

InterSystems IRIS for Health

Technology Brief



**FHIR-Based Healthcare Application Development.
Advanced Interoperability.
Unprecedented Speed.
Massive Scale.**



Introduction

InterSystems IRIS for Health™, a comprehensive, cloud-first data platform, along with its predecessors, is the basis for healthcare software that manages over a billion health records globally. Its use is broad and deep across the healthcare industry, and our own healthcare product offerings are based on it.

InterSystems IRIS for Health, an extension of the InterSystems IRIS Data Platform is at the heart of commercial healthcare solutions such as Epic's EMR and Solventum Encompass 360 (formerly 3M). Clinical laboratory companies such as Sonic, Labcorp, and Bioreference have based their internal management systems on our platform. It's used by startups and established healthcare delivery and medical technology organizations, payers, pharmaceutical companies and others as a foundation for innovation through its connectivity, interoperability, and application development capabilities.

The Core of a Smart Data Fabric


Recently, smart *data fabrics* have emerged as an IT architectural approach that goes beyond previous designs used to gain access to data silos across an organization. Smart data fabrics combine technology to integrate, transform, and harmonize data from disparate sources on demand, and leverage that data in applications through a wide range of analytics capabilities.

For an increasing number of organizations and developers, InterSystems IRIS for Health forms the core of a smart data fabric that includes analytics and data exploration, business intelligence, natural language processing, complex business rules, and machine learning. Developers, researchers, and organizations across the healthcare continuum are using it to gain new insights from data and power a transformation where patients are at the center of every service provided.

InterSystems IRIS for Health

InterSystems IRIS for Health provides many of the capabilities required to implement real-time, smart data fabric architectures in a single product, eliminating the need to implement, integrate, and maintain dozens of different technologies. Key components of InterSystems IRIS for Health include:

- A new approach to data and application integration
- Seamless healthcare interoperability via HL7® FHIR® and other standards
- Advanced FHIR capabilities
- Embedded advanced analytics tools
- Support for application development



Data and Application Integration: A New Approach

Nearly every application today needs to connect to other systems, whether to simply transfer data, or to offer new services that aggregate information from multiple sources, support users with consistent data from different systems, or build seamless composite end-to-end workflows to automate and optimize processes.

Connect and Collect

InterSystems IRIS for Health provides a low-code approach to data and application integration that supports connect and collect approaches. Connect enables new applications to access information on demand, without creating additional copies of the data. Of course, data can be stored within the InterSystems IRIS for Health multi-model, multi-workload database (collect), and applications can leverage both approaches for the highest performance and resource efficiency.

High Performance, Multi-Model, Multi-Workload Database Management

At the core of InterSystems IRIS for Health is an ultra-high performance, multi-model, multi-workload database management engine that supports vertical and horizontal scalability. It ingests, processes, and stores transactional data at high rates while simultaneously processing high volume analytic workloads involving historical and real-time data (including ACID-compliant transactions).

Multi-Model

Within the InterSystems IRIS for Health database, data is stored once and can be accessed as tables, objects, documents, key-value pairs, or multidimensional arrays without duplicating data or executing performance-killing mapping between models. All access methods can be used simultaneously on the same data with full concurrency. This pure approach to multi-model database management allows developers to use the most appropriate model types for their applications within a single environment.

Multi-Workload

InterSystems IRIS for Health is optimized for real-time applications that require high-throughput data ingestion with high-performance, concurrent analytics at scale. It is deployed in mission-critical healthcare applications that must have the capacity to ingest thousands or millions of records per second while simultaneously querying the incoming data in real time.

Proven Scalability Technology

InterSystems IRIS for Health enables applications to scale efficiently to meet the needs of your application – horizontally through sharding and our Enterprise Cache Protocol, and vertically through parallel SQL queries. Whether it's high volumes of concurrent users, a massive dataset to analyze, or both at the same time, InterSystems IRIS for Health scales out to a distributed architecture on affordable hardware to reduce total cost of ownership.

Vertical Scalability

Vertical scaling takes advantage of bigger, multicore machines through the efficient and fully automated use of parallelization, which enables organizations to right-size infrastructure resources in the cloud to achieve optimal price-performance ratios. However, big, multicore machines can quickly become cost prohibitive, both in the cloud and on premises.

Horizontal Scalability


InterSystems IRIS for Health addresses this “big machine cost” issue with InterSystems Enterprise Cache Protocol (ECP). ECP enables horizontal scaling of the number of users by caching data on application servers. User queries are satisfied from the local application server cache, if possible, retrieving data from the data server only if necessary. ECP automatically synchronizes the data and is entirely transparent to users and applications, providing superior performance and resource efficiency as workloads increase.

InterSystems IRIS for Health also enables horizontal scaling of data through sharding — splitting up very large databases onto multiple machines. Queries are run in parallel on each shard, and the results are aggregated before being returned to the user.

Sharding and ECP are transparently combined so applications can handle large volumes of data and high volumes of compute workloads efficiently and independently.

Flexible Scalability for Your Business

You don't have to make a large investment in InterSystems IRIS for Health to get the performance and scalability your business needs. Flexible licensing and deployment options, and our Smart Health Data Services in the cloud, enable you to address current requirements with the technology you need now, and easily scale as your needs change.



Healthcare Interoperability

In fragmented healthcare systems, data is stored in multiple sources and in a variety of formats and standards, creating silos that are difficult to integrate. Even when sources make data available via an API, those APIs aren't compatible with one another. To ensure you can leverage all this data in your solutions, InterSystems IRIS for Health offers extensive development tooling for healthcare interoperability standards and templates.

These include:

- HL7® FHIR® (STU3, R4)
- HL7® V2
- IHE Profiles, including XDS.b, XCA, PIX, PDQ, and MHD
- CDA/C-CDA Documents
- DICOM
- X12

An individual's health record typically needs to be brought together from different systems using different standards, so developers often need to transform data from one format to another. InterSystems IRIS for Health provides pre-built extensible data transformations between modern and legacy data representations. For example, using out-of-the box InterSystems IRIS for Health features, you can:

- Transform an HL7 V2 message from one schema version to another, including HL7 FHIR
- Produce HL7 V2 messages from a CDA document
- Transform a portion of a CDA document into HL7 FHIR resources



Advanced FHIR Capabilities

InterSystems support for HL7 FHIR is broad and deep. InterSystems IRIS for Health provides the building blocks needed to work with FHIR data and develop FHIR applications, including:

- A base FHIR server supporting standard FHIR REST APIs
- A FHIR resource repository fully supporting the FHIR Standard
- Data access authorization via the standard OAuth protocol
- Built-in data transformations between FHIR and other healthcare interoperability standards such as HL7 V2 and CDA
- Support for FHIR-based IHE Profiles, including PIXm, PDQm, and MHD
- Bulk FHIR
- FHIR SQL Builder

The InterSystems IRIS for Health FHIR repository offers full read-write capabilities, receiving or sending FHIR resources or bundles via the FHIR RESTful API in JSON or XML formats. This allows applications built on the latest technologies to use FHIR data — new or mapped from legacy systems — for patient care, quality improvement, research, and other use cases.

Bulk FHIR

One of the most useful capabilities that have been added to the FHIR Standard is Bulk FHIR import and export. Bulk FHIR is important for several reasons:

- **Efficiency:** Bulk FHIR enables the retrieval of large datasets of FHIR resources in a single request, which can significantly improve the efficiency of data exchange and management.
- **Population health:** Bulk FHIR supports population health use cases by enabling the analysis of large datasets to identify trends, patterns, and gaps in care.
- **Research:** Bulk FHIR supports research use cases by enabling the retrieval of large volumes of data for analysis and study.
- **Reporting:** Bulk FHIR supports quality reporting and compliance use cases by enabling the retrieval of large quantities of data for reporting purposes.
- **Interoperability:** The Bulk FHIR API facilitates interoperability by enabling the exchange of large quantities of data between different healthcare systems, which can support care coordination and continuity of care.

InterSystems FHIR server not only supports Bulk FHIR export but includes an intuitive Bulk FHIR client for easily exporting FHIR data from external FHIR servers. For FHIR servers that do not support the Bulk FHIR API, InterSystems offers pure FHIR fetch, which enables customers to get data from FHIR servers that support the operator \$everything.

FHIR SQL Builder

FHIR is foundational to health data exchange but is poorly suited to analytics. FHIR data is represented as a complex directed graph, which is difficult to model in a relational schema. Nevertheless, business intelligence tools require a relational representation of data to build cubes, dashboards, and reports. FHIR data must be transformed, “projected”, to make this relational representation possible.

Why is FHIR analytics so important in healthcare? Because it is the foundation of population health management, decision support, and operational/quality improvement programs. InterSystems IRIS for Health FHIR SQL Builder transforms the FHIR directed graph into a relational schema for access by analytics applications.

InterSystems recognizes that there is no one relational projection all customers would agree on. The IRIS for Health FHIR SQL Builder guides users toward the optimal projection for each use case.

The FHIR SQL Builder :

- Analyzes the FHIR graph, and reports on the number of resources, links between resources, and number of elements for each resource
- Gets user input on what resources, links, and elements are important for building their schema
- Projects the FHIR data to a relational schema of the customer's choice, so there is no need to move the data
- The FHIR SQL Builder takes advantage of the unique capability of InterSystems IRIS for Health FHIR Repository

Note: InterSystems IRIS for Health Advanced Server is required



Analytics

Embedded and Open Analytics


InterSystems IRIS for Health provides a range of powerful, built-in analytics capabilities that run against data stored in its database. Users gain rapid access to insights with the ability to run a wide variety of analytics directly within applications close to the data. Because data in FHIR format is by nature nonrelational, we offer an option — FHIR SQL Builder — that gives data scientists access to it using SQL. Data managers do not need to translate and move the data in the FHIR repository into a separate relational database.

InterSystems IRIS for Health also supports third-party analytics tools, so data scientists can continue to use whatever works best for them.


Machine Learning for SQL Developers

InterSystems IntegratedML[®] brings the power of machine learning to SQL developers. With three simple SQL statements, users can create and train ML models on their data and then use those models to make predictions on unseen data in SQL-based applications. This turnkey tool dramatically increases the productivity of data teams and enables data scientists to focus on only the most complex problems, without having to worry about data access or model deployment.

The Machine Learning Process



Automating the Machine Learning Process With IntegratedML



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

Vector Search

Vector search leverages advanced machine learning models to transform data into high-dimensional vectors, enabling the retrieval of information based on meaning rather than simple keyword matching. This method provides faster, more accurate results, making it particularly useful in scenarios like personalized recommendations, complex healthcare queries, and natural language search. By allowing organizations to search and explore data in more intuitive ways, vector search empowers them to uncover deeper insights with enhanced precision and efficiency.

Note: InterSystems IRIS for Health Advanced Server is required

Natural Language Processing

Built-in capabilities for NLP and text exploration can deliver insights from unstructured data directly within applications. NLP is often used with InterSystems IRIS for Health to create new ML features from text such as clinical notes, that can in turn be used to develop and train ML models. InterSystems text exploration technology—also available as open source—is unique in that it uses a bottom-up approach, discovering concepts and relationships within the text itself.

Business Intelligence (BI)

InterSystems IRIS for Health provides fully integrated support for business intelligence modeling, analysis, and end-user dashboards. An InterSystems IRIS for Health BI model runs directly on transactional data and any other data that might be needed. InterSystems IRIS for Health integration capabilities ensure that its database is always up-to-date, and eliminates the need for extract, transform, and load processing to bring in data from external sources. Drag and drop analysis editing enables non-technical users to examine the data at any level, performing complex queries with ease. InterSystems IRIS for Health dashboards provide a way to display live business metrics and give users options to explore and change what is displayed.

InterSystems IRIS for Health BI uses selective Cube Build, which makes it much faster to add measures and dimensions to a build, without bringing down the cube. By eliminating the need to rebuild the cube each time, hours to days of time can be saved depending on the size of the dataset.

Note: InterSystems IRIS for Health Advanced Server is required

Adaptive Analytics

InterSystems IRIS for Health Adaptive Analytics is an add-on component that provides business users with self-service analytics capabilities to visualize, analyze, and interrogate data from multiple sources in a consistent format. Its semantic layer and drag-and-drop data modeling capabilities allow business users to interactively explore the data to make timely and accurate business decisions.

Note: InterSystems IRIS for Health Advanced Server is required

Columnar Storage

Columnar Storage is an option for IRIS for Health SQL tables that offers order-of-magnitude faster analytical queries compared to traditional row storage. Such queries typically aggregate data over very large tables and involve filters and groupings on one or more columns. Through laying out the table data by column rather than by row (which works best for transactions on a handful of rows at a time), we can dramatically reduce the amount of I/O required to run such queries, and exploit modern chipset-level optimizations called SIMD (Single Instruction Multiple Data) to further improve performance as part of vectorized query processing. That means that, in addition to I/O reductions, bottlenecks such as unpacking values and memory access no longer constrain performance, and operations are only bound by the number of CPU cores the system can access.


If you are implementing a use case where you're running queries across large volumes of data in a data warehouse, data lake, or data lakehouse; or if you're running analytical queries on live application data (a Hybrid Transaction-Analytical or Translytical use case), Columnar Storage will help you achieve excellent query performance.

Note: InterSystems IRIS for Health Advanced Server is required

Reporting

InterSystems IRIS for Health enables the creation of pixel-perfect forms and reports in a variety of formats. It supports the scheduling, exporting, and embedding of reports in customer and partner applications.

Note: InterSystems IRIS for Health Advanced Server is required



Application Development Support

Healthcare application developers need to understand complex regulations and privacy rules, and make sure the software they create complies with them.

Developers need to keep up with new interoperability standards – like FHIR – while still managing to historical standards. And dealing with healthcare data, terminology, and data models requires specialized knowledge. Using a healthcare-specific data management platform such as InterSystems IRIS for Health accelerates development.

Your Language. Your Choice.



Java



C#/.NET



Node.js



Python



ObjectScript



Unified, Low-Code Development Environment

InterSystems IRIS for Health includes a unified graphical and code-based environment that streamlines development and maintenance of sophisticated data- and analytics-intensive healthcare applications that connect data and application silos.

Developers can use graphical and drag-and-drop editors to design integration flows and business processes, to incorporate business rules and human workflow, and to define data and message transformations. InterSystems IRIS for Health seamlessly manages all connection states, connection adapters, message queues, and payloads between itself and external applications and systems.

The platform provides a consistent representation of diverse programming models, programming interfaces, and data formats to simplify development of applications that access and share data across the care continuum. Its trace capability enables developers to track and see the behavior of messages to and from the application, simplifying debugging and diagnosis, lowering development costs, and accelerating time to market.

Client-Side Development

InterSystems IRIS for Health supports client-side development using many popular development technologies, including Java, C#/.NET, Node.js, Python, and ObjectScript, for the greatest flexibility in supporting a wide range of developers and existing applications. ObjectScript is a high-performance, flexible object programming language optimized for developing complex data- and analytics-intensive applications with InterSystems IRIS for Health. ObjectScript classes can be exposed through built-in libraries to Java, .NET, C++, JavaScript, and many other languages.

Server-Side Development

Server-side development provides the highest performance since the applications run close to the data. InterSystems IRIS for Health supports server-side application development with Python and InterSystems ObjectScript. Developers can build applications in either or both programming languages, choosing whichever language is best for the application. Python and ObjectScript code execute within the InterSystems IRIS for Health kernel on the server for extremely high performance.

Message Management

Messages are automatically saved and can be easily audited and visually tracked. This eliminates the need to develop additional application logic to monitor inbound/outbound traffic, queues, and message volumes; for persisting historical message content; for message resending, rerouting, alerting, and event logging; or even for coordinating multiple simultaneous threads of executing process logic.

Full Life-Cycle API Management

InterSystems IRIS for Health provides full life-cycle API management capabilities that support discovering, consuming, routing, throttling, securing, logging, monitoring, and monetizing APIs to support a modern microservices approach to development.

InterSystems API Manager (IAM) enables organizations to monitor and control traffic to and from web-based APIs, including FHIR APIs. They can configure API management using a web-based user interface and API calls, which simplifies remote deployments. The more distributed the environment, the more critical it becomes to properly govern and monitor API traffic. IAM enables all traffic to be routed through a centralized gateway and forwards API requests to appropriate target nodes.

This enables organizations to:

- Monitor all of their API traffic in a central spot
- Plan, document, and update the list of APIs they are using and the servers that provide them
- Identify issues before they become critical
- Control API traffic by throttling throughput, configuring allowed payload sizes, whitelisting and blacklisting IP addresses and domains, and quickly taking endpoints into maintenance mode
- Onboard internal and external developers by providing interactive API documentation through a dedicated and customizable developer portal
- Secure APIs in a central location

Extensibility

Integration capabilities are flexible and extensible. In addition to its built-in integration capabilities, InterSystems IRIS supports the incorporation of existing integration components written in Java, .NET, and Python, enabling developers to build and integrate custom inbound and outbound adapters that can be called at runtime and can send messages to other components.



Flexible Deployment

InterSystems IRIS for Health is available as a managed service (Data Platform as a Service) and can be deployed on all major cloud platforms, in private clouds, on premises, and in multi-cloud and hybrid environments.

It offers the following advantages:

- Eliminates lock-in to a single cloud provider
- Runs on standard hardware to support on-premises, private, and hybrid deployments, with no custom configurations required
- Seamlessly supports all deployment types through a single API, without requiring any modifications
- Continuously leverages the optimizations released by cloud providers and hardware vendors



Smart Health Data Services

InterSystems Smart Health Data Services let you access much of the power of InterSystems IRIS for Health database, interoperability, and analytics technology as smaller, flexible, cloud-centric managed services. See the table below for a list of services and uses. They are available from major public cloud marketplaces, major public cloud marketplaces.

PRODUCT	DESCRIPTION	BEST FOR
InterSystems IRIS for Health	A comprehensive cloud-based software development platform for data management, with interoperability support for all major healthcare standards, including FHIR, and analytics	Developing intelligent, scalable, standards-based health solutions where connectivity to a variety of source and destination systems matters
InterSystems FHIR Server	A smart data service delivering an enterprise-grade FHIR server to provide secure, scalable data sharing and storage for healthcare data	Easily adding FHIR data storage and management to an application
InterSystems FHIR Transformation Service	Fully managed, automated cloud services to transform HL7 V2 and CDA messages to FHIR	Populating a FHIR server from data sources that do not “speak FHIR”

InterSystems Smart Health Data Services



Security

InterSystems IRIS for Health provides a strong, flexible, consistent, and high-performance security infrastructure while minimizing its burden on application performance. This security architecture is based on authentication, authorization, auditing, and database encryption.

- **Authentication:** InterSystems supports multiple authentication mechanisms, including two-factor authentication.
- **Authorization:** Using our System Management Portal, system administrators can easily assign and manage role- and application based resource access privileges.
- **Auditing:** InterSystems products record all system and application events in an append-only log, which can be queried using SQL or a reporting tool.
- **Database encryption:** InterSystems IRIS for Health encrypts data at rest and data in motion. To protect entire databases, it offers block-level encryption.