

APACHE PHOENIX: TRIALS AND TRIBULATIONS OF BUILDING AN OPEN SOURCE, DISTRIBUTED DATABASE

This talk will focus on the journey Apache Phoenix[1] has taken from a one-man implemented tool for easing the access of HBase data, to a collaborative, multi-company, multi-project distributed relational database layer.

- Apache Phoenix is a relational database layer on top of Apache HBase that supports:
 - SQL 92 for querying, DML, and DDL
 - Mix of OLTP and OLAP workloads
 - Compilation of SQL statements into parallelized, native store calls
 - Client embedded or remotable JDBC driver
 - Multi-tenancy through views and tenant-specific connections
 - Multiple secondary index implementations optimized for different workloads
 - Flashback queries leveraging immutability of underlying store
- Leverages many other projects as component parts:
 - Apache HBase[2], an LSM tree implementation of a distributed KeyValue store based on Google Bigtable[3].
 - Apache Calcite[4], a query planner and cost-based optimizer framework based on Volcano optimizer[5], which opens the door for federation across other Calcite consumers such as
 - Apache Hive[6] for HDFS data
 - Apache Drill[7] for JSON data
 - Apache Kylin[8] for materialized views
 - Standard RDBMS sources
 - Apache Calcite Avatica[9], a remotable JDBC driver plugin
 - Cask Tephra[10], a snapshot-isolation transaction provider for HBase.
 - Apache HTrace[11], a distributed tracing framework for monitoring and management

[1] Apache Phoenix: <http://phoenix.apache.org/>

[2] Apache HBase: <http://hbase.apache.org/>

[3] Bigtable: A Distributed Storage System for Structured Data: <http://research.google.com/archive/bigtable.html>

[4] Apache Calcite: <https://github.com/apache/incubator-calcite/blob/master/README.md>

[5] Volcano optimizer: <http://www.seas.upenn.edu/~zives/03s/cis650/P209.PDF>

[6] Apache Hive: <https://hive.apache.org/>

[7] Apache Drill: <http://drill.apache.org/>

[8] Apache Kylin: <http://kylin.incubator.apache.org/>

[9] Apache Calcite Avatica: <https://github.com/apache/incubator-calcite/tree/master/avatica>

[10] Cask Tephra: <http://tephra.io/>

[11] Apache HTrace: <http://htrace.incubator.apache.org/>



James is an architect at [salesforce.com](https://www.salesforce.com) in the Big Data Group. He founded the Apache Phoenix project and leads the development effort. Prior to working at Salesforce.com, James worked at BEA Systems on projects such as a federated query processing system and an event driven programming platform and has worked at various other start-ups in the computer industry over the past 20+ years.