



Summary

CUSTOMER

Major international investment bank

CHALLENGE

Address limitations of existing Hadoop data lake, which, unable to support real-time data streams, adequately scale, or meet analytical needs, slowed the response time of the bank's customer management application.

OUTCOME

Seamless integration of InterSystems IRIS data platform with existing data lake enables advanced querying, ability to scale and handle large data volumes, and high-speed real-time data processing. Bank's analytic needs are met across business lines for portfolio analysis, risk, and compliance.

Major International Investment Bank + InterSystems

InterSystems IRIS Complements Existing Data Lake

Enables Real-Time Capabilities, Advanced Analytics, and Scalability

When one of the world's largest investment banks was impacted by limitations with its Hadoop data lake, it implemented InterSystems IRIS® data platform as a dynamic data layer between the data lake and its production applications. The result not only addressed the bank's growing list of requirements, but also saved money by leveraging existing hardware.

With a high performance transactional-analytic database engine, a complete integration platform and a rich set of embedded analytics capabilities, InterSystems IRIS integrates seamlessly with the data lake to combine and process historical data with current transactions at the rate of 50 megabytes per second and with a response time under 100 milliseconds. And while enabling more than 100 concurrent connections, it meets the bank's analytic needs across business lines for such use cases as portfolio analysis, risk, and compliance.

Problem: Limited Analytics, Slow Response

Problems with the petabyte-sized Hadoop data lake impacted the bank's customer management application, which it used to query customer assets and demonstrate that trades met regulatory requirements. The application was slow, responding to queries in seconds when it should have responded in milliseconds.



“INTERSYSTEMS IRIS DATA PLATFORM IS A WONDERFUL PRODUCT THAT GIVES THE SCALE AND PERFORMANCE OF AN IN-MEMORY DATABASE AT A MUCH LOWER COST FOR VERY LARGE MULTI-TERABYTE DATASETS.”

Vice president at the bank

The petabyte-sized Hadoop data lake, while acceptable for historical analysis, had a mass of limitations, including its inability to support real-time data streams, scale out efficiently, or support the organization’s analytical needs. It could handle only simple queries, so the bank was hampered in its analysis reporting. What’s more, the data lake could not store previous, frequently used query results, resulting in inefficiencies and duplication of effort.

Ideally, the bank needed an application that could complement the existing data lake while also accommodating real-time data, facilitating more sophisticated queries, and handling large volumes of data at a low cost.

Solution: Seamless Integration of InterSystems IRIS

InterSystems IRIS, which easily integrates with existing systems and applications, could meet the bank’s performance requirements at scale. The solution was to add it to the existing configuration to support the complex queries the former architecture could not.

The new architecture includes a data access layer, which, feeding on the data lake, links to the InterSystems IRIS data platform via a message bus. Storing a few days of query history, InterSystems IRIS, in turn, feeds the application programming interface for data analysis. What’s more, it combines shared-nothing and shared-everything architectures for enhanced performance.

The data platform’s cloud-friendly architecture and its ability to run and scale on commodity hardware meant the bank could leverage its existing hardware infrastructure.

“InterSystems IRIS data platform is a wonderful product that gives the scale and performance of an in-memory database at a much lower cost for very large multi-terabyte datasets,” says a vice president at the bank.

