

A New Data Platform for Transactional-Analytic Applications

InterSystems IRIS Data Platform

Technology Brief



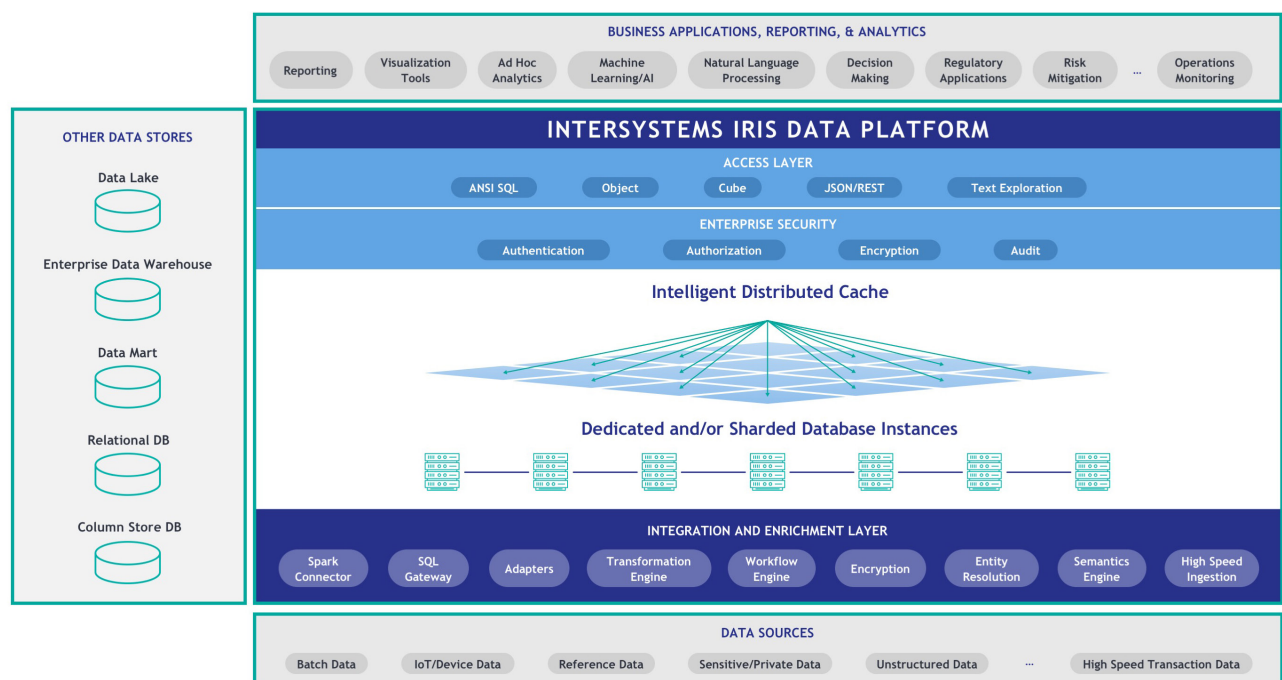
InterSystems IRIS: A New Data Platform for Transactional-Analytic Applications

Introduction

Historically, transactional applications (dealing with real-time data) and analytic applications (dealing with historical and other non real-time data) were separate entities. But as organizations wanted to make use — in real time — of the insights gleaned from their data, a new breed of solutions emerged. Transactional-analytic applications simultaneously support the rapid ingestion of data from a variety of sources and sophisticated analytics.

InterSystems IRIS® data platform is designed for development and execution of high performance transactional-analytic applications. It combines in-memory performance with highly optimized disk storage and unique intelligent distributed data aware caching that eliminates the need to duplicate entire data sets in memory or on disk. Applications built on InterSystems IRIS can ingest large sets of transactional data — while simultaneously executing complex analytics, including machine learning, business rules, natural language processing, BI, and SQL queries — with better performance and higher reliability at a lower cost. This paper describes some of the major technical advantages of InterSystems IRIS.

InterSystems IRIS Reference Architecture



InterSystems IRIS data platform provides:

- High performance, concurrent transactional-analytic processing to support real-time use cases.
- Support for executing machine learning, business rules, natural language processing, BI, and SQL queries in real time.
- A complete interoperability platform to integrate with existing applications and data sources.
- Ability to scale up and scale out with superior resource utilization and lower total cost of ownership (TCO).
- Multi-model database management support for object, relational and document technologies.
- Intelligent distributed data aware caching which allows workloads to be distributed among multiple nodes using commodity hardware.
- Interfaces to all prevalent programming languages, including Java, Node.js, .NET, Python, and others.
- Highly customizable data access and update layers that can be adapted individually to meet the application performance demands.
- Optimized interfaces for JDBC, ODBC, Apache Spark, and others.

Easy Integration with Existing Applications

InterSystems IRIS uses versatile data structures that can easily work with data as relational, object, or document models. All data can be accessed and manipulated in any appropriate model at any time without additional coding. For example, the InterSystems IRIS relational model is fully ANSI SQL compliant, and has adapters that allow easy integration with other relational products.

The result? InterSystems IRIS integrates with existing data warehouses and data lakes. It can be used as a data accelerator to make the best use of existing investments in Apache Spark, Amazon S3, and other technologies, allowing for real-time data ingestion and highly efficient analytics of large data sets.

InterSystems IRIS also integrates with various machine-learning libraries via the Predictive Model Markup Language (PMML). PMML models can be natively incorporated into the database for high performance. What's more, InterSystems IRIS has a built-in text analytics engine and can also be easily integrated with other text analytics solutions via the Apache Unstructured Information Management Architecture (UIMA) standard.

Intelligent Distributed Data-Aware Caching

Not only does InterSystems IRIS support simultaneous high-performance data ingestion and analytics, but it can do so at scale, thanks to its unique, intelligent distributed data-aware caching capabilities. In a multi-node setup of InterSystems IRIS, each node can be aware of what other nodes hold. Scenarios of improved application performance include:

- In a sharded environment, InterSystems IRIS significantly improves performance and reliability for complex cross-shard joins.
- In a multi-node cluster, it allows query workloads to be distributed among all nodes in the cluster.
- For concurrent transactional-analytic workloads, it ensures that analytic workload processing does not affect transaction/ingestion performance or reliability.

Integrated Machine Learning (ML)

InterSystems IntegratedML is an embedded feature of InterSystems IRIS. It allows application developers to automatically create and train machine learning (ML) models without requiring specialized data science skills. IntegratedML makes it possible for organizations that do not yet have data scientists on staff to develop and deploy machine learning models. And it can improve the effectiveness of organizations with teams of skilled data scientists by automating much of the tedious work involved in data wrangling, feature engineering, and model building, freeing data scientists to work on higher value tasks.

Deep integration with InterSystems IRIS enables applications to seamlessly execute these ML models dynamically in response to real-time events and transactions, without extracting or moving any models or data.

Comparison with Other Technologies

	InterSystems IRIS	OLAP Data Warehouse	Distributed In-Memory DBMS	OLTP DBMS
Architecture	Distributed; with storage; not limited by RAM	Column store	Distributed; limited by total RAM	Not distributed
Performance/ Scaling	<ul style="list-style-type: none"> • Transaction, query, and user workloads can be distributed and partitioned horizontally • High, consistent performance (milliseconds) on large partitioned data using less infrastructure which results in lower costs • Optimized for in-memory processing, but not constrained by memory; elastic caching tier spills over to disk as needed, resulting in lower costs and higher reliability 	<ul style="list-style-type: none"> • All data must be replicated on all machines • Inefficient scaling for queries • Performance is limited per machine • Need shared storage with flash storage preferred 	<ul style="list-style-type: none"> • High costs (servers and RAM) to scale • Stability problems when memory is exhausted • High performance dependent on in-memory capacity 	<ul style="list-style-type: none"> • Workloads must fit on a single machine (<1 TB) • Inefficient scaling • Query performance and concurrency limited

Performance Metrics

In customer proofs of concept and production implementations, InterSystems IRIS routinely outperforms competitive technologies for processing transactions and analytic queries – including in-memory, column store data warehouse, and open source technologies – using fewer hardware resources.

In one head-to-head comparison, InterSystems IRIS is more than **30 times faster** at querying real-time transactional data than the leading commercial in-memory database, while consuming fewer system resources.

Visit InterSystems.com/SpeedTest for full details and to run the open source performance comparisons for yourself.

**INTERSYSTEMS IRIS IS
TESTED TO BE MORE
THAN 30 TIMES FASTER
THAN THE LEADING
COMMERCIAL
IN-MEMORY DBMS**

Easy Deployment and Provisioning

InterSystems IRIS also includes capabilities that make these applications easier to build, deploy, and provision, including flexible interoperability, cloud neutrality, and a robust security model.

Built for Cloud Neutrality

InterSystems IRIS is supported on all major cloud providers and offers the following advantages:

- Eliminates lock-in to a single cloud provider.
- Runs on standard hardware (e.g. Amazon EC2 instances) to support on-premises, private, and hybrid deployments with no custom configurations required.
- Provides one single API, seamlessly supporting various cloud and hybrid deployments without requiring any modifications.
- Continuously leverages the optimizations released by the various cloud providers and hardware vendors.

Security

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 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 SQL security with granularity of a single row. This is called row-level security. 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).

Low Total Cost of Ownership

The unique architectural approach of InterSystems IRIS delivers lower total cost of ownership for customers, based in part because it:

- Does not require that all the data fit in memory.
- Runs on commodity hardware and a smaller footprint than competing products.
- Provides built-in high-availability capabilities.
- Is highly reliable, scalable, enterprise-proven, and requires less maintenance and administration support.

InterSystems: An Industry Leader

The industry analyst firm Gartner recognized InterSystems as a Leader in its 2019 Magic Quadrant for Operational Database Management Systems (OPDBMS).

Gartner Peer Insights Customers' Choice

InterSystems has achieved the highest overall ranking among all of the vendors in the Gartner Peer Insights Customers' Choice for OPDBMS for the second year in a row. The Gartner Peer Insights Customers' Choice distinction recognizes vendors and products that are highly rated by their customers. The survey data represents a top-level synthesis of vendor software products that are most valued by IT enterprise professionals.



Conclusion

InterSystems IRIS is an ideal platform for real-time, transactional-analytic applications that require high performance and the ability to efficiently scale to handle high data ingest rates and large data sets. It enables rapid ingestion of multiple forms of data from a variety of sources, while simultaneously supporting high-performance analytics processing — at scale, with robust security and a lower total cost of ownership.

At InterSystems, we're all-in as a partner, and we'll do whatever it takes for our clients to be successful. Going the extra mile is our policy. Our best people are on the front lines 24/7 to provide our clients with immediate help.

InterSystems was named a March 2020 Gartner Peer Insights Customers' Choice for Operational Database Management Systems (OPDBMS). The GARTNER PEER INSIGHTS CUSTOMERS' CHOICE badge is a trademark and service mark of Gartner, Inc., and/or its affiliates, and is used herein with permission. All rights reserved. Gartner Peer Insights Customers' Choice constitute the subjective opinions of individual end-user reviews, ratings, and data applied against a documented methodology; they neither represent the views of, nor constitute an endorsement by, Gartner or its affiliates. <https://www.gartner.com/reviews/customers-choice/operational-dbms>

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.

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

Gartner®

Read the report at
[InterSystems.com/MQ](https://www.gartner.com/MQ)

**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.**

