

NewSQL Cloud Database / DaaS



CMS: Elastic, Scalable, Highly-Available



# ScaleDB

“ScaleDB looks to be the MySQL version of Oracle RAC”

--- Brian Aker, Drizzle Architect, Sun Microsystems

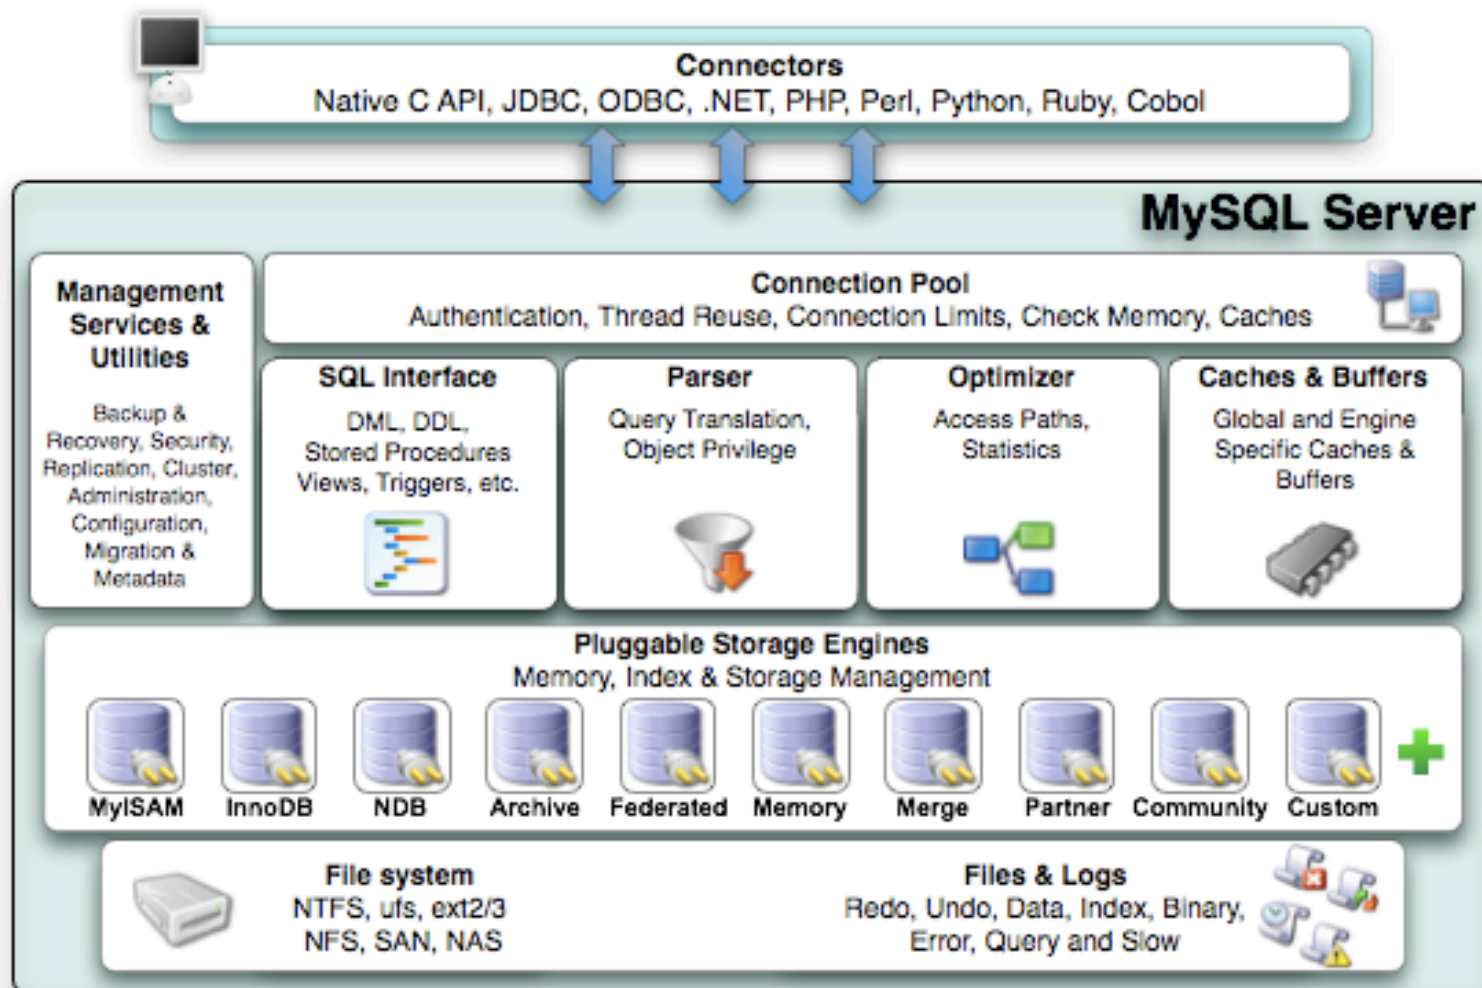
Xin Jia

Mar 05, 2012

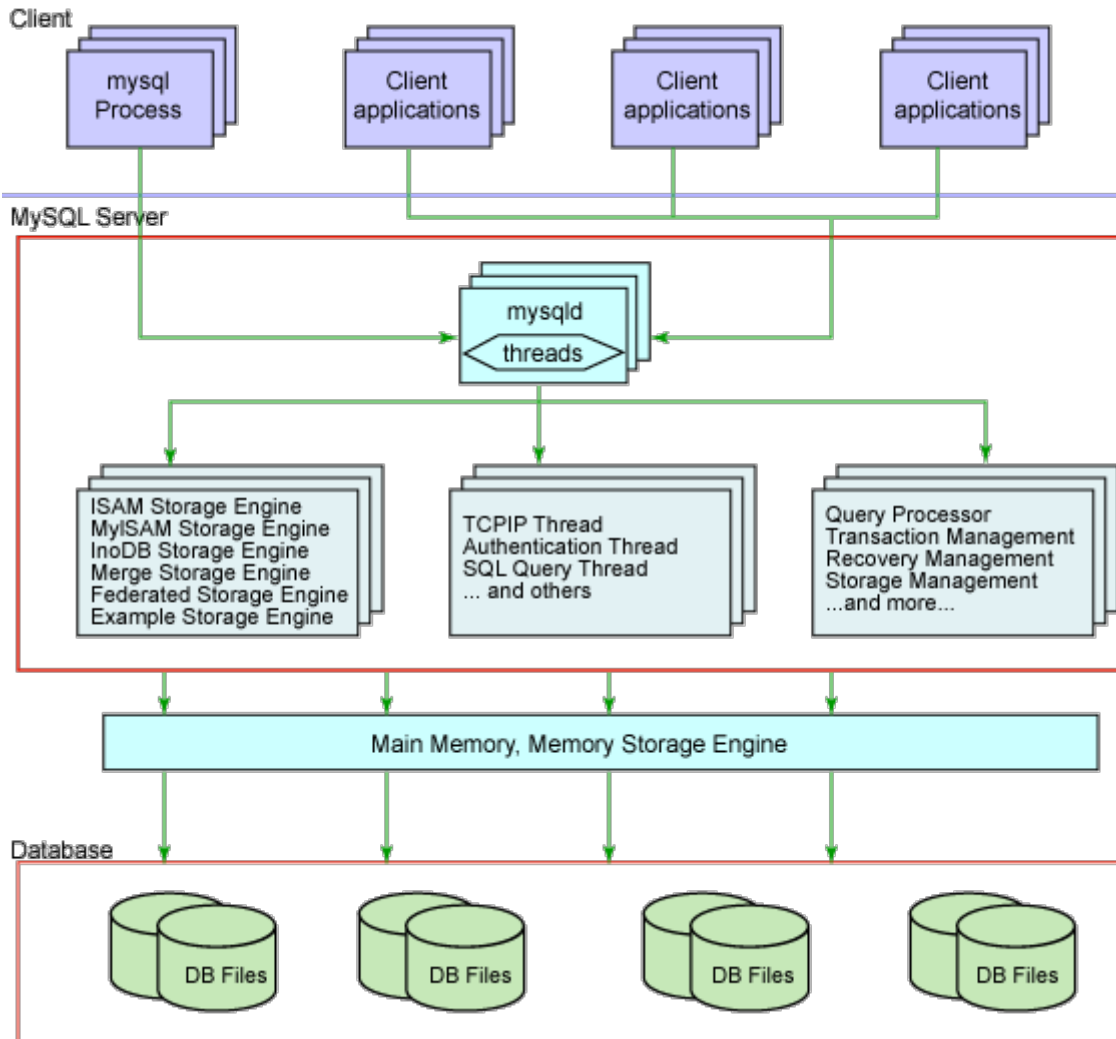
# What's ScaleDB

- ScaleDB is a **pluggable storage engine** for **MySQL**. It turns your MySQL application into an enterprise-class, **highly-available, clustered database** that **scales elastically** in a public cloud, private cloud, or on premise.
- ScaleDB is a NewSQL pioneer, delivering the advantages of both SQL and NoSQL, while targeting the **cloud**.

# MySQL Architecture(1)



# MySQL Architecture(2)



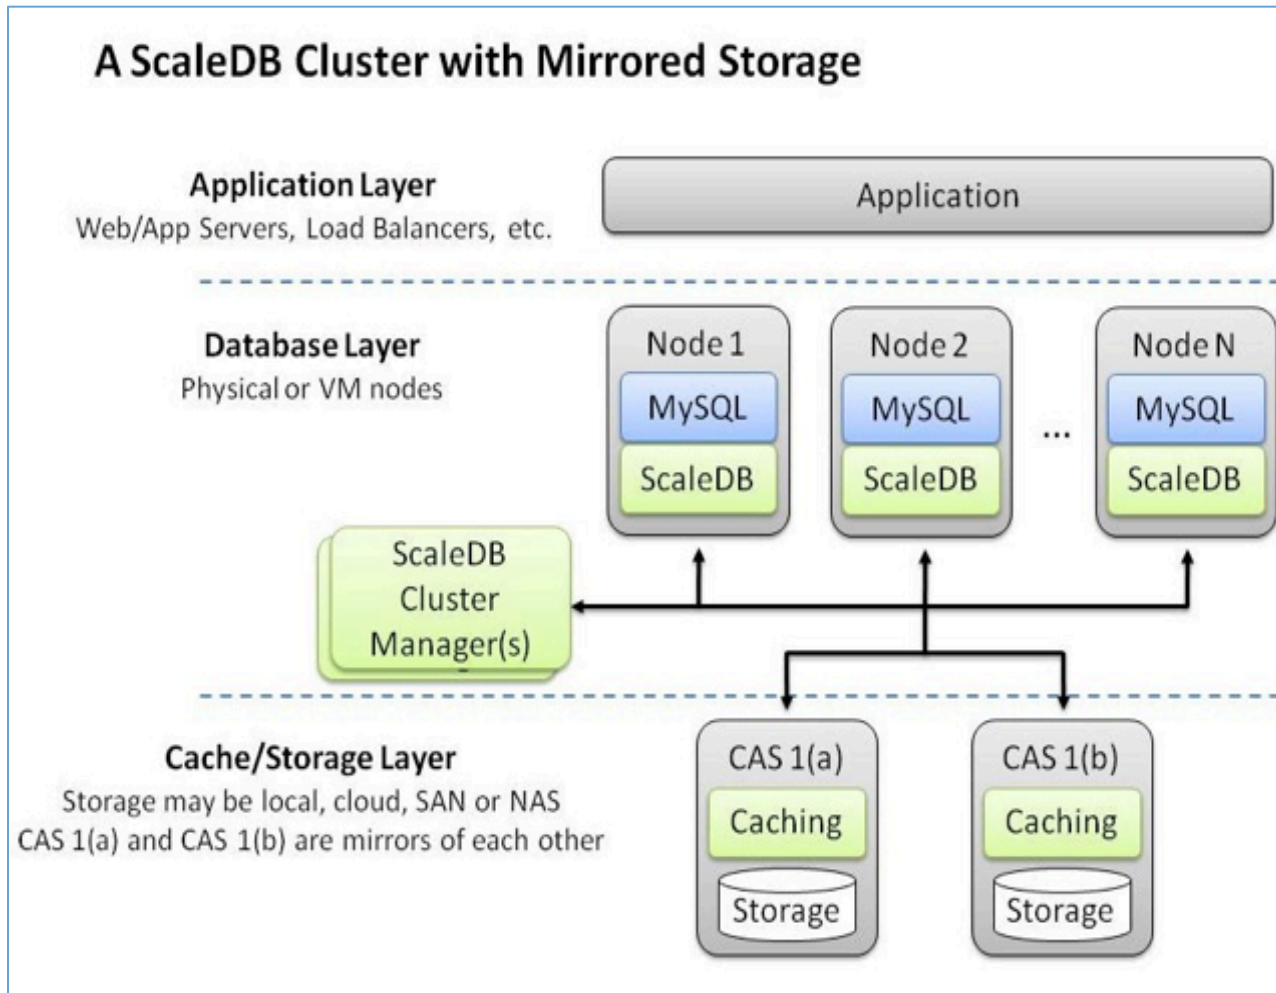
Bottleneck:

1. Disk Speed (I/O)
2. CPU
3. Capacity (users,data)
4. Performance

ScaleDB improvements:

1. Large numbers of concurrent users
2. Large data sets
3. Large numbers of tables queries
4. Load balancing
5. Cloud elasticity
6. Lower TCO

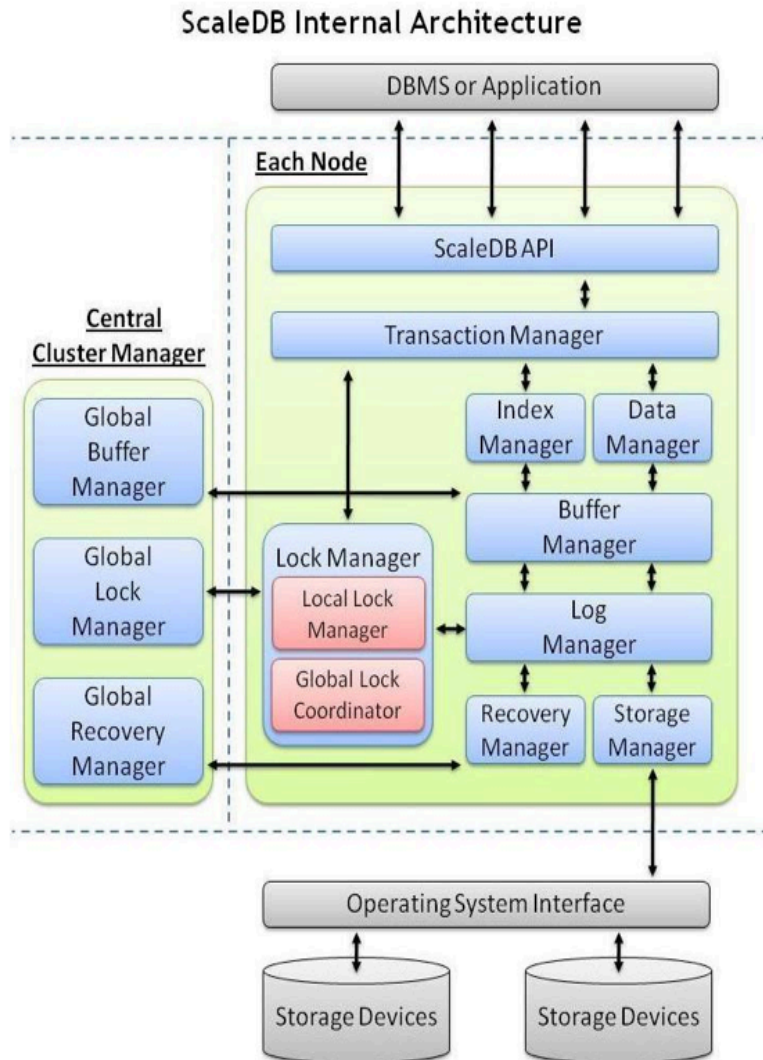
# ScaleDB cluster with Mirrored Storage



What's New:

1. Shared-Disk Clustering
2. Multi-Table Indexing
3. CAS (Cache Accelerator Server)
4. Locality

# ScaleDB Internal Architecture



## The Performance Challenges

### 1. Data contention

- Locality to reduce data movement (certain nodes in the database specialize in certain data)

### 2. Sharing data via the disk

- Shared cache done through CAS

### 3. Maximize local cache

- Improving cache hit ratio using CAS

### 4. Network traffic for messaging

- Locality eliminate synchronizing messaging

### 5. Network traffic for large chunks of data

- Local processing by CAS to send results only

# Features Comparison

| Feature           | Shared-Disk    | Shared-Nothing | ScaleDB  |
|-------------------|----------------|----------------|----------|
| Elasticity        | Yes            | No             | Yes      |
| High-Availability | Yes            | No             | Yes      |
| Ease of Use       | Medium         | Easy           | Easy     |
| Ease of Scaling   | Easy           | Hard           | Easy     |
| Software Costs    | High           | Low            | Low      |
| Storage Costs     | High (SAN/NAS) | Low (PC)       | Low (PC) |
| Performance       | Medium         | High           | High     |

# Features Summary

- **High-Availability**: Node failure affect less
- **Flexibility**: Share workload
- **High Performance**: MTI (Multi-Table Index)
- **Seamless scalability**: Add nodes & storage with shutting down
- **Less TCO**(Total Cost of Ownership):

| Innovative high-speed indexing | High-performance transaction processing      |
|--------------------------------|--|
| Plug-and-Cluster™ simplicity   | Graceful fault-tolerance                     |
| Automatic data recovery        | Row-level locking                            |
| ACID compliance                | Multi-node concurrency control               |
| Shared-everything architecture | Eliminates the requirement to partition data |

# Different with H-Store

- ScaleDB
  - Locality
  - CAS other than disk
  - Paired CAS on different nodes
  - Data contention lock
  - Cluster Manager
- H-Store
  - CTA
  - Main Memory
  - Partition Data replica on different nodes
  - Multi-partition txns
  - Site Coordinator

# Answer questions

- 1. Does the system support distributed transactions?
  - No.
- 2. Is the system open-source?
  - NO. (Except for connectors for MySQL)
- 3. Does the system support stored procedures?
  - NO.

# Sources