

Query Application for 3.9 Petabytes of Data

Bloomberg

client

Company: Bloomberg
Industry: Financial Services
HQ: New York City, New York
Size: 15,500 employees
Revenue: \$9 billion

technology

- Apache Cassandra™
- Scala
- Amazon Web Services

about kpi

- 4+ years experience with DataStax/Cassandra
- DataStax/Cassandra Strategic Partner
- Cassandra, DSE, Spark, Solr, Java, Scala, Chef expertise
- 200+ enterprise customers
- Onshore and Offshore consulting

KPI delivers custom Scala application on Apache Cassandra™ that allows users to query 3.9 petabytes of data

About Bloomberg

Bloomberg L.P. provides financial software tools such as an analytics and equity trading platform, data services, and news to financial companies and organizations. Connecting decision makers to a dynamic network of information, people and ideas, Bloomberg delivers business and financial information, news and insight around the world.

Bloomberg's Business Needs

Bloomberg has a 3.9 petabyte data store and needed a metadata query application to meet regulatory and compliance needs.

This data store contains data generated at Bloomberg terminals by their customers and includes both structured and unstructured data such as emails, PDFs etc.

Bloomberg's Selection Process

Bloomberg selected Apache Cassandra™ for its high availability for its global user base.

Bloomberg selected KPI for their proven Cassandra expertise and their expertise in building Java/Scala applications on the Apache Cassandra™ platform.

What KPI Delivered

KPI delivered a custom Scala application on Apache Cassandra™ that created metadata for the 3.9 petabytes of data and allowed users to query the metadata. The application was developed and deployed on Amazon Web Services.

KPI also did extensive performance testing and tuning using a billion object test database. KPI developed the test database, conducted the tests and tuned both the application and Apache Cassandra™ to meet Bloomberg's stringent performance and scalability requirements.

Business Benefits

- Enabled Bloomberg to query data to meet compliance requirements
- An application that is available 24x7 to Bloomberg's global user base
- An application that scales to meet Bloomberg's performance and scalability requirements

