S&P Global
Market Intelligence

InterSystems' re-platforming strategy leads to new GenAI functionality in IRIS data platform

Analysts - James Curtis

Publication date: Wednesday, August 14 2024

Introduction

InterSystems continues to see steady adoption of its IRIS data platform, particularly with the heightened interest in all things generative AI. Initially rolled out in 2019, IRIS represents the company's next-generation data platform. IRIS required significant development, including a modernization of some core elements of InterSystems' previous platform, Caché. The strategic move seems to be paying off and is enabling the company to incorporate a host of GenAI functionality for developers, enterprises and buyers of its technology. Specifically, new GenAI functionality includes support for vectors as a foundational data type, opening the door for vector search, retrieval-augmented generation and similar scenarios. InterSystems has also embedded GenAI as a coding co-pilot to assist developers and drive data transformation efforts.

The Take

InterSystems has shown that its technology has staying power, perhaps because the company has been iterating on its core database technology for more than four decades in demanding markets, with a record of positive revenue. The rollout of IRIS in 2019 was strategically significant because of its unified architecture that can accommodate disparate data types that are stored uniformly and are accessible from a common data plane. A build-out architecture where the company has added vectors as a data type is one example. Vectors are opening up GenAI and other use cases for its client base, and specifically for the broader healthcare space — a big GenAI adoption vertical — in which InterSystems has traction. Financial services and supply chain/logistics may also present GenAI adoption opportunities for the company. Awareness continues as a factor with InterSystems for more widespread adoption, but for now, the company seems content to let its technology do the talking.

Context

Based in Cambridge, Mass., InterSystems has been around for the better part of 46 years. The company was founded in 1978 by Terry Ragon, who remains the CEO. Never having taken in outside funding, InterSystems is privately held and has been consistently profitable since 1979. In 2023, revenue surpassed $1 billion, and the company reports there are roughly 2,000 employees.

InterSystems has built a loyal customer following with its technology platform. It serves the financial services and logistics/supply chain verticals, touting that it processes more than two billion equity trades daily and tracks more than 20 million shipping containers globally. However, InterSystems' primary foothold is within the healthcare space, where an estimated 75% of its revenue is derived from, and the company claims that it manages some one billion patient records.

Strategy and positioning

The company's core offering consists of the InterSystems IRIS platform, which serves as the basis for many of the company's products, such as IRIS for Health and IRIS Adaptive Analytics, as well as a collection of cloud services, such as IRIS Cloud SQL and IRIS Cloud IntegratedML. Enterprises can also add cloud applications to IRIS to enable specific use cases, such as for aggregating and manipulating health records, but also for supply chain and asset management.

At its core, IRIS is a database platform capable of processing operational and analytical workloads (what we refer to as hybrid operational and analytic processing, or HOAP). The database platform is also multi-model, in that it can store and process many data types (e.g., relational, time series, documents, vectors, etc.). InterSystems, however, positions IRIS as a technology platform powering the company's Smart Data Fabric. In this case, the data fabric serves as an enabler that exposes a collection of services and functionality — some of which InterSystems has built and some developed by partners — for users to access, govern and manipulate the data within the IRIS platform. The "smart" moniker of the data fabric has to do with the embedded AI and GenAI functionality (more on this later) that the company has added to IRIS.

In June, InterSystems rolled out Data Fabric Studio, which functions as a front-end, self-service interface for accessing the services and capabilities that are included with the Smart Data Fabric. Adoption of Data Fabric Studio is expected to widen for IRIS customers because it reduces adoption hurdles by providing capabilities such as a click-and-point interface, as well as a low-code development environment for developers.

Technology

While the Smart Data Fabric functions as a sort of clearing house for IRIS services and access, the underlying technology of IRIS is worth noting, as it forms the basis to deliver hybrid (operational and analytic) processing, including the ability to store and process multiple data models.

A key element of IRIS is its common data plane, a scalable virtual data space. Underneath is an ordered key-value database. Data, regardless of data type, is stored once as multi-dimensional arrays within the common data plane. The arrays include multi-key subscripts that enable fast lookups and scans based on the built indexes. Access to the data is based on a rich framework that IRIS leverages, called projections. Projections map a data type with its underlying storage form such that applications can consume the data in the form that makes the most sense for an application's specific use case.

Because IRIS employs a multi-compute architecture based on a common data plane, maintaining data consistency and performance becomes paramount. One way IRIS addresses this is with a cloud-scale coherent cache, or what InterSystems calls its Enterprise Cache Protocol.

Using a scale-out cache, the data can be distributed along with a system topology, which functions independently of the application. The benefit is that each application maintains a logical view of all the data for optimized performance. IRIS also includes a resilient architecture such that additional compute instances can be used as a recovery mechanism if needed.

Regarding specific GenAI updates and capabilities, InterSystems is tackling this in a couple of ways. One way is to assist developers when creating GenAI applications, which takes advantage of IRIS' platform interoperability. Recently, vectors were added as a new data type that can be stored and accessed just like any other data type within IRIS. Vectors open up a variety of new use scenarios, such as semantic search and retrieval-augmented generation scenarios. InterSystems cites one customer combining natural language processing, vector search and large language models (LLMs) for antibody discovery. InterSystems does not plan to build its own foundation models, either LLMs or embedding models. The company notes that roughly half of its GenAI customers use OpenAI models today, but other models can be used; the company takes a "bring your own model" strategy.

GenAI is also being integrated within IRIS to assist users. For instance, InterSystems has a forthcoming co-pilot to help integration engineers with their data transformation efforts. The co-pilot was fine-tuned on a model based on the company's propriety Data Transformation Language. From a drag-and-drop interface, users determine the transformation — from one type to another — and the co-pilot generates the appropriate code in the background and allows the user to adjust it as necessary. Further, IRIS can also integrate third-party co-pilots — for instance, Microsoft Corp.'s GitHub Copilot for Python.

Competition

Traditionally, InterSystems has competed with a cadre of vendors targeting hybrid workloads. While hybrid can be a differentiator, it is becoming less so as many vendors are able to process operational and analytic workloads within a single environment. A few vendors are worth noting, including Oracle Corp., IBM Corp., Microsoft, Google, SAP SE, SingleStore, FairCom, PingCAP, MariaDB and Percona. There are also several NoSQL vendors, such as DataStax, MongoDB Inc., Aerospike, Couchbase Inc. and Redis, all of which can accommodate hybrid workloads.

However, with the trend toward developing platforms centered on data, we believe that InterSystems' strongest competition is likely to come from the cloud providers, specifically Amazon Web Services, Microsoft (Azure) and Google Cloud, which all offer a litany of database services that address operational and analytical workloads via individual, isolated and specialized database services. They also integrate GenAI functionality, including co-pilots and proprietary-built LLMs and models. Another notable vendor is Oracle, which has built and positioned the Oracle Database as a "converged" database because it accepts multiple data types and development styles and enables multiple workloads; it runs on the company's cloud platform, Oracle Cloud Infrastructure. Two other vendors prominent in the data platform space are Databricks and Snowflake Inc. Known for its lakehouse architecture, Databricks positions itself as offering a data intelligence platform, whereas Snowflake promotes its AI data cloud.

SWOT Analysis

Strengths	Weaknesses
IRIS' common data plane, particularly how it stores data and makes it accessible, is innovative. Vector support was recently added, but other specialized types used by some customers can also be added, highlighting IRIS' build-out architectural design. Other notable architecture elements, like a distributed cache, built-in data integration, and	The company's profile is somewhat small compared with many of its competitors, perhaps in part because of its partnership strategy and deep presence in the healthcare space, compared with other verticals.

embedded analytics and AI, are also noteworthy.	
Opportunities	**Threats**
While InterSystems plays in other verticals besides healthcare, healthcare has high potential and interest in adopting GenAI. Given that the company receives the bulk of its revenue from healthcare, there is a tremendous opportunity to continue to build out specific functionality to help healthcare companies and institutions move forward with GenAI.	There is increased pressure with the predominant cloud platform vendors — AWS, Microsoft Azure and Google Cloud — that provide a collection of data and GenAI services, as well as functionality that competes directly with InterSystems' IRIS platform. This, combined with the fact that the company's profile is much smaller, makes for a highly competitive market.

Source: 451 Research.