

Managing Data in the High-Speed World of Financial Services: InterSystems IRIS Data Platform

Technology Brief





A LEADING GLOBAL BANK IMPROVED DATA THROUGHPUT 500%, SUBSTANTIALLY REDUCED LATENCY, AND LOWERED OPERATING COSTS BY 75% COMPARED WITH ITS PREVIOUS IN-MEMORY DBMS, ALL WITHOUT A SINGLE INCIDENT SINCE ITS INITIAL IMPLEMENTATION.

InterSystems IRIS Data Platform for Financial Services

InterSystems IRIS® data platform is an ideal technology for high performance, mission-critical financial services applications. It is optimized to accommodate both very high transactional workloads and a high volume of analytic queries simultaneously, without incident or performance degradation, even when message rates and analytic workloads spike unexpectedly. It provides greater resource efficiency and lowers costs compared with other technologies, and is non-disruptive to existing data management infrastructure.

At the core of InterSystems IRIS is a multi-model database that delivers fast transactional and analytic performance without sacrificing scalability, reliability, or security. It handles relational, object, document, key-value, and multi-dimensional data in a common, persistent storage tier without any replication of the data.

High Performance for Transactional Workloads with Built-in Persistence

InterSystems IRIS features a high-performance database that provides transactional performance equal to or better than in-memory databases, with built-in persistence. It is the permanent data store, and it is always current. Data is not lost when a machine is turned off, eliminating the need for any database recovery or rebuilding efforts.

InterSystems IRIS excels at data ingest, due in part to an advanced data engine that allows efficient and compact storage of data in a rich data structure. By using an efficient, multi-dimensional data model with sparse storage techniques instead of two-dimensional tables, data access and updates are accomplished faster, using fewer resources and less disk capacity. In addition to traditional TCP/IP access APIs, it provides in-memory, in-process APIs, further increasing ingest performance.

High Performance for Analytic Workloads

InterSystems IRIS also provides a wide range of analytic capabilities, including full SQL support, enabling financial organizations to use their existing SQL-based applications with few or no changes. Since the database stores data in efficient multi-dimensional structures, SQL applications achieve better performance than with traditional relational databases. It provides native support for other data paradigms, including objects, documents, key-value data, and unstructured data.



Its high performance for analytic workloads also derives from its unique approach to real-time indexing – transactional bitmap indexing – which enables extremely fast performance on complex queries on live transactional data. Most databases use traditional indexes that maintain a list of IDs for values in rows or objects. In contrast, bitmap indexes contain a separate bitmap for each possible value of a column or property. The advantage of bitmap indexes is that complex queries can be processed by performing Boolean operations without searching through the entire database, increasing response times for queries that search large volumes of data by 100 times or more.

However, with traditional bitmap indexing, updating the index can be slow. Such indexes can also require a significant amount of storage, which has limited their use with databases, especially if they need to analyze real-time transactional data. The unique transactional bitmap indexing in InterSystems IRIS leverages multi-dimensional data structures to eliminate these two problems. Consequently, updating the bitmaps is fast, often faster even than traditional indexes, and advanced compression techniques significantly reduce storage requirements. The result is ultra-fast bitmaps that can search very large data sets and enable analytic queries to incorporate “live” data with high performance at scale.

InterSystems IRIS also includes an embedded machine learning (ML) runtime engine that processes machine learning models natively within the database engine to support fraud detection, portfolio modeling, compliance, and other real-time predictive and prescriptive use cases with the highest performance.

Finally, InterSystems IRIS allows users to perform sophisticated analytic queries on very large data sets, including performing joins that can access data distributed on disparate nodes or shards, with extremely high performance and efficiency, and without making multiple copies of the data. The result is that organizations can include more data in their analytics to support downstream applications and ad hoc analytic queries to gain new and more accurate insights from the data, with high performance and low resource utilization.

This capability is also ideal for many back-office applications that require processing very large data sets within reduced processing windows using fewer compute resources.

And Both, Concurrently

InterSystems IRIS provides the highest levels of performance for both transactional and analytic workloads concurrently, at high scale, without compromising performance for either type of workload. Since rising order volumes increase both the transactional and analytic workloads on the system, the data platform must scale to handle increasing workloads without experiencing performance or availability issues, especially during periods of market volatility. InterSystems IRIS provides several features to support these demands:



THE RESULT IS ULTRA-FAST BITMAPS THAT CAN SEARCH VERY LARGE DATA SETS AND ENABLE ANALYTIC QUERIES TO INCORPORATE “LIVE” DATA WITH HIGH PERFORMANCE AT SCALE.

1. Enterprise Cache Protocol. To achieve superior multi-workload performance and scalability, InterSystems has developed a unique technology, Enterprise Cache Protocol (ECP). A key benefit of ECP is that resources can be scaled independently based on workload type (i.e. transaction processing or analytic queries) and load. ECP optimally coordinates the flow of data across a multi-server environment from ingestion to consumption. It enables full SQL access to all the data in the environment without replicating or broadcasting it.

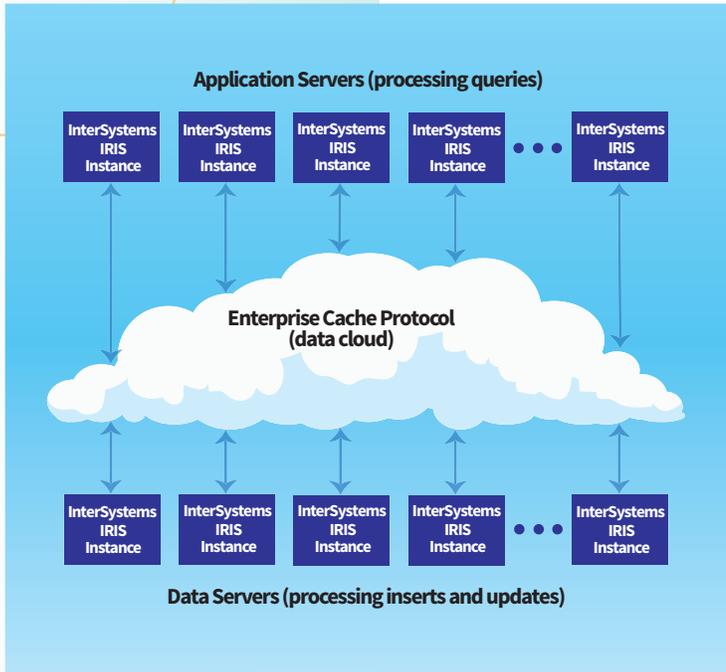


Figure 1: Most databases can be tuned either for inserts and updates, or for processing queries. Enterprise Cache Protocol is a unique technology that optimizes processing for both kinds of workloads in a single environment.

ECP enables the servers in a distributed system to function as both application servers and data servers, and to dynamically access remote databases as if they were local. Primary ownership of the data only needs to be held by a small percentage of the servers in the system. If analytic requirements increase, more application servers can be added instantly. If disk throughput becomes a bottleneck, more data servers can be added and the database becomes logically partitioned.

Each node in the distributed system can operate on data that resides in its own disk system or on data transferred to it from another data server by ECP. When a client makes a request for data, the application server will attempt to satisfy the request from its local cache. If the data is not local, it will request it from the remote data server, and the data is

then cached on the local application server where it becomes available to all applications running on that server. ECP automatically manages cache consistency and coherency across the network.

Using ECP is transparent and requires no application changes or specialized techniques. Applications simply treat the entire database as if it were local.

2. Sophisticated memory management techniques. InterSystems IRIS uses sophisticated memory management techniques to ensure high performance and availability. Most in-memory databases rely on the operating system for memory management, and thus can exhaust memory resources during periods of high workloads. InterSystems IRIS requests a continuous chunk of memory from the operating system at startup, and then intelligently manages and optimizes the memory, independently of the OS. It is possible (and expected) that during periods of high workloads, the allocated memory may be completely consumed. If that occurs, InterSystems IRIS dynamically releases the least recently used data, freeing memory, and continues operating seamlessly. If some required data is not available in memory, it simply retrieves it from disk.

3. Cache optimization. Unlike most other databases, which maintain a separate cache for each process running on a machine, InterSystems IRIS maintains a single cache for each machine, and allows processes running in their own memory address space to access the data. Since multiple clients are able to share a single cache, only one copy of the data needs to be maintained for each machine, resulting in reduced storage requirements, reduced network I/O, and superior scalability. For example, in an InterSystems IRIS-based system of 250 machines, each with eight cores, only 250 caches need to communicate with one another to maintain cache coherence. In contrast, systems that require a separate cache for each core would need to coordinate among 2,000 caches.

The performance and scalability benefits of InterSystems IRIS are dramatic, enabling organizations to efficiently process transactional and analytic workloads concurrently, without compromising either type, using a single platform, with the highest levels of performance and reliability, even when transaction volumes spike.

| | InterSystems | In-Memory DBMS |
|---|--------------|----------------|
| High performance for transactions and analytics | ● | ● |
| Automatic built-in durability | ● | ● |
| No hard scalability limit | ● | ● |
| No restart delays | ● | ● |
| Low total cost of ownership | ● | ● |

Figure 2: Comparing InterSystems IRIS with in-memory databases

Lower Total Cost of Ownership

InterSystems IRIS provides a single, consistent architecture for working with multiple data types and for processing transactional and analytic workloads, eliminating the need to learn or integrate multiple technologies and approaches. It also provides a highly intuitive application development and maintenance environment, speeding implementations, simplifying application maintenance, and enabling firms to accomplish more with fewer resources while lowering development and maintenance costs.

Its highly efficient use of system resources reduces the amount of necessary infrastructure, especially compared with in-memory databases, significantly lowering infrastructure costs.

Using InterSystems technology, one financial services firm is experiencing a 75% reduction in operating costs while increasing throughput, performance, and reliability, compared with its previous in-memory DBMS-based system.

Security

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

- Offers a strong, consistent, and high-performance security infrastructure for applications
- Meets certification standards
- Makes it easy for developers to build security features into applications
- Places a minimal burden on performance and operations
- Ensures that InterSystems IRIS can operate effectively as part of a secure environment and that other applications and InterSystems IRIS can work together well
- 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.
- Managed key encryption protects information against unauthorized viewing.

In addition to its general security, InterSystems IRIS provides security with granularity of a single row. With row-level security, each row holds a list of authorized viewers, which can be either users or roles.

InterSystems IRIS also supports the use of SSL/TLS and provides tools for a public key infrastructure (PKI).

Customer Examples

InterSystems technology is used to power a wide range of mission-critical, high-performance front-, middle-, and back-office applications throughout the financial services industry. For example, it is used to process up to 15% of the worldwide equity trades every day; one firm is using it to power its entire global post trade management function; and a major brokerage firm relies on it to power its entire suite of 3,000 back office processes supporting 11 million client accounts.

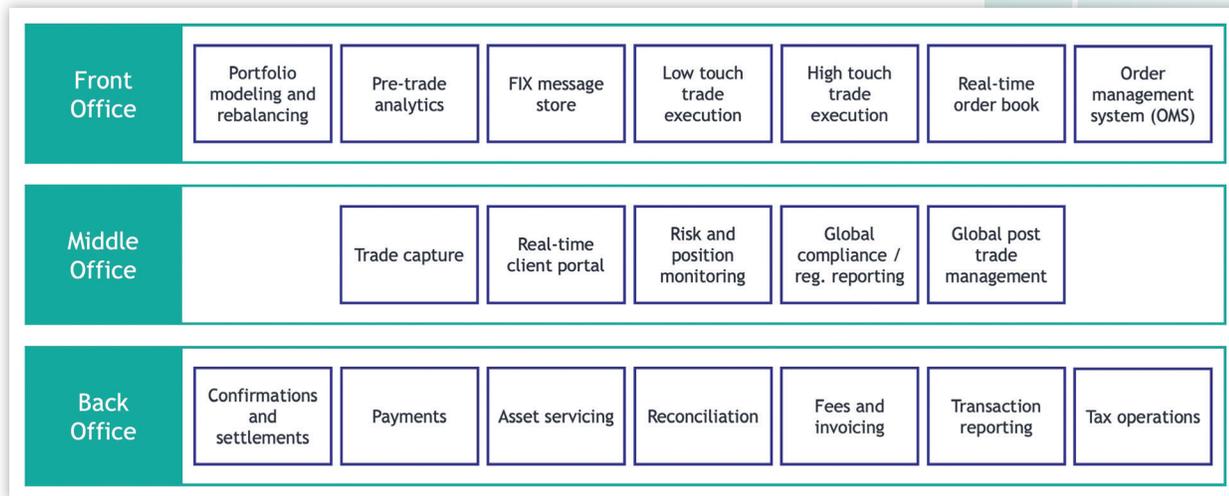


Figure 3: Selected InterSystems IRIS data platform production implementations at capital markets firms

High-Performance Transaction Order Management System (OMS)

A leading global sell-side investment bank with 50,000 employees and more than \$1 trillion in assets under management had conflicting requirements for its high-performance order management system to reliably handle increasing trade order volumes without incident, while lowering infrastructure costs. The globally distributed system must accommodate billions of daily orders, with sufficient headroom to handle unexpected spikes during periods of market volatility. At the same time, the system must service analytic queries from more than 300 applications across the organization.

The bank had previously used an in-memory DBMS-based system that experienced serious performance and availability issues during periods of peak loads. Our customer migrated the application to InterSystems technology, resulting in 3-5x increase in throughput, 10x increase in performance, with a 75% reduction in operational costs. In addition, the bank is now able to include more order data in analytic queries, compared with its previous in-memory based system.

With the InterSystems-based implementation, there is no single point of failure. Each functional component is linearly scalable, and it has successfully handled all trading days – without a single performance or availability issue.

INTERSYSTEMS IRIS HAS PROVEN RELIABILITY AND SCALABILITY UNDER THE MOST STRENUOUS MARKET CONDITIONS, INCLUDING CURRENT VOLATILITY WITHOUT THE NEED TO OVERBUILD OR INCREASE THE HARDWARE FOOTPRINT, MAINTAINING OVER 1 BILLION TRANSACTIONS PER DAY WITH PEAKS NEAR 100,000 TRANSACTIONS/SECOND.



A Leader in 2019 Gartner
Magic Quadrant for OPDBMS

Gartner[®]

Read the report at
[InterSystems.com/MQ](https://www.gartner.com/doc/reprints?id=338286&tid=646&cid=392217)

**INTERSYSTEMS WAS
RECOGNIZED FOR ITS
ABILITY TO EXECUTE
AND COMPLETENESS OF
VISION.**

**INTERSYSTEMS
WAS RECOGNIZED
IN TRADITIONAL
TRANSACTIONS,
DISTRIBUTED VARIABLE,
EVENT PROCESSING/
DATA IN MOTION,
AND AUGMENTED
TRANSACTIONS USE
CASES.**

Source: Gartner, Magic Quadrant for
Operational Database Management Systems
(25 November 2019)

Gartner does not endorse any vendor, product or service depicted in its research publications, and does not advise technology users to select only those vendors with the highest ratings or other designation. Gartner research publications consist of the opinions of Gartner's research organization and should not be construed as statements of fact. Gartner disclaims all warranties, express or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose.

GARTNER is a registered trademark and service mark of Gartner, Inc. and/or its affiliates in the U.S. and internationally, and is used herein with permission. All rights reserved.

Analytics on Historical Data Lake and Real-Time Transactional Data

A different sell-side investment bank with more than 35,000 employees and more than \$800 billion in assets under management had requirements to analyze its historical transactional data and respond to end-user queries in milliseconds. The bank had imported 70 years' worth of historical data into a data lake in Hadoop, but it was not able to meet all of its requirements using the analysis capabilities in Hadoop alone.

While their existing data warehouse was able to support simple SQL queries, it was unable to perform the necessary SQL queries to correlate data (e.g., via joins) to identify new relationships and insights. Many of the queries exceeded the bank's strict performance service-level agreement (SLA) of 200 milliseconds.

Our customer implemented InterSystems technology to complement its data lake in Hadoop. As a result, it is now able to perform complex SQL queries faster, including inner, outer, cross, and cross-network joins. Traditional SQL queries now execute five times faster on average, and all queries execute in less than 200 milliseconds, meeting the firm's SLAs. In addition, more data, including the real-time transactional data, is now included in the analytics to deliver new and more accurate insights.

Conclusion

More and more financial services firms are turning to InterSystems IRIS to keep pace with the escalating requirements that are challenging the industry.

It is ideal for high-performance front- and middle-office applications that must process both transactions and analytic workloads concurrently, at very high scale, with the highest levels of reliability, and with a low total cost of ownership.

It is also ideal for back-office applications that require processing large very large data sets within reduced timeframes using fewer hardware resources.

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