

Oracle Exadata for Consolidated Business Growth

November 8, 2012

Ken Wood

Troy Ligon



nielsen
.....



nielsen
.....


Copyright © 2012 The Nielsen Company. Confidential and proprietary.

The Nielsen Company



Nielsen is a leading global information and measurement company that provides clients with a comprehensive understanding of consumers and consumer behavior

Headquartered in New York, USA
Presence in more than 100 countries
4 major datacenters, with numerous satellite datacenters



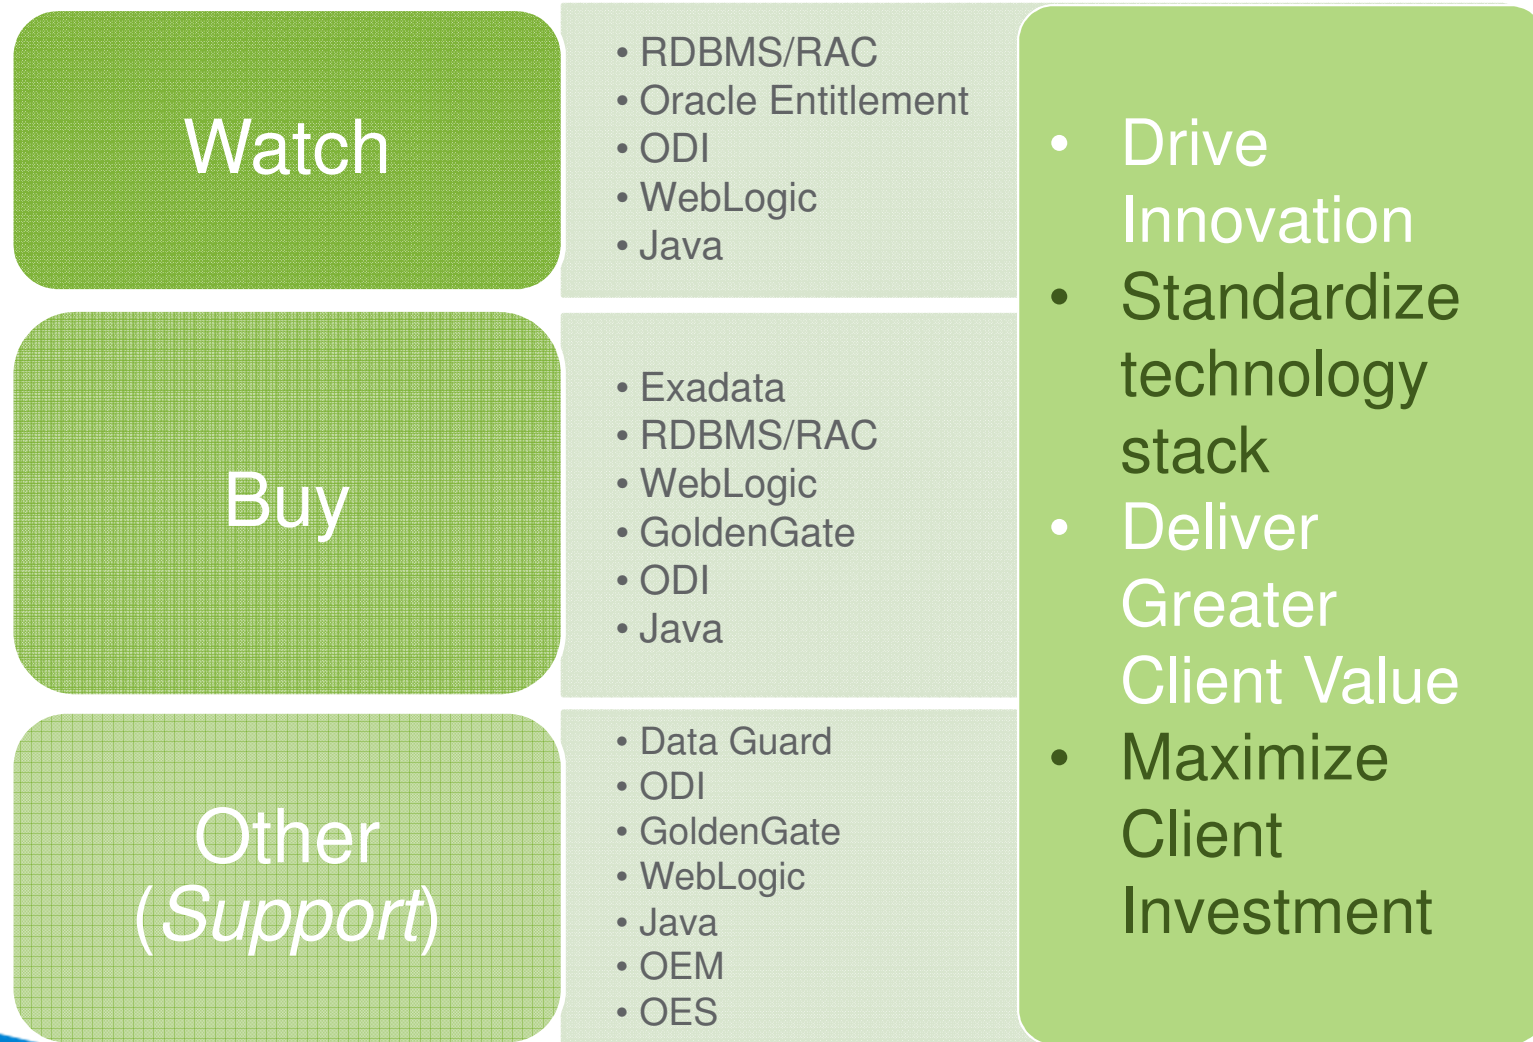
What Consumers Watch

Provides viewership data and analytics primarily to the media and advertising industries across television, online, and mobile screens

What Consumers Buy

Provides retail transactional measurement data, consumer behavior information, and analytics primarily to businesses in the consumer packaged goods industry

Nielsen/Oracle Relationship



Project Background – One Global Reference Data System

GOAL

To consolidate all regional reference data systems and implement a single global platform supporting all Nielsen company functionality. Consolidation of consumer reference and location specific data today and media reference data in the future.

NEED

A characteristic management platform supporting single view of unique products (with/with out unique identifiers) or locations globally.

BENEFIT

The global reference data system enables Consolidation, Standardization, and a common user experience

OGRDS – Key Capability Requirements

The Characteristics Management platform must

- Handle significant numbers of attributes coupled with a high frequency of growth and change.
- Variation of categories including varied characteristics based on category type and region.
- Enable centralized reference data and process governance across business domains and geographies.
- Manage characteristics, descriptions, translations, and properties to the appropriate business entities through a single workflow process.

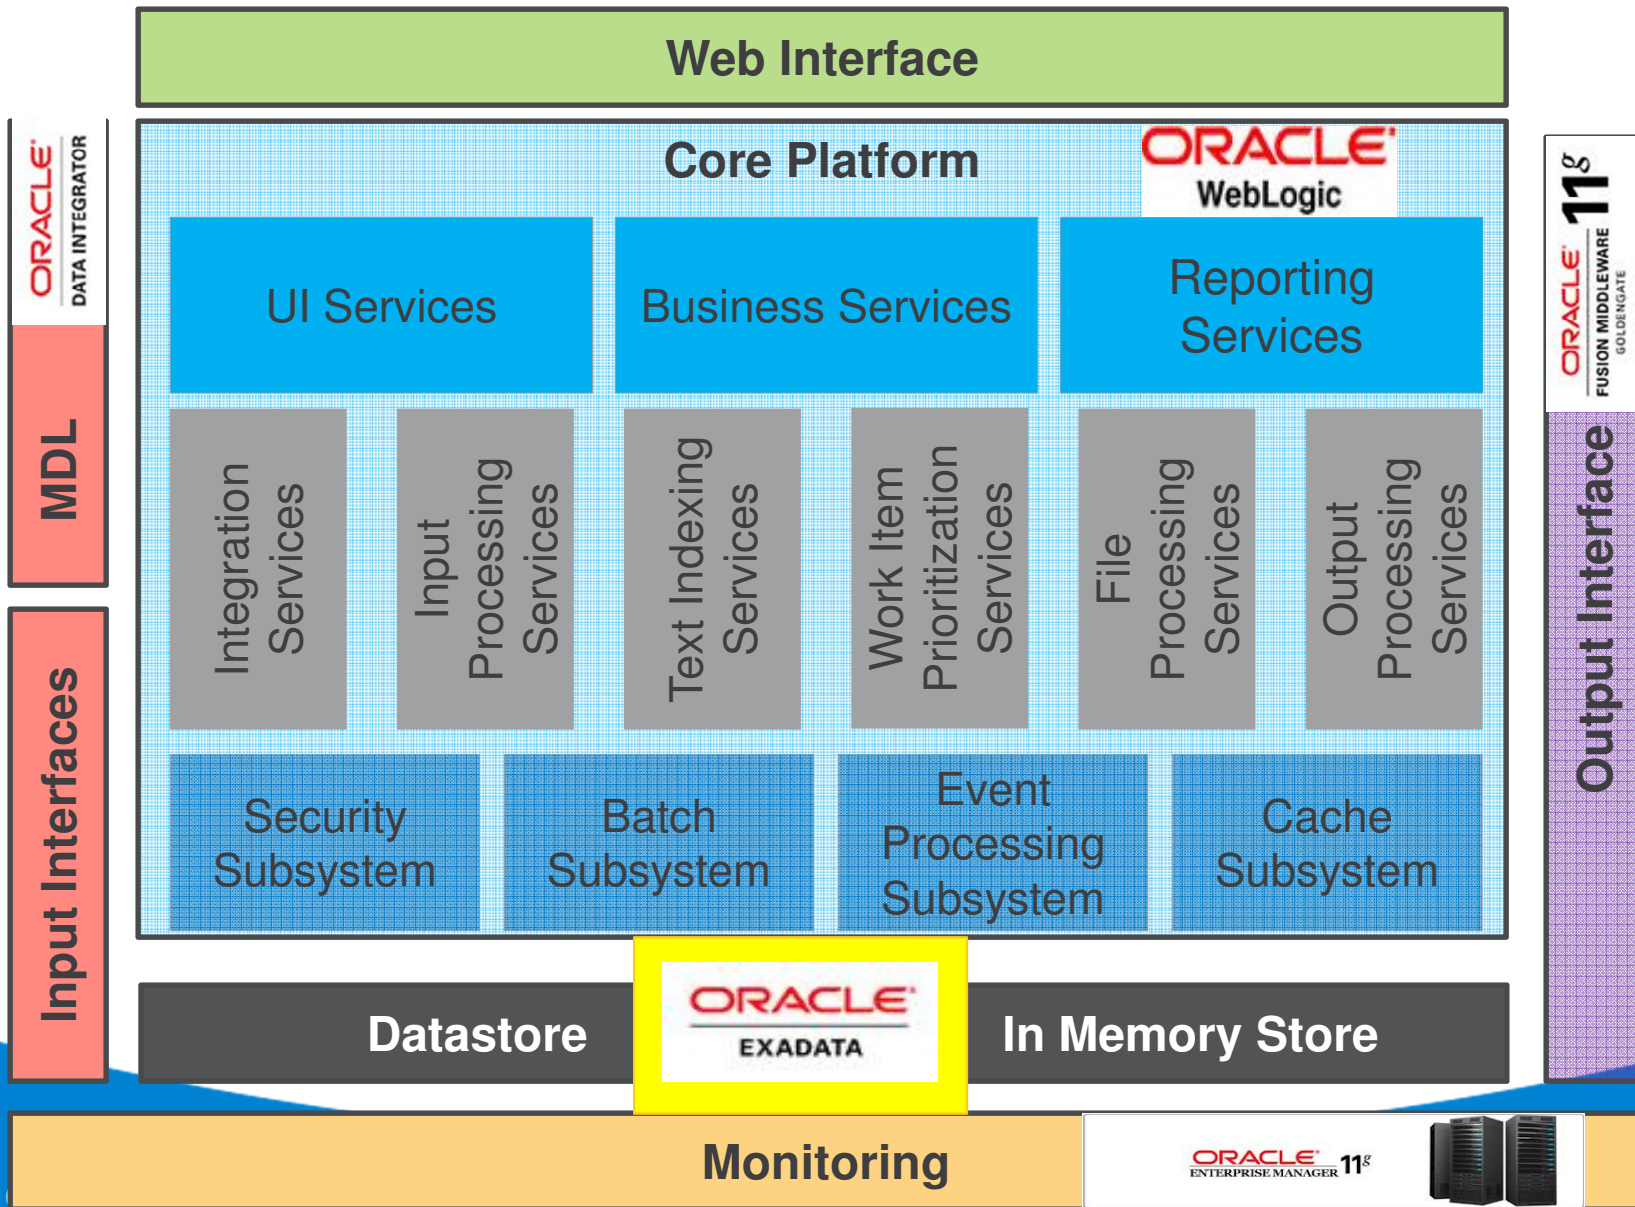
Consequently utilizes a Key Value Pair data structure which

- Must meet global OLTP coding speeds.
- Handles mixed workload.

The hardware of choice must support this...

5

Logical Architecture



Why Exadata?

- ✓ High data volume and processing requirements
 - Key Value Pair Model
- ✓ Scalability - Growth for BAU and New Countries
- ✓ Support for Mixed workload (OLTP, Batch, Report Pub)
- ✓ Flexibility for change in consumer behavior
- ✓ Easier Maintenance and Administration Support
- ✓ Data Center Consolidation / Virtualization Support

Exadata POC Summary

Objective

Compare performance and scalability of Exadata to Sun/M5000 (10g)

Results

- ❑ Incremental **load time decreased by 60%** in Exadata compared to 10g
- ❑ Sample jobs **completed in 40% lesser time** on Exadata compared to 10g
 - Exadata **295X faster** on individual search queries
 - Query scaling test – **Avg. response time up to 50X faster on Exadata**

Recommendations

Positive results – further testing under *mixed workload* required for definitive decision

In production - by the numbers..

2012		Expected growth in 2013
Number of users	850	+ 950
Concurrent users	550	+ 250
Data volumes	4 TB	+ 4 TB
Daily reports	150	+ 150*
Daily batches	16,500	+ 4000*
Queries per day	48.5 million	+ 15 million*
IOPS	177,000 @ 0.7ms	NA
Cell offload efficiency	87%	NA
*Estimated		

Cycle time – Before and After (Exadata + other factors)

	Legacy (RAC + flat data model)	OGRDS (Exadata + KVP model)	Difference
Processing Updates (Records / Hour)			
Historical Compare	176,891	854,362	5 times faster
Bulk Score	7,691	21,156	3 times faster
Bulk Address Standardization	300	150,000	500 times faster
Report (in Minutes)			
Time for Report completion	5 – 240	1 – 14	5 to 17 times faster
Advance to Next Page / Record (Seconds)			
Historical Review	1.8	3.8	2 times slower
Visual Match	1.0	2.5	2.5 times slower
Manual Match	2.0	4.3	2 times slower

Results indicate user experience, **cannot be isolated for Exadata**

AWR supports Exadata 5x faster across the board

Lessons Learned

Technology


- Leave Flash Cache alone
- Dev/QA Systems use Exadata
- Set Resource Governor as a last resort
- HCC does not work well with frequently updated data
- Review indices on a case by case basis

Business

- Demanding model (KVP) and mixed workload = fun tuning exercises
- Upgrade to next Configuration Level vs Extension Packs
- Partition / Cluster your Exadata
- Design for Patching / Simple downtimes
- Leverage Oracle Advanced Customer Support

Answers Retail Edition – Exadata Summary

- ✓ Simplified Architecture
 - Full Rack replaces 2 IBM P770 servers and some processing on Netezza
- ✓ High availability and redundancy
- ✓ Capacity for future organic growth and new business
- ✓ Supports 24x7 ETL processing with no impact to online reporting



Thank you!

