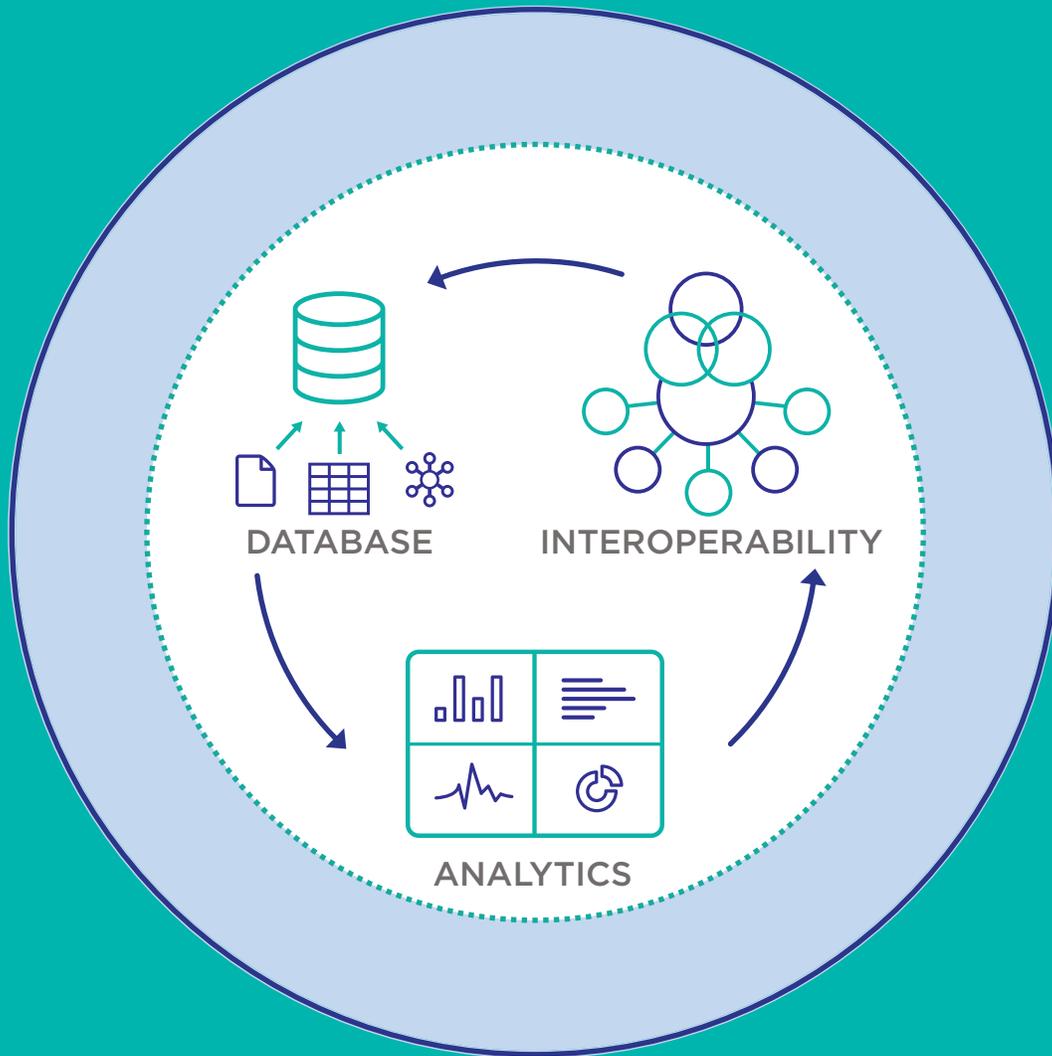

InterSystems IRIS Data Platform

For Current Caché and Ensemble Customers



INTUITIVE
|
RELIABLE
|
INTEROPERABLE
|
SCALABLE

Contents

Introduction	3
Performance & Scalability	
Performance	4
Scalability	5
Database Development	
Document Database Model	6
Optimized Client-Side Programming	6
API Management	7
Analytics	
Business Intelligence	8
Machine Learning	9
Text Analytics	10
Third-Party Business Intelligence Connectors	11
Connected Applications	
Interoperability	12
Java Business Hosts	13
Managed File Transfer	13
Deployment	
Support for Container Technology	14
Cloud Support	15
References	15

Introduction

InterSystems IRIS™ is the latest evolution of our data management software and provides our customers with powerful technology upon which to build modern, high-performance, data- and analytics-intensive applications.

Today's applications are becoming smarter and faster; more API driven; more integrated; incorporate more and different kinds of analytics; and leverage cloud and containers.

For organizations to survive and grow, their technology must allow them to keep pace with these trends, as competitors are introducing new applications and revising their existing ones to offer new capabilities. InterSystems IRIS is the data platform software that empowers organizations to capture new market opportunities and remain competitive and relevant.

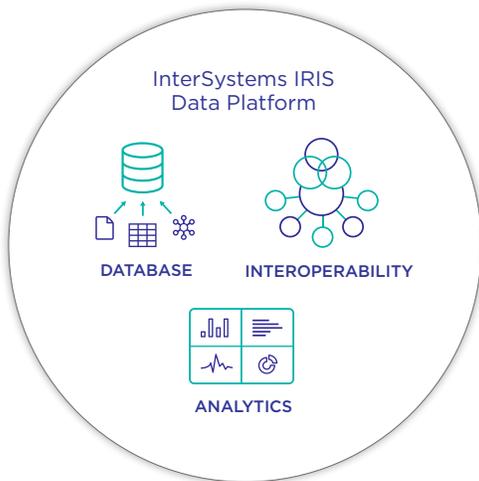
This document is aimed at current InterSystems Caché® and InterSystems Ensemble® customers, and presents the new functionality offered with InterSystems IRIS to stimulate ideas for revising and refreshing current applications and creating new ones, by leveraging new features, more data, new kinds of analytics, and greater workloads.

Upgrading to InterSystems IRIS

While InterSystems is committed to supporting Caché and Ensemble for our existing customers and partners, we plan to develop new functionality primarily for InterSystems IRIS.

Given the importance and urgency customers have expressed for migrating Caché and Ensemble applications to InterSystems IRIS, InterSystems has provided functionality and material that makes this transition as easy as possible, including supporting in-place conversion of applications, beginning with InterSystems IRIS 2019.1.1.

This document provides only a summary of new InterSystems IRIS features, with limited details. For additional information about each of these features, please contact your InterSystems Sales Engineer.



Performance & Scalability



Performance

InterSystems IRIS delivers significant performance and efficiency improvements compared to Caché and Ensemble. Queries run on InterSystems IRIS, in general, are at least 25% faster than the same queries run on Caché or Ensemble, and often much faster, based on real-world tests using existing applications.¹

As such, InterSystems IRIS enables applications to analyze more data from more sources faster, which in turn means producing more accurate real-time insights and actions.

For example, InterSystems IRIS can be used to analyze fraudulent transactions or execute compliance rules on a stock trade using machine learning algorithms and business intelligence in real time, not after the transaction or trade has gone through.

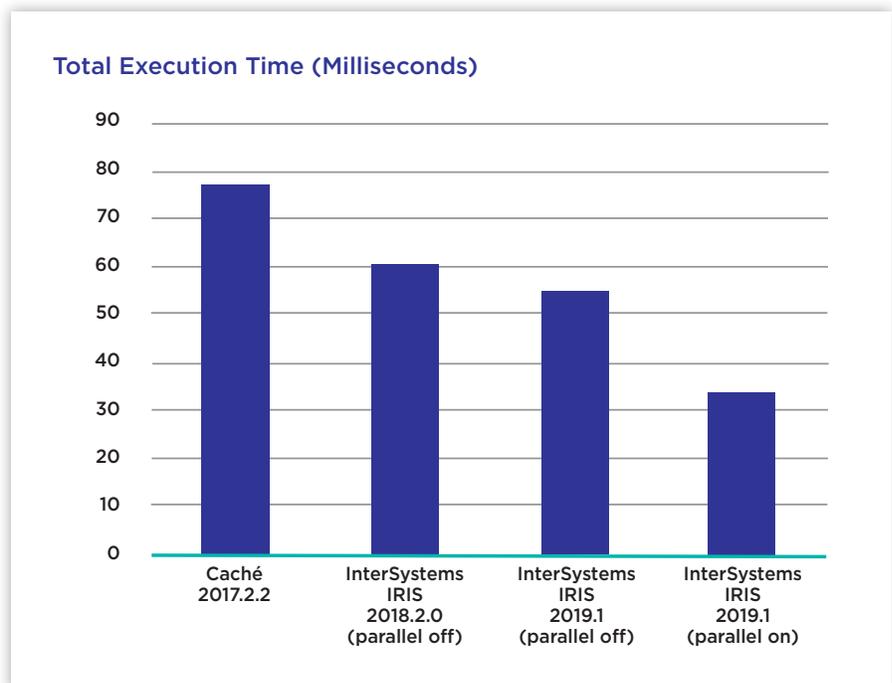
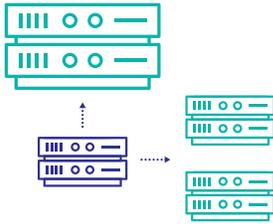


Figure 1: Time to execute complex query

¹www.intersystems.com/resources/detail/sql-performance-benchmark-of-intersystems-iris-data-platform



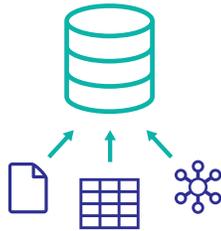
Scalability

InterSystems IRIS introduces significant scalability enhancements.

When scaling vertically, InterSystems IRIS improves the use of memory and multiple CPUs for each individual query through improved and automated parallel processing as well as lower-level engine optimizations.

When there is a need to scale horizontally, InterSystems IRIS introduces the ability to split a large database table; such that portions of that table are stored in independent and separate instances of InterSystems IRIS, through a mechanism called sharding. Each of these “data nodes” has its own CPU and memory resources at its disposal. Queries on sharded tables are executed in parallel on each node, dramatically reducing the time to produce a query result.

Even for applications with datasets that are currently small, we are seeing a common trend towards growing data volumes and the ability to turn more data into meaningful actions in real time.



Document Database Model

InterSystems IRIS can now define a new type of database model – the document database model.

A document, in this context, is a container holding an arbitrary set of properties and collection of properties. Typically, mobile applications retrieve data from the user interface and send this to the database in a JSON document for persisting into the database.

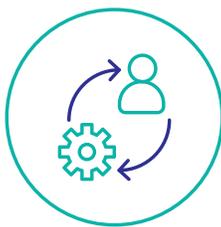
Building applications that utilize the new document database model in InterSystems IRIS is a quick way of building solutions. InterSystems IRIS also provides a REST API for creating, reading, updating, and deleting documents from a document database so that you do not need to build an API for your solution.

Traditionally data is mapped to a SQL table schema or an object model schema, where the data values are checked for validity (e.g. date format, or that required fields are supplied, etc.) before being saved; this is not the case with document databases.

In this aptly named “schema-free” approach, the document is saved as a stream, and just given an ID. No validation is typically performed. This allows for a future version of the same document to have new properties included in the document as enhancements are made to the application. (Note: typically data verification is performed on the client before a document is submitted to the server).

Documents inserted into an InterSystems IRIS database can be indexed and projected as SQL tables so traditional SQL queries can be used to search for collections of relevant document data.

Optimized Client-Side Programming



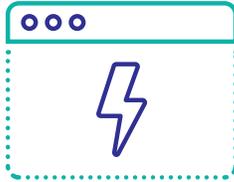
Caché and Ensemble solution builders have always had a broad set of technology choices to build their client-side interfaces and mid-tier applications. These different language environments then connect to the back-end Caché or Ensemble applications via drivers to retrieve or post data.

InterSystems IRIS rationalizes these language connectivity capabilities and, more importantly, improves the performance of the communication between client and server.

Supported languages include Java, .NET, Node.js, and Python, giving developers a choice of popular technologies upon which to base their clients. External programs can connect to InterSystems IRIS using a choice of connectivity options; for example: xDBC and ADO.NET, Hibernate and Entity frameworks, XEP, or Native/Direct access.

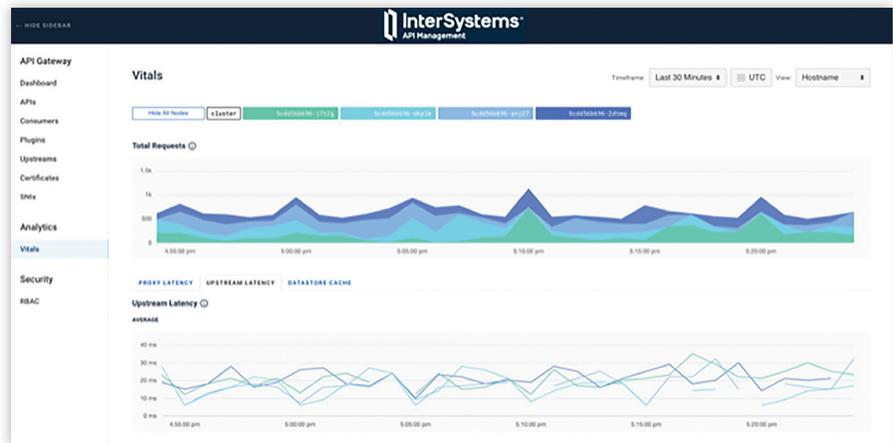
InterSystems IRIS also supports mobile client development using JSON and REST web frameworks like Angular or REACT.

Many application design patterns leverage Application Programming Interfaces (APIs). InterSystems IRIS brings you the following:



API Management

InterSystems IRIS introduced an API management gateway that allows solution developers to manage their collection of REST APIs developed and hosted within InterSystems IRIS.



The first release of the API management gateway provides authentication, encryption, protection against Distributed Denial of Service (DDoS) attacks, as well as API filtering.

The API management gateway also provides logging, request monitoring and management, analysis of API use, and support for API versioning and “canary” testing, which is used prior to releasing new API versions.

Caché and Ensemble have had support for REST for some time; InterSystems IRIS introduces support for an API-first approach to development, which means the developer’s starting point is to design the REST API endpoints using the OpenAPI specification. The subsequent tasks are the back-end programming code that actually implements the functionality of the service in InterSystems IRIS, and any front-end GUI that is consuming the service (if required). This approach ensures effective agreement between the front-end and back-end teams as the API becomes a pivotal point for communicating and implementing functionality changes.

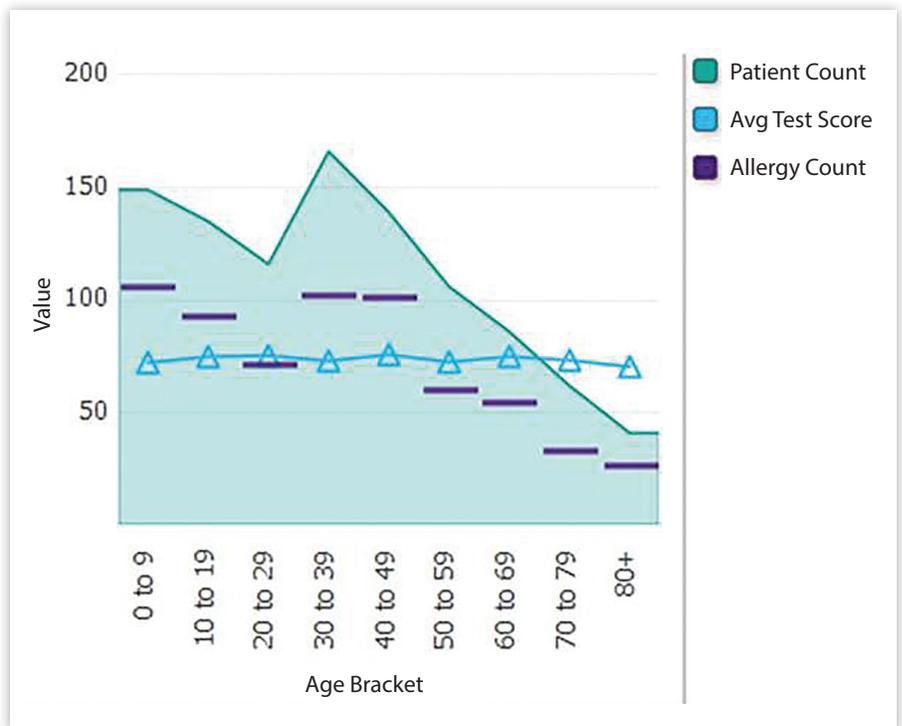
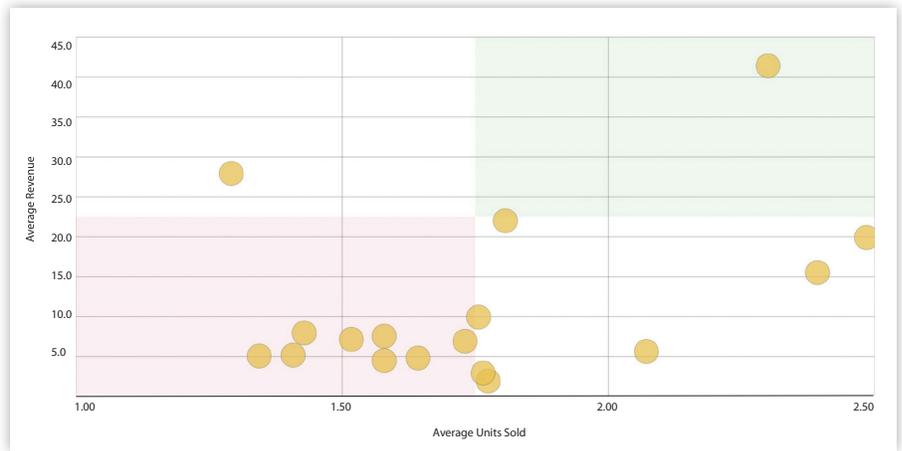
InterSystems IRIS also introduces additional capabilities that allow developers to build user interfaces against InterSystems IRIS REST APIs through the adoption of the OpenAPI specification.

You can now generate the OpenAPI specification for any REST APIs defined in InterSystems IRIS. This specification can be loaded into tools such as Swagger and other open source applications for developers to quickly develop solutions that use APIs.



Business Intelligence

Caché and Ensemble support the ability to model OLAP cubes using DeepSee, which provides the ability to embed real-time business intelligence and answer multi-dimensional analytical queries swiftly. ***With InterSystems IRIS, the performance of business intelligence queries have been enhanced, and the business intelligence capabilities are now included by default at no additional cost.***





Machine Learning

There is rising demand to incorporate advanced analytics including machine learning (ML) and artificial intelligence (AI) into solutions.

Analysts predict that applications that do not offer some level of AI will continue to lose market share to those applications that do. This is a trend occurring across all industries.

To defend against this competitive trend and incorporate AI and ML into solutions, large historical and real-time volumes of data need to be captured, integrated, cleansed, and made available to engines such as Apache Spark, and other environments in which data scientists build models.

The unique interoperability and scalability features of InterSystems IRIS can integrate, aggregate, normalize, and cleanse data from multiple sources, while also providing a distributed, scalable architecture to AI and ML modeling engines.

InterSystems IRIS is also ideal for performing feature engineering tasks that transform the data into formats and structures required by machine learning tools.

The InterSystems IRIS Apache Spark connector provides high throughput connectivity between Apache Spark and InterSystems IRIS. Using knowledge about a given configuration's data locations, and any implemented database shards, InterSystems IRIS is able to intelligently distribute computational workloads, enabling parallel processing of queries across multiple instances of InterSystems IRIS. ***The outcome is a more efficient use of system resources, increased throughput, reduced query times, and an improved user experience for data scientists.***

Machine learning models developed by data scientists using external tools can be seamlessly incorporated into InterSystems IRIS productions, where they are executed natively within the database via the Predictive Modeling Markup Language (PMML) industry standard. These models can be embedded into (real-time) business processes and run with real-time data, ***transforming your existing productions into AI-enabled solutions.***



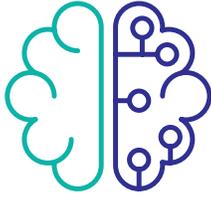
Text Analytics

InterSystems IRIS includes Natural Language Processing (NLP) capabilities for exploring unstructured data and extracting specific information and insight from it. It is differentiated from other approaches in that it uses a unique “bottom up” approach, discovering concepts and their context within the text itself, leveraging a deep understanding of the language, rather than a top-down view in a particular domain. There is no need for the tedious – and limiting – process of predefining dictionaries or ontologies.

This powerful capability, when used in applications directly or combined with traditional business intelligence capabilities, gives you the ability to analyze all the data, not just the structured data.

With InterSystems IRIS, NLP is now included by default at no additional cost. InterSystems IRIS also includes support for the Unified Information Management Architecture (UIMA) standard. This allows solution developers to embed InterSystems, third-party and custom NLP components in a single pipeline and easily manage and leverage the output in InterSystems IRIS applications.



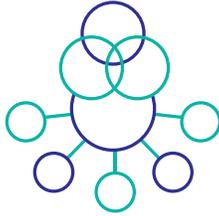


Third-Party Business Intelligence (BI) Connectors

InterSystems IRIS not only provides application developers with a choice of programming languages at the front end, but also a choice of third-party business intelligence tools to visualize and explore the data held within InterSystems IRIS. These integrations combine the advanced self-service data visualization capabilities of popular BI tools with the high-performance data access offered by InterSystems IRIS.

Microsoft (Power BI), Tableau, and Qlik are all making dedicated, high-performance connectors for InterSystems IRIS available through their channels.

Connected Applications



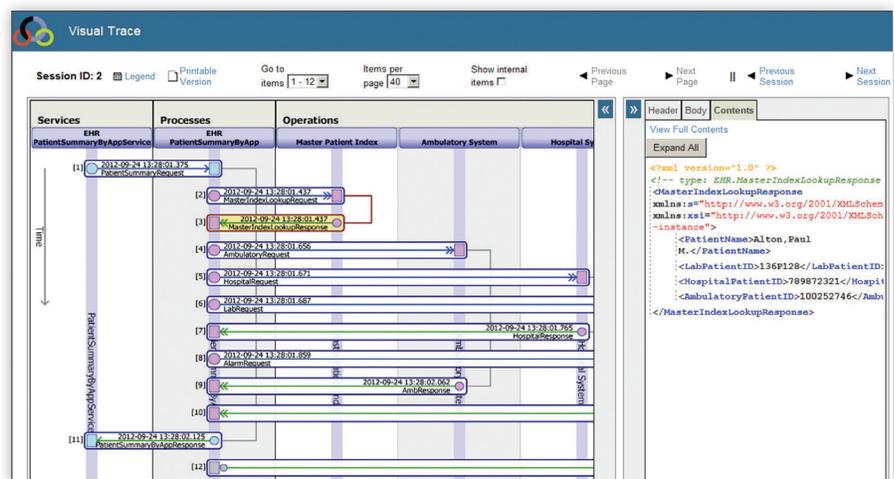
Interoperability

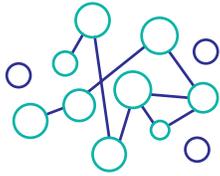
Nearly every application today needs to connect to other systems, whether simply to transfer data, offer new services that aggregate information from multiple sources or build seamless composite end to end workflows to automate processes. Therefore, we've incorporated Ensemble's integration "engine" and development framework for connected applications into InterSystems IRIS.

InterSystems IRIS takes a "low-code" approach to development that reduces the time it takes to deliver integrated solutions. It provides graphical and drag-and-drop editors for designing business processes, incorporating business rules and human workflow, and defining data and message transformations.

InterSystems IRIS manages all connection states, connection adapters (e.g. TCP, FTP, relational databases, REST), message queuing and payloads between IRIS and external applications and systems.

Messages are automatically saved and can be easily audited and visually tracked. Software developers do not need to develop programs that monitor inbound/outbound traffic, queues or message volumes. Nor do they need to worry about persisting with historical message content, message re-sending, re-routing, alerting, event logging, or even coordinating multiple simultaneous threads of executing process logic.



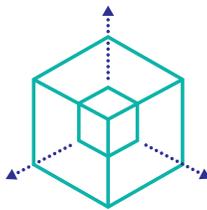
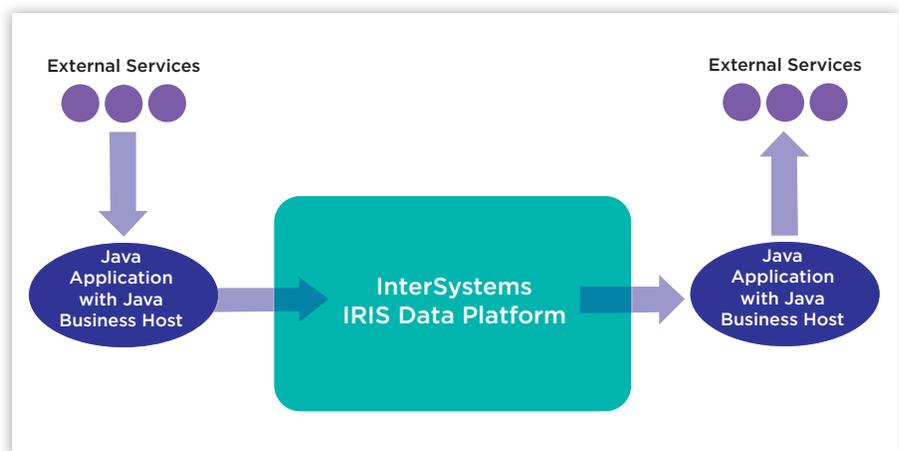


Java Business Hosts

Given the ever-changing nature of standards and communication protocols, specialized adapters may be needed from time to time.

Java Business Hosts expand Java developers' ability to leverage existing Java libraries and code to build new InterSystems IRIS interoperability components using Java. Once done, developers can import and use these Java-based components directly within the InterSystems IRIS integration platform.

Java Business Hosts greatly expand the library of applications and systems with which InterSystems IRIS can integrate and communicate.



Managed File Transfer

Managed File Transfer (MFT) is a general term referring to the features provided by popular cloud-based storage providers (e.g. DropBox, Box, or kiteworks) that allow customers to securely transfer a file from one party to another without needing to worry about or install certificates used in the encryption process.

Support for MFT saves developers from having to build their own mechanisms and certificate key management framework for the secure exchange of sensitive information.

InterSystems IRIS introduces connectors for DropBox, Box, and kiteworks so that when one party securely transfers a document to another party, InterSystems IRIS is aware of the event and has access to the contents of the document for processing.



Support for Container Technology

In accordance with today's rapidly increasing adoption of cloud infrastructure, secure packaged deployments, and the need to deploy solutions quickly and easily inside a container, InterSystems IRIS is supported and is also provided in the industry-standard OCI-compliant Docker container.

A container allows you to package your solution's software and all its required setup and configuration (including InterSystems IRIS software, your application code any dependencies, and other libraries needed) into an immutable "Docker image" that can be deployed and run anywhere.

Deployments using Docker images eliminate issues with incorrect configuration or missing libraries as applications are moved between environments.

Once an application is packaged into a container image, developers and testers can quickly and easily deploy and instantiate that image locally or on cloud infrastructure for interim testing or demonstration, and then easily destroy that instance once done.

Container images can be moved from one server (e.g. Development) to another (e.g. Test) or pushed to the cloud instances that can "start" the container image. Any image that is started is a copy of that image and is referred to as the running container.

Container images and containers are a critical component of modern infrastructure that enable easier software development life cycles, as organizations seek to quickly and incrementally introduce software enhancements and fixes on a daily basis for their solutions.



Cloud Support

InterSystems IRIS provides support for cloud deployments, and includes the InterSystems Cloud Manager (ICM). ICM is a tool that is used to automatically:

- **Create (provision) servers** in Microsoft Azure, Amazon Web Services, Google Cloud Platform. Configurations can include single-server architectures, mirrored server architectures or architectures with multiple hosts; and
- **Deploy software applications** into the provisioned hardware based on a configuration file that specifies a Docker image of a particular solution.

ICM can be used to quickly and automatically provision infrastructure in cloud or on-prem environments and to deploy applications into that environment.

InterSystems IRIS containers can also work with third-party provisioning tools such as Kubernetes. InterSystems is actively developing a series of scripts that will automate much of the work involved with using Kubernetes to deploy InterSystems IRIS solutions.

References

For additional information, please contact your InterSystems Sales Engineer, or visit InterSystems Learning Services at <http://learning.intersystems.com> and review the following suggested material:

[Introducing InterSystems IRIS Data Platform](#)

[What is InterSystems IRIS?](#)

[Learn InterSystems IRIS Data Platform](#)

InterSystems IRIS Data Platform Comparison

CAPABILITY	InterSystems IRIS	Ensemble	Caché
Cloud, On-Premise, Hybrid Deployment	✓		
Containerized Delivery	✓		
Full Lifecycle API Management	✓		
Optimized Client Side Programming Using xDBC, XEP, Native	✓		
Horizontal Data Scaling (Sharding)	✓		
3 rd Party BI Connectors (Tableau, PowerBI, Qlik)	✓		
Managed File Transfer	✓		
UIMA Integration	✓		
Spark Connector + AI Enablement Via PMML	✓		
Connectivity Adapters and Data Transformations	✓	✓	
Messaging Engine With Guaranteed Message Delivery	✓	✓	
Composite Business Process Orchestration	✓	✓	
Graphical Editor For Business Process Orchestration	✓	✓	
Multi-Workload Database (Transactions + Analytics)	✓	✓	✓
Multi-Model Database	✓	✓	✓
Business Intelligence (DeepSee)	✓	✓	✓
Text Analytics (iKnow)	✓	✓	✓
Row Level Security, Authentication, Authorization, and Auditing	✓	✓	✓

