
ORACLE®

ORACLE DATABASE 11g

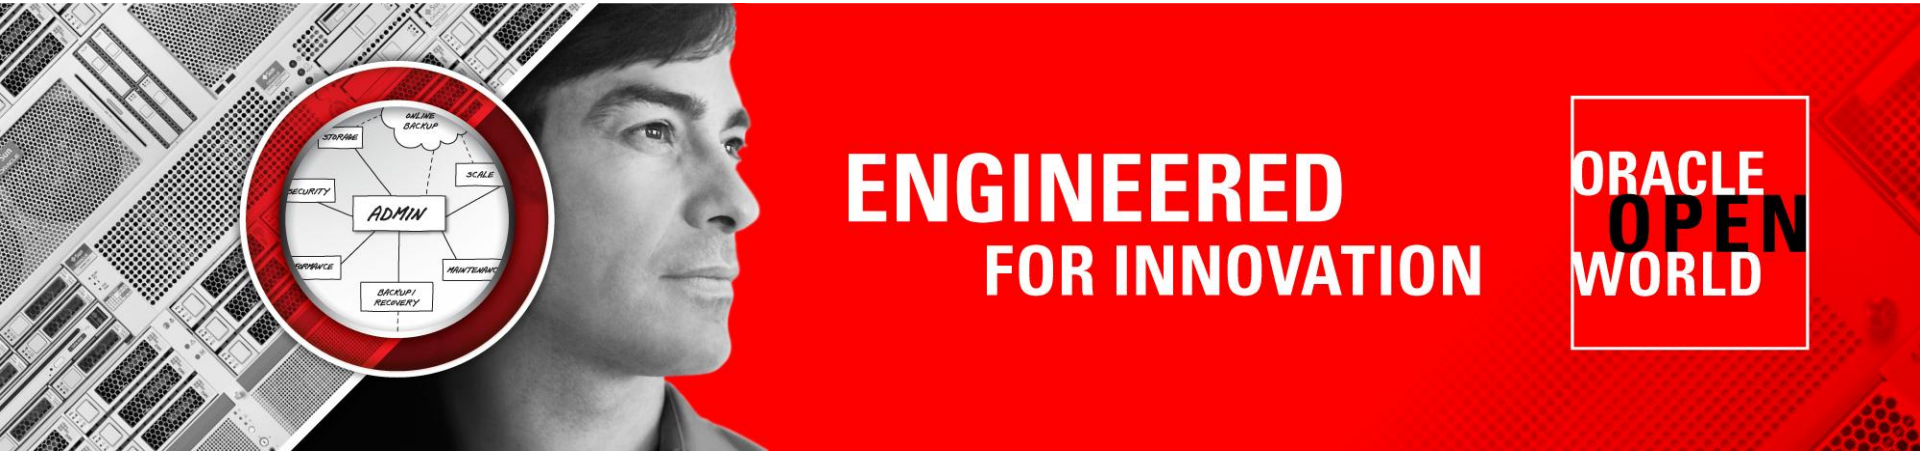

ORACLE®

Oracle Advanced Compression – Reduce Storage, Reduce Costs, Increase Performance

Gregg Christman -- Senior Product Manager

Presenting with

SHARP

DELL

ORACLE OPEN WORLD

ENGINEERED FOR INNOVATION

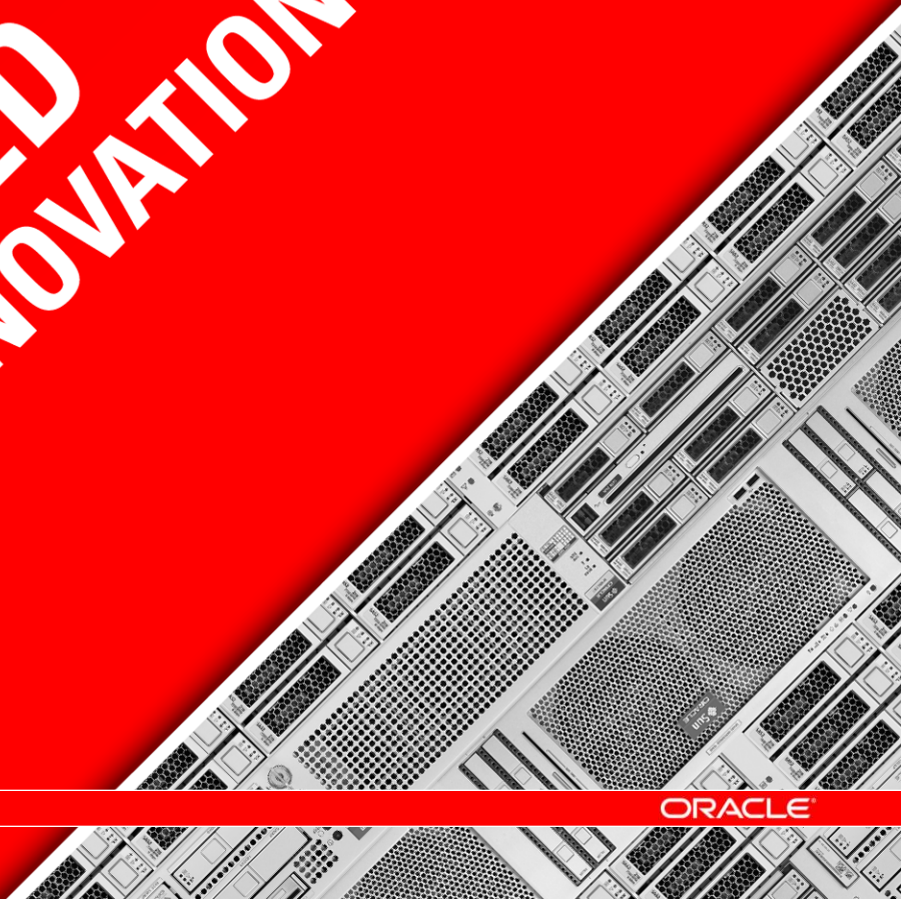

ORACLE

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

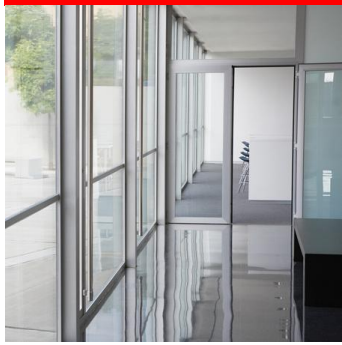

- Data Growth Challenges
- Advanced Compression Overview
- Customer Experience's
- Questions and Answers

Data Growth Challenges

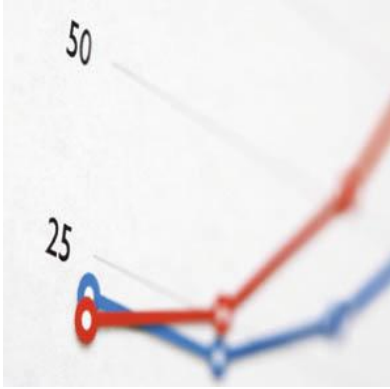

- **Explosion in Data Volumes**
  - Government Regulations (Sarbanes-Oxley, etc)
  - User Generated Content (Web 2.0)
  - Application Consolidation
- **IT Managers Must Support Larger Volumes of Data with Limited Technology Budgets**
  - Need to optimize storage consumption
  - Also maintain acceptable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

Oracle Advanced Compression Option

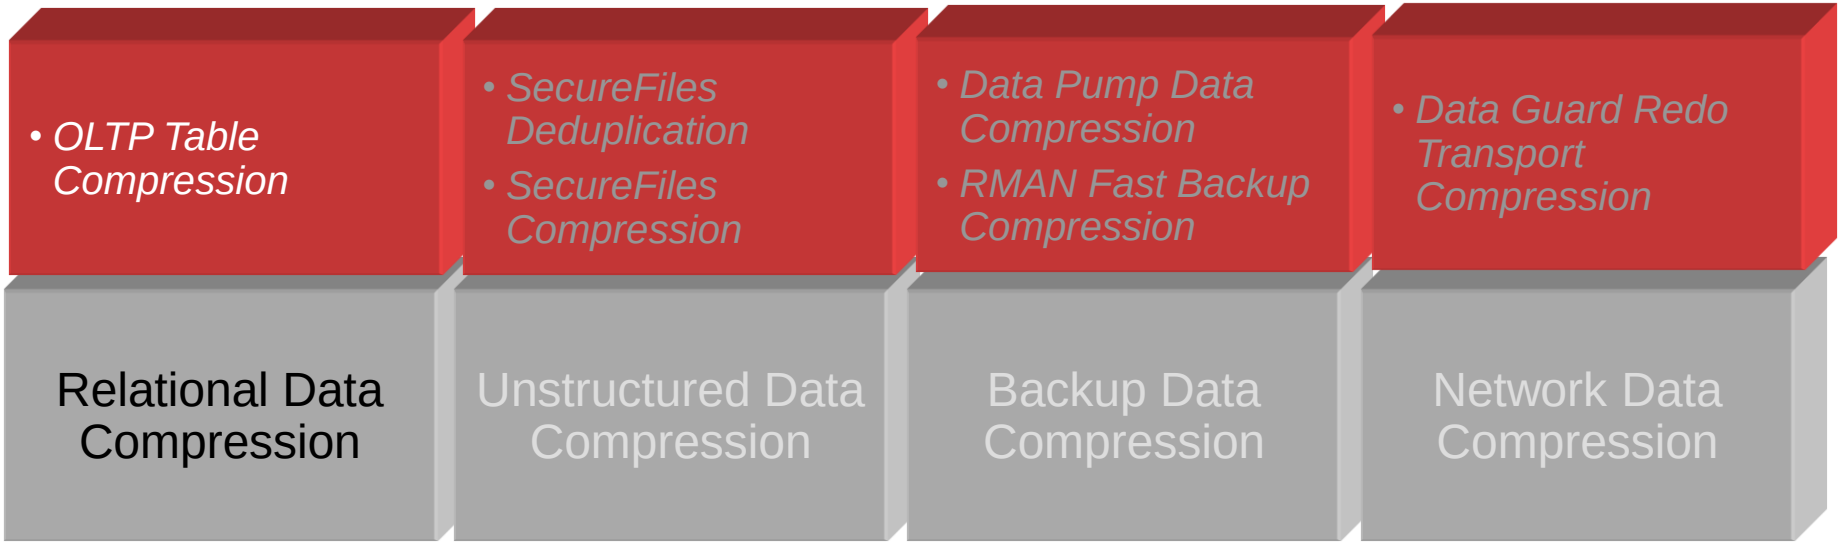

Basic Table Compression

- **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

Oracle Advanced Compression Option

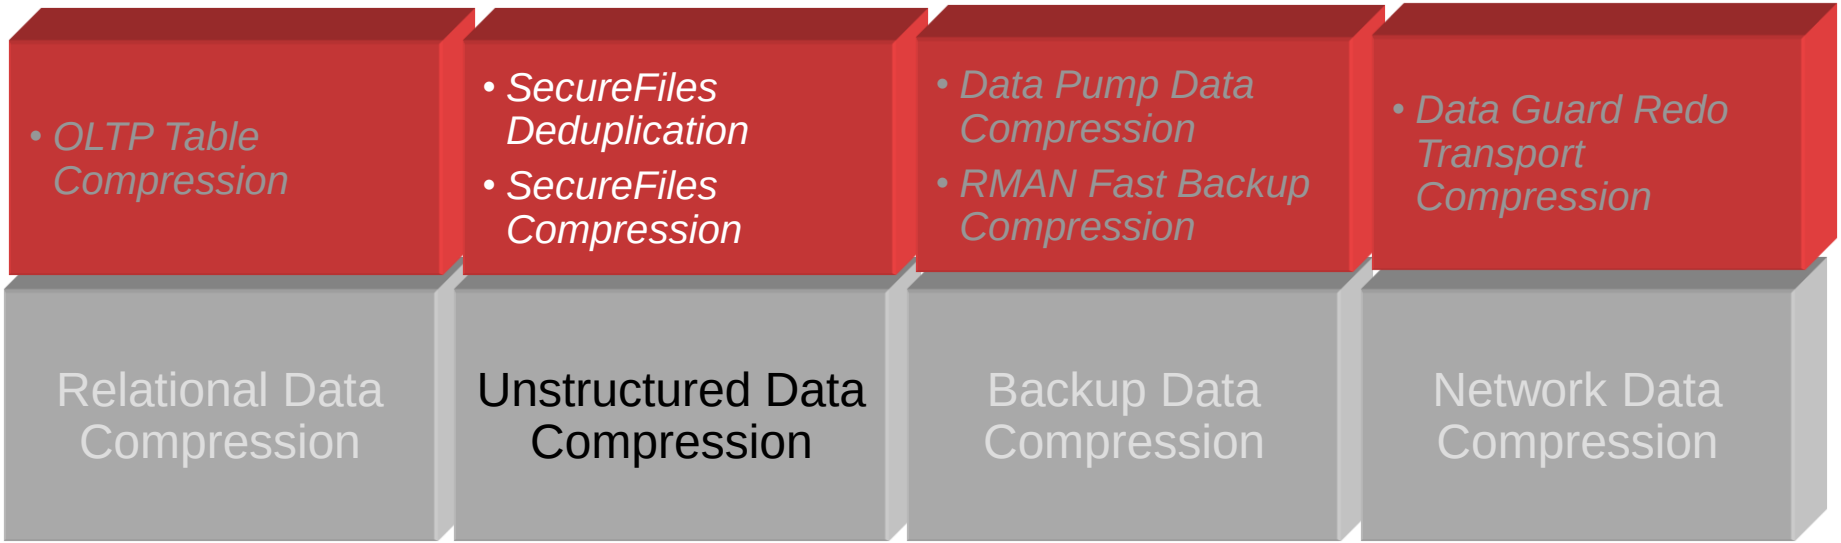

SecureFiles Deduplication/Compression

- **Enables storage of a single physical image for duplicate data**
  - Significantly reduces space consumption
  - Dramatically improves writes and copy operations
- **No adverse impact on read operations**
  - May actually improve read performance for cache data
- **With SecureFiles compression, typical files such as documents or XML files, experience a reduction of 2x to 3x times in size**
  - Automatically avoids compressing data that would not benefit from compression
  - Specially useful for content management, email applications and data archival applications

Oracle Advanced Compression Option

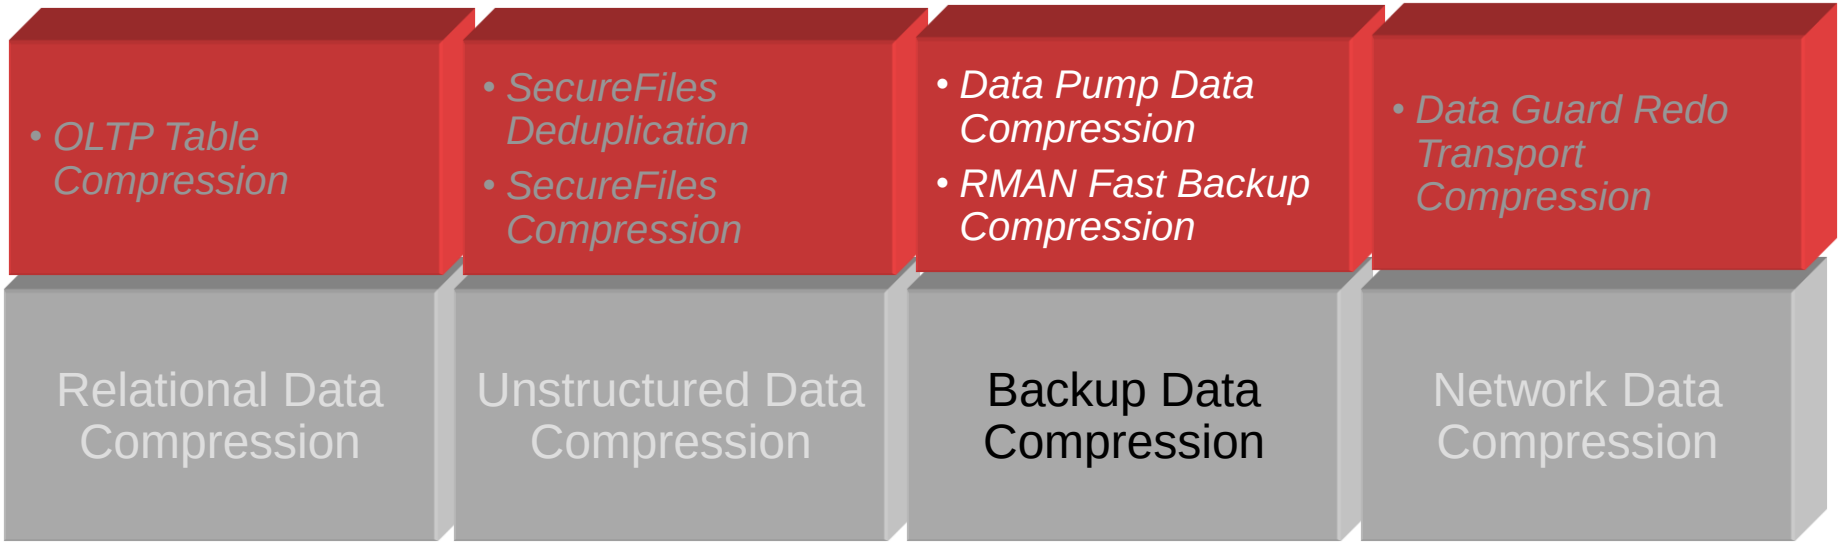

Data Pump Compression

- **Metadata compression since Oracle Database 10*g***
- **Oracle Database 11*g* extends compression to 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 Data Pump functionality available on compressed files

RMAN Backup Compression

- **RMAN compression can 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 (New in 11.2)**

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

    Balance between CPU usage and compression ratio
  - **Compression LEVEL HIGH (New in 11.2)**

    Best compression ratio and highest CPU utilization

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

Oracle Advanced Compression Option

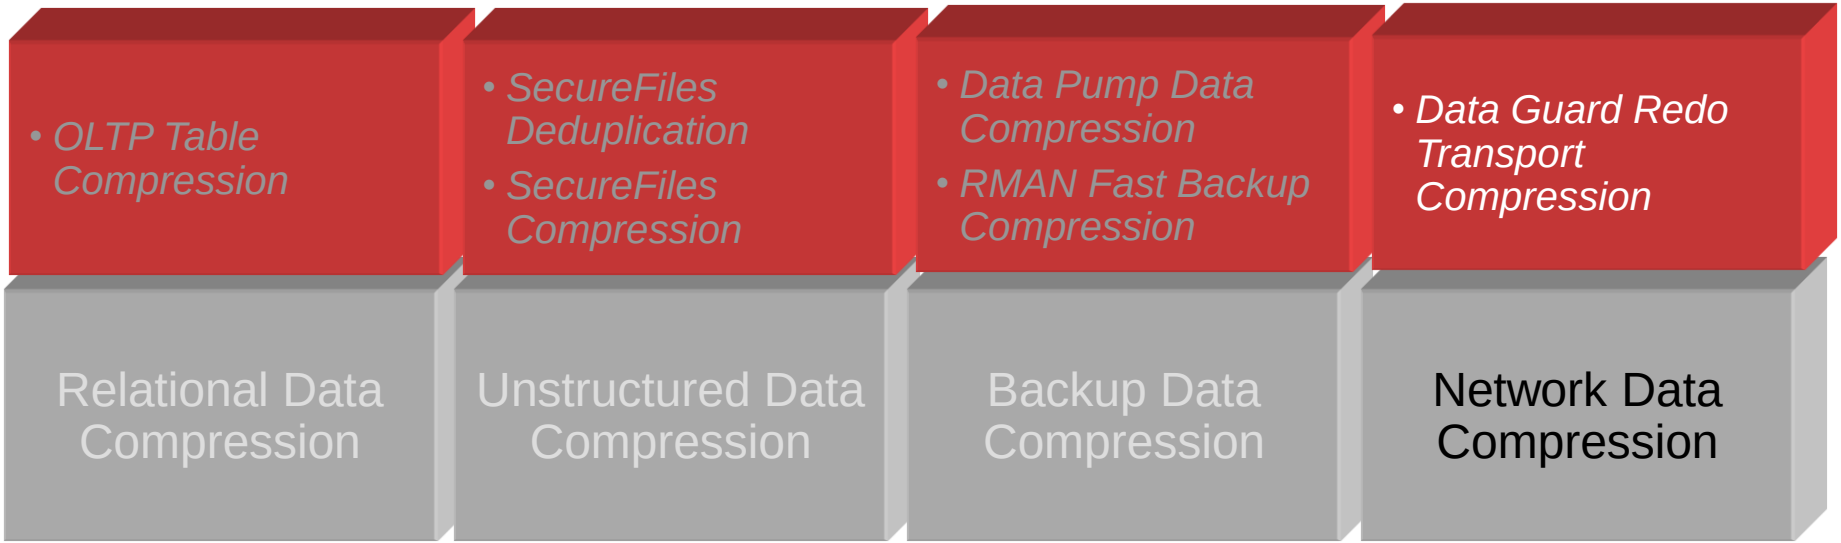

Network Compression

- **Data Guard Redo Transport Services**
- **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%

Customer Experience's

Presenting with

Sharp Healthcare

Oracle Advanced Compression

Kim Nguyen, DBA

Sharp Healthcare

- Location: San Diego, California
- Not-for-profit health care delivery system
- 4 acute-care hospitals, 4 specialty hospitals, 2 medical groups, 1 health plan, plus full spectrum of other facilities and services
- 2,600 physicians
  - More than 1000 physicians in medical groups:
    - Sharp Rees-Stealy and Sharp Community Medical Group
- More than 14,000 employees

Sharp Healthcare

- 2007 Malcolm Baldrige National Quality Award Winner
- Ranked 13 nationally in Modern Healthcare and # 1 in California for “Most integrated Healthcare Networks”
- 2 acute hospitals received prestigious MAGNET designation
  - Excellence in nursing practices and quality patient care
- RDBMS: Oracle, MySQL, SQL Server, Cache
  - 350 Production and non Production databases
  - Over 100 TB of data

Sharp Healthcare

- **Background**
  - Sharp Healthcare's SQL Server data warehouse was rapidly growing
  - Strategic decision to migrate from SQL Server to Oracle began in fall of 2008
  - New Cerner EMR (electronic medical record) application would feed new Oracle data warehouse in addition to other enterprise applications using Informatica as the ETL
  - New reporting requirements to meet Meaningful Use criteria
  - Oracle 10.2.0.4 using 2.5TB of storage and rapidly increasing

Sharp Healthcare

- **Implementation and Migration**
  - Upgraded to 11.1.0.7 to use Oracle's tablespace level Advanced Compression in summer 2009
  - Tests on small tablespaces resulted in 2:1 compression ratio
  - Testing of creation of tables and load took same amount of time as uncompressed tablespaces
  - General Oracle issues
    - Datapump import was very lengthy
    - Tables moved through database link

Sharp Healthcare

- **Implementation and Migration Cont'd**
  - After successful completion of overall testing, go live plan was developed to use advanced compression for all user data and index tablespaces
  - Feb 2010 development and test environments upgraded to 11.2.0.1 along with AIX 6.1.4
  - New production hardware purchased
    - Same migration plan was tested again and implemented
    - Datapump import successfully finishes with Oracle 11gR2
  - 2.75TB of disk space compressed to 1.5TB during migration

Sharp Healthcare

- **Application Details**
  - ETL: Informatica 9
  - Reporting: Cognos 10.1
  - Critical business intelligence reporting and patient care studies used enterprise wide
  - Comply with ‘Meaningful Use’ for electronic health records ARRA/HITECH Act

Sharp Healthcare

- **Performance Statistics**

  Sharp Healthcare's 3 largest

  electronic health and medical record applications

  – Over 5 million rows loaded nightly
  – Average load time 5-6 hours with jobs running concurrently over night for each large application
  – Inserted into ~270 Dims and ~70 Fact tables
  – User report query times remained consistent after compression
  – Update performance remained consistent after compression
  – Compression used concurrently with partitioning for increased performance

Sharp Healthcare

- **Specifications Today**
  - Oracle 11.2.0.2
  - IBM Power 6 P560
  - AIX 6.1.4 (64bit)
  - 16 processors, 256 GB of RAM
  - 4.5TB of EVA 8400 storage used, and growing…
  - 2 non prod environments

Sharp Healthcare

- **Outcome**
  - Overall 2:1 compression ratio
  - Significantly reduced storage requirements
  - Allowed us to build a second non-production environment for analyst testing also using compression with Oracle's Datapump tool
  - No performance issues due to advanced compression
  - $$$ saved in purchasing storage and time associated to managing storage

Oracle Advanced Compression
A Case Study at Dell

Dmitry Donetskov, Andrey Dubnik, Alan Goodall,
Vrishab Kakade, Alexander Vetrov, Phillip Woodruff,
Raj Alagirisamy

Advanced Compression – Business Case Target System

OMEGA	Dell Order Fulfillment system for back office	Business critical system impacting manufacturing	Goal and expected numbers: OLTP Table Compression **50-70%**
Currently supporting EMEA and APJ	Going Global soon with LATAM and Americas	Oracle E-Business Suite R 12.1.3	Index Compression **30-40 %**
16-node RAC database / 11gR2	OEL 5.4 64-bit / 48GB RAM per Node	32TB Storage	Overall Compression **40%**

Results from Testing in Non-Production

Test databases:

All tables and Indexes compressed; Running fine for the past 4 months

- Table compression: 50-70%
- 2.5 TB to 1.2 TB
- 10 non-prod databases: total storage reduction: 10 TB
- 15 % to 33% performance Improvement in some reporting queries
- 50-60% reduction in physical reads

Archive DB (OMEGAPA):

- OLTP table compression running fine for the past few months
- 1 TB to 300 GB
- 70% compression
- Tablespace level

Conclusion

Enabling secure file storage in row with medium or high compression provided excellent compression rates.

Index Compression and Test with Larger Block Size

Index compression: Some real figures on space saving achieved from compressing indexes in one of our test environments

Online Time ~ 100 min, index size 33Gb -> 18Gb

Offline Time ~ 40 min, index size 17Gb -> 11Gb

So, there is somewhat **30-40% compression ratio for indexes.** The most tricky part about compressing them is choosing an optimal prefix key length size.

Larger database block size is recommended for OLTP compression for better space saving. Some real figures – the RA_CUST_TRX_LINE_GL_DIST_ALL table with size 454GB was compressed to 158GB with the 8KB block size and to 138GB (i.e. 20GB less) with the 32KB block size.
We considered several methods to compress existing data

- by moving a table and rebuilding indexes unfortunately it requires considerable downtime
- CTAS – works great but we did not use much
  →does not save time comparing to moving a table

→adds complexity of renaming tables and DDL preparation for indexes

- with the online redefinition option, it comes with a set of bugs/limitations (Bug 8507266, Note 1142595.1, Note 1237593.1, Note 1233204.1 – the most notable one, Note 1116785.1) and we ruled it out as being a too risky one for a heavily loaded 24/7 system.
- by updating each row to same values – did not work for us, basically it did not trigger compression

Reporting Queries – Performance Improvement

Test	BU	Run Time (hh:mi:ss)	Rows	Consistent gets	Physical reads
Single - nocomp	CN	20:30:54	340	64,149,097	12,909,861
Parallel - nocomp	CN	2:00:03	340	98,043,933	64,984,875
Parallel - comp	CN	2:03:11	340	129,060,967	38,890,960
Single - nocomp	DE	15:14:41	103	35,514,363	9,904,162
Parallel - nocomp	DE	1:14:40	103	77,439,792	62,156,661
Parallel - comp	DE	0:59:47	103	66,273,649	29,687,811
Single - nocomp	FR	11:54:19	96	30,578,814	8,175,152
Parallel - nocomp	FR	1:05:03	96	74,089,454	61,614,674
Parallel - comp	FR	0:55:14	96	59,574,892	28,200,993
Single - nocomp	GB	18:24:51	135	67,664,501	14,866,912
Parallel - nocomp	GB	1:29:10	135	82,370,708	62,882,690
Parallel - comp	GB	1:28:52	135	78,241,447	31,964,625
Single - nocomp	XM	1:37:46	254	9,437,017	1,064,806
Parallel - nocomp	XM	0:32:50	254	63,916,791	59,832,077
Parallel - comp	XM	0:21:47	254	36,377,869	22,806,990

15% to 33% performance Improvement in some reporting queries
50-60% reduction in Physical reads

Secure File Compression - A Sample Table

```
-- case 5b
create table xxeom.xmlpr_log_c5b  (
 "NODE_ID"        number,
 "RETMSG"         varchar2 (255)  default 'OK',
 "NAME"           varchar2 (1000), -- not null enable,
 "KEY"            varchar2 (1000), -- not null enable,
 "CREATED_DATE"    date default sysdate, --  not null enable,
 "MSG_BODY"       clob,
 "LOG"            clob)
 lob ("MSG_BODY") store as securefile (tablespace xxdell_ts_interface  compress medium enable
storage in row)
 lob ("LOG")      store as securefile (tablespace xxdell_ts_interface compress medium enable
storage in row)
 partition by range (created_date)
 interval(numtodsinterval(7, 'DAY'))
 (partition p0 values less than (to_date('03-01-2011', 'DD-MM-YYYY')))
 pctfree 0 initrans 11 compress for oltp
 tablespace xxdell_ts_interface;
```

Secure File Compression - Data from Test Cases

XMLPR_LOG	Case 1a	Case 1b	Case 2a	Case 2b	Case 3b	Case 4a	Case 4b	Case 4c	Case 5b
Table Compression	No	No	Yes	Yes	No	No	No	No	Yes
SecureFile Compression	No	No	No	No	Medium	Low	Medium	High	Medium
SecureFile Storage in Row	Disabled	Enabled	Disabled	Enabled	Disabled	Enabled	Enabled	Enabled	Enabled
Table, MB	24	136	24	136	32	158	141	180	144
LOB, MB	2721	2376	2728	2375	1070	648	378	291	368
Sum	2745	2512	2752	2511	1102	806	519	471	512
CPU used by this session	30626	16166	30331	20148	32297	15005	17444	19956	18,994
HSC OLTP positive compression	0	0	3723	8771	0	0	0	0	11,634
HSC OLTP negative compression	0	0	0	7037	0	0	0	0	4,198
securefile uncompressed bytes	0	0	0	0	1,714,986,772	1,981,901,630	1,981,901,630	1,981,901,630	1,981,901,630
securefile compressed bytes	0	0	0	0	255,540,663	444,899,559	327,403,306	269,075,197	327,403,306
redo size	2,332,731,056	2,114,374,552	2,362,608,560	2,316,044,740	816,602,908	523,484,564	382,866,640	317,254,156	542,032,600

Observations/Comments from Secure File Compression Tests

LOB storage options:

High CPU consumption for each case with the disabled storage in row for LOB segments regardless the table compression level or the LOB compression level. Cases 1a, 2a, 3b.
Looks like internal Oracle mechanism requires much additional CPU resources to handle LOBs with the disable storage in row option.
Higher storage consumption with the disabled storage in row for LOB segments under same circumstances (it probably can be tuned with the CHUNK parameter?). Case 2a vs. 2b.

Table compression options:

Data in columns outside LOBs are compressed at ratio 32 to 24 i.e. 4 to 3, not a great one (cases 2a and 3b).
So, compressing the table is not very much beneficial comparing to the size of LOB segments but could cause much higher redo (cases 4b and 5b - took ~10% more CPU and generated ~40% more redo).

LOB compression options:

The compression of LOBs is not a great CPU overhead. Most of CPU seems to be consumed by other operations. Cases 2a and 3b.
Difference in gain for the compression between low and medium levels is several times greater than between the medium and high compression levels i.e. the medium one is quite close to the high one.
The high compression level requires only ~15% more CPU and provides ~20% less redo and ~10% less segment storage.

Production Implementation

OMEGA-Production

OLTP - Table Compression

3 tables fully compressed	**Compression Ratio: 5.5 : 1**
13 other tables enabled for compression (1 TB) (Compression data not available at this time)	Rest of the Implementation – in phases including Secure file compression
No degradation in transaction throughput	Overall compression expected: 40% (This is due to the fact this particular system has many tables with 256 columns or above)

Conclusion

The test results are quite encouraging especially for Tables with LOBs (with Secure file compression)

With proper Index compression we aim to achieve at least 40% storage reduction for this particular system

So far no degradation seen in transactions

In fact, some of the reports have better throughput times (up to 33%)

With additional enhancements from Oracle, we hope to see much better compression ratios and thereby significant storage cost savings as well as overall performance improvement (due to reduced I/O)

To learn more about how Dell can help you drive an Efficient Enterprise visit:

- Dell solutions booth #3109 (West) and Dell Storage and Networking booth #409 (South)
- Michael Dell's keynote on Tuesday at 8 a.m.
- One of Dell's 20+ conference sessions

www.dell.com/oracle

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®