


ORACLE®

OTN Developer Day: Oracle Big Data



ORACLE®

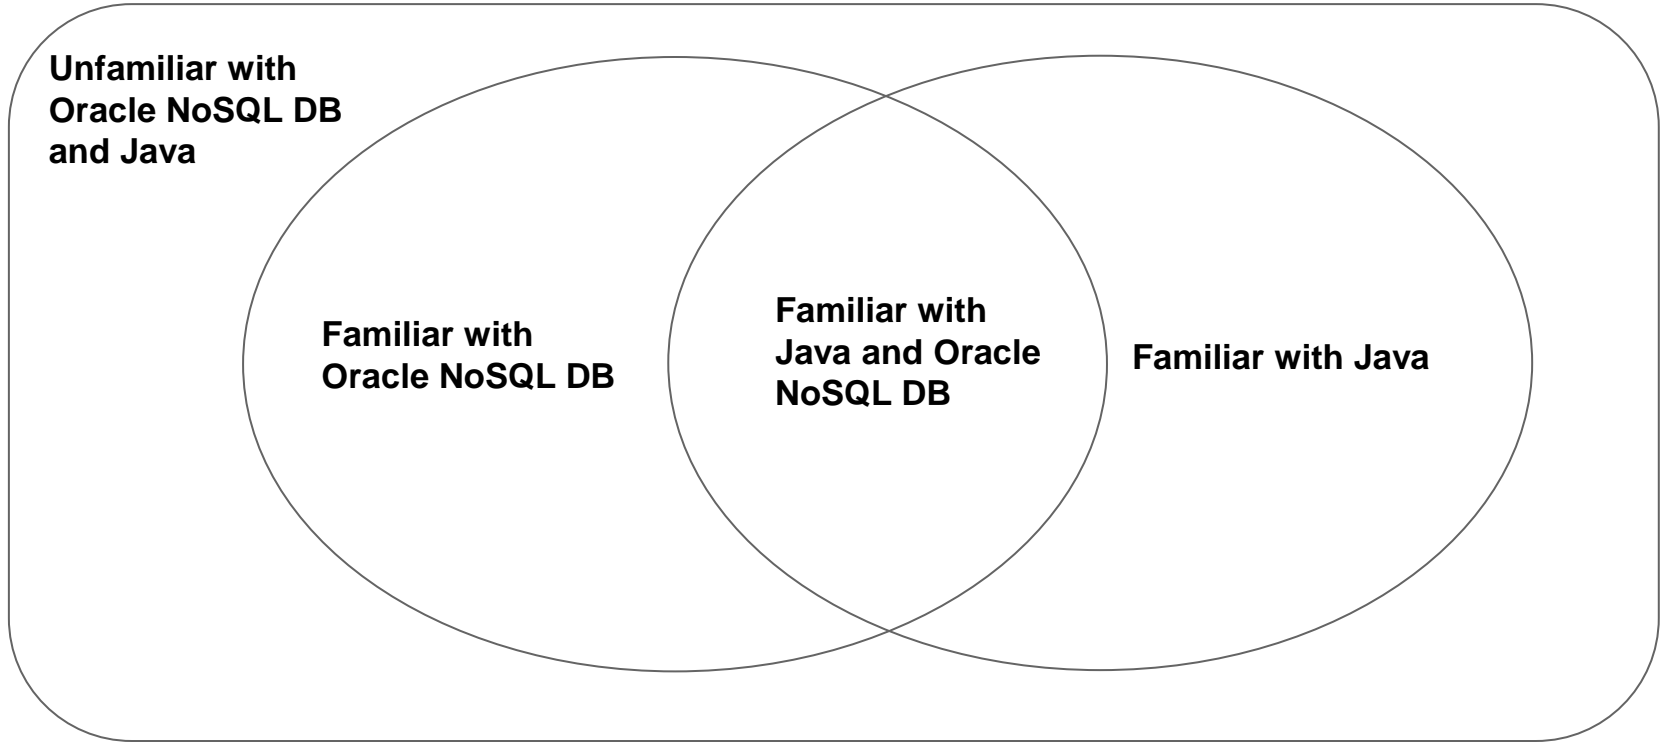
Data Management in Oracle NoSQL Database

Anuj Sahni, Principal Product Manager

Agenda

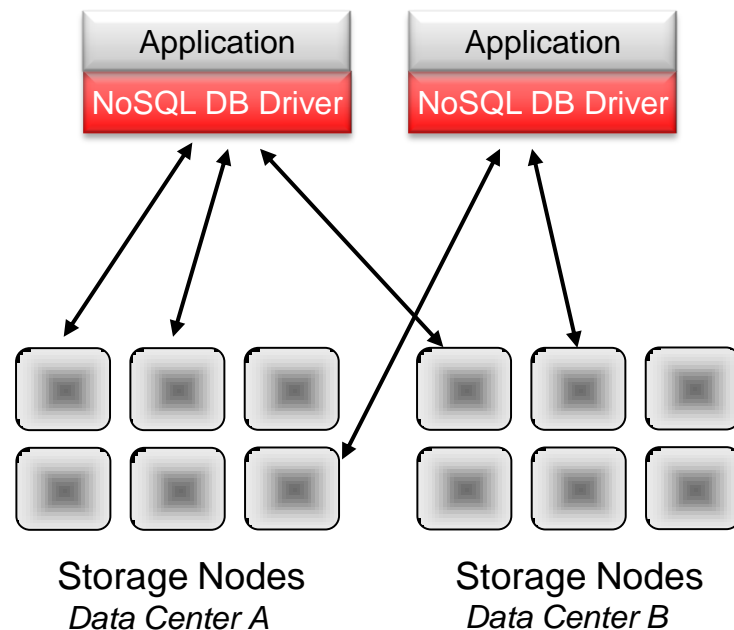
- Oracle NoSQL Database
 - Overview
 - Architecture
 - Benchmark
- What's new ?
- Hands-on Lab

Poll: Who is familiar with Oracle NoSQL DB and/or Java ?

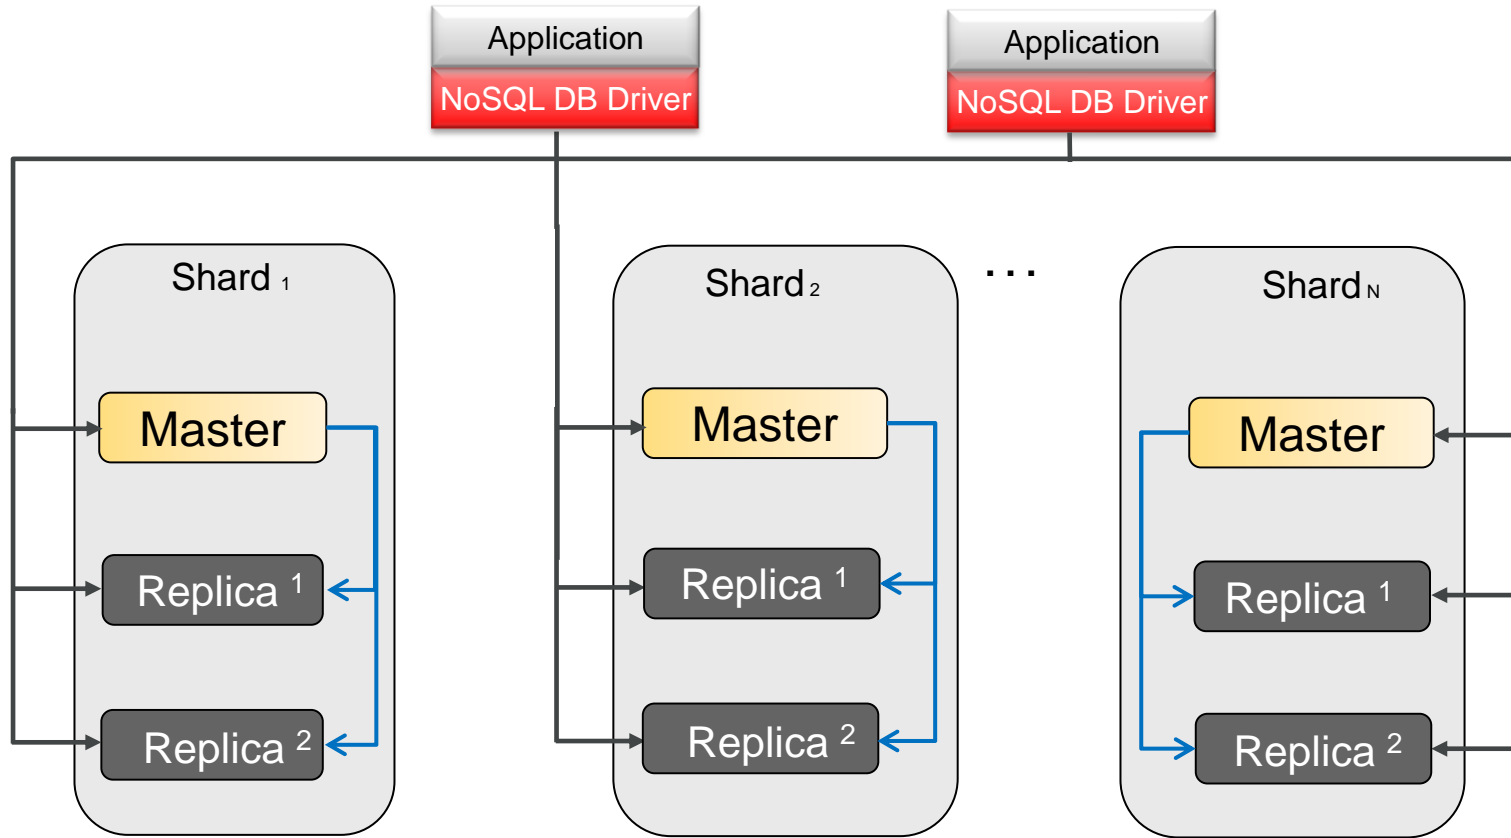


Oracle NoSQL – A Distributed Key-Value Database

- **Simple Data Model**
 - Key-value pair with major+minor-key paradigm
 - CRUD operations plus range scans
- **Scalability**
 - Dynamic data partitioning and distribution
 - Optimized data access via intelligent driver
- **High availability**
 - Resilient to partition failures
 - No single point of failure
- **Transparent load balancing**
 - Reads from master or replicas
 - Driver is network topology & latency aware
- **Elastic Expansion**
 - Online addition/removal of storage nodes and automatic data redistribution



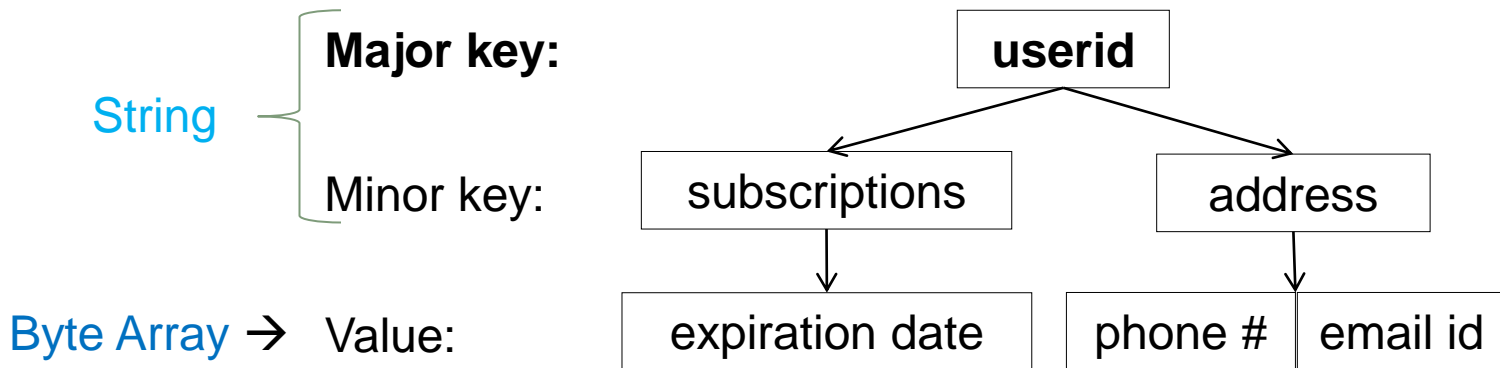
Architecture – The Application's Perspective



Oracle NoSQL Database

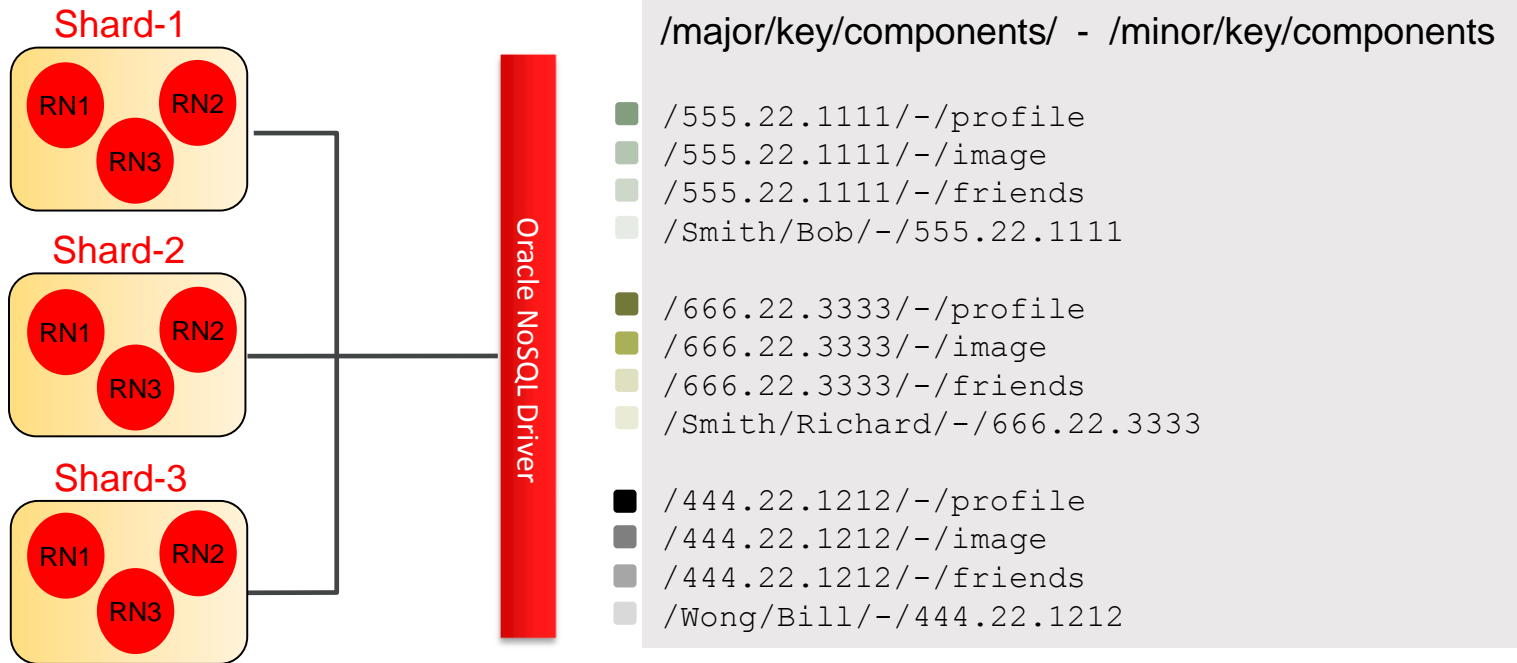
Simple Data Model

- 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

Major-Minor Key Paradigm



Oracle NoSQL Database Use Cases

Use Cases

High-throughput event processing

Customer profile management

Click-through data processing

Sensor & statistics data capture

Social networks

Personalization

Mobile application backend infrastructure

Authentication & Content management

Archiving

SIMPLE QUERIES

DYNAMIC SCHEMA

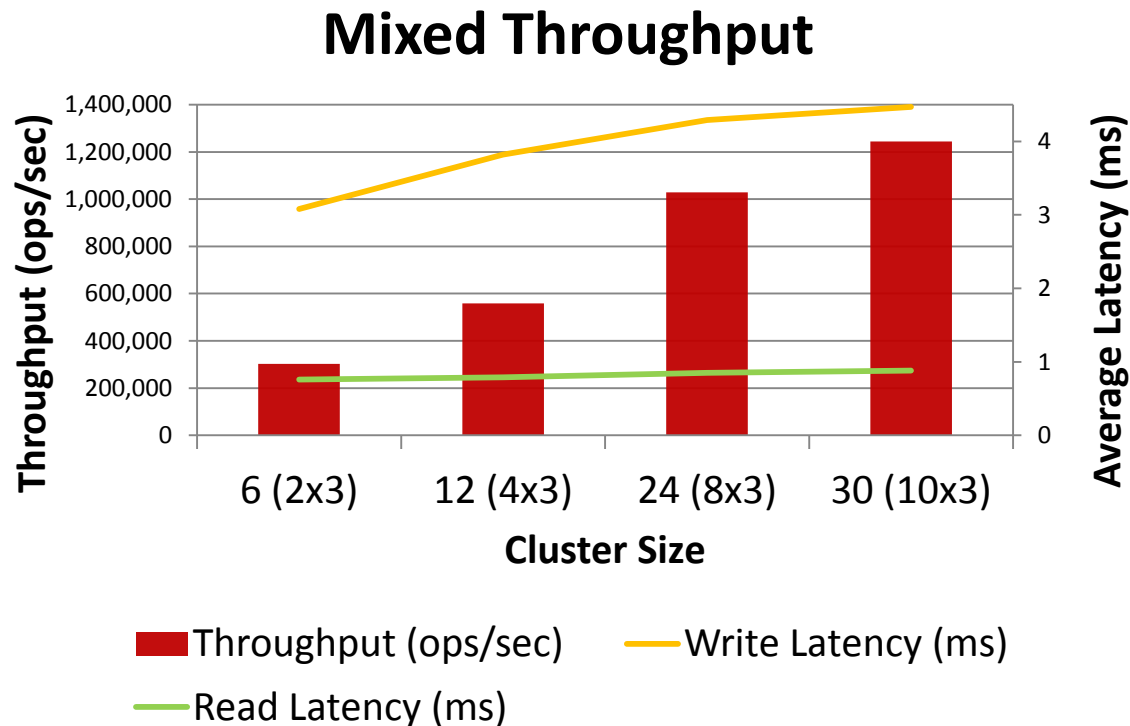
HIGH VOLUME

REAL TIME DATA ACCESS

“Last mile” problems

Benchmark Results

- 1.25M ops/sec
- 2 billion records
- 2 TB of data
- 95% read, 5% update
- Low latency
- High Scalability

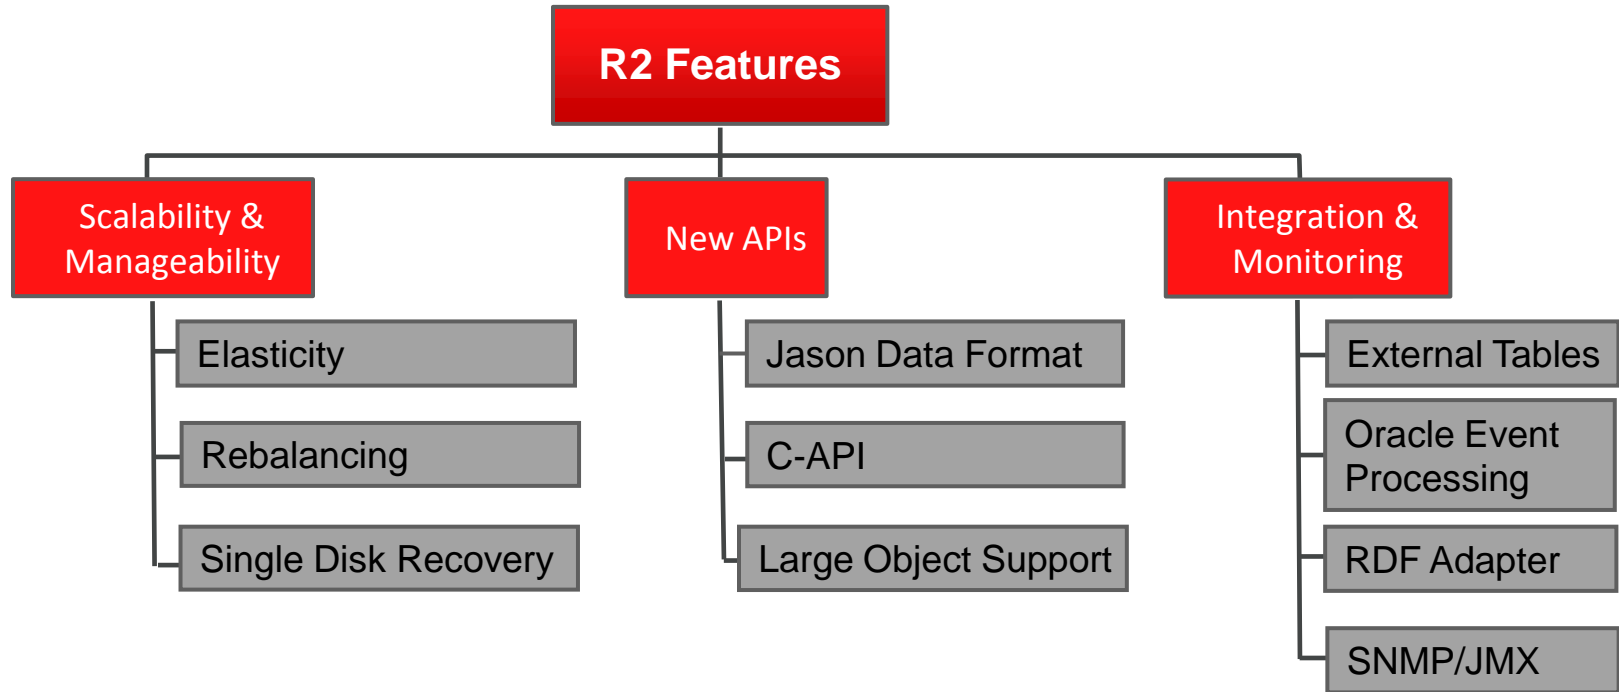


Agenda

- Oracle NoSQL Database
 - Overview
 - Architecture
 - Benchmark
- What's new ?
- Hands-on Lab

What's New?

Release 2 Feature Summary

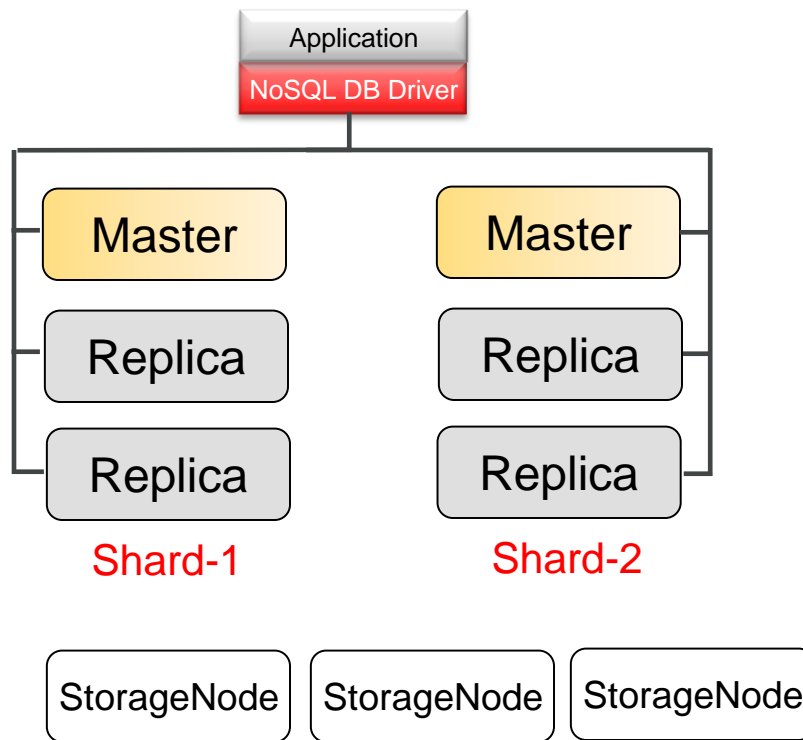


Elasticity

On-Demand Cluster Expansion

On Demand

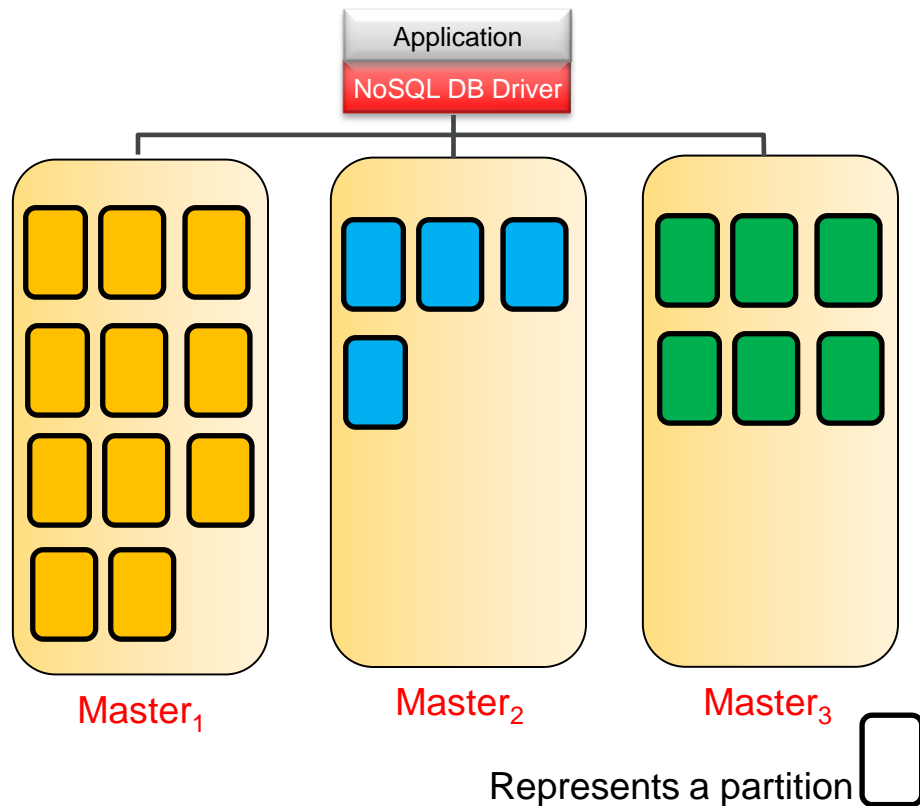
- Increase Data Capacity
 - Add more storage nodes
 - New shards automatically created
- Increase Data Throughput
 - More shards = better write throughput
 - More replicas/shard = better read throughput



Rebalance an Unbalanced Store

Improve Performance

- Replication nodes move from over-utilized to under-utilized storage nodes
- Number of shards and replication factor remain unchanged



JSON Data Format

Avro based Serialization/Deserialization

- Why Avro?
 - Compact, highly efficient serialization
 - Synergy with Hadoop
 - Multiple binding options (JSON, Generic, POJO)
- Schema
 - DDL allows schema creation through Avro JSON definition
 - Supports serialization from/to JSON strings
- Schema evolution
 - Easy to use mechanism for schema evolution
 - Schema versions can be opaque to readers

Support for Large Objects

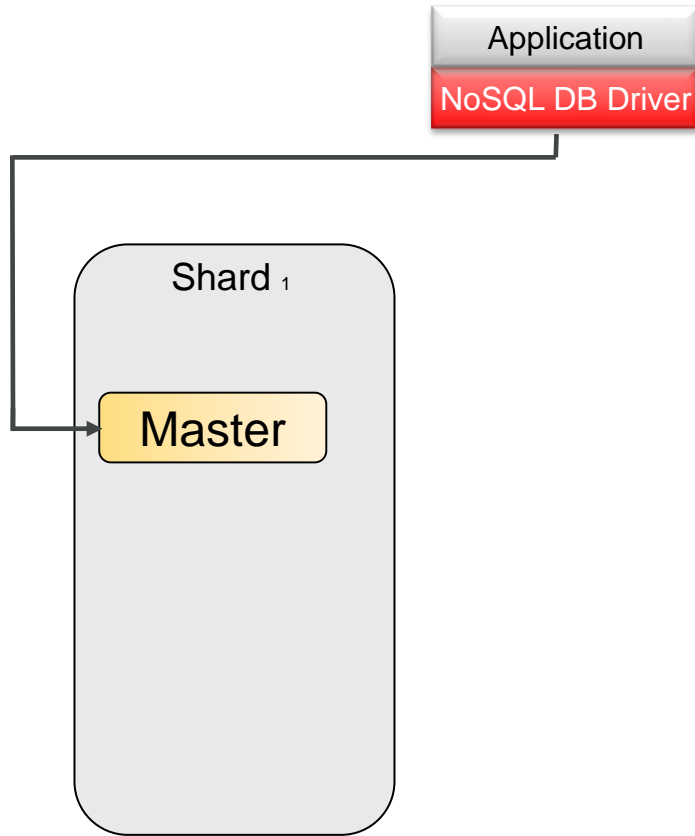
- Efficient storage and retrieval of large objects
- Client side streaming interface for low memory consumption
- Server side splitting and distribution of object chunks across nodes for better read/write latency

Integration with Oracle Products

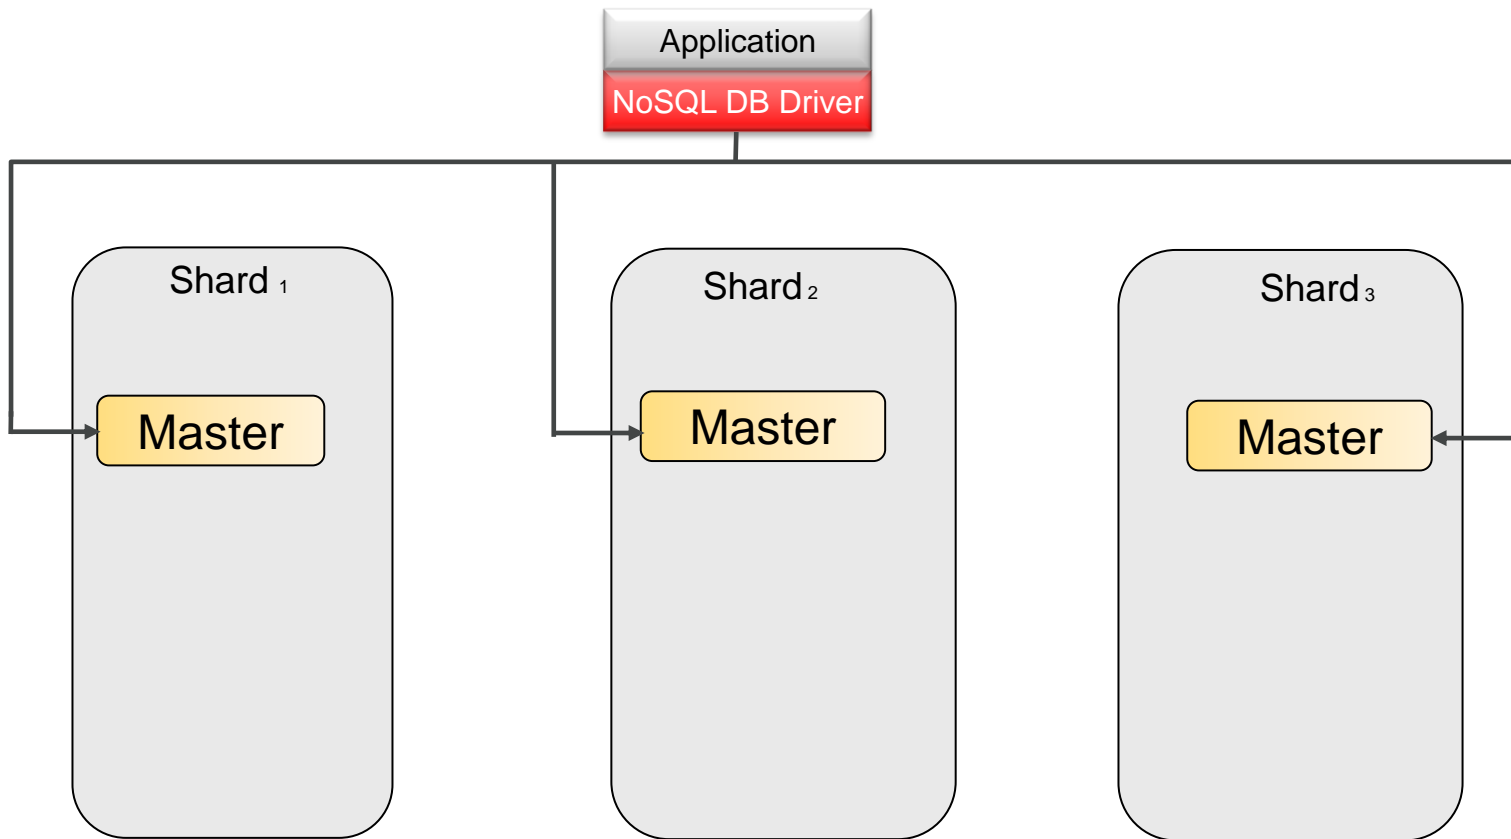
- Database External Tables
 - Access NoSQL data directly from Oracle
 - Available in the Enterprise Edition
- Oracle Event Processing (OEP)
 - NoSQL cartridge for Oracle Event Processing
 - Java serialization utilized for values
- Oracle Semantic Graph
 - RDF Jena adapter

Hands on Lab: Elastic Expansion and Rebalancing

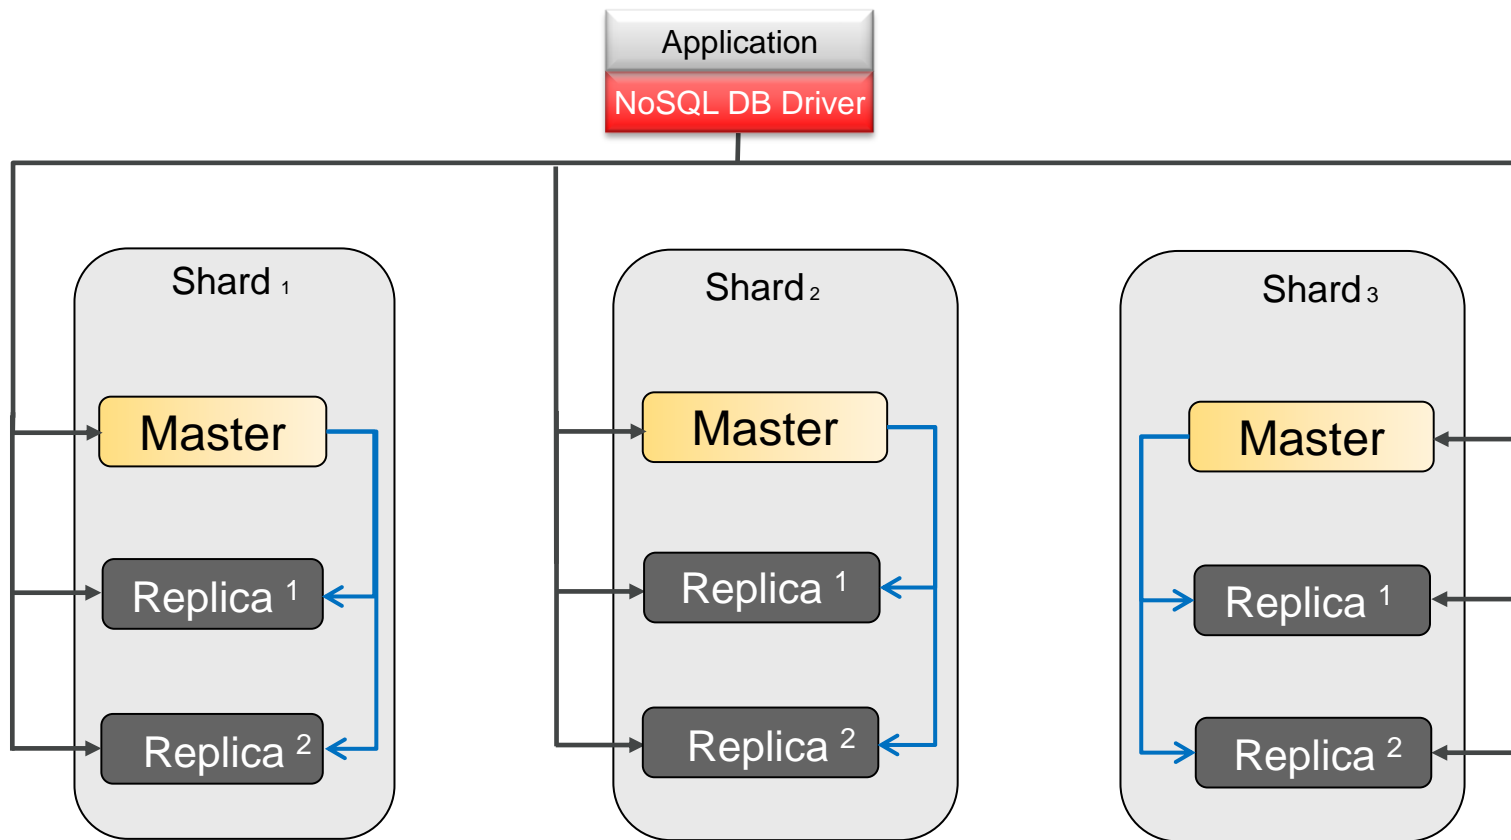
Elastic Expansion and Rebalancing (1x1)



Elastic Expansion and Rebalancing (3x1)



Elastic Expansion and Rebalancing (3x3)



Hands-on Lab

Key takeaways...

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

Quick Survey

- Oracle NoSQL Database Release 2.0 Feedback
<https://www.surveymonkey.com/s/NoSQLDB-2>
- Oracle NoSQL Database Feature Survey
<https://www.surveymonkey.com/s/NoSQLDB-Futures>

APPENDIX

Oracle NoSQL DB Resources

Support

- Support via OTN forums and Oracle Support process
- OTN Forum:
 - [Forum Home](#) » [Big Data](#) » [NoSQL Database](#)
 - forums.oracle.com/forums/forum.jspa?forumID=1388
- Oracle.com:
 - www.oracle.com/us/products/database/nosql/overview/index.html
- OTN (including documentation and download):
 - www.oracle.com/technetwork/products/nosqldb/overview/index.html

Oracle NoSQL DB Resources

Documentation

- On OTN and in download
 - docs.oracle.com/cd/NOSQL/html/index.html
- Getting Started Guides
- Programmatic API
- Installation & Release Notes
- FAQ

Oracle NoSQL DB Resources

Internal

- bigdata.us.oracle.com
- nosqldatabase.us.oracle.com or “Oracle NoSQL Database” Beehive workspace
- NoSQL_DB_help_ww@oracle.com
- bigdata_ww@oracle.com
- Prod Mgmt: dave.segleau@oracle.com, anuj.sahni@oracle.com
- Dev Mgmt: ashok.joshi@oracle.com, dave.rubin@oracle.com
- Prod Marketing: peter.jeffcock@oracle.com

What we have been testing

- YCSB-based QA/benchmarking
 - Key \approx 10 bytes, Data = 1108 bytes
- Configurations of 6-30 nodes
 - Typical Replication Factor of 3 (master + 2 replicas)
 - 200m records per shard, 2 billion records in total
 - 2 RNs per SN
 - Used SSDs, and two of them per host
- Minimal I/O overhead
 - B+Tree fits in memory => one I/O per record read
 - Writes are buffered + log structured storage system == fast write throughput