
ORACLE®

ORACLE®

Oracle E-Business Suite Performance with Real Application Clusters & Advanced Compression

Lester Gutierrez, Uday Moogala
Applications Performance Group

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.

Agenda

- Optimize E-Business Suite on RAC
  - The Configuration bit
  - The Affinity bit
  - Triage & Tuning
- Optimize E-Business Suite with Advanced Compression
  - Introduction
  - Performance Testing
  - Oracle's Production System live & with 3X Storage Savings
  - Lessons Learnt & Best Practices
- Q & A


Optimize EBS on RAC
SOFTWARE.
HARDWARE.
COMPLETE.
TAPE
~AC ⎓DC
STANDBY
XSCF
ORACLE

Optimize E-Business Suite on RAC

- The Configuration bit…
  - RDBMS
    - Be at 10.2.0.4+ , 11.1.0.7+, 11.2.0.1+
    - Use ASM, ASSM & AUTOALLOCATE extent LMT's
    - Affinitize PX processes to each node
      - 10g - Set instance specific values for INSTANCE_GROUPS, PARALLEL_INSTANCE_GROUP
      - 11gR1 - Can be done as above or via services
      - 11gR2 - PARALLEL_FORCE_LOCAL =true
    - Buffer Cache +10%

Optimize E-Business Suite on RAC

- The Configuration bit…
  - Applications
    - Follow EBS RAC conversion MOS documents for patch and parameter info
      - 466219.1 11gr1 RAC with EBS R12
      - 823587.1 11gr2 RAC with EBS R12
    - Use Parallel Concurrent Processing
    - Use services
      - Apps workloads can be defined as services- batch/web/forms/ad-hoc; group related modules
      - 12.1.3 more detailed info for module/action

Optimize E-Business Suite on RAC

- The Configuration bit…
  - HW
    - Interconnect– Private; GigE or Infiniband
    - Network Configuration– Use Jumbo frames
      - `ifconfig <adapter> mtu 9000`
    - IO subsystem sized for performance, not just storage; write optimized
    - CPU utilization, run queue lengths & multi-threading considerations

Optimize E-Business Suite on RAC

- The Affinity bit …
  - Application Affinity:
    - What is it? Is it required? Why?
    - Goal: Direct some workload components to certain nodes to maximize scalability by reducing inter-node communication & synchronization work

Optimize E-Business Suite on RAC

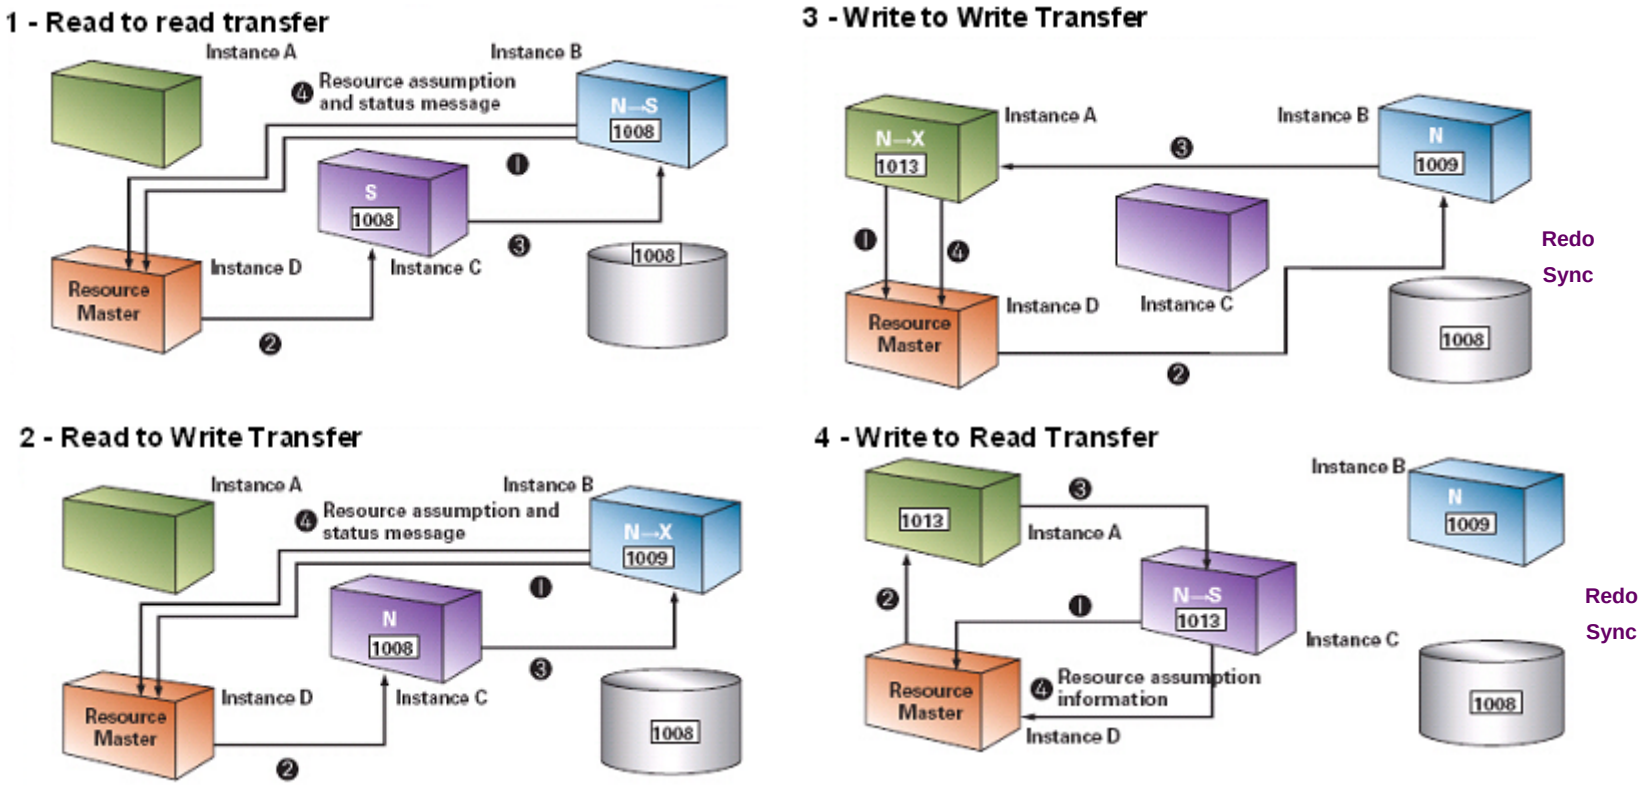

- This all happens very fast and is optimized with every version, but what happens when concurrent sessions on multiple nodes are inserting, updating, querying & committing the same blocks repeatedly ?

Optimize E-Business Suite on RAC

- The Affinity bit … Concurrent Manager
  - Use Parallel Concurrent Processing with 1 – 1 correspondence of CM nodes to RAC nodes
    - Set <s_cp_twotask> per CM node to specific RAC instance
      - Define primary/secondary CM nodes for failover
    - Use dedicated CM queues with specialization rules (inclusion/exclusion) for programs with high scalability needs
    - Group by application, program, user, request_type; examples:
      - OM, Receivables Super User, Create Accounting
      - WF background engines Node to ITEM_TYPE assignment
    - 12.1.3 introduced option to define node affinity at the program level – Concurrent -> Program -> Define

Optimize E-Business Suite on RAC

- The Affinity bit … Applications Tier
  - Forms
    - Set Profile Option "Database Instance" at App or Resp level which can be tied to two_task value or to a service
  - Self Service
    - Set Profile Options “App%Agent” to web tier hosts configured for specific services …or
    - Set Profile Option "Applications Database ID" to node specific DBC file name

Optimize E-Business Suite on RAC

- The Affinity bit …Examples
  - Oracle's Internal Production System
    - 16TB DB 4 Node Cluster, 11.1.0.7, 15K concurrent peak users

Node 1	Node 2	Node 3	Node 4
Order Management	Payroll	Financials –	Accounts Receivable
Install Base	HR (with self-service)	GL, AP, PO, FA	Project Accounting
Shipping	Benefits	Self-Service –	Imaging
Education	OTL (self-service)	iProc, Expenses	
Contracts		DBI	

Optimize E-Business Suite on RAC

- The Affinity bit …Examples
  - Large AR Telco system: **Batch Affinity**
    - *Node 1*: AutoInvoice,Create Accounting-Invoices
    - *Node 2*: AutoReceipt, Remittance, Create Accounting- Receipts
  - Order Management Benchmark: **Batch Affinity**
    - *Node 1*: Order Management, Shipping, WF background- OE Item Type
    - *Node 2*: Financials, Procurement, WF for PO item types
  - Incentive Comp system: **No Affinity**
    - *4 node cluster*, commodity HW; nodes too small to handle required OIC batch workload in one node
    - Scaled to customer's needs with RAC to required throughputs, BUT …
    - Needed custom partitioning, and several development and custom code optimizations, index changes/drops and other tuning

Optimize E-Business Suite on RAC

- Triage & Tuning
  - Some things are more expensive in RAC
    - Logical IO's in RAC must do more work; More LIO's, more RAC tax
      - Full table scans (gc cr multi block request)
      - Unselective Range Scans
      - High block concurrency work (CR clone - Undo work)
        - Right growing or unselective indexes
        - Can also happen at table level, but less frequent
    - Excessive parsing & locking
    - How much tax ?
      - *Lower Tax bracket* : Tuned SQL + Some level of App Affinity
      - *Higher Tax bracket*: Inefficient SQL and/or
        High inter-instance demands for same blocks (especially DML)

Optimize E-Business Suite on RAC

- Triage & Tuning
  - RAC focused AWR Analysis
    - Interconnect & Messaging Performance
    - RAC Wait Events
    - Top SQLs by Cluster Waits
    - Segment Statistics
      - *Global Cache Buffer Busy*
      - *CR blocks received*
      - *Current blocks received*
  - Need data from all instances
    - Global reports added in 11gR2 –
      - awrgrpt.sql and awrgdrpt.sql (cluster-wide base and diff reports)
    - 11gR1 - ADDM RAC report via DBMS_ADDM

Optimize E-Business Suite on RAC

- Triage & Tuning – AWR RAC Analysis
  - Interconnect Messaging Performance
    - Baseline & Monitor AWR
      - *Global Cache Load Profile*– "Estimated Interconnect Traffic"
      - *Global Cache and Enqueue Services - Messaging Statistics*
        - "Average message sent queue time on ksxp" – IPC efficiency; > 2ms = interconnect/CPU capacity/issue
    - CPU: Out of juice ? – Long run queues & CPU spikes; LMS (by default scheduled real time in 11g) involved in
      - All commits: as other instances must ACK the SCN propagation (BOC) - (could see longer log file sync waits)
        - bug fixes for BOC in 11.1.0.7 & 10.2.0.4
        - 2 new stats : broadcast on commit waits/time
      - Inter-instance messaging (ges/gcs) and block shipping/receiving

Optimize E-Business Suite on RAC

- Triage & Tuning – AWR RAC Analysis
  - Interconnect Messaging Performance

RAC Statistics

	Begin	End
Number of Instances:	4	4

Global Cache Load Profile

Eg) Network Traffic received = 232.31*8k=1.9 Mb/sec

	Per Second	Per Transaction
Global Cache blocks received:	232.31	6.57
Global Cache blocks served:	473.15	13.38
GCS/GES messages received:	7,437.34	210.35
GCS/GES messages sent:	7,447.23	210.63
DBWR Fusion writes:	3.29	
Estd Interconnect traffic (KB)	8,550.82	

Global Cache and Enqueue Services - Messaging Statistics

Avg message sent queue time (ms):	0.2
Avg message sent queue time on ksxp (ms):	0.6
Avg message received queue time (ms):	0.0
Avg GCS message process time (ms):	0.1
Avg GES message process time (ms):	0.0
% of direct sent messages:	73.91
% of indirect sent messages:	25.15
% of flow controlled messages:	0.94

Optimize E-Business Suite on RAC

- Triage & Tuning
  - Common Top RAC Wait Events
    - *gc cr/current [2/3] way*
    - *Gc cr/current grant*
    - *gc cr/current block busy*
    - *gc buffer busy %*
  - Related events
    - *Log file sync*
    - *buffer busy*
    - *enq: HW/FB – contention*

Optimize E-Business Suite on RAC

- Triage & Tuning
  - SQL ordered by Cluster Waits

Top 5 Timed Foreground Events

Event	Waits	Time(s)	Avg wait (ms)	% DB time	Wait Class
DB CPU		83,660		46.73	
gc buffer busy acquire	3,495,424	23,219	7	12.97	Cluster
gc cr block busy	7,368,841	16,071	2	8.98	Cluster
gc current block busy	627,037	11,137	18	6.22	Cluster
db file sequential read	3,060,429	8,277	3	4.62	User I/O

SQL ordered by Cluster Wait Time

- %Ela is Cluster Wait Time as percentage of SQL Elapsed Time
- %Total is Cluster Wait Time as percentage of Total Cluster Wait Time
- Only SQL with Cluster Wait Time > .005 seconds is reported
- Total Cluster Wait Time (s): 176
- Captured Segments account for 77.4% of Total

Cluster Wait Time (s)	%Ela	%Total	Elapsed Time(s)	CPU Time(s)	Executions	SQL Id	SQL Module	SQL Text
145.10	6.52	82.40	2,226.34	1,107.17	4	5t39uchjqpyfm	XLAACCUP	BEGIN xla_accounting_pkg.unit_...
56.34	36.31	32.00	155.16	4.78	10	9m0fuv7qx6pa0	XLAACCUP	INSERT INTO AR_XLA_LINES_EXTRA...
33.43	45.26	18.98	73.86	1.48	4	0tsz27uj32p77	XLAACCUP	UPDATE GL_BC_PACKETS SET GROUP...
17.73	69.51	10.07	25.52	0.56	4	59y8ffvpxndw5	XLAACCUP	UPDATE GL_JE_BATCHES JEB SET S...
6.36	2.43	3.61	261.19	111.68	10	cyj399dwznajq	XLAACCUP	UPDATE /*+ INDEX(ra_cust_trx_l...
5.10	1.62	2.90	315.14	69.04	10	8q6dh8mqr8j47	XLAACCUP	INSERT INTO AR_XLA_LINES_EXTRA...
3.19	23.35	1.81	13.67	3.82	10	dqh0sx2ra584c	XLAACCUP	UPDATE (SELECT /*+ leading(tmp...
2.12	0.29	1.20	740.10	576.82	10	96cqwbvtkyqk5	XLAACCUP	BEGIN :1 := XLA_00222_AAD_S_00...
2.06	6.36	1.17	32.37	3.43	1	cfz686a6qp0kq		select o.obj#, u.name, o.nam...
1.57	12.77	0.89	12.29	1.54	10	78qrvpzszm6j7	XLAACCUP	INSERT INTO XLA_AE_HEADERS (A...
1.52	2.26	0.87	67.57	31.55	10	b5tq6435p6j2h	XLAACCUP	INSERT INTO XLA_DISTRIBUTION_L...
1.43	5.47	0.81	26.22	16.23	10	70upqfsmsq66a	XLAACCUP	INSERT INTO XLA_AE_LINES (AE_...

Optimize E-Business Suite on RAC

- Triage & Tuning - Segment Statistics

Segments by Global Cache Buffer Busy

- % of Capture shows % of GC Buffer Busy for each top segment compared
- with GC Buffer Busy for all segments captured by the Snapshot

Owner	Tablespace Name	Object Name	Subobject Name	Obj. Type	GC Buffer Busy	% of Capture
APPLSYS	APPS_TS_TX_IDX	WF_ITEMS_N3		INDEX	279,973	19.54
MRP	APPS_TS_INTERFACE	MRP_RELIEF_INTERFACE_U1		INDEX	178,510	12.46
INV	APPS_TS_TX_DATA	MTL_RESERVATIONS		TABLE	150,959	10.53
ASO	APPS_TS_QUEUES	ASO_ORDER_FEEDBACK_T		TABLE	123,129	8.59
INV	APPS_TS_TX_IDX	MTL_MATERIAL_TRANSACTIONS_U2		INDEX	53,330	3.72

Segments by CR Blocks Received

- Total CR Blocks Received: 25,023,824
- Captured Segments account for 51.8% of Total

Owner	Tablespace Name	Object Name	Subobject Name	Obj. Type	CR Blocks Received	%Total
INV	APPS_TS_TX_DATA	MTL_RESERVATIONS		TABLE	5,725,582	22.88
INV	APPS_TS_INTERFACE	MTL_TRANSACTIONS_INTERFACE		TABLE	1,612,625	6.44
WSH	APPS_TS_TX_DATA	WSH_DELIVERY_DETAILS		TABLE	1,578,776	6.31
ONT	APPS_TS_TX_DATA	OE_ORDER_LINES_ALL		TABLE	426,800	1.71
INV	APPS_TS_TX_IDX	MTL_RESERVATIONS_N2		INDEX	400,771	1.60

Optimize E-Business Suite on RAC

- Triage & Tuning

  - RAC AWR & ASH Analysis – How ?
    1. Verify SQL execution is optimized – 10046 trace/SQL Monitoring
    2. Collect AWR data from all instances
    3. Verify IO, CPU utilization and Interconnect messaging performance are within acceptable limits
    4. Drill down on cluster wait events + “SQL order by Cluster Waits”
    5. Correlate SQL executions on all instances accessing/modifying top segments in Segment Statistics by
       - Global Cache Buffer Busy & Buffer Busy Waits
       - CR/Current Blocks Received
       - Logical Reads

Optimize E-Business Suite on RAC

- Triage & Tuning
  - Common Remedial Actions
    - For the Nth time! ….SQL Tuning & App Affinity
    - Avoid over-indexing: consolidate/drop redundant
    - Right growing indexes: global hash partition or use larger sequence caches
    - Unselective indexes: append sequence driven key
    - Table partitioning
    - Workload staggering & smaller commit cycles
    - Avoid ordered sequences- 10g:DFS lock handle wait bug:5209859

Optimize E-Business Suite on RAC

- Triage & Tuning - Remedial actions by issue

Wait or Symptom	Causes	Remedial Actions
gc cr multi block request	Large Full Table Scans	Tune SQL, Review App Affinity; If *gc cr multiblock* timeouts are high check UDP receive buffer size
enq: HW/FB - contention + gc current grant	High insert concurrency; too frequent space allocations	Use ASSM with autoallocate or large uniform extents are used for the objects with significant growth

Optimize E-Business Suite on RAC

- Triage & Tuning - Remedial actions by issue

Wait or Symptom	Causes	Remedial Actions
gc cr block 2/3 way and or gc cr/current grant	lots of LIO's , blocks not in local cache and not locally mastered. Lots of blocks in other caches	Review app affinity, and review SQL doing lot's of LIO's ..SQL tuning these is your best bet
gc%congested/failure/ retry	Interconnect related problem and/or CPU bound nodes. Rare.	ping/netstat and related tools to check interconnect performance; workload mgmt

Optimize E-Business Suite on RAC

- Triage & Tuning - Remedial actions by issue
  - ***gc buffer busy***

Causes	Remedial Actions
• Data/Index blocks accessed at a higher rate than the time it takes to load and pin buffers in the cache • GC Buffer Busy > accesses to buffer made between a remote read starts and completes will queue up on this • Pin time= (time to read block into cache) + (time to modify/process it) • Busy time = (Avg. pin time) * (# users waiting ahead of me) • Queues can build up on multiple nodes	- Check AWR Buffer Wait section. For 'data block' class, usually SQL tuning issue: bad execution plan with excessive/repeating LIOs of a set of blocks coupled with concurrent sessions across nodes executing bad code (i.e. mjc) - AWR & trace will identify specific SQLs, objects and plans - Fix execution plans - Review App Affinity

Optimize E-Business Suite on RAC

- Triage & Tuning - Remedial actions by issue
  - ***gc cr/current block busy***
    - Block received but LMS delayed send as block was not ready. Wait times mostly block flush (log file sync) or defer time (cu blocks) on the serving node

Causes	Remedial Actions
• Results from high concurrency and block contention • Concurrent updates to the same blocks • Inserting of monotonically increasing primary keys or same values to indexes at high rates • Worsened by high IO times or system load	• Check log file sync time & DBWR IO performance on other instances • Check segment stats, Cluster wait SQL • For right growing, PK/UK driven idxs > hash partition or raise sequence cache • For unselective indexes > drop/consolidate or append sequence driven or node-affinity key • Review App affinity • For tables, review partitioning if options

Optimize EBS with Advanced Compression

SOFTWARE.
HARDWARE.
COMPLETE.

ORACLE

Advanced Compression – Introduction

- Introduced in Oracle Database 9i Release 2
  - Compression during bulk load operations (Direct Load, CTAS)
  - Data modified using conventional DML not compressed
- Oracle Database 11*g* extends table compression for OLTP data
  - *Support for conventional DML Operations (INSERT, UPDATE)*
  - New algorithm significantly reduces write overhead
    - Batched compression minimizes impact on OLTP transactions
  - No impact on reads
    - Reads data directly in its compressed format
    - Reads may actually see improved performance due to fewer I/Os and enhanced memory efficiency

Advanced Compression – Introduction
OLTP Table Compression Process

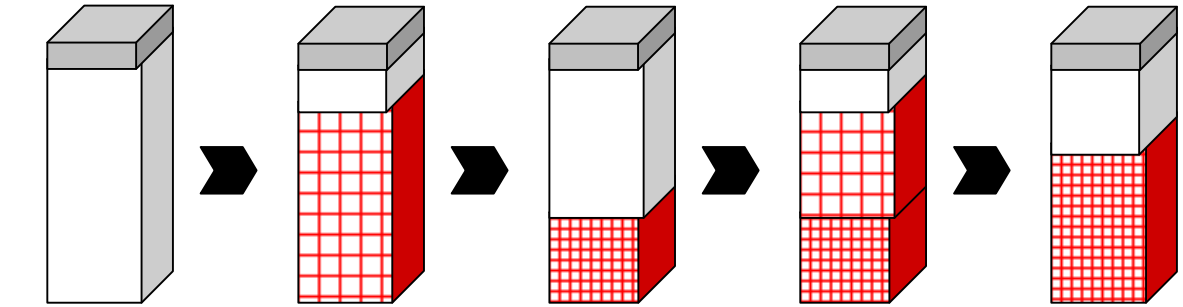

Empty Block	Initially Uncompressed Block	Compressed Block	Partially Compressed Block	Compressed Block

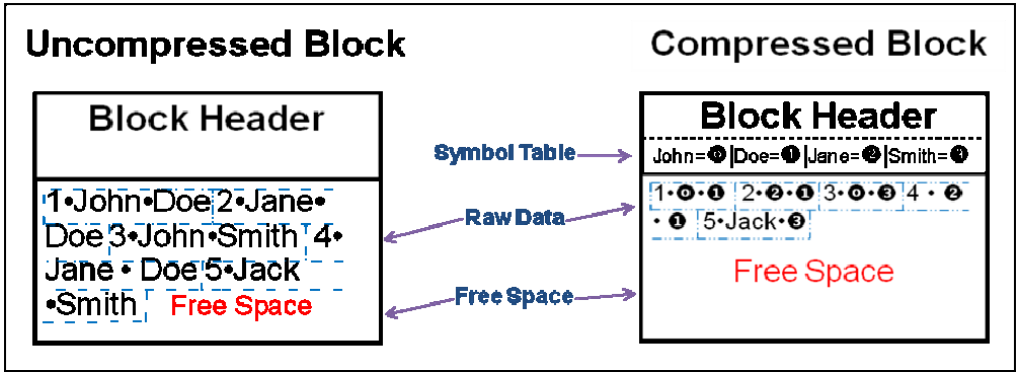

Advanced Compression – Introduction
Block-Level *Batch* Compression

- 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

Advanced Compression – Introduction
Index Compression

- Compress multi-column indexes to reduce the storage overhead of repeated values
- Stores more keys per leaf block
- Less I/O,better performance of range scans

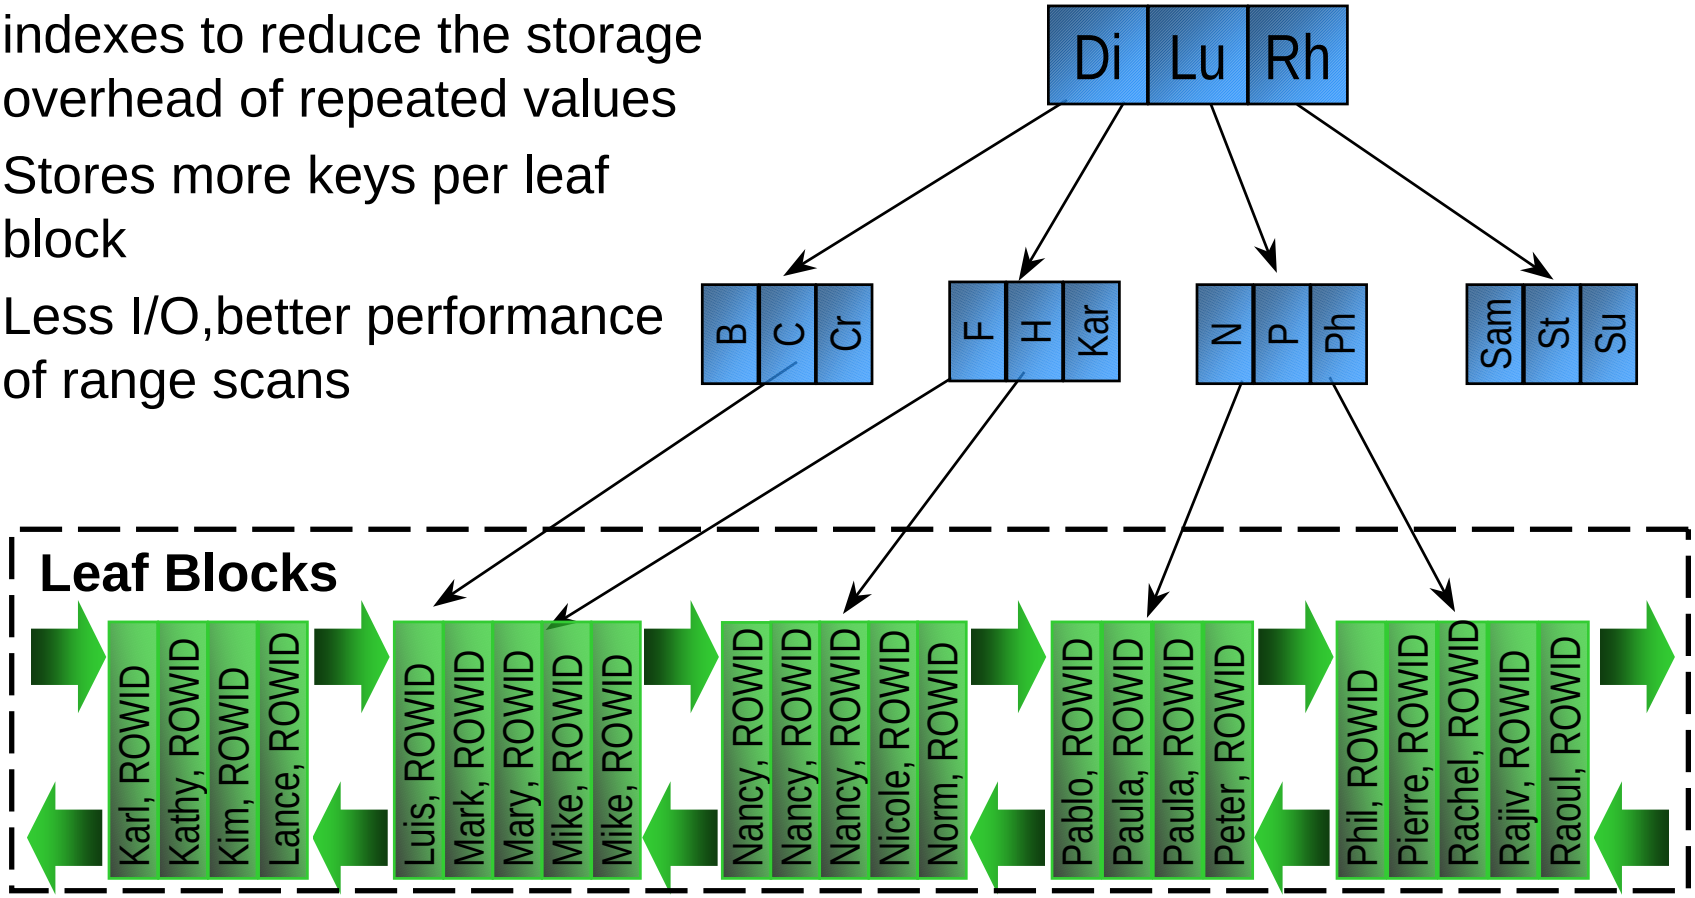

Table Compression Advisor
Estimate Potential Storage Savings

- 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)

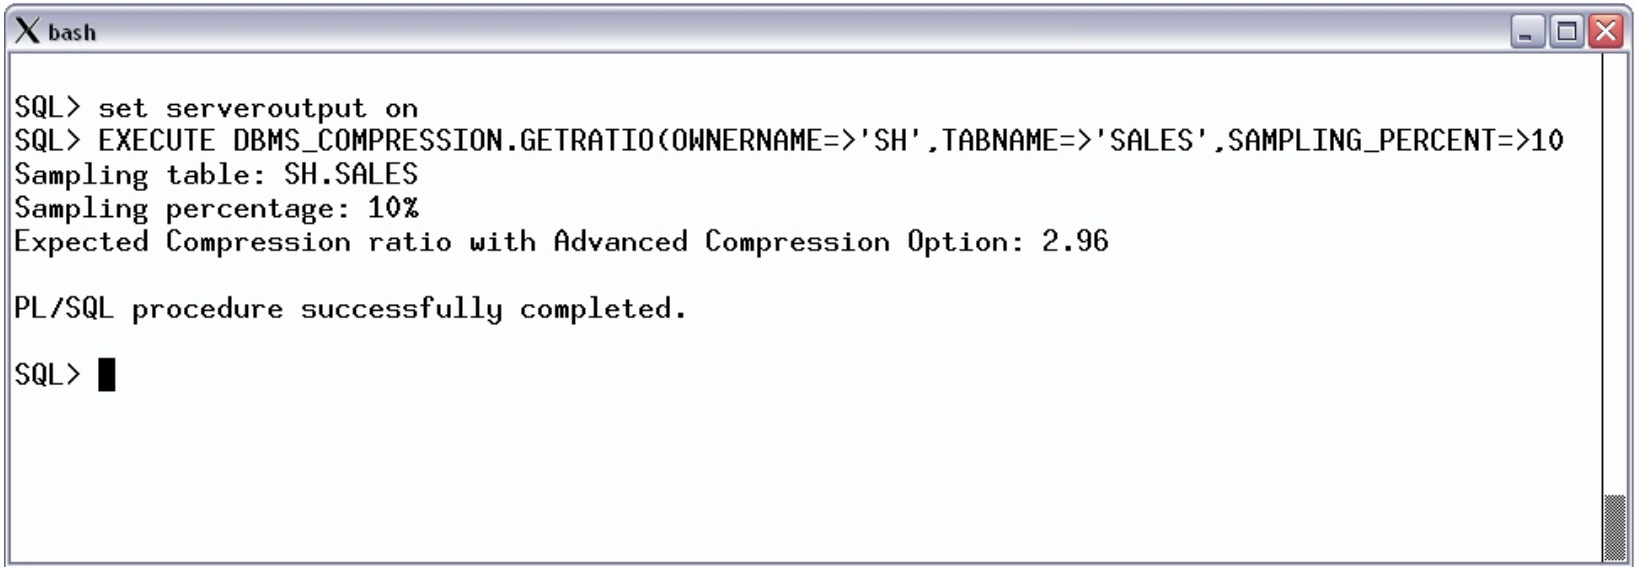

** http://www.oracle.com/technology/products/database/compression/compression-advisor.html*

Oracle 11g Advanced Compression Option (ACO)

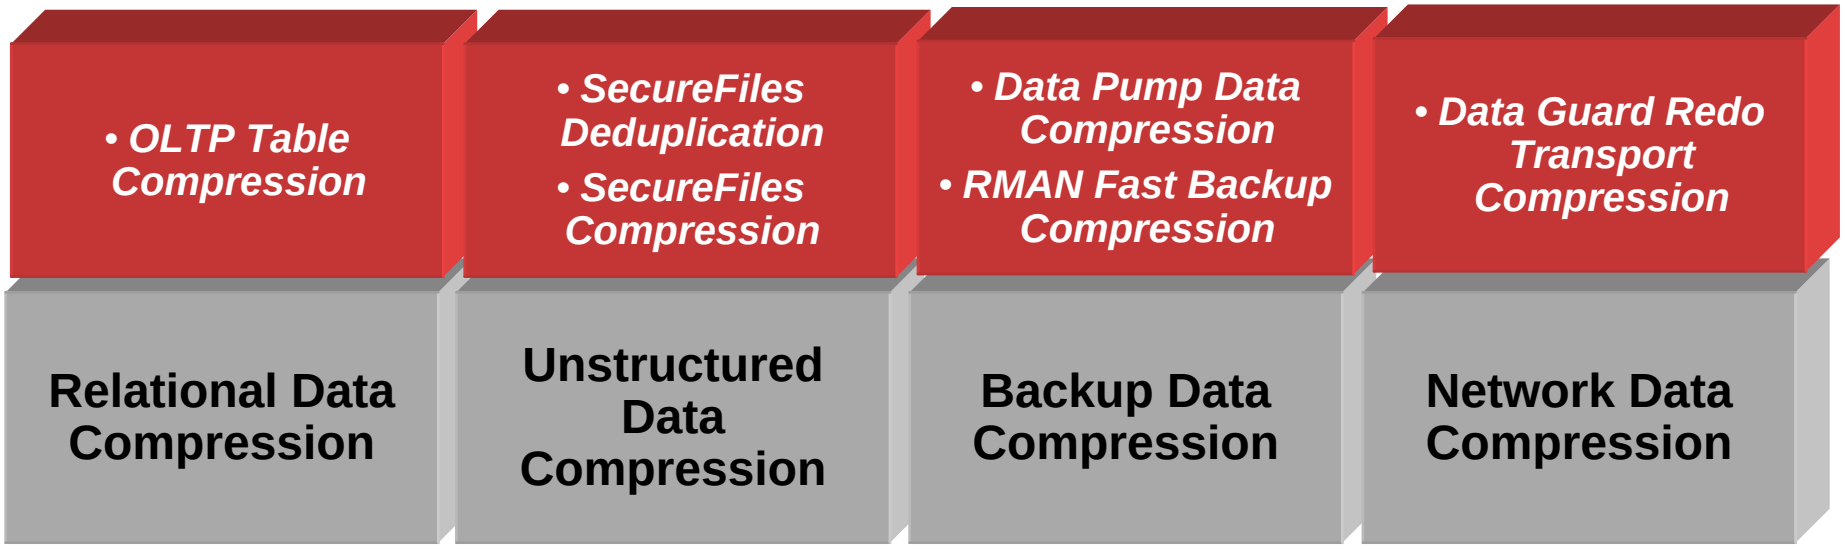

- Reduces resource requirements and costs
  - Storage System
  - Network Bandwidth
  - Memory Usage

ORACLE DATABASE 11g

Hybrid Columnar Compression (HCC)

Two New Features in Exadata V2

Warehouse Compression

- 10x average storage savings
- 10x reduction in Scan IO

Optimized for Speed

Archive Compression

- 15x average storage savings
  - Up to 70x on some data
- For cold or historical data

Optimized for Space

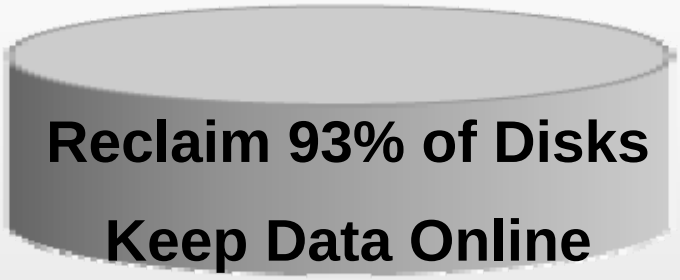

Can mix OLTP and hybrid columnar compression by partition for ILM

Oracle E-Business Suite
Advanced Compression
Performance Testing

Test Scenarios

- Batch Flows
  - Payroll Process
  - Order-To-Cash
- Online Flows
  - Self Service 100 user test of 7 Self Service actions (related to expenses, timecards and payslips)
  - 120 user test of 16 Purchasing actions (related to creating, approving and searching PO’s and invoices)

Results – Storage Savings

- Enabled OLTP Table Compression
  - 200 largest tables were selected
  - Tables reorganized before compression

Description	Storage Utilized	Storage Saved	Percent Savings
Uncompressed Tables	21.9 GB	N/A	N/A
Uncompressed Tables After Reorganization	21.8 GB	110 MB	0.5
Tables with OLTP Table Compression	6.9 GB*	14.8 GB	68

* 3.1 Compression Ratio

Test Results - Observations

- The online workload showed up to 30% improvement with up to 6% increase in CPU consumption
- The Order-To-Cash Batch runtime increased by 0.43% and CPU usage decreased by 7%
- The Payroll Process Batch runtime reduced by 17% with an additional 3.7% CPU usage
- 38 of the largest tables could not be compressed due to table with more than 255 columns or that include LONG columns, which are Advanced Compression restrictions. They consume 1.4 GB.

Results – Performance

- Purchasing CPU – user+sys
  - Average CPU consumption increased by 0.7%

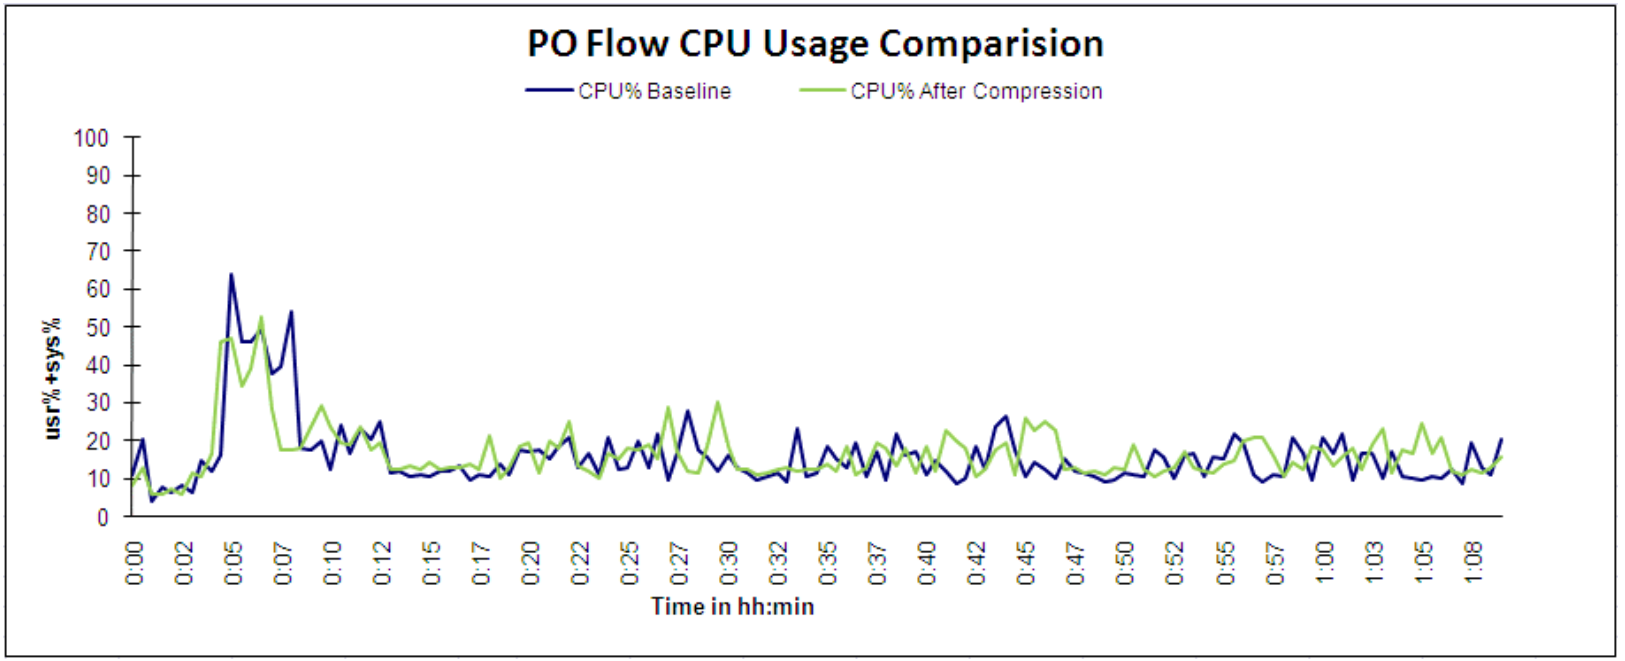

Results – Performance

- Payroll CPU – user+sys
  - Average CPU consumption increased by 3.7%

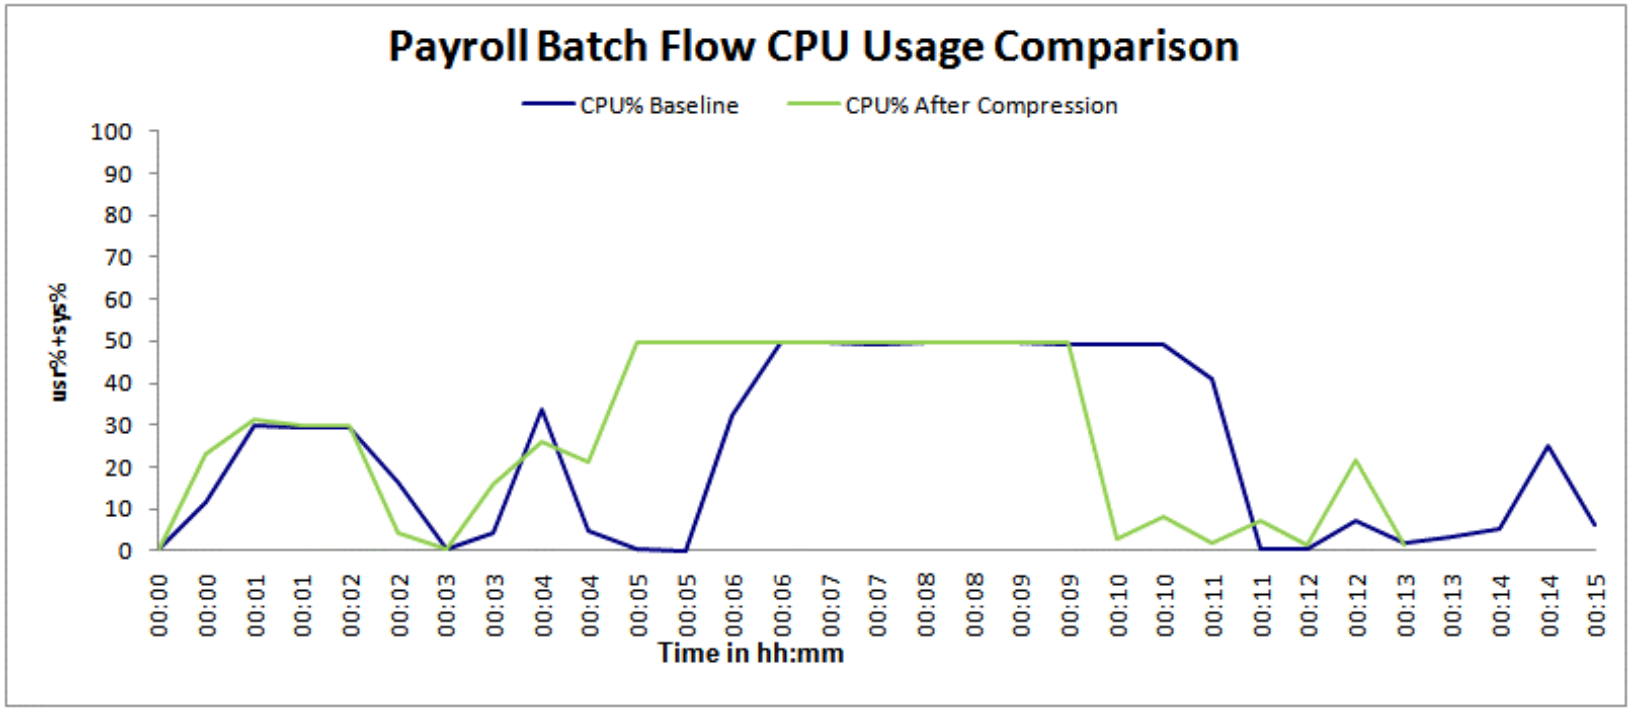

Advanced Compression Results In Oracle's Production System (GSI)

Oracle's Internal E-Business Suite DB
- Global Single Instance (GSI)

- Major modules implemented
  - HR - Payroll, Benefits, OTL
  - Financials - GL, AP, AR, FA, CM, PN, Self Service
  - Order Management, Shipping
  - Projects, Contracts
- Supports 90,000 employees and business groups
  - Over 15,000 concurrent sessions (peak)
  - 500,000 concurrent jobs completed per week (peak)
- 18 TB database growing by 500 GB per quarter
  - ~1.6 Billion rows in XLA Distribution Links
  - ~1 Billion rows Subledger Journal Entry Lines
  - ~900 Million rows in GL Journal Entry Lines
  - ~220 Million rows in Invoice Distributions
  - ~216 Million in AP Accounting Lines

43

Oracle's Internal E-Business Suite DB
- Global Single Instance (GSI)

Compressed:

~260 tables

~1600 indexes

~28 LOBs using 11g SecureFiles

- 1st phase, compressed 100 largest segments, which accounted for 40% space usage
- 2nd phase, compressed tables that accounted for most IO

44

Oracle's Internal E-Business Suite DB
- Global Single Instance (GSI)

231 TB
of realized storage savings!

- In primary, standby and test systems
- Additional benefits were also accrued in dev clones and backups

45

Oracle's Internal E-Business Suite DB
- Global Single Instance (GSI)

- **Average overall storage savings 3x**
  - Table compression 4x
  - Index compression 2x
  - LOB compression 2.3x

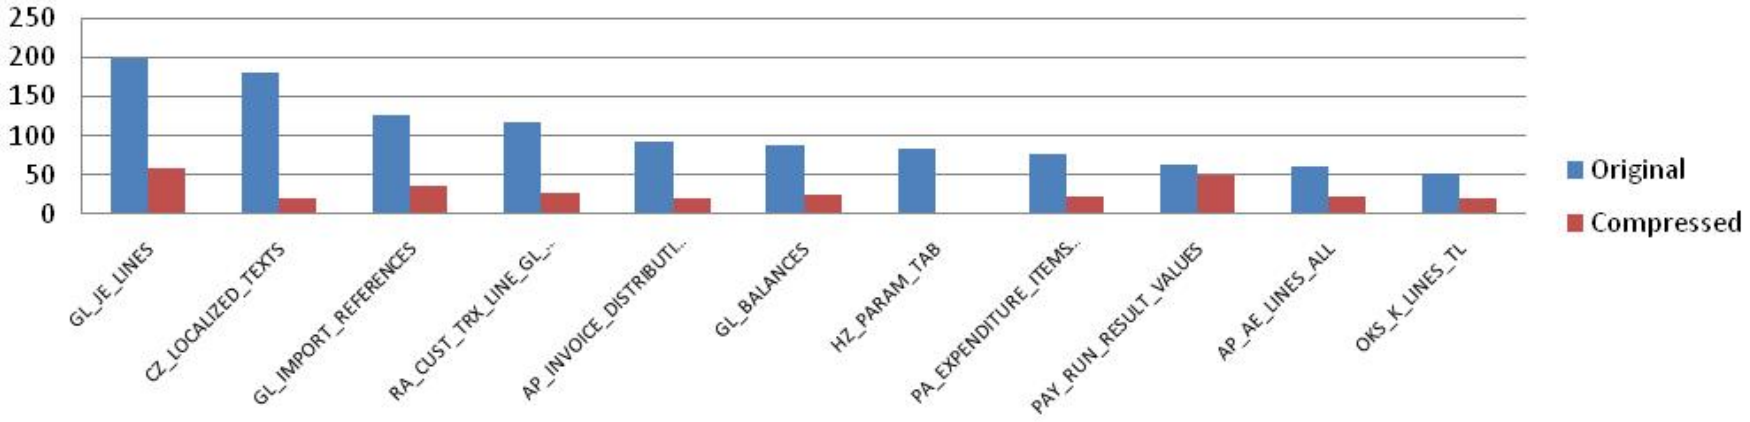

46

Oracle's Internal E-Business Suite DB
- Global Single Instance (GSI) with Exadata HCC

- Archive High
  - Overall: 23x
  - Largest Table: 52x
- Query High
  - Overall: 15x
  - Largest Table: 35x
- OLTP Compression
  - Overall: 3.6
  - Largest Table: 4x

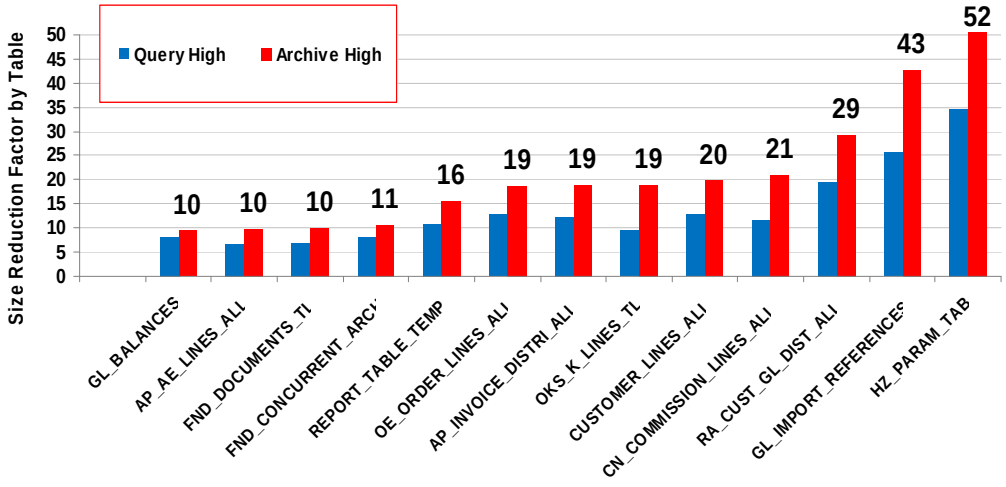

Up to 52x reduction in table size for EBS

47

Lessons Learnt & Best Practices

Lessons Learnt & Best Practices

- **Table Compression**
  - Focus on your largest tables and
  - Tables that account for most IO
  - Consider excluding high transaction tables such as FND_CONCURRENT_REQUESTS
  - Sort the read only tables to achieve best compression

Lessons Learnt & Best Practices

- **Index Compression**
  - Apply compression to the subset of indexes that provide significant space reduction rather than applying index compression on all indexes
  - Compress indexes with repeating key values to achieve best compression ratio
  - To see optimum number of columns to compress
    - Analyze the index
    - Look at opt_cmpr_count and opt_cmpr_pctsave in index_stats table
  - All index partitions should be compressed
  - Do not include columns that provide uniqueness to multi-column index i.e., compress leading columns that can have repeated values

Lessons Learnt & Best Practices

- **Patches**
  - Review and apply recommended Advanced Compression Option and DML/Space Layer patches. MOS Doc: 1061366.1

- **ITL Contention**
  - In case of concurrency issues, *INITRANS* and *PCTFREE* can be increased
  - INITRANS should only be raised if significant ITL wait events are observed in AWR/SQL Trace data

Lessons Learnt & Best Practices

- **Expect some SQL plan regressions**
  - Establish SQL Plan Baselines
    - Recommended best practice anyway
  - Consider using SQL Profiles or SQL Plan baseline while issue is being fixed

- **DB Block Checking**
  - Increased overhead observed for customers running with db_block_checking set to TRUE
    - more rows fit into one block -> more data will be checked for corruption resulting in higher CPU consumption

Lessons Learnt & Best Practices

- **SecureFiles LOB Compression and Deduplication**
  - LOB data should be stored using the SecureFiles feature
  - SecureFiles Compression and Deduplication can reduce storage for LOBs significantly
- **RMAN: use ACO to compress your backups**
- **Data Pump: use ACO to compress your Data Pump exports**

Lessons Learnt & Best Practices
- Information Lifecycle Management (ILM)

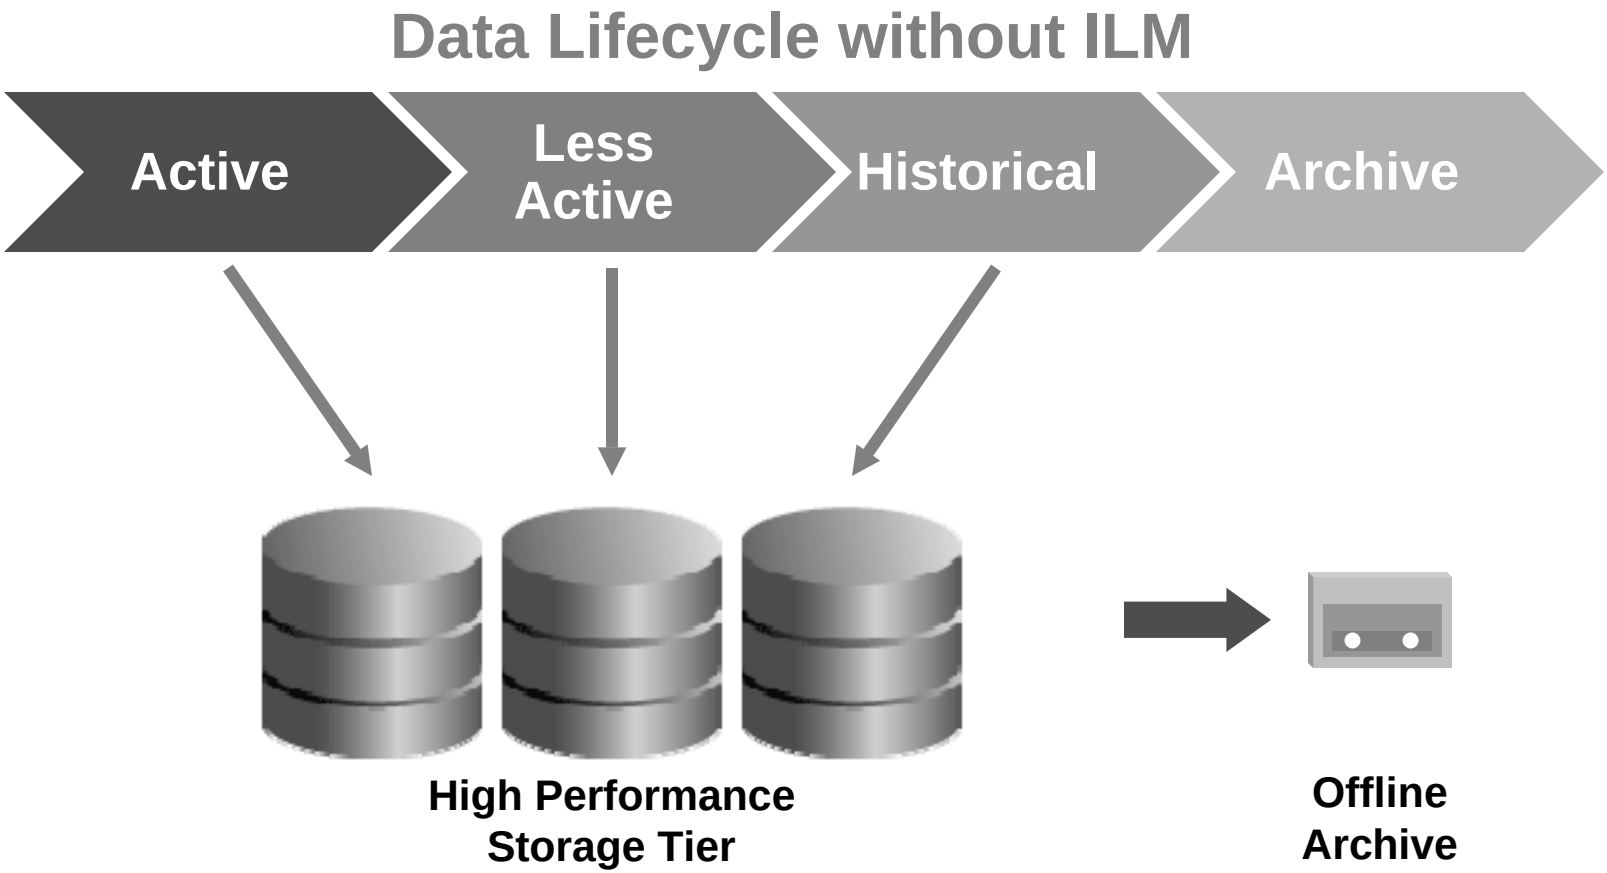

Lessons Learnt & Best Practices
- Information Lifecycle Management (ILM)

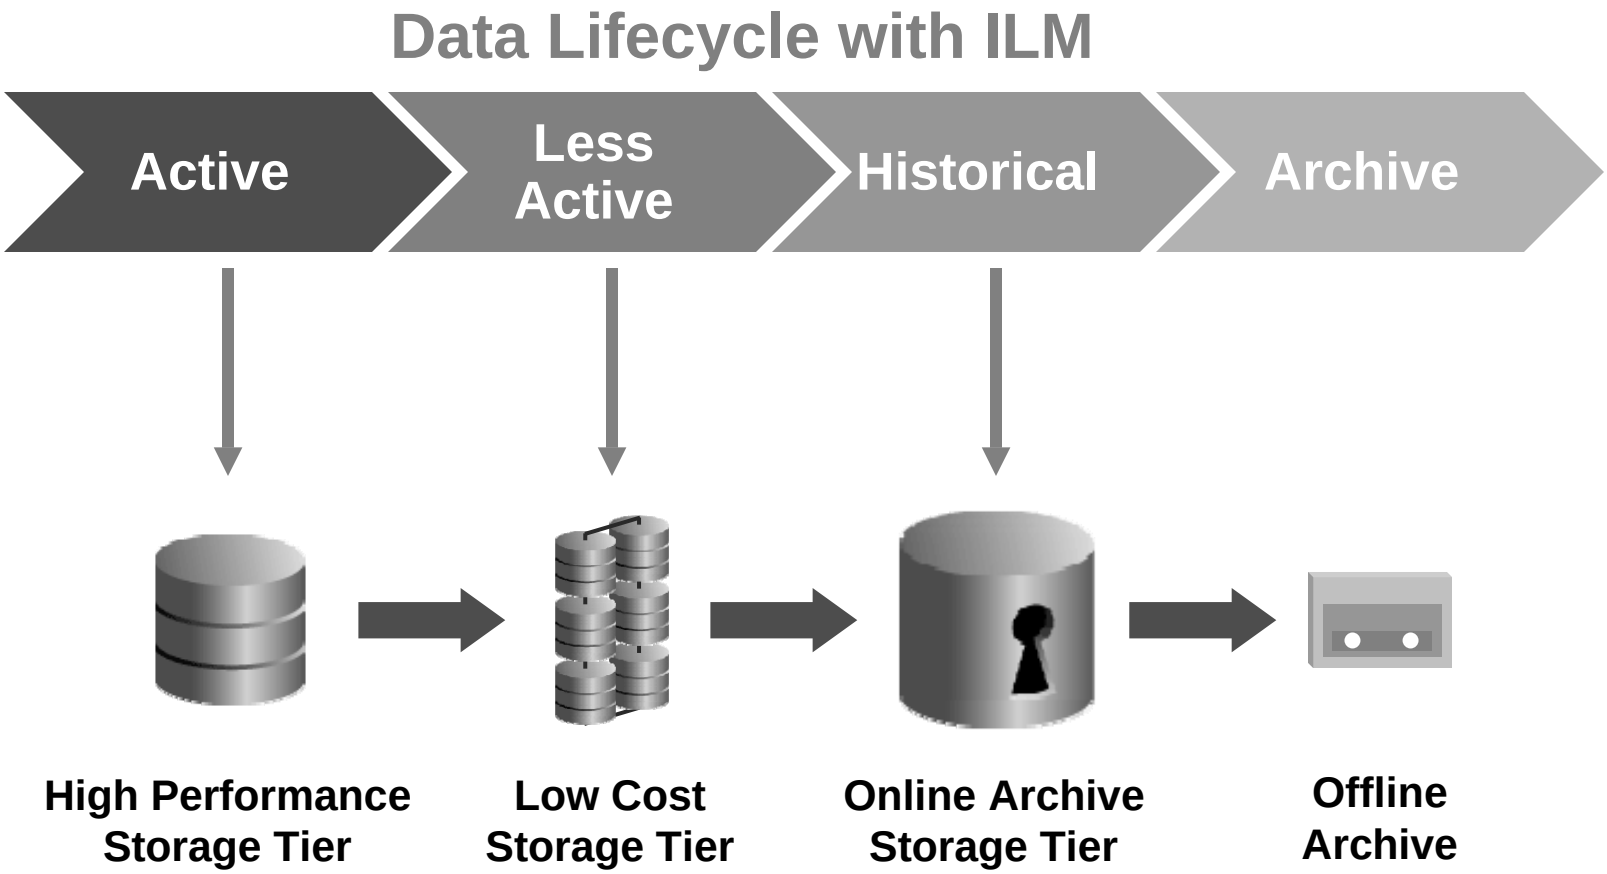

Lessons Learnt & Best Practices
- Information Lifecycle Management (ILM)

- **ILM with Partitioning + ACO**
  - Can partition large EBS tables by CREATION_DATE
    - Interval partitioning can be used for automatic partition creation
  - Keep indexes GLOBAL
  - Monthly/Quarterly/Yearly compress partitions holding older/dormant data
    - Online partition maintenance operation
    - Can move older data to lower tier storage

Lessons Learnt & Best Practices
- Information Lifecycle Management (ILM)

- On going deployment with a large customer
  - Special case of 10x compression of Invoices Distributions
  - Lot more data because compression in one block resulting in high block contention for high volume batch updates
- Proposed solution
  - Partition the table using Interval partitioning on creation_date column and create global indexes
  - Compress partitions with non-current data only
  - Reduced contention
  - Other alternatives could be increasing PCTFREE and also increase INITRANS
  - Can still reap ACO benefits for vast majority of data

Lessons Learnt & Best Practices

ILM and Data Archiving Strategies with Exadata HCC

- OLTP Applications
  - Table Partitioning
  - Heavily accessed data
    - Partitions using OLTP Table Compression
  - Cold or historical data
    - Partitions using Online Archival Compression
- Data Warehouses
  - Table Partitioning
  - Heavily accessed data
    - Partitions using Warehouse Compression
  - Cold or historical data
    - Partitions using Online Archival Compression

Optimize E-Business Suite with Advanced Compression - Summary

- Comprehensive data compression capabilities for all types of data
  - Structured, Unstructured, Backup, Network
- Reduces storage consumption by 2 to 4 times
- Improves read performance
- Enhances memory, buffer cache utilization
- Complete application transparency
- Benefits diverse application workloads
- White paper on ACO with E-Business Suite Release 12 at MOS Document 1110648.1

Thank You

Oracle Products Available Online

ORACLE STORE

Oracle Store

Buy Oracle license and support online today at
oracle.com/store

SOFTWARE. HARDWARE. COMPLETE.

We encourage you to use the newly minted corporate tagline “Software. Hardware. Complete.” at the end of all your presentations. This message should replace any reference to our previous corporate tagline “Oracle Is the Information Company.”

Oracle OpenWorld

Latin America 2010

December 7–9, 2010

Oracle OpenWorld

Beijing 2010

December 13–16, 2010


Q&A