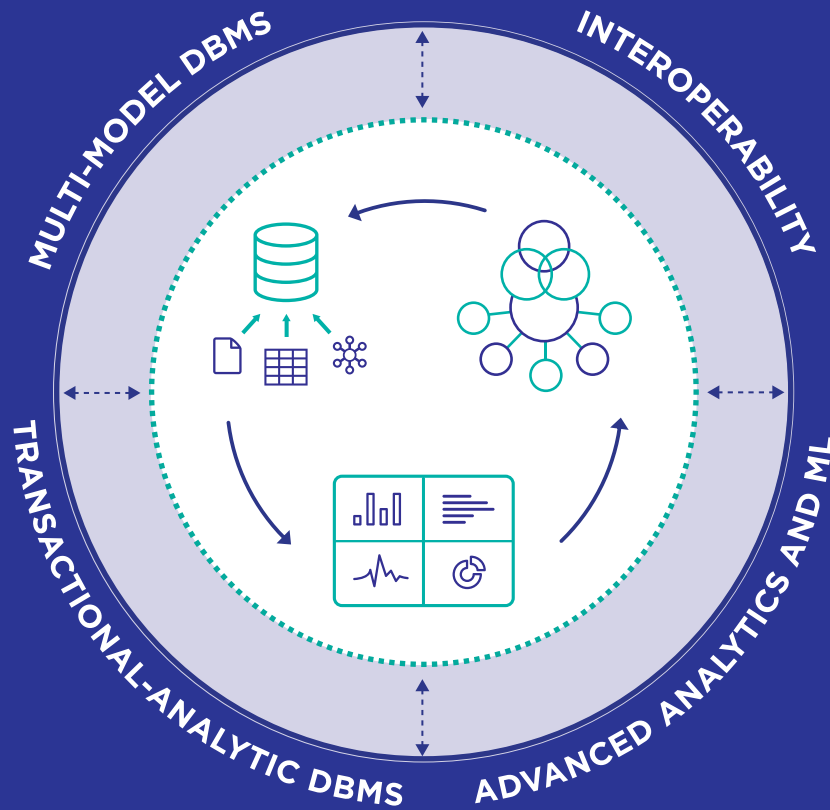


Migrate to the Next-Generation Data Platform: InterSystems IRIS and InterSystems IRIS for Health

For Current Caché and Ensemble Customers



INTUITIVE
|
RELIABLE
|
INTEROPERABLE
|
SCALABLE

Contents

Introduction	3
Performance and Scale: Run Faster, Bigger, and Leaner	4
Deployment: Cloud, On-Premises, or Hybrid	6
Development: API-First, in the Language of Your Choice	8
Analytics and AI: Build Smarter Applications	10
Integration and Interoperability: Field Connected Applications	14
Healthcare-Specific Functionality: InterSystems IRIS for Health	15
References	15

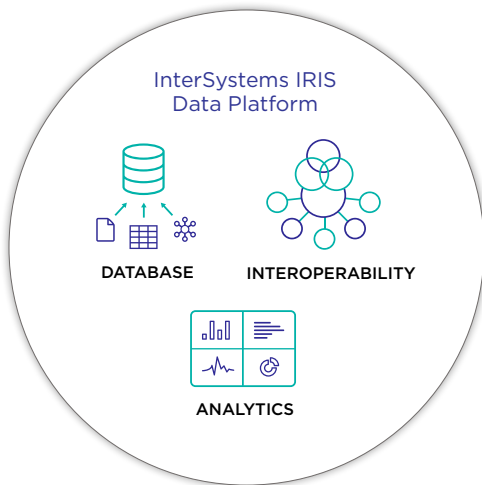
Introduction

InterSystems IRIS® and **InterSystems IRIS for Health™** are the newest generation of our data management software, developed to provide our customers with powerful technology upon which they can build modern, high-performance data- and analytics-intensive applications.

Today's applications are becoming smarter, faster, more API-driven, and more integrated. They incorporate more and different kinds of analytics and leverage cloud and container technology.

For organizations to survive and grow, their technology must enable them to keep pace with these trends as competitors introduce new applications and revise their existing ones to offer new capabilities. InterSystems IRIS is the data platform software that empowers organizations across industries to capture new market opportunities and remain competitive and relevant. InterSystems IRIS for Health extends InterSystems IRIS to help organizations harness the complexity of connecting, managing, and transforming healthcare data.

This document was created to inform current InterSystems Caché® and InterSystems Ensemble® customers about new functionality offered in InterSystems IRIS and InterSystems IRIS for Health. It is meant to help you generate ideas and plans for revising and refreshing your current applications and creating new ones by leveraging new features and new kinds of analytics for even greater performance and scalability.



For the sake of brevity, throughout this document we will use “InterSystems IRIS” to refer to both InterSystems IRIS and InterSystems IRIS for Health. All sections are applicable to both products, and there is a section at the end with healthcare-specific information about migrating to InterSystems IRIS for Health.

Migrating to InterSystems IRIS

InterSystems IRIS was first introduced in 2018 and is now well established in the market, with hundreds of customers in production. While InterSystems® is committed to supporting Caché and Ensemble for our existing customers and partners, we are focused on InterSystems IRIS for developing new functionality.

Because customers have communicated the importance and urgency of migrating Caché and Ensemble applications to InterSystems IRIS, InterSystems provides tools to make this transition as easy as possible, with support for in-place migration of applications.

This document summarizes important new features and capabilities available with InterSystems IRIS. For additional information about each of these features, please review the documentation and training material, and contact your InterSystems account team.



Performance

InterSystems IRIS delivers significant performance and efficiency improvements compared with Caché and Ensemble. Applications accessing Globals directly will benefit from many scale and performance enhancements at the kernel level, including node tables, which provide fully transparent optimization for cached data blocks. Tests have shown that node tables can speed up random global access and reverse Global traversal by 20% and up to 75%, respectively.

The InterSystems IRIS SQL engine has also incorporated many significant enhancements since Caché 2018.1, including improved and fully automated parallelization, optimizations to code generation, and advanced use of statistics to bring you the best query plans. SQL queries run on InterSystems IRIS in general are at least 25% faster—often significantly faster—than the same queries run on Caché or Ensemble on the same hardware, based on tests performed using a real-world application.¹

InterSystems IRIS enables applications to analyze more data from more sources faster, which in turn means 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, business intelligence, and business rules in real time rather than after the transaction or trade has executed.

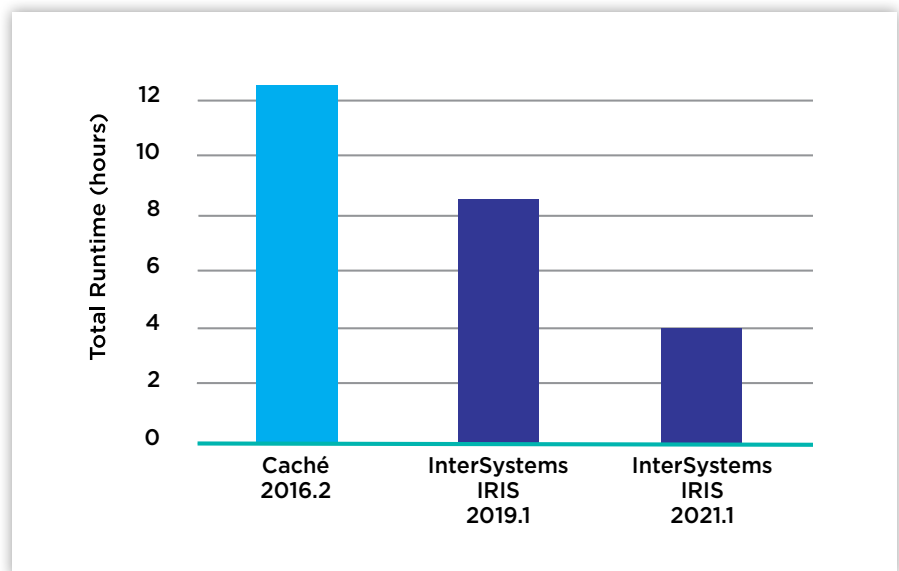



Figure 1: Times to execute complex queries (Lower is faster)

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



Scalability

InterSystems IRIS introduces significant scalability enhancements.

Vertical scalability is all about making the best use of the available system resources and has always been an InterSystems strength. With InterSystems IRIS, we continue to invest in making sure it can leverage all available memory, CPU, and IO bandwidth on the most advanced hardware platforms through enhancements at the kernel level, including more efficient journaling and dejournaling.

When scaling SQL applications, InterSystems IRIS improves the use of memory and multiple CPUs for individual queries through improved and automated parallel processing. The work queue manager infrastructure supporting this functionality on InterSystems IRIS is also available for ObjectScript developers, for increased convenience, robustness, and efficiency for applications that relied on manual use of the JOB command on Caché and Ensemble.

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 separate, independent 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 required to produce a query result.

This new capability complements InterSystems Enterprise Cache Protocol (ECP) functionality, which enables applications to scale the compute tier by adding ECP application servers, thus building a distributed cache to serve very high user counts efficiently. InterSystems IRIS improves ECP in a number of ways, especially for the most demanding deployments supporting tens of thousands of concurrent users.

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



Support for Container Technology

InterSystems IRIS is available as an OCI-compliant Docker container, and supports customer-built containers including InterSystems IRIS as well. Containers are being adopted rapidly for cloud deployments, and they also provide portability and DevOps advantages for on-premise and device-based deployments.

A container allows you to package your solution's software and required setup and configuration (including InterSystems IRIS software, your application code any dependencies, and other libraries) into an immutable "image" that can be deployed and run on any OCI-compliant engine.

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 a cloud infrastructure for interim testing or demonstration. Once they are done, they can then easily destroy that instance.

Container images can be moved from one server (such as development) to another (such as testing) or pushed to the cloud instances that can "start" a container based on that image.

Container images and containers are a critical component of modern DevOps workflows and infrastructure. They enable easier software development lifecycles as organizations seek to quickly and incrementally introduce software enhancements and fixes for their solutions on a daily basis.

The InterSystems Kubernetes Operator

Kubernetes is an open-source platform for managing containerized workloads and services. It facilitates both declarative configuration and automation and has a large, rapidly growing ecosystem. Kubernetes services, support, and tools are widely available. It is available as a managed service from every major public cloud.

The InterSystems Kubernetes Operator (IKO) packages InterSystems IRIS-specific knowledge and best practices into an easy-to-use automated tool. IKO extends the Kubernetes API with a custom resource representing an InterSystems IRIS sharded cluster. This resource, called an IrisCluster, can be deployed on any Kubernetes platform on which the IKO is installed.

InterSystems IRIS support for container technology, and the IKO in particular, enable a broad range of DevOps practices, from simplifying test setups to deploying fully automated CI/CD pipelines.



Cloud Support


InterSystems IRIS provides support for cloud deployments and includes the InterSystems Cloud Manager. Cloud Manager is a tool that is used to automatically do the following:

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

Cloud Manager can be used to quickly and automatically provision infrastructure in cloud or on-premises environments and deploy applications into that environment. Its “all-in” approach enables InterSystems customers who are new to cloud and container technology to get started quickly and deploy their solutions using best-practice configurations.

Development: API-First, in the Language of Your Choice

API Management



InterSystems IRIS includes InterSystems API Manager, a full lifecycle API management platform that allows solution developers to manage their collection of REST APIs developed and hosted within InterSystems IRIS.



InterSystems API Manager provides authentication, encryption, protection against distributed denial-of-service attacks, and API filtering.

It also provides logging, request monitoring and management, analysis of API use, and support for API versioning and “canary” testing, which is used before new API versions are released. API Manager has exceptionally high usage among InterSystems IRIS customers and has been a key driver for existing Caché and Ensemble customers to migrate to InterSystems IRIS.

Caché and Ensemble have provided 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 applications for developers to quickly develop solutions that use APIs.



Enhanced Client- and Server-Side Programming

Caché and Ensemble solution builders have always had a broad set of technology choices to build their client-side interfaces and middleware 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 they can base their client applications. External programs can connect to InterSystems IRIS using a variety of connectivity options, including ODBC, JDBC, ADO.NET, Hibernate, Entity Framework, XEP, or native direct access.

InterSystems IRIS also supports mobile client development using JSON and REST web frameworks such as Angular and React.

Recent releases of InterSystems IRIS have also significantly strengthened support to implement server-side business logic in Java, .NET, Python, and R through External Language Gateways and invoke it directly from ObjectScript, SQL, or Interoperability business processes. The gateway technology enables solution developers to work in their technology of choice at every level of their application, with guaranteed highest throughput and robustness.

Analytics and AI: Build Smarter Applications



InterSystems IRIS provides many different types of analytic capabilities to support the varied requirements of our customers and partners. InterSystems IntegratedML™ and natural language processing (NLP) are available to all InterSystems IRIS customers, along with direct access through popular third-party business intelligence (BI) tools via ODBC and JDBC. Additional capabilities—including InterSystems Reports, InterSystems IRIS Business Intelligence, and InterSystems IRIS® Adaptive Analytics—are available as options with InterSystems IRIS.

Adaptive Analytics

Adaptive Analytics is an optional extension that makes InterSystems IRIS even more powerful by providing a business-oriented virtual data model layer between InterSystems IRIS and popular BI and artificial intelligence (AI) client tools. It enables data stewards and business users to more easily explore and analyze data without relying on IT.

Adaptive Analytics includes an intuitive user interface for developing a data model in the form of a “virtual cube” where data can be organized, calculated measures can be consistently defined, and data fields can be clearly named. Having a centralized common data model enables enterprises to solve the problem of differing definitions and calculations, so that end users have a single, consistent view of business metrics and data characterization.

Key capabilities of Adaptive Analytics include the following:

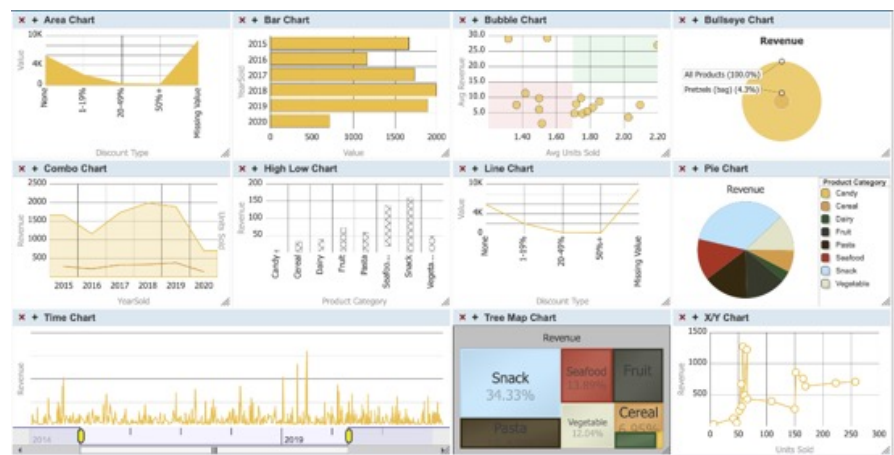
- Data stays in InterSystems IRIS, eliminating the complexity and performance issues associated with moving or copying data.
- Data stewards can use the Adaptive Analytics modeler to make data accessible for business users—without having to expose complex data structures, tables, or relationships.
- Changes to the data model are published as virtual cubes without disrupting users—so there is no need to wait for the rebuilding of cubes.
- Analytics users can employ the BI tools of their choice, such as Microsoft Excel, PowerBI, and Tableau, and access the same online analytical processing (OLAP) model.
- Adaptive Analytics uses the full breadth of data stored within InterSystems IRIS through live connectivity rather than partial content or stale data extracts.
- Adaptive Analytics provides a single layer to govern data access and protects sensitive data from unauthorized access.



Business Intelligence

There is rising demand to incorporate advanced analytics including Caché and Ensemble support the ability to model OLAP cubes using InterSystems IRIS Business Intelligence (formerly known as InterSystems DeepSee), which provides the ability to embed real-time business intelligence and answer multidimensional analytical queries quickly. The availability of this OLAP engine inside the data platform means application partners can include analyses and dashboards in their applications without requiring additional server- or client-side installations. This speeds and simplifies development and deployment.

With InterSystems IRIS, the performance of BI queries has been enhanced, and the new Selective Build feature enables you to change your dimensional models without the downtime traditionally associated with such redesign work.



Machine Learning



There is increasing demand to incorporate advanced analytics, including machine learning (ML) and AI, into solutions.

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

Building accurate ML models requires the unique skills of data scientists, one of the most sought-after job titles in the marketplace today. And if acquiring data science talent weren't hard enough, embedding or *operationalizing* the ML models into production applications is challenging as well, as illustrated by various analysts reporting that more than half of models built never make it into production.

Both these challenges are addressed through InterSystems IntegratedML, an automated ML capability embedded within InterSystems IRIS. IntegratedML provides an easy-to-use SQL interface so application developers can quickly build and test ML models without any upfront data science training. Leveraging these models within an application requires

only a simple SQL function to be called, with InterSystems IRIS taking care of all the wiring and model management that normally complicate this operationalization task.

InterSystems IntegratedML changes the game for in-application ML, empowering application developers to build and embed ML models with unprecedented speed and ease of use.

IntegratedML leverages additional compute resources to perform the feature engineering and iterative optimizations that would traditionally fall upon data scientists, but it does not replace them. It helps data teams build initial models quickly so the predictive value of a dataset for a given problem can be assessed with very little effort. Your skilled data scientists can then select the hard problems that require their skills and experience to take model accuracy to the next level, so in fact IntegratedML increases the throughput of your teams by making efficient use of available skills.


To enable data scientists to work more efficiently with large datasets managed by InterSystems IRIS, it has an Apache Spark connector that provides high throughput connectivity between Apache Spark and InterSystems IRIS. Apache Spark is commonly used by data scientists to build data and ML pipelines when working with large set of data, so this throughput is paramount. 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 the data platform software.

The outcome is a more efficient use of system resources, increased throughput, reduced query times, and an improved user experience for data scientists.

IntegratedML by default leverages InterSystems own AutoML engine, which is based on well-known open source packages such as TensorFlow, SciKit-learn, and others. Users can also opt for third-party engines by H2O and DataRobot, two leading providers of AI, which are available out-of-the-box, following the same simple SQL-based API described above.

InterSystems also partners with DataRobot to enable customers to tap into DataRobot's cloud-based platform for Enterprise AI and keep those efforts integrated with their InterSystems IRIS-based applications.

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 engine via the PMML (Predictive Model Markup Language) industry standard. These models can be embedded into real-time business processes and run with real-time data with high performance, transforming your existing productions into AI-enabled solutions.



Text Analytics

InterSystems IRIS includes NLP capabilities for exploring unstructured data and extracting specific information and insights from it. It is different from other methods in that it uses a unique bottom-up approach, discovering concepts and their context within the text itself and leveraging a deep understanding of the language, rather than taking a top-down approach focused on 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 BI capabilities, gives you the ability to analyze data from many different sources, not just structured data.

With InterSystems IRIS, NLP is now included by default at no additional cost. InterSystems IRIS also includes support for the UIMA (Unstructured Information Management Architecture) 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.

InterSystems Reports




InterSystems IRIS includes InterSystems Reports, a robust modern reporting solution for customers who need to distribute and view reports and forms. InterSystems Reports includes the following:

- Embedded operational reporting, which can be customized by both report developers and end users.
- Pixel-perfect formatting that lets you develop highly specific form grids or other special layout elements for invoices, documents, and forms.
- Banded layouts that provide structure for aggregated and detailed data.
- Exact positioning of headers, footers, aggregations, detailed data, images, and sub-reports.
- A variety of page report types.
- Large-scale dynamic report scheduling and distribution, including export to PDF, XLS, HTML, XML, and other file formats; printing; and archiving for regulatory compliance.

Integration and Interoperability: Field-Connected Applications

Interoperability

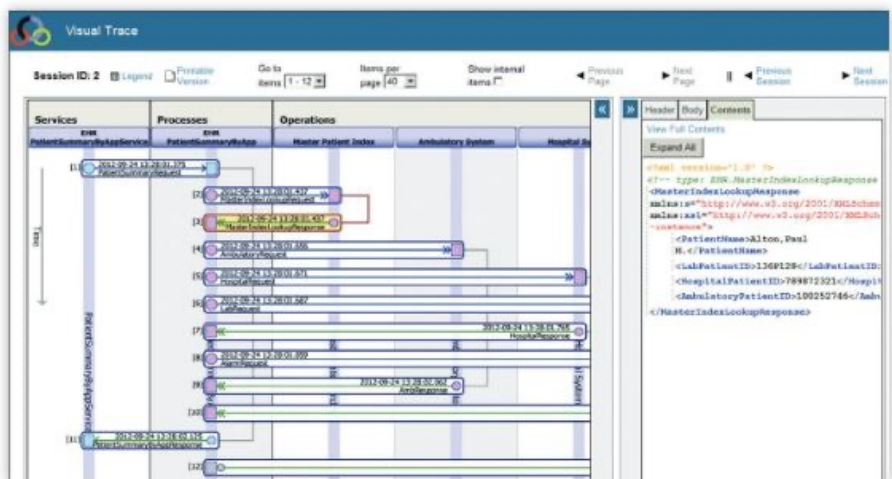



Nearly every application today needs to connect to other systems, whether simply to transfer data, offer new services that aggregate information from multiple sources, support business users with consistent data from different systems, or build seamless composite end-to-end workflows to automate processes.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

InterSystems IRIS provides a "low-code" approach to interoperability 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 includes a number of new and enhanced interfaces not available in Ensemble that conveniently help users manage the most complex setups.

InterSystems IRIS manages all connection states, connection adapters (including TCP, FTP, relational databases, and REST), message queuing, and payloads between InterSystems 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 historical message content; message resending, rerouting, alerting, and event logging; or even coordinating multiple simultaneous threads of executing process logic.





Production Extensions

Given the ever-changing nature of standards and communications protocols, specialized adapters may be needed to connect to external sources and systems.

The InterSystems IRIS Production EXtension (PEX) framework expands Java and .NET developers' ability to leverage existing libraries and code to build new InterSystems IRIS interoperability components using the Java and .NET programming languages.

PEX provides flexible connections across business services, processes, and operations that are implemented in PEX-supported languages or in InterSystems ObjectScript. In addition, PEX can be used to develop inbound and outbound adapters. The PEX framework allows application developers to create an entire production in Java or .NET or to create a production that has a mix of Java, .NET, or ObjectScript components. Once they are integrated, the production components written in Java and .NET are called at runtime and use the PEX framework to send messages to other components in the production.

PEX greatly expands the library of applications and systems with which InterSystems IRIS can integrate and communicate.

Healthcare-Specific Functionality: InterSystems IRIS for Health




InterSystems IRIS for Health extends InterSystems IRIS with the built-in ability to transform healthcare data from one format to another. InterSystems supports HL7 FHIR Release 4 (R4), the latest healthcare standard for interoperability. FHIR R4 enables developers to create innovative healthcare applications to unify data inside and outside the hospital. InterSystems FHIR support includes FHIR profiling and packages, FHIRPath, and FHIR R4 transformations from older healthcare standards, such as HL7v2. The HL7 Migration Tooling supports migrating transformation logic from other interface engines, including Cloverleaf, DataGate, and eGate to Health Connect.

InterSystems also provides X12 schema support, including element validation for improved data quality, which opens up opportunities for billing application support. InterSystems can route and transform X12 messages using new X12 business services and business operations.

InterSystems IRIS for Health also supports a variety of IHE profiles.

References



For additional information, please contact your InterSystems Sales Engineer or visit [InterSystems.com/Migrate](https://www.intersystems.com/Migrate).

Compare InterSystems Product Capabilities

CAPABILITY	InterSystems IRIS for Health	InterSystems IRIS	Ensemble	Caché
Normalized Healthcare Message Model	✓			
Pre-built Transformations between Healthcare Messaging	✓			
Advanced FHIR Development	✓			
FHIR Resource Repository	✓			
Enhanced Standards Support – FHIR and IHE	✓			
Healthcare Interface Engineer Productivity Tools	✓			
Flexible Deployment: On-Premises, Cloud, Multi-Cloud, or Hybrid	✓	✓		
Software Delivery with Containers, Kubernetes	✓	✓		
Enhanced Client-Side Development, ODBC, JDBC, Java, .NET, Python, and Node.JS	✓	✓		
Machine Learning with IntegratedML	✓	✓		
Self-Service Analytics with Adaptive Analytics	✓	✓		
Apache Spark Connector and Data Transformations	✓	✓		
Full Lifecycle API Management	✓	✓		
Pixel-perfect Reporting	✓	✓		
Horizontal Scaling for Data and Compute	✓	✓		
Graphical Editor for Business Process Orchestration	✓	✓	✓	
Composite Business Process Orchestration	✓	✓	✓	
Messaging Engine with Guaranteed Message Delivery	✓	✓	✓	
Connectivity Adapters and Data Transformations	✓	✓	✓	✓
Text Analytics	✓	✓	✓	✓
Embedded Business Intelligence	✓	✓	✓	✓
Row-Level Security, Authentication, Authorization, and Auditing	✓	✓	✓	✓
Multi-Model Database	✓	✓	✓	✓
Multi-Workload Database (Transactions and Analytics)	✓	✓	✓	✓

