



www.InterSystems.com

One Memorial Drive, Cambridge, MA 02142, USA

Tel: +1 800 753 2571


Email: Info@InterSystems.com

InterSystems IRIS

The company

InterSystems Corporation is a privately-owned company that was founded in the USA in 1978 and it opened its first office in the UK in 1986. Revenues were nearly \$700 million in 2018. Apart from North America it has offices throughout Europe (including Russia and Israel), across both South America and South East Asia, and in Australia, South Africa and the Gulf, in a total of over 20 countries. It also has distributors around the globe.

Historically the company has largely relied on implementation partners and ISVs for its sales, with a relatively small direct sales force. While the company's partner base remains a major strength InterSystems has also expanded its direct sales force to support large corporations that require a direct relationship and because of the HealthShare and TrakCare solutions, which provide healthcare applications for the US (HealthShare) and elsewhere (TrakCare).



The image in this Mutable Quadrant is derived from 13 high level metrics, the more the image covers a section the better. Execution metrics relate to the company, Technology to the product, Creativity to both technical and business innovation and Scale covers the potential business and market impact.

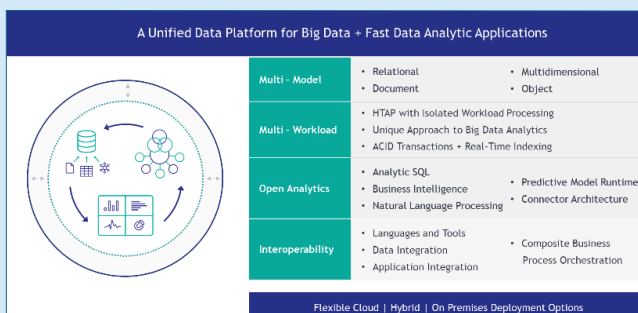


Figure 1 – An overview of InterSystems IRIS

What is it?

InterSystems IRIS Data Platform™ was announced in late 2017 and made generally available shortly thereafter. It is an evolution of the company's Caché database, Ensemble interoperability platform, and its various analytics offerings for both structured and unstructured data, which are all available in this single platform. In this context, it is worth commenting that InterSystems IRIS is the company's fourth generation of database technology; Caché was first introduced in 1997. InterSystems IRIS is available for on-premises, cloud-based, and hybrid deployments.

The major elements of InterSystems IRIS – illustrated in **Figure 1** – are a horizontally and vertically scalable multi-model transactional-analytic database with full ACID compliance and immediate consistency; scalable and distributed application server(s); a development environment (Eclipse-based) supporting several languages; a business process layer with a rules engine, workflow and process orchestration; and specific capabilities to support analytics against both structured (including integration with Apache Spark) and unstructured (including natural language processing) data, as well as integration with streaming environments such as Apache Kafka.

How does it work?

InterSystems IRIS contains a multi-model database that stores data in multi-dimensional arrays. Models supported are relational, object (with full persistence, polymorphism, inheritance and so forth, and with no requirement for object-relational mappings), document (XML, JSON), and multi-dimensional, and you can implement any number of these within the same environment, with full interoperability across these projections, and without any duplication of data. Thus, for example, other models can also be implemented on top of the arrays within InterSystems IRIS.

Architecture	★★★★★
Concurrent analytics	★★★★★½
Criticality	★★★★★½
Ease of use	★★★★★

Flexibility	★★★★★½
In-process analytics	★★★★★½
Performance	★★★★★
Scale	★★★★★

“We view InterSystems (IRIS) as a powerful and comprehensive platform for development. InterSystems provides all the capabilities we need to meet our business demands, in one seamless environment.”
CFAO

Where InterSystems IRIS differs from most other vendors is that competitive products typically have a row store and a column store to support transactions and analytics respectively. This is not necessary in InterSystems IRIS because you can read an array either vertically or horizontally. We consider this to be a major advantage for InterSystems IRIS.

InterSystems IRIS is able to process transactions, index the incoming data in real time, and perform analytics on the real-time data and non-real-time data (that is, historical data and reference data) at scale.

InterSystems IRIS scales to accommodate large workloads and data sets on commodity

hardware. **Figure 2**

illustrates how both the database and application servers scale out horizontally. From the point of view of hybrid applications, the most important thing is how you maintain both performance and consistency. This is enabled by means of

InterSystems Enterprise Cache Protocol (ECP). This decouples the execution of application code from the process of persisting the data, which allows application servers to scale horizontally, independently of the number of shard servers, and to distribute workloads automatically. It is possible to configure the environment so that some nodes are designated as query servers and others as transactional or hybrid servers, to cater for different workload resource requirements. Data ingestion can be done in parallel directly to each shard server, providing high-performance ingestion for streaming data. InterSystems IRIS also allows analytic queries to be pushed down to partitioned or sharded data tables, further increasing performance and resource efficiency.

There are a variety of specific query capabilities provided by InterSystems IRIS. To begin with, there is an Apache Spark connector

for leveraging machine learning and predictive modelling within the IRIS environment. The connector operates in parallel, with high-speed connections from each of the shard servers shown in **Figure 2** into a Spark cluster. Predictive models developed in R, IBM SPSS and other environments are executed natively within InterSystems IRIS through support for PMML (predictive modelling mark-up language).

To facilitate the embedding of real-time business intelligence into operational applications InterSystems IRIS includes a designer for creating dashboards and extending operational applications; an analysis component, employed by business users to explore and display relevant data; and an architect component, which is used to define your data model. From a performance perspective a major feature is that the software can make direct use of graphics processors for pipelining and so forth. An ODBC interface is available to support third-party visualisation products such as Tableau, Power BI and Qlik. For unstructured data InterSystems IRIS includes natural language capabilities and it supports Apache UIMA (unified information management architecture) for processing natural language processing (NLP) pipelines.

Why should you care?

Hybrid real-time use cases are becoming increasingly common in financial services (fraud detection), retail (customer personalisation), manufacturing/IoT (detecting errors, predictive maintenance), and other industries. InterSystems IRIS not only caters to these environments but it is going to be more cost efficient than many of its rivals, thanks to the fact that you only need to store your data once and in only one place. And unlike some other multi-model databases you get full interoperability across all data models without any duplication of data.

The bottom line

We are long-time fans of InterSystems products and that hasn't changed with the introduction of IRIS. We think there is a compelling argument for shortlisting InterSystems if you are considering a new hybrid transactional/analytic application deployment.

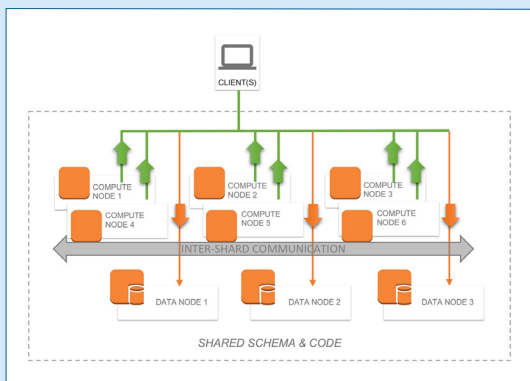


Figure 2 – Showing horizontal scaling

[FOR FURTHER INFORMATION AND RESEARCH CLICK HERE](#)