



Oracle NoSQL Database Overview

Marie-Anne Neimat, VP Development

June 14, 2012


ORACLE®




Agenda

- **Big Data Overview**
- Oracle NoSQL Database
 - Architecture
 - Technical Overview
 - Benchmark Results
 - Usage scenarios

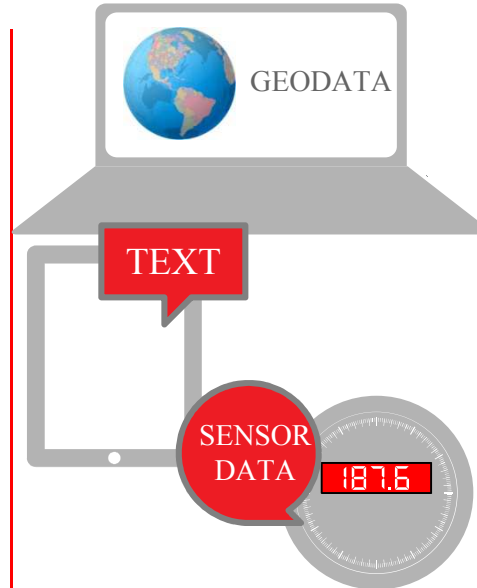
What is Big Data ?




VOLUME



VELOCITY




VARIETY



VALUE



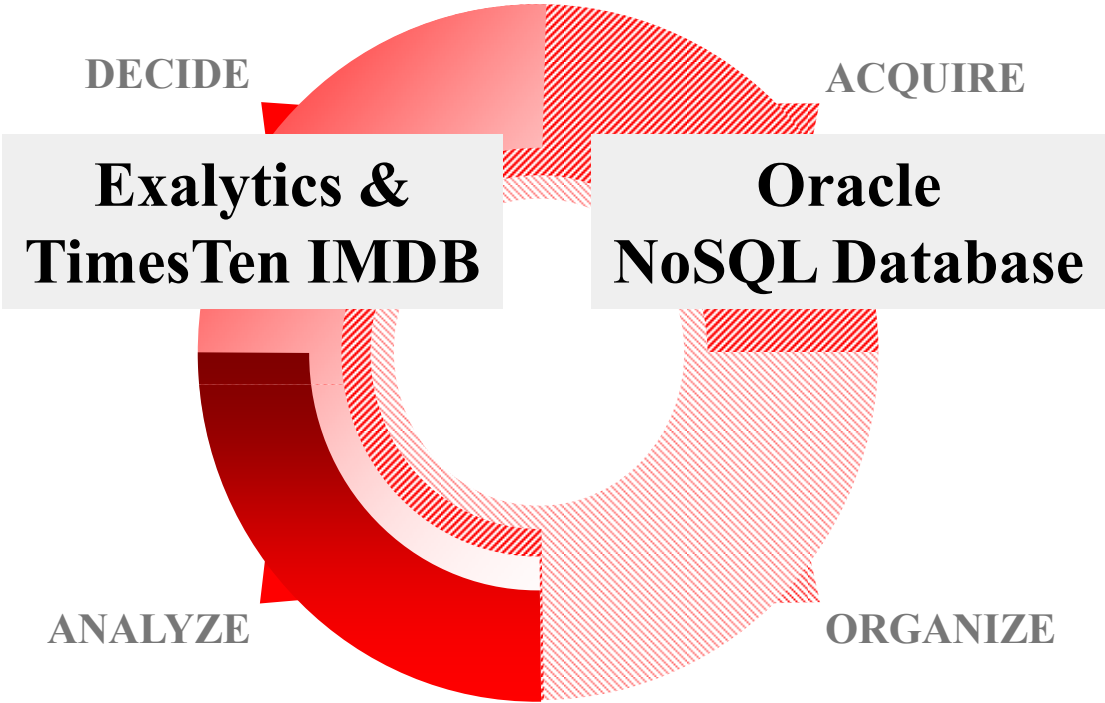
Big Data Lifecycle



**Make
Better
Decisions
Using
Big Data**



Big Data Lifecycle



**Make
Better
Decisions
Using
Big Data**



Big Data Use Cases

Today's Challenge	New Data	What's Possible
Healthcare Expensive office visits	Remote patient monitoring	Preventive care, reduced hospitalization
Manufacturing In-person support	Product sensors	Automated diagnosis, support
Location-Based Services Based on home zip code	Real time location data	Geo-advertising, traffic, local search
Utilities Complex Distribution Grid	Detailed consumption statistics	Increased availability, reduced cost, tiered metering plans
Retail One size fits all marketing	Social media	Sentiment analysis segmentation



Why is Big Data important?

US HEALTH CARE

Increase industry value per year by

\$300 B

US RETAIL

Increase net margin by

60+%

MANUFACTURING

Decrease dev., assembly costs by

-50%

GLOBAL PERSONAL LOCATION DATA

Increase service provider revenue by

\$100 B

EUROPE PUBLIC SECTOR ADMIN

Increase industry value per year by

€250 B

“In a big data world, a competitor that fails to sufficiently develop its capabilities will be left behind.”

Source: * McKinsey Global Institute: Big Data – The next frontier for innovation, competition and productivity (May 2011)

ORACLE®



Oracle NoSQL Database for realtime data management

Hadoop Distributed File System (HDFS)	Oracle NoSQL Database
File System	Database
Parallel scanning	Indexed storage
No inherent structure	Simple data structure
Bulk writes	High volume random reads and writes
Batch Oriented	Real-Time



Agenda

- Big Data Overview
- **Oracle NoSQL Database**
 - Architecture
 - Technical Overview
 - Benchmark Results



Oracle NoSQL Database Use Cases

Customer profile management

Social networks

Personalization

Click-through data processing

High-throughput event processing

Sensor & statistics data capture

Authentication & Content management

Mobile application backend infrastructure

Archiving



SIMPLE QUERIES

DYNAMIC SCHEMA

HIGH VOLUME DATA

ORACLE®

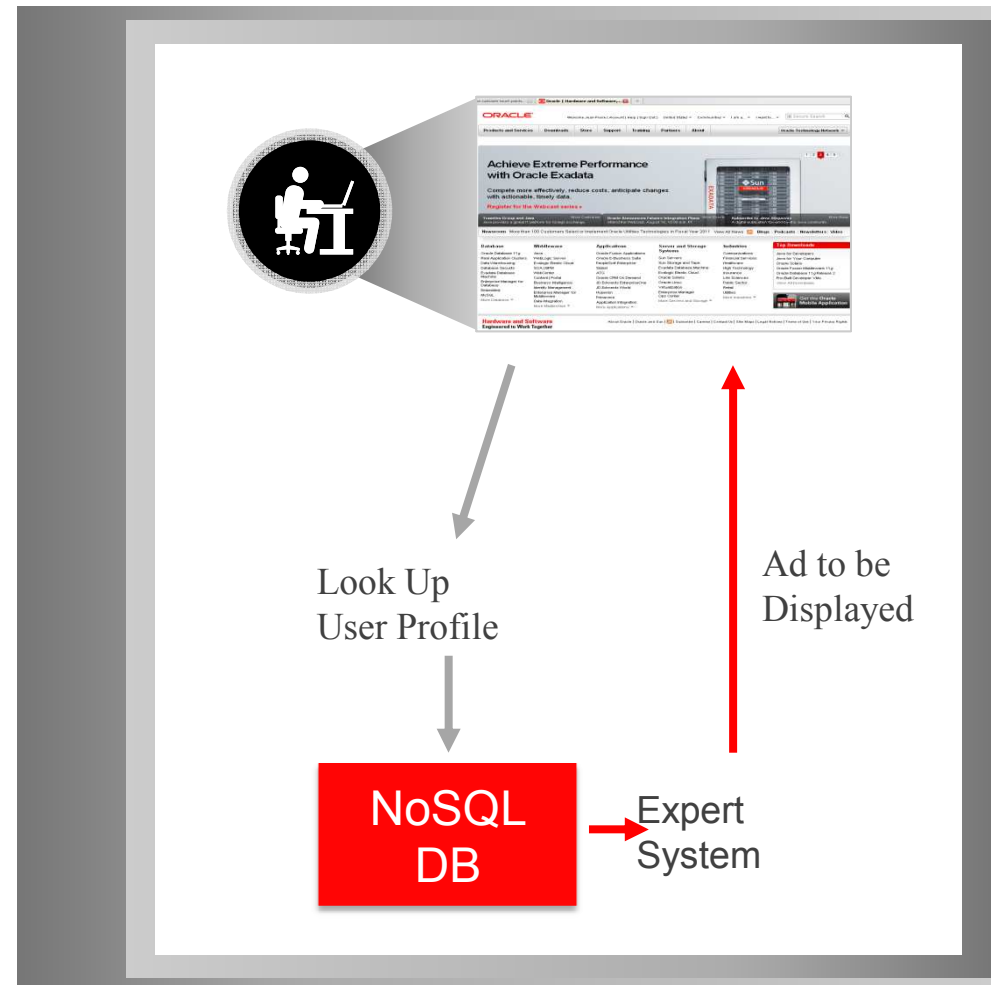
Example: On-line Ads and Content

Scenario

- User clicks on a web page
- Cookie used to look up user profile
- Store new profile if required
- User profile sent to expert system
- Appropriate ad placed on page

Important Attributes

- Fast response time and high throughput
- High availability and scalability
- Flexible data model



Goal: Optimize Internet Advertising to Increase Revenues

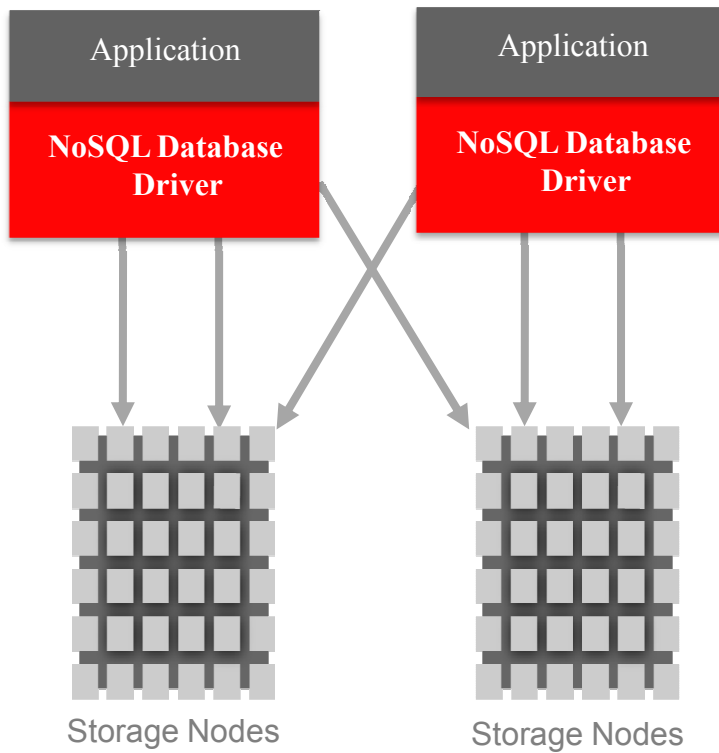


Customer-Driven Requirements

- Terabytes to petabytes of unstructured or semi-structured data
- No single point of failure
- Cost effective, distributed storage
- Scalable on commodity hardware
- Fast, predictable response to simple queries
- Fast, reliable transactions
- Simple administration, enterprise support
- **Commercial-grade** NoSQL solution
 - Real 24x7 support
 - Real database expertise
 - Large vendor & dedicated resources building & testing the software

Oracle NoSQL Database

A Distributed, Scalable Key-Value Database



Simple Data Model

Small, distributed footprint

Highly scalable, available

Transparent load balancing

Integrates with Oracle Stack

ORACLE®

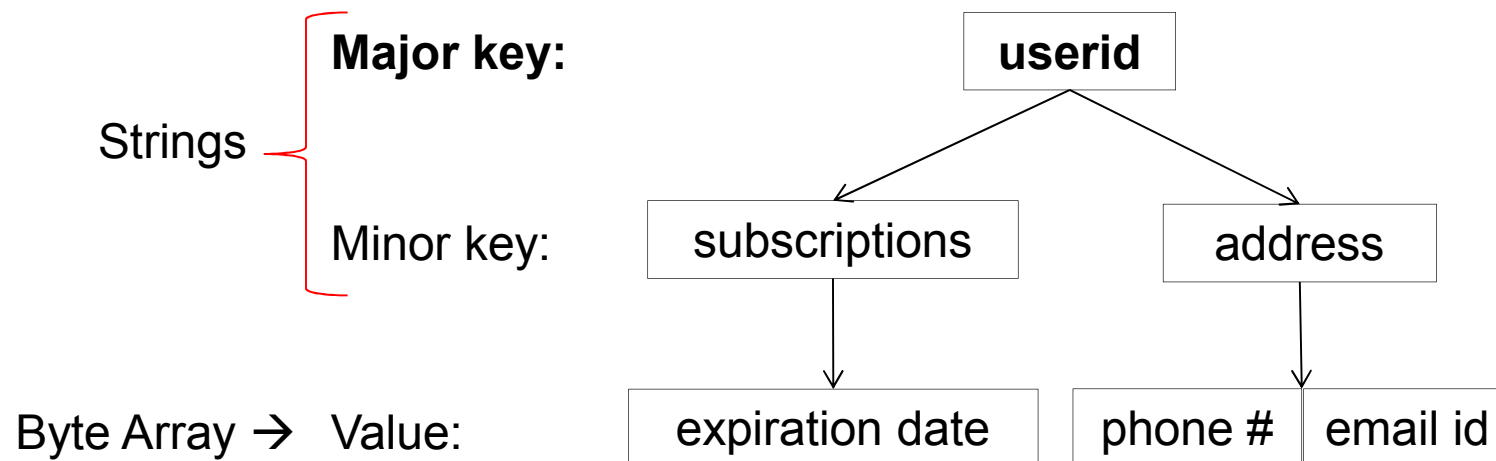


Architecture Summary

- Scalability
 - Dynamic data partitioning and distribution
 - Optimized data access via intelligent driver
- High availability
 - One or more replicas
 - Resilient to partition master failures
 - No single point of failure
 - Disaster recovery through location of replicas
- Transparent load balancing
 - Reads from master or replicas
 - Driver is network topology & latency aware
- Optimized for high throughput and low latency

Simple Data Model

- Simple data model – key-value pair (major+minor-key paradigm)
- Simple operations – read/insert/update/delete, RMW support
- Scope of transaction – records within a major key, single API call
- Unordered scan of all data (non-transactional)





Simple Data Model

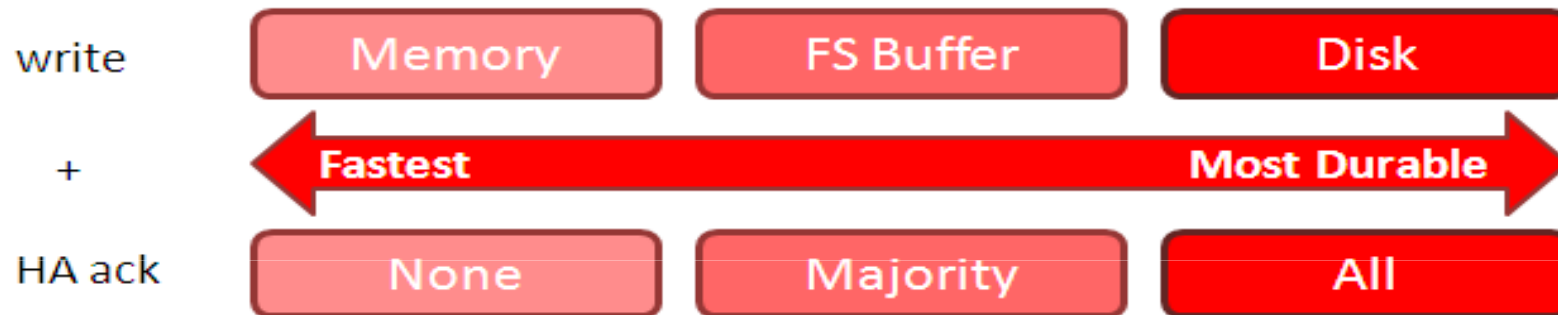
ACID Transactions

- ACID transactions by default
- Transaction Scope
 - Single API call
 - All records must have the same major key
 - Support for multiple operations within a transaction
- Can be relaxed for increased performance on a per-operation basis

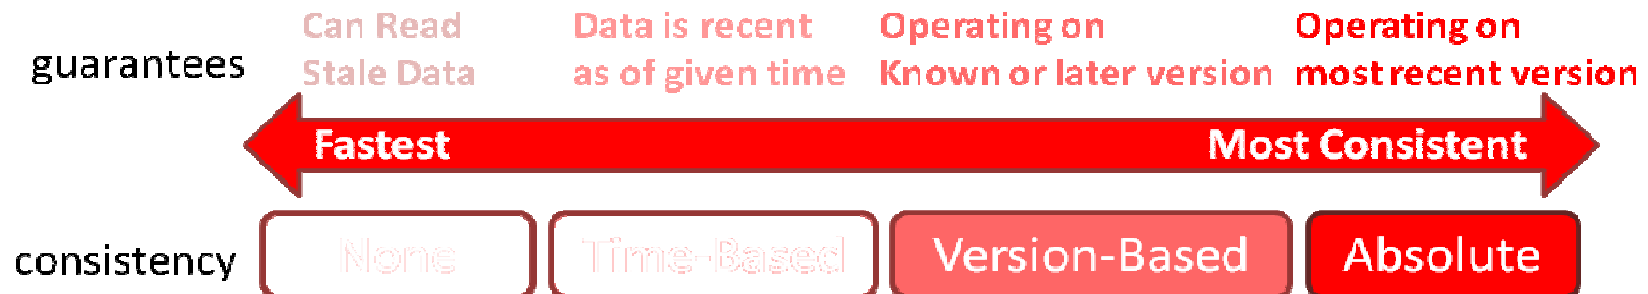
Simple Data Model

ACID Transactions – Configurability

- Configurable **Durability** Policy

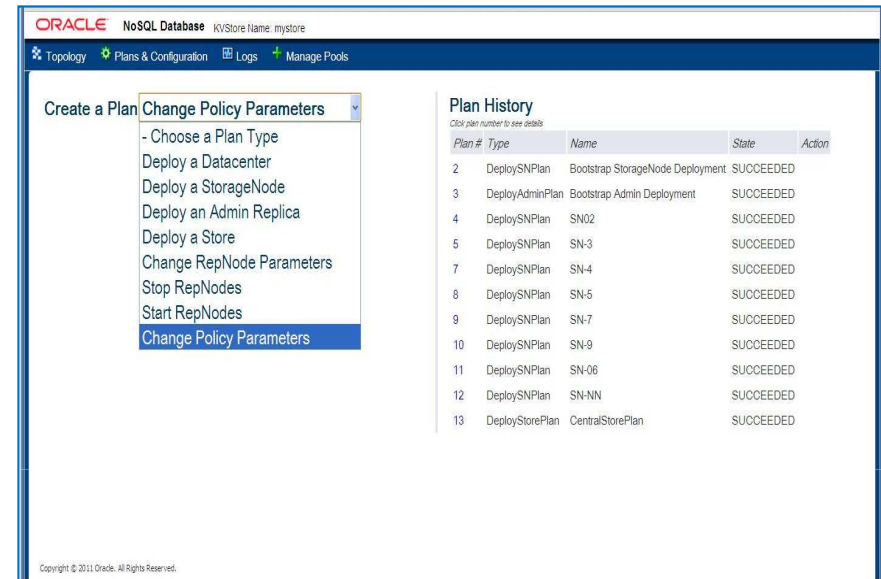


- Configurable **Consistency** Policy



Simple Administration

- Web-based console and CLI commands
- Manages and Monitors
 - Topology
 - Configuration changes
 - Load: Number of operations, data size
 - Performance: Latency, throughput. Min, max, average, trailing, ...
 - Events: Failover, recovery, load distribution
 - Alerts: Failure, poor performance, ...



The screenshot displays the Oracle NoSQL Database Administration Console interface. The top navigation bar includes 'ORACLE NoSQL Database' and 'KVStore Name: mystore'. Below the navigation bar, there are tabs for 'Topology', 'Plans & Configuration', 'Logs', and 'Manage Pools'. The main content area is divided into two sections. On the left, under 'Create a Plan', a dropdown menu is open, showing options: '- Choose a Plan Type', 'Deploy a Datacenter', 'Deploy a StorageNode', 'Deploy an Admin Replica', 'Deploy a Store', 'Change RepNode Parameters', 'Stop RepNodes', 'Start RepNodes', and 'Change Policy Parameters' (which is highlighted). On the right, the 'Plan History' section shows a table with columns 'Plan #', 'Type', 'Name', 'State', and 'Action'. The table contains 13 rows of data, all with a 'SUCCEEDED' state.

Plan #	Type	Name	State	Action
2	DeploySNPlan	Bootstrap StorageNode Deployment	SUCCEEDED	
3	DeployAdminPlan	Bootstrap Admin Deployment	SUCCEEDED	
4	DeploySNPlan	SN02	SUCCEEDED	
5	DeploySNPlan	SN-3	SUCCEEDED	
7	DeploySNPlan	SN-4	SUCCEEDED	
8	DeploySNPlan	SN-5	SUCCEEDED	
9	DeploySNPlan	SN-7	SUCCEEDED	
10	DeploySNPlan	SN-9	SUCCEEDED	
11	DeploySNPlan	SN-06	SUCCEEDED	
12	DeploySNPlan	SN-NN	SUCCEEDED	
13	DeployStorePlan	CentralStorePlan	SUCCEEDED	

Oracle NoSQL Database Differentiation

Integrates seamlessly with Oracle Stack (ODI, CEP, OLH)

Commercial Grade Software and Support

- **General Purpose**
- **Reliable** – Based on proven Berkeley DB JE HA
- **Easy** to Install & Configure

Scalability and Availability

- **Intelligent** Oracle NoSQL DB Driver
 - Evenly distributes Data
 - Sends operation to fastest node
 - Bounded network hops for all operations
- **Automatic** replication and failover

Simple Data Model

- **Simple** Major + Minor key and Value data structure
- **ACID** transactions
- **Configurable** consistency and durability

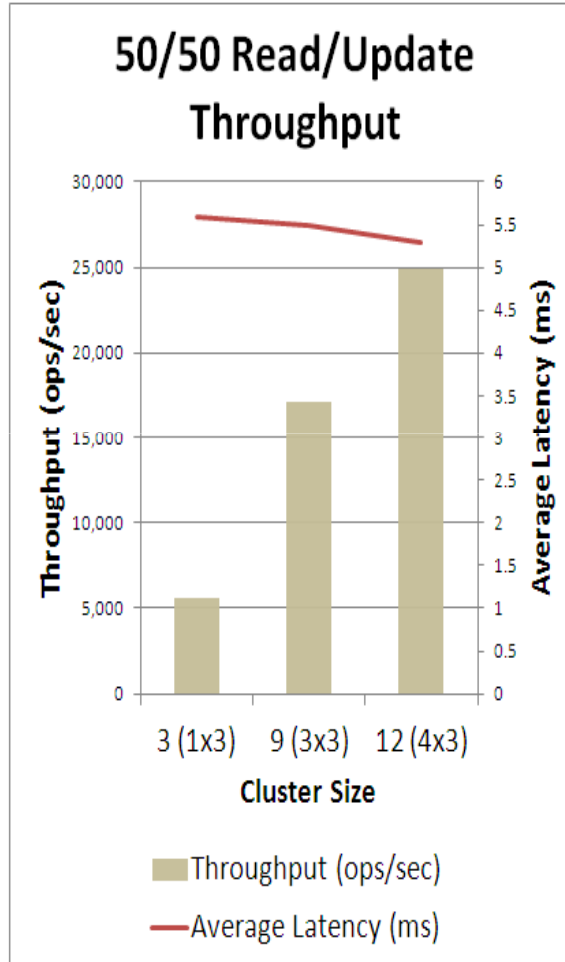

Simple Administration

- **Web-based** Console and **CLI** commands
- **Manages and Monitors:**
 - Topology
 - Load
 - Performance
 - Events
 - Alerts

ORACLE®

Benchmarking

Results @ Cisco



- 1.6 billion records
- 94K insert/sec
- 25K read/update/sec
- Low latency
- Linear scalability



Oracle NoSQL Database

- ✓ Easy to use, easy to manage
- ✓ Scalable, Available, Predictable Latency
- ✓ A NoSQL Database from a vendor you trust



Hardware and Software

ORACLE®

Engineered to Work Together

ORACLE®