

# Consolidating Risk Management and Compliance Silos in Financial Services

Technology Guide



---

## Introduction

For years, financial institutions have been implementing disparate risk management and compliance applications to meet regulatory requirements and reduce operational risk.

Recently, however, many organizations have begun unifying these separate applications via a single, comprehensive, enterprisewide platform. Doing so allows them to gain a more accurate and complete view of their enterprise data, identify operational and compliance risks and suspicious activity faster and more accurately, comply with new and changing regulations, and simplify overall management and maintenance.

This paper explains how InterSystems IRIS™ data platform complements organizations' existing data management environments and their risk and compliance applications to provide a unified, panoramic view of all enterprise data and an ultra-high-performance scale-out processing layer.

## InterSystems IRIS Data Platform for Consolidated Risk Management and Regulatory Compliance

InterSystems IRIS is an ideal technology to help firms successfully consolidate their risk management and regulatory compliance applications onto a single platform.

It seamlessly complements an organization's existing data management infrastructure, including legacy applications, data warehouses, big-data technologies, and data lakes. Organizations can obtain a unified, panoramic, secure view of data from multiple sources across the organization via a real-time distributed caching layer.

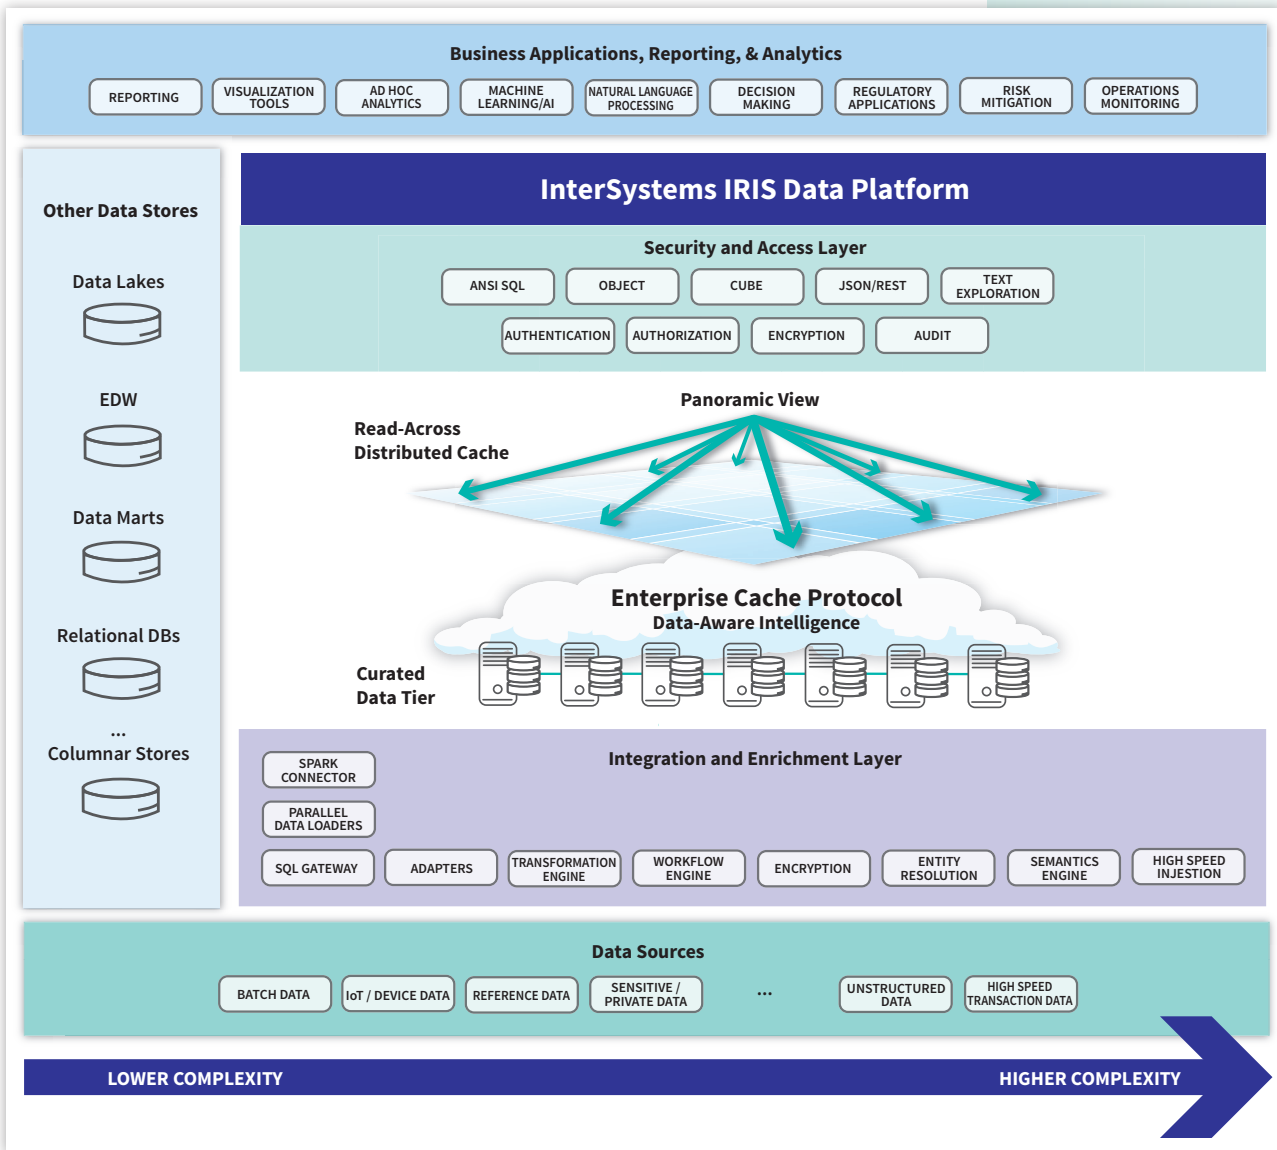


Figure 1: InterSystems IRIS data platform for consolidated risk and compliance: reference architecture

InterSystems IRIS provides an ultra-high-performance, scale-out processing layer to perform complex batch and real-time processing tasks on large, distributed data sets. It can perform complex multi-table joins on sharded data without requiring co-sharding,<sup>1</sup> replicating data, or performing network broadcasts. This enables organizations to reliably and quickly process queries that would otherwise take hours or time out before completion, at a lower operational cost.

<sup>1</sup> Co-sharded data refers to data that is partitioned on a common key.

---

InterSystems IRIS lets organizations gain a consistent view across all their enterprise data, gain more accurate and timely business intelligence, ensure better compliance with regulations, and respond more quickly to unanticipated questions from financial regulators and compliance analysts. At the same time, it provides the role-based access, security, encryption, governance, and lineage capabilities businesses need. Its horizontally scalable technology supports scale-out, sharded architectures to manage and analyze very large data sets using low-cost distributed processing and storage nodes.

**InterSystems IRIS executes fast and efficient complex processing tasks on distributed data sets and can perform multi-table joins without requiring co-sharding, replicating data, or performing network broadcasts. The result is dramatically higher performance and reliability, lower operational costs, and the ability to quickly process queries that would otherwise time out or take hours to complete.**

## HTAP Multi-Model Database

At the core of InterSystems IRIS is a proven, enterprise-grade, distributed, multi-model HTAP (hybrid transaction/analytical processing) database that is designed to work with large sets of heterogeneous data. It ingests, stores, and indexes large volumes of transactional data at very high rates to support real-time analytical use cases.

It provides the flexibility to store dissimilar source data in the most appropriate format. The data is stored once and can be represented as relational data, objects, multi-dimensional arrays, key-value pairs, documents, and other data types. This eliminates the need to duplicate data or map between different representations (e.g., object-relational mapping), for superior performance and efficiency. It natively supports sharded, distributed scale-out architectures, providing a cost-effective data platform for working with large data sets using commodity resources. It provides strong enterprise-level security measures, integration with Kerberos and LDAP, role-based access control, and encryption for data both in transit and at rest.

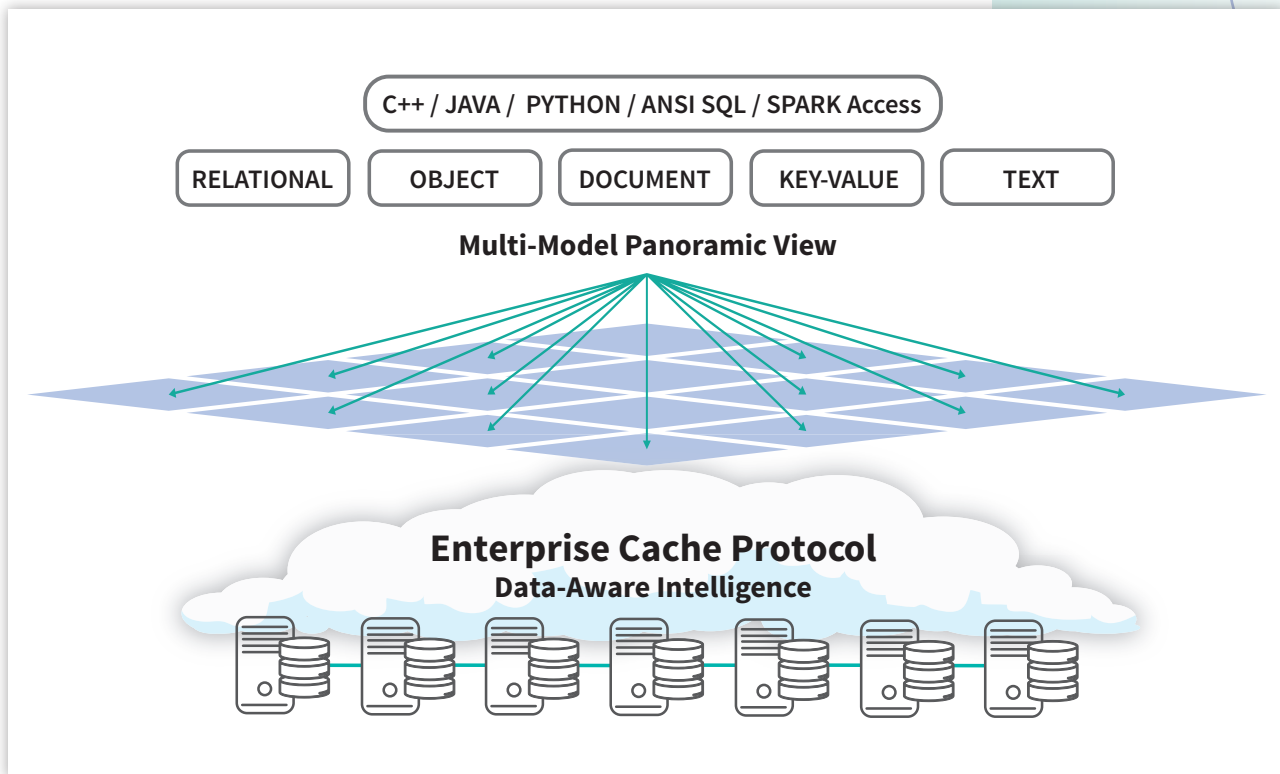


Figure 2: InterSystems IRIS multi-model hybrid transaction/analytical processing database

## Enterprise Cache Protocol

InterSystems provides powerful capabilities for reliable, high-performance, distributed, multi-workload (transaction and analytic) processing at very high scale. This comes from a unique technology, Enterprise Cache Protocol (ECP), which has been implemented, optimized, and used successfully in thousands of mission-critical production environments.

ECP is an integral capability of InterSystems IRIS. It quickly and reliably returns responses to queries on distributed data sets regardless of how the data is organized. It natively supports distributed, sharded architectures. Complex joins are processed locally rather than broadcast across the network, eliminating the latencies and time-outs typically associated with broadcast joins while increasing performance and reducing operational costs. ECP makes it possible for regulators and compliance analysts to quickly obtain accurate results from ad hoc queries without expensive and time-consuming pre-processing or replication of data.

Using ECP is intuitive. It requires no application changes or specialized techniques: applications and processing tasks simply treat the data as if it were local. The performance and scalability benefits are dramatic, enabling organizations to answer questions, correlate information, and identify patterns in distributed data sets that are not co-sharded with performance and reliability, and at a significantly lower cost.

---

## Integration With Apache Spark

Many of the business and regulatory drivers for risk management and compliance now require intraday and near-real-time reporting and visibility, driving organizations to increasingly employ higher-performance computing techniques to reduce latency. As a result, Apache Spark, with its in-memory architecture, is being evaluated and adopted by some organizations.

InterSystems IRIS provides fast and efficient parallel connectivity with Spark and is a high-performance, seamless complement to architectures already in use at many financial institutions. The shard-aware integration and enrichment layer supports batch inserts and individual inserts, allowing very large data sets to be ingested into InterSystems data shards quickly.

InterSystems technology integrates directly with Spark via a shard-aware native Spark connector. The connector presents InterSystems data shards as native partitions for the highest performance. This deep integration enables organizations to leverage InterSystems seamlessly and optimize Spark queries for improved performance.

## Concurrent Transactional and Analytic Processing Using Real-Time Data

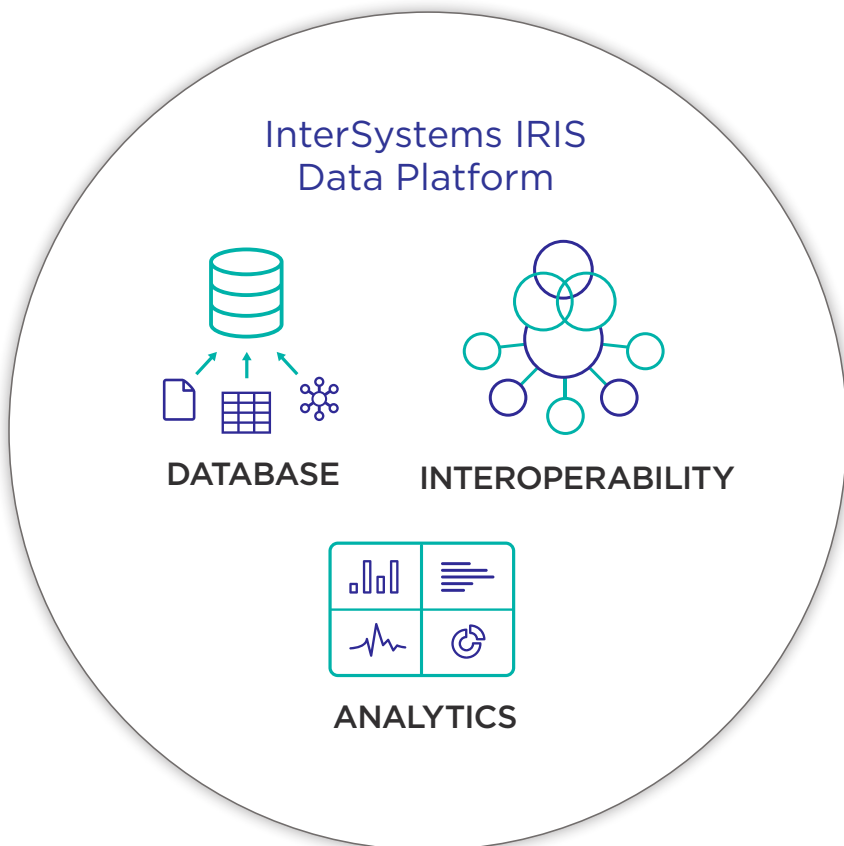
Whereas some risk and compliance use cases call for data to be analyzed in batches, others require real-time data to be analyzed with historical or reference data. For example, they may require transaction monitoring and filtering, intraday or pre-trade liquidity calculations, and other real-time and low-latency applications.

InterSystems IRIS is ideal for handling both real-time and batch requirements. It is optimized to process both very large transactional workloads and a high volume of analytic queries on transactional data and historical or other batch data simultaneously, without compromising performance.

## Interoperability

InterSystems IRIS provides connectivity to a wide range of applications and data sources. It also includes a built-in adapter library that provides connectivity and data transformations for traditional industry standards, protocols, and technologies, such as REST, SOAP, HTTP/S, Financial Information Exchange, Kafka, and Java Message Service. InterSystems IRIS offers a SQL gateway that can access and present data and metadata in common databases and data warehouses, including Oracle, Sybase, and DB2 – as well as SQL-on-Hadoop engines such as Hive and Impala – to InterSystems client applications as native tables. In addition, the data platform provides a workflow engine, encryption, entity resolution, and the ability to apply transformations.

InterSystems IRIS provides comprehensive, unified access to third-party reporting, analytics, and visualization tools – as well as risk, compliance, and other applications – that are already in use. It provides ANSI SQL support with time-tested, proven SQL optimizations on fully sharded scale-out data architectures, as well as the flexibility to support object, OLAP, and JSON/REST access to the data.



---

## Security

**In financial services applications, security is absolutely paramount. InterSystems IRIS provides a simple, unified security architecture with the following features:**

- It offers a strong, consistent high-performance security infrastructure for applications;
- It meets certification standards;
- It makes it easy for developers to build security features into applications;
- It places a minimal burden on performance and operations;
- It ensures that InterSystems IRIS can operate effectively as part of a secure environment and work well with other applications; and
- It provides infrastructure for policy management and enforcement.

**The security is based on authentication, authorization, auditing, and database encryption:**


- Authentication verifies the identity of all users;
- Authorization ensures that users can access the resources that they need, and no others;
- Auditing keeps a log of predefined system and application-specific events; and
- Managed key encryption helps prevent unauthorized viewing of information.

In addition to its general security features, InterSystems IRIS provides row-level, user-, and role-based SQL security for managing data access. It also supports the use of SSL/TLS and provides tools for a public key infrastructure.

## Integrated Analysis of Unstructured Text

Unstructured data, including free text in emails, documents, text messages, master agreements, Suspicious Activity Reports, and external data from blogs and tweets, can provide valuable insight to help banks reduce risk and identify suspicious or fraudulent behavior.

Natural language processing (NLP) is an integral component within InterSystems IRIS. Its native NLP technology uses a unique bottom-up approach to analyze text based on what is contained within the text itself. It can work with customer-defined dictionaries and ontologies, and it provides native embedded semantic analysis capabilities for analyzing patterns and correlations in unstructured data. For unstructured data, InterSystems IRIS also supports data exploration; signal detection and trend analysis; content-based profiling and clustering; and information extraction, categorization, and mapping. These capabilities can be useful in summarizing and contextualizing large amounts of free text for various compliance and surveillance initiatives.



## Data Lineage and Data Governance

Effective data lineage – the ability to describe the source of the data and how it changes as it moves through the data pipeline – and data governance are critical for risk and compliance initiatives. For example, the Securities and Exchange Commission’s Rule 613 (Consolidated Audit Trail) requires organizations to collect and accurately identify every order, cancellation, modification, and trade execution for all exchange-listed equities and options across all U.S. markets.

Different applications that perform different functions may store different representations of the data or modify the data (e.g., breaking a large initial order into smaller child orders for execution). Compliance analysts and regulators must have confidence in the original data sources and the processes and transformations that are applied. InterSystems IRIS provides strong support for multiple data types, including both object and SQL schemas, and its flexible metadata capabilities allow the application of proper data lineage and governance.




## Conclusion

InterSystems IRIS data platform offers a powerful and seamless complement to financial institutions' existing infrastructures to deliver a secure, panoramic view of all of their data assets enterprise-wide. It provides an ultra-high-performance, distributed scale-out processing layer for handling a range of complex batch and real-time tasks that are required for consolidating risk management and regulatory compliance solutions into a comprehensive, unified platform.

By incorporating InterSystems IRIS into their consolidation initiatives, financial institutions can more quickly and cost-effectively ask more questions of their enterprise data in applications, data lakes, warehouses, and other data sources; gain more accurate and timely intelligence into their businesses; ensure better compliance with industry regulations; and respond quicker to questions from financial regulators and compliance analysts, thus reducing operational and regulatory risk.

InterSystems is the engine behind the world's most important applications. In healthcare, finance, government, and other sectors where lives and livelihoods are at stake, InterSystems is the power behind what matters. Founded in 1978, InterSystems is a privately held company headquartered in Cambridge, Massachusetts (USA), with offices worldwide, and its software products are used daily by millions of people in more than 80 countries.

For more information, visit [\*\*InterSystems.com/Financial\*\*](https://www.intersystems.com/Financial)



The power behind what matters.

