



Client Overview

- The client is a leading Italian bank in operation for more than 130 years with more than 300 branches spread across Italy, Ireland, India and Romania.
- The bank offers a wide range of financial and banking products to meet its customers' needs, mixing the benefits of specialization with tailor-made solutions.
- The bank has always been at the forefront of innovation and has been part of many new initiatives – Among the first to provide online trading for brokers, part of the first European ecommerce initiative by Microsoft & Yahoo etc.



Business Requirements

- The bank was looking for a scalable real-time solution that would analyze data from all workloads and provide operational intelligence, to proactively reduce server downtime
- As part of this initiative, the bank wanted to mine server logs, at real-time, to enable accelerated troubleshooting, RCA and action plan to solve server performance issues



Key Challenges

- Integration of Apache Storm, chosen for its real-time data processing capabilities, with Cloudera Enterprise CDH 5
- Design of solution architecture that would fetch and push logs to Cloudera without additional coding required in all existing applications that generate the logs
- Implementation and performance tuning of the integrated application to provide real-time insights



Our Approach

- We studied the system first to understand the infrastructure setup, expected growth in the number of servers and data volume pertaining to the business requirements.
- We designed the architecture such that fetching logs has no impact on existing applications that generate them and suits new applications to be developed



Our Solutions

- **Apache Flume** daemons were leveraged to fetch and push server logs to **Apache Storm** through **Kafka** messaging queue
- The data processing was done at real-time in Apache Storm and the processed data was loaded into **Hbase NoSQL database**. **D3.js** visualization was built by the bank on top of this processed data
- The raw data from Apache Storm was pushed to **Apache SOLR** to enable Admins perform text searches and gain insights, at real-time



Results

- The entire application was built from the ground up in under **2 weeks**
- The application was tested with production grade data and tuned for performance to get real-time insights from server logs generated in **420 servers**



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