Brought to you by Informa Tech

Publication date:
19 Aug 2024
Author(s):
Bradley Shimmin, Chief Analyst, AI & Data Analytics

The art of unlocking value through smart data fabrics – Lessons from InterSystems

Omdia view

Summary

InterSystems, the silent data giant behind more than one billion healthcare records, two billion daily equity trades, and more than two million tracked shipping containers, is stepping into the limelight. The 45-year-old data management, integration, and analytics player recently invited several analysts to its headquarters in Cambridge, Massachusetts, to take stock of the company as it seeks to expand its reach both within and beyond its core vertical markets. In this report, Omdia discusses several of the company's efforts to extend its data platform into the age of Generative AI (GenAI).

Why this matters

The market for modern data platforms has undergone some unique changes over the past two years as we move further into the age of GenAI. Prior to 2023, data and analytics players promoted basic ideas like breaking down data silos and creating and democratizing data-driven insights.

Perhaps unsurprisingly, in the GenAI era, ushered in by OpenAI ChatGPT in late 2022, those messages remain and are in a way more important than ever. GenAI demands data—lots of data from across the business. It also does a truly fantastic job of democratizing access to data for users across that same business.

In short, GenAI prioritizes the value of data and accentuates access to that data, represented within GenAI models themselves. Now more than ever, companies need a data platform capable of unifying disparate data types, query engines, and storage paradigms. Such perennial requirements, along with data security and privacy, will remain relevant regardless of the shape AI takes, now or in the future.

No one expected, however, that the role of data itself would change in the age of GenAI. Data, of course, remains a vital source of corporate truth. However, its value no longer rests solely within the confines of a database of one sort or another. Corporate truth is moving into GenAI models themselves.

Specifically, data increasingly serves as a "representation" of corporate knowledge and insights, even its history and values. Moving beyond basic relational tables and discrete data points, these broader entities live as vector embeddings or graph node/edge triads and manifest as semantic search results fed into large language models (LLMs) like OpenAI GPT-4 via cloud-borne API or Meta Llama 3.1 directly on premises.

Why this is important could fill a book, headlined by broadly impactful use cases like retrieval augmented generation (RAG), a technique that feeds timely, contextual information into a GenAI model as a means of improving model output quality across numerous measures such as truthfulness and relevancy.

This brings us back to InterSystems. Any data platform player capable of enabling companies to live in both worlds, the world of discrete and representational data, will be well situated competitively. Full stop. Within the remainder of this report, Omdia will take a close look at some of the ways InterSystems intends to use its smart data fabric to bring these two worlds together for its customers.

Today's challenges as tomorrow's opportunities

At present, InterSystems maintains a fleet of vertical solutions that run on top of a single data platform, InterSystems IRIS, which itself breaks down into several vertical and horizontal components, with IRIS Adaptive Analytics as one example.

Given that InterSystems has been in business for more than thirty years, it may come as a surprise that, like Apple before it, the company completely rebuilt this core platform in 2019. Its goal at that time was to modernize the IRIS infrastructure, adding cloud services and focusing on high performance machine learning (ML) use cases. Using this new foundation, InterSystems provided native support for several data models:

- Time series
- Object
- Time series
- Key-value
- Dataframe
- Cube
- Vector/tensor

The addition of vector/tensor support is interesting. Looking back, it appears as a highly prescient move, anticipating the rise of GenAI and the associated demand for RAG, which relies upon support for vector data. Rather, InterSystems simply added this functionality to enable columnar table support within their core data plane.

This enabled them to better support traditional, predictive AI use cases (which utilize vectors/tensors), but it also positioned the company to hit the ground running in 2023, when the need for RAG as a means of grounding GenAI models really started picking up steam.

A grand unified field theory

Reworking its core data platform was not limited in scope to the addition of new data types. The objective was to modernize the IRIS platform, such that it could handle multiple types of analytic workloads like ML closer to the data—a philosophy underscored by the company's strong commitment to Python as a native server-side language within the platform.

Such support for Python alongside SQL is not a new idea. Most modern data platforms actively combine ML and analytics together. For InterSystems, which has built up over the decades a large community of customer/partner developers (many still relying on the company's legacy ObjectScript development tooling), the addition of more modern development tools and practices is an absolute must.

InterSystems must show its more than 1,000 customers and 800 partners that their commitment to the IRIS platform can carry them into the future without risk or undue expense. Nowhere is this more true than with development languages.

Why are programming languages becoming so important? The data and analytics marketplace is rapidly transitioning away from discrete tasks like business intelligence (BI) or ML that live either on-premises or in the cloud.

Instead, the market has begun focusing on a higher level of abstraction, moving up from language-specific tasks and discrete databases to focus on enterprise-wide data quality, security, and privacy challenges that depend upon breaking down data silos and doing away with fragile and brittle data integration pipelines (see **Figure 1**).

Figure 1: When it comes to AI, data quality trumps all other areas of investment

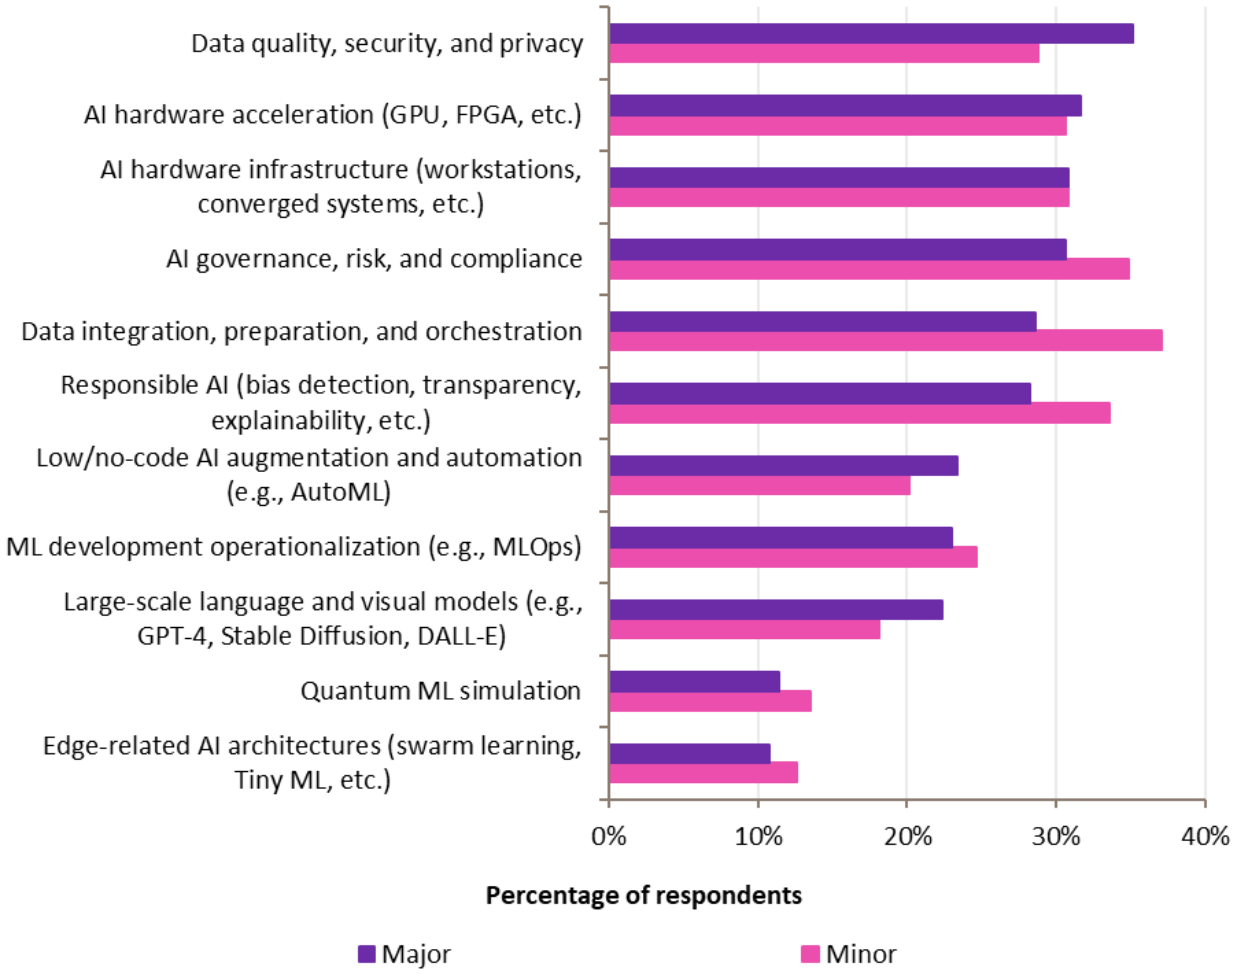

Notes: n=1,695

Source: Omdia

Solving these well-understood problems demands several Herculean feats of strength.

- Creating a universal view of all corporate data assets, usually through the collection and analysis of huge amounts of metadata
- Unifying both transactional and analytical use cases—even in real time
- Supporting multiple data modalities (e.g., vector data next to document data)
- Fusing deployment paradigms (cloud, premises, and edge)
- Delivering direct value at the application layer (i.e., industry solutions)
- Creating a conjoined compute and data engine capable of speaking the programming languages of business—whatever those may be—at any given time
- Democratizing access to the resulting data assets across the business without incurring risk or cost

- Orchestrating all of the above through smart automation and augmentation (often fueled by AI)

Put these together, and you get what the current marketplace terms as a data fabric. With the rapid adoption of GenAI within data and analytics products, a more accurate term may be “smart data fabric” in that they broadly use GenAI and other types of analytics to improve human decision-making, optimize tasks, secure intelligent programmatic actions, and democratize access to data itself.

Returning to the importance of programming languages, the rise of smart data fabrics is rapidly changing the way enterprises build and maintain software. The growing importance of embedded analytical data within business applications and the highly adaptive and composable nature of GenAI itself are pushing enterprises away from ideas like embedded BI dashboards.

Instead, companies are writing apps using full-stack programming languages (Java, TypeScript, et al.) that incorporate functionality that normally would require data specialists working in SQL or Python. Increasingly, developers are not even writing code but instead using low/no-code tooling (itself fueled by GenAI) to generate both the user experience and all underlying functionality.

At present, InterSystems embraces this philosophy entirely, baking AI into their platform as a feature rather than offering AI capabilities as a separate product. Likewise, the company is actively pushing to bring in more full-stack programming languages. Python and SQL are currently available server-side for use side by side. And, in looking to the future, the company intends to bring in more full-stack languages like TypeScript as first-class platform citizens. This is already underway, driven by existing customers like healthcare giant, Epic, which is starting to use TypeScript alongside and even in place of some legacy ObjectScript code.

What’s next

As with many of its business app-oriented competitors (Salesforce, SAP, and Oracle), InterSystems understands the importance of garnering support from the developer community. Corporate buying decisions increasingly turn on speed to market. And that speed depends entirely on the underlying tools used by developers (both in-house and through partners) to integrate new functionality without getting bogged down managing technical debt.

Achieving a proper balance between innovation and maintenance is an absolute must for high volume, data-intensive, mission-critical use cases, especially for customers solving complex issues like supply chain optimization or operating within highly regulated markets such as finance and healthcare—three key verticals for InterSystems.

Taking the time (and the risk) to invest in new tooling, let alone pausing long enough to reinvent the underlying data platform without breaking what came before, can be a tricky business, especially for well-established companies. Fortunately, InterSystems operates as a privately held company, led by one individual with an eye on the long game, not on quarterly margins.

This makes InterSystems a unique player in the broader data and analytics marketplace—one less likely to cut costs for temporary financial gain and more willing to focus on delivering customer value, adopting new capabilities when and where they are deemed valuable.

Technologically, this means InterSystems will not chase tangential opportunities such as entering the enterprise data cataloging marketplace, leaving that to established market players such as Collibra. As we have seen with its adoption of vectors as a native data type and inclusion of ML workloads directly inside its data platform, InterSystems will explore new technologies when it has one capable of delivering end-to-end business solutions for high volume, mission-critical data.

A principled approach should serve the company well as it seeks to move beyond healthcare, finance, and supply chain markets to explore new opportunities, such as the public sector. Omdia believes that to reach new markets, InterSystems will need to focus on the following objectives.

Scale up and out: As InterSystems expands its market presence, the company must address its lack of visibility in the industry. InterSystems enjoys the support of several large and notable customers (M-Tech, UST, Epic Systems, Les Voies Navigables de France, et al.). But with approximately 1,000 customers across three markets, rivals could position the company as being resource-constrained, unable to enter new markets or capture the long tail within its core areas of competency.

Go open: InterSystems wants to grow its developer ecosystem. At present, this community actually outpaces the company's customer count, featuring more than 19,000 Developer Hub members. These members have built more than 950 App Gallery solutions, which are available as open source. But there is more InterSystems can do, such as creating a compelling open source plan of investment that prioritizes the creation of open source industry solutions for new markets—solutions that can be readily extended by partners to help the company scale out and up more aggressively.

Pick up the pace. InterSystems is ahead of many market players in tackling emerging opportunities, as with its interest in supporting TypeScript. However, the company could move more quickly in adopting rapidly maturing data technologies that target data fabric requirements, especially open table formats such as Apache Iceberg and Databrick's Delta format.

Considered together, these and other similar areas of investment will help InterSystems thrive over the long haul by enriching the company's secret sauce, namely its multi-model, multilingual core data engine. There is a good reason the majority of medical records flow through InterSystems software. All that remains is for InterSystems to demonstrate value beyond the company's comfort zone. Based on the information shared with Omdia at InterSystems' inaugural analyst summit in Boston, the company is off to a great start.

Appendix

Further reading

Generative AI Innovation Tracker – 2024 Data (July 2024)

Generative AI Enterprise Survey: Early Adopters Survey Data Tool (July 2024)

Market Landscape: Vector databases powering Generative AI (July 2024)

Author

Bradley Shimmin, Chief Analyst, AI & Data Analytics

askananalyst@omdia.com

Citation policy

Request external citation and usage of Omdia research and data via citations@omdia.com.

Omdia consulting

We hope that this analysis will help you make informed and imaginative business decisions. If you have further requirements, Omdia's consulting team may be able to help you. For more information about Omdia's consulting capabilities, please contact us directly at consulting@omdia.com.

Copyright notice and disclaimer

The Omdia research, data and information referenced herein (the "Omdia Materials") are the copyrighted property of Informa Tech and its subsidiaries or affiliates (together "Informa Tech") or its third party data providers and represent data, research, opinions, or viewpoints published by Informa Tech, and are not representations of fact.

The Omdia Materials reflect information and opinions from the original publication date and not from the date of this document. The information and opinions expressed in the Omdia Materials are subject to change without notice and Informa Tech does not have any duty or responsibility to update the Omdia Materials or this publication as a result.

Omdia Materials are delivered on an "as-is" and "as-available" basis. No representation or warranty, express or implied, is made as to the fairness, accuracy, completeness, or correctness of the information, opinions, and conclusions contained in Omdia Materials.

To the maximum extent permitted by law, Informa Tech and its affiliates, officers, directors, employees, agents, and third party data providers disclaim any liability (including, without limitation, any liability arising from fault or negligence) as to the accuracy or completeness or use of the Omdia Materials. Informa Tech will not, under any circumstance whatsoever, be liable for any trading, investment, commercial, or other decisions based on or made in reliance of the Omdia Materials.

CONTACT US

omdia.com

askananalyst@omdia.com