



Client Overview

- Mobile engagement platform solutions provider helping enterprises drive their marketing outreach



Status Quo

- The client's platform powers mass mobile interactions and time bound campaigns permitting billions of transactions – 150mn/day on average
- The data aggregates were stored in database and undergo ETL process to derive timely performance and integrity reports
- Achieve maximum reach and responsiveness by designing efficient data storage and processing architecture



Business Challenge

- The existing architecture used MySQL for storing messages, which limited insertion rate. The server required heavy IO and computation power for ETL processes
- PostgreSQL server used for reporting and aggregating logs was unable to lend access to real time reporting data
- The system required a horizontal data storage scale up, reduced ETL process time and support real time access to reporting data



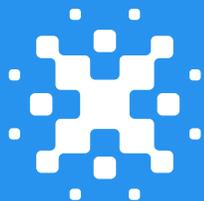
Technology

- Hadoop, Oozie, SQOOP, HIVE, HDFS, HBase, Phoenix



Business Impact

- Enhanced the ETL process by enforcing a centralized Big Data Hadoop ecosystem
- Merged the log database into Hadoop clusters. This enabled distributed processing
- Data required for real time reporting was generated using Hive tables
- Removed the data insertion bottleneck by introducing direct inserts into Hbase
- Hbase on top of HDFS and Phoenix implementation loaded real time data into Hbase and made accessing reports in real time effortless
- Reduced data processing times from 3 days to 3 hours



INDIA

Chennai | Bengaluru | Mumbai
Toll-free: 1800-123-1191

USA

Cupertino | Princeton | Boston
Toll-free: +1-888-207-5969

UK

London

SINGAPORE

+65 9630 7959

MALAYSIA

Kuala Lumpur
+60 (3) 2298 8465