

hosted by  Alibaba Group 阿里巴巴集团  APACHE
HBASE

AntsDB

MySQL Compatibility for HBase

August 17, 2018

Content

- 01 About Us
- 02 Scenario
- 03 Architecture
- 04 Complementary to HBase

01 About Us

- ✧ Founder of AntsDB
- ✧ Love Hadoop and HBase but a heavy RDB user in the past
- ✧ Love programming for extreme performance
- ✧ Previously Software Architecture at Informatica, DataMirror, founder of BDI Systems.


- ❖ Open source project at <https://github.com/waterguo/antsdb>
- ❖ Problem to resolve: NoSQL is taking over the world but majority of the data applications are still using RDB
- ❖ Database virtualization software
- ❖ Attempt to bring MySQL compatibility to HBase

02 Scenario

Virtual Telephone Number

- ✧ Telecom industry
- ✧ Rigid latency requirement: 20ms max
- ✧ Data is growing fast
- ✧ Built on top of relational database – MySQL

Solution Architecture



What we like HBase

- ✧ Linear scalability
- ✧ Hadoop ecosystem
- ✧ Good random access performance

What we were struggling

- ✧ Latency spike
- ✧ Transaction
- ✧ Join performance

Why don't cache hot data

- Operational data = hot data
- Historical data = majority of big data
- Daily hot data = 1-5%
- Weekly hot data = 5-10%
- Monthly hot data = 10-15%


Silver bullet - caching

- ✧ Large amount local caching using SSD – 500 GB
- ✧ Local data fetch -> sub-millisecond latency
- ✧ Distributed transaction -> local transaction
- ✧ Distributed joins -> local joins

03 AntsDB Architecture

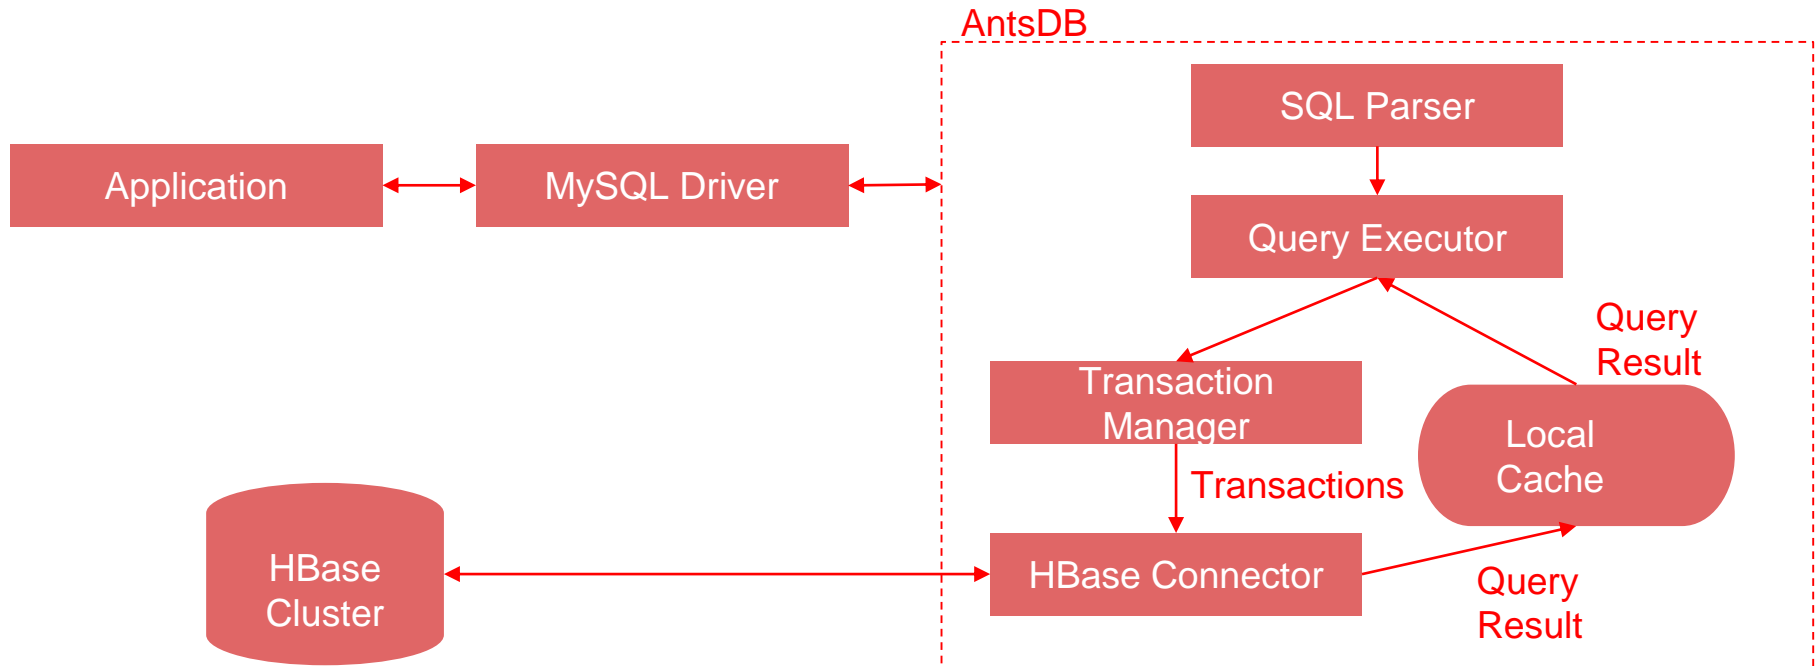
Layers

hosted by



Architecture


Subtitle Text



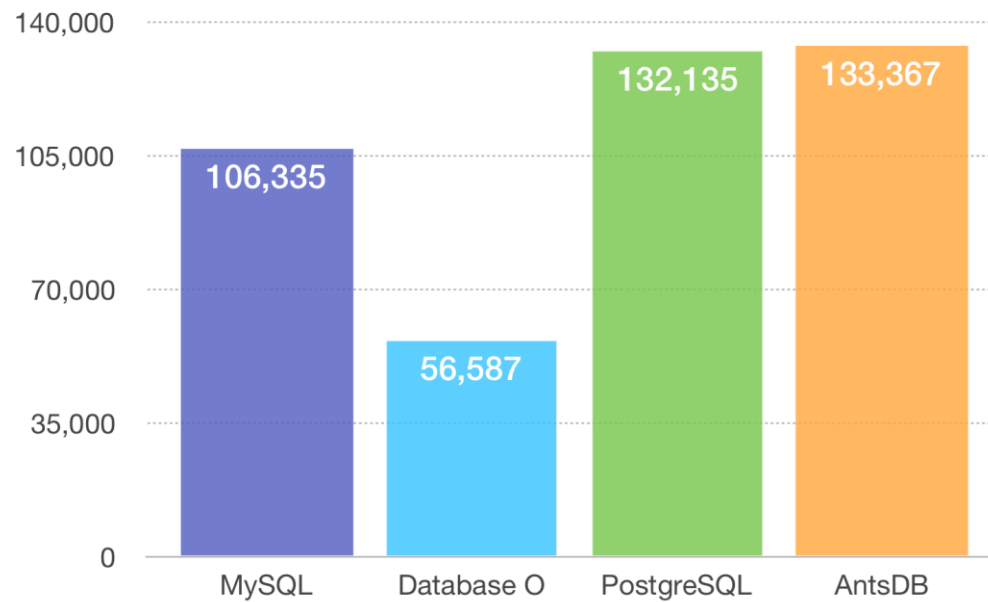
Fight the GC

- Java Unsafe - skip list, cache, WAL, row locking, cursor
- Atomic operation for row locks
- Small heap size - less than 4G
- Escape analysis
- Scan performance : 1 million rows / thread / second
- Insert performance: 200K inserts / thread / second

Performance - YCSB write



Performance - TPCC



04 Complementary to HBase

MySQL Compatibility

- MySQL protocol - JDBC, ODBC, PHP, Python, Native
- Core SQL syntax
- MySQL data types
- MySQL functions - 45
- Low latency transactions
- Low latency joins
- Indexes - Unique, non-unique, full-text

MySQL Behaviour Simulation

- Locking
- Commit and rollback
- MVCC
- Quirks such as 0000-00-00 00:00:00 datetime, zerofill integers
- Millisecond latency

Compatible applications

- MySQL command line utilities
- MySQL Workbench
- DBeaver
- SquirrelSQL
- MediaWiki
- SonarQube
- PHPMyAdmin
- BechmarkSQL - TPC-C benchmarking
- YCSB

Next step



High availability



Cluster



Low latency CP system

hosted by  **Alibaba** Group
阿里巴巴集团

APACHE
HBASE 

Thanks