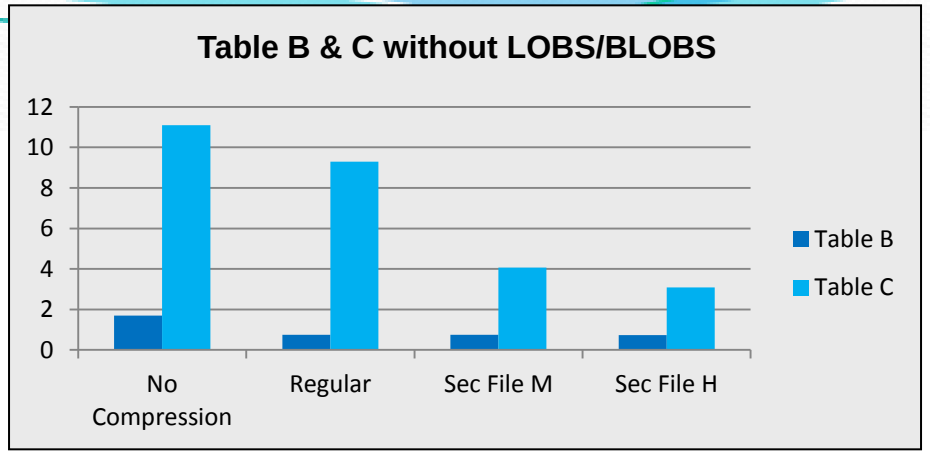

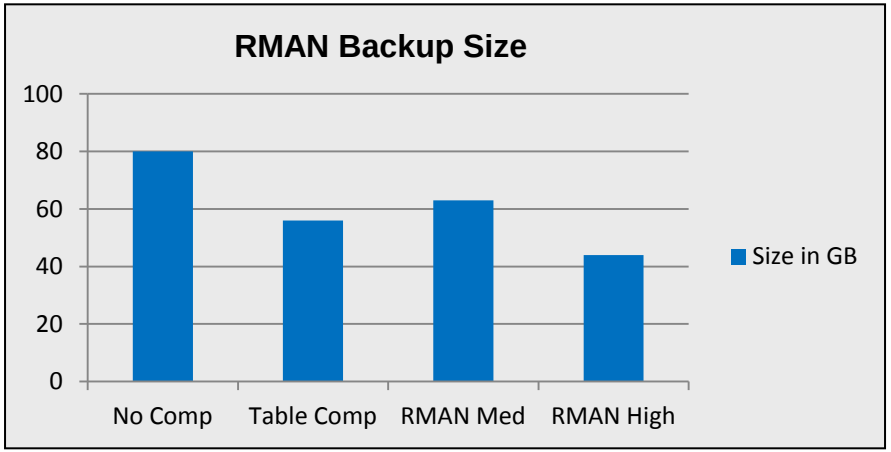

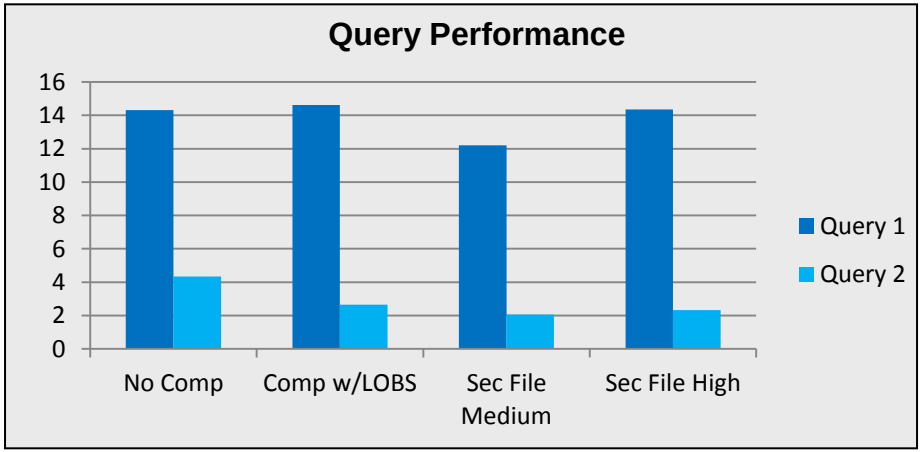

Conclusion

- Significant reduction in storage utilization
- No performance impact
- Huge space save using Secure Files, above 80%
- Reduced TCOA
- No significant space save using EXPDP
- Easy to Implement

Summary

- **Comprehensive data compression capabilities for all types of data**
  - Structured, Unstructured, Backup, Network
- **Reduces storage consumption by 2x to 4x**
- **Improves read performance**
- **Enhances memory, buffer cache utilization**
- **Complete application transparency**
- **Benefits diverse application workloads**

Q&A

Hardware and Software

ORACLE

Engineered to Work Together


ORACLE®