



Rationalizing Oracle, MySQL & NoSQL

NoCOUG – Nov 2017

Ritesh Chhajer, Core Data Platforms, PayPal

About Speaker, Ritesh Chhajer

- Director, Core Data Platforms at PayPal in San Jose, CA
- 17 years of industry experience (CNBC TV18, Yahoo, PayPal)
- Loves solving data challenges at scale
- Blog: <http://compuritesh.tumblr.com/>

Stewards of Data

- Traditional Definition of Data
: body of facts; information
- Modern Interpretation of Data
: new oil; new currency

“If Data is the new currency then databases are essentially the place where company’s substantial wealth is stored”

Choosing the right data store

Quality Attributes

- Business use case: SOR, Transient, Persistent, Durability
- Functionality: ACID, Secondary Index, Concurrency, Serialization, Consistency, DLM
- Availability: Node failure, Data loss, Read from replicas, Auto failover, Load balancing, DR strategy, A/A
- Performance: Latency, Throughput, Optimization knobs
- Scalability: Vertically Scale up, Horizontally Scale out, Read scale, Write scale, Sharding
- Query Expressiveness: SQL Support
- Operations: Vendor support, Product maturity, Monitoring Capabilities, Online production changes
- Security: Authentication, Authorization, Encryption, Auditing
- TCO: Licensing cost, Cost of storage, Development cost

RDBMS vs NoSQL

RDBMS	NoSQL
Fully ACID Compliant	Limited ACID
Trends towards availability and consistency	Trends towards partition tolerance
Persistent data store, Guaranteed Consistency	Heavily Sharded, Eventually Consistent
Good fit for transactional databases	Good fit for distributed databases
Data stored in rows and columns	K-V pairs, documents, graphs, wide column stores
Recommended for system of record	Recommended for system of engagement
Data access through SQL	Data access through Key
Limited Scalability, Manual Sharding	Highly scalable with auto sharding
Predefined schema for structured data	Dynamic schema for structured, unstructured and semi-structured data
Each business attribute is a column	Each business attribute is a key

RDBMS vs NoSQL

RDBMS	NoSQL
Database → Tables → Rows	Database → Collection → Documents
CREATE TABLE EMP (name varchar2(100), age int(2), sex char(6));	db.emp.save ({"name":"Mike","age":39,"sex":"male"});
INSERT INTO EMP (name,age,sex) values ('Mike', '39', 'male');	
Vertically Scale up	Horizontally Scale Out
Usually one data copy per DB	Multiple data copies due to replication factor
Well Instrumented. Relies and thrives on vendor support.	Relies on DEV expertise and community support
Mix of open source (MySQL, Postgress) and commercial (Oracle, MS-SQL)	Mostly all Open Source (MongoDB, Cassandra, Hbase)
TCO High on CAPEX	TCO High on OPEX

MySQL

Benefits

- High on Developer Productivity, Easy and quick to setup
- Small to Medium footprint ACID compliant relational and persistent metadata
- Free, Fast, Light weight, Open Source
- Independent of infrastructure expertise
- Active Open Source Development Community
- Cloud friendly
- Popular with Web Applications
- Amenable to Automation
- Flexible and easy to setup replication

Case Study: MySQL to Oracle Migration

MySQL Challenges

- DDL blocking reads, downtime releases on a daily basis ^{*1}
- Suboptimal mutexes and locks, rogue SQL bringing down the application ^{*2}
- Single threaded replication, unable to support heavy writes ^{*2}
- Limited scalability on single node, unable to handle large data size
- 3 weeks to go live, not enough time to design and implement sharding

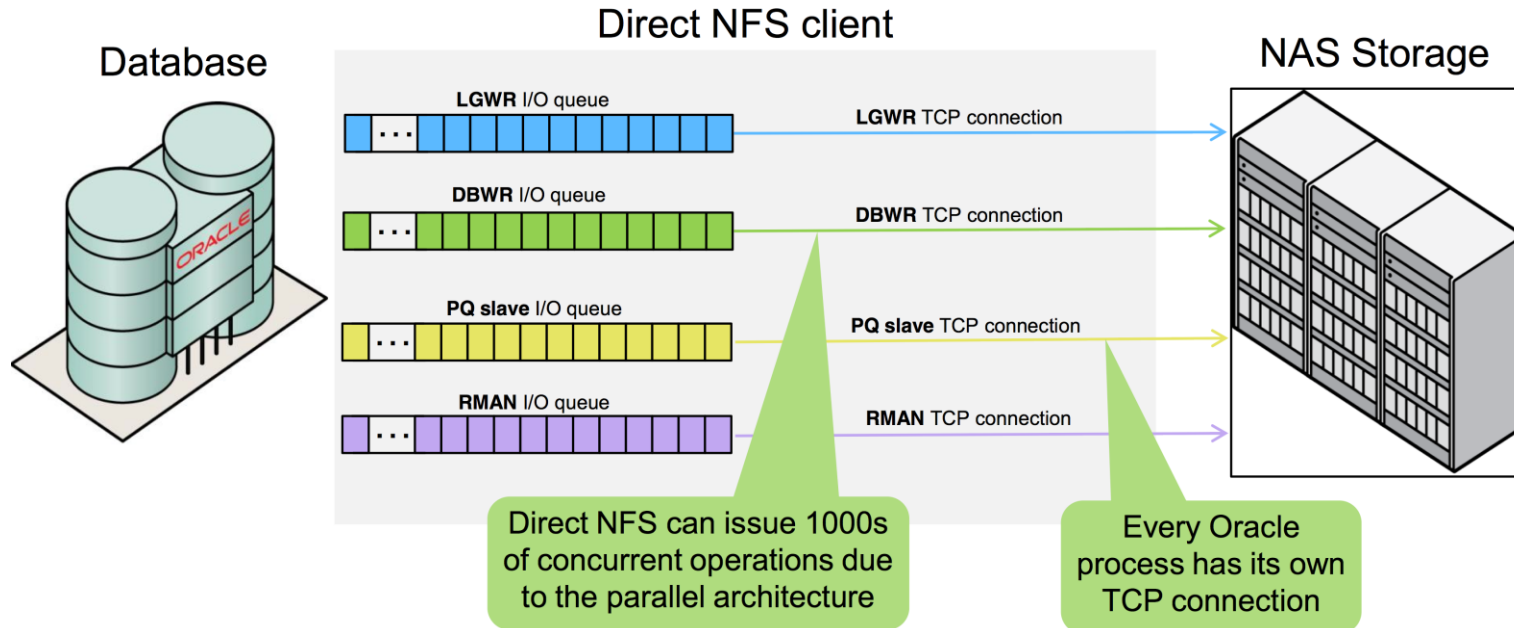
Oracle Benefits

- Readers don't block writers and writers don't block readers.
- HA with data consistency, no replication issues
- Write Scalability, Distributed reads for load balancing
- Well instrumented with detailed data dictionary for performance monitoring and tuning
- Scalable storage, Improved IO throughput through dNFS

*1 – Pt-online-schema was not integrated with internal CI/CD process for executing online changes

*2 – MySQL 5.5 was in use. Backups running with FTWRL

Oracle dNFS



Benefits

- Provides multi-pathing to NFS
- Automatic Failover and load balancing
- Free (Not a licensed feature)
- Support for parallel IO
- 2x gain in read throughput by addition of NIC card

Oracle

Limitations (Oracle RAC)

- Hardware setup is rack dependent.
- All RAC nodes to be setup within same cabinet, connected to same switch
- Infrastructure expertise dependency
- Custom configurations need special cabling diagrams/instructions to data center teams.
- Annual Licensing Cost

Oracle vs MySQL

Feature	Oracle	MySQL
ACID	✓	✓
Standard SQL Support	✓	✓
Open Source	✗	✓
Easy to install and manage	✗	✓
Independent of infrastructure expertise	✗	✓
Popular with Web Apps	✗	✓
VLDB	✓	✗
Reporting Analytics	✓	✗
Enterprise Data warehouse	✓	✗
HA with data consistency	✓	✗
Resource Manager	✓	✗
Instrumentation	✓	✗
Low Cost	✗	✓

Summary

- Polyglot Persistence: RDBMS vs NoSQL is not a zero sum game
- No one size fits all. Every organization may have a different use case.
- Don't blame the technology for a failed project and vice versa
- Co-existence and healthy competition
- Existentialism perspective in Database World

Passing Thoughts



**Did Someone
Move Your
Cheese?**



Comments & Questions

<https://www.linkedin.com/in/riteshchhajer/>