Celent.

InterSystems IRIS Data Platform

Enabling Banks to Improve Analytics and Decision Intelligence Through Better Data

Craig Focardi and Colin Kerr

June 1, 2025

Contents

Why It Matters

Banks are among the most data-intensive institutions—almost everything a bank does is based on digital processing and data analytics, including advanced analytics and models. **For banks, their data *is the business*.**

In recent years, the banking industry has experienced several simple but momentous trends that drive significant investment in data platforms and management:

- Banking is increasingly digital and data dependent. Not only is more data used in processing, but more data is collected and created by banks—and increasingly from new sources.
- Digital banking strategy is also increasingly cloud-first or must at least integrate on-premise data with data deployed or distributed in the cloud.
- The insights gleaned from that data, and the potential to use that data in advanced analytics and AI, can be a game changer for banks. Data and AI are the building blocks for future banking products and services, and the quality and performance of quant and AI models becomes more critical.

However, banks still face significant challenges to manage data effectively. Celent research[1] also shows that many banks have either started the journey or plan to implement a data fabric solution to address problems in managing data. It is in this context that Celent has written this Solution Brief about InterSystems IRIS data platform.

Get the Data House in Order

47% of banks ranked **advanced analytics and AI** as a top investment priority for 2025.

Source: Celent Dimensions Survey 2025

The potential of AI at scale has driven a realization that many banks may not have their data houses in order. Celent's latest Dimensions survey and research show that AI tops the agenda for technology priorities, but the journey to enable AI at scale—and to support a new generation of intelligent banking products—has become increasingly complex to manage. Individual systems being built, bought, or acquired over many years have created a messy architecture of old and new applications and data stores. Moreover, the data management challenge is significantly greater for banks operating in multiple geographies compared to a single country or region.

Despite all the investments banks have made in data platforms and data management over the past 10 or so years, data issues remain pervasive! Figure 1 shows an excerpt of Celent's Decision Management and Data Fabric Survey 2025 regarding retail and commercial banking technology investment priorities. Integrating new data sources, reducing the latency of data, and improving data quality are the top three issues faced by banks of all sizes.

---

[1] Celent Decision Intelligence and Data Fabric Survey, Celent, March 2025.

Figure 1: Bank Concerns About Data Used in Decision-Making Models and Analytics
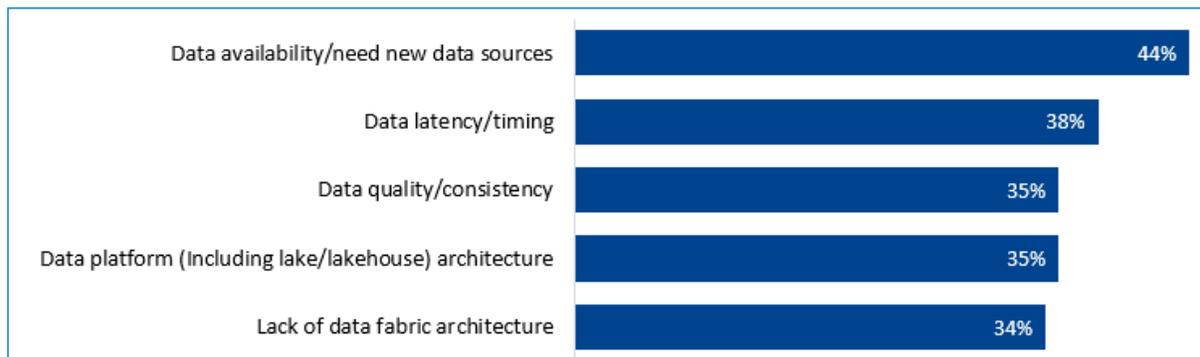
Source: Celent Decision Intelligence and Data Survey 2025.

Note that, in addition to data-specific challenges, banks also cite legacy data architecture as a top barrier to enhancing decision intelligence processes and models. These warehouses and data lakes/lakehouses allowed central management of data assets, but in Celent's view, it has been challenging to deliver the intended business value that requires distribution of the right data to the right business process, at the right time. However, many centralized warehouses remain in place. As a part of Celent's research, a senior banker from a top-tier UK bank noted that 70% of quantitative model performance issues at that institution are data related. This includes gaps in source data, incorrect data sources, and data transformation errors. These issues impact not only AI and model performance but also enterprise data management quality objectives. Furthermore, these issues are symptomatic of the challenges in moving to modern data architecture while still needing to run and grow the business on a day-to-day basis.

60% of tier 1 banks see legacy data architecture as a barrier to improving decision intelligence.

Source: Celent Survey 2025

Many of these challenges could likely be addressed by a data fabric architecture that can deliver high quality data to the point of need at the right time. What does that mean? A data lake or warehouse stores data centrally. Conversely, a data fabric can weave together data from multiple sources (including traditional warehouses) and deliver it to multiple business units, applications, and data scientists at the right time. Digital transformation can last multiple phases over several years, but a unified data fabric should be nonintrusive and support such transformations. This architecture is designed to connect and deliver data precisely to the point of need at the right time, ensuring that organizations can make timely and informed decisions even while underlying infrastructure changes over time.

> **66** 70% of quantitative model performance issues are data related.
>
> Celent interview with a senior banker from tier 1 UK bank

Company and Solutions Overview

InterSystems is a privately held vendor of cloud- and hybrid-based data platform technology that provides unified database management, analytics, integration, and data services capabilities.

Company Context

InterSystems serves many industries, including financial services, healthcare, manufacturing, supply chain, and public sector. Based in Boston, Massachusetts, the company has 37 global office locations from which it serves thousands of enterprise clients in over 80 countries.

Table 1: InterSystems Snapshot as of 12/31/2024

Year Founded	1978
Headquarters	Boston, Massachusetts, USA
Ownership	Privately funded
Annual Revenue	>$1bn
Client Demographics	>40,000 deployments across >1,000 enterprise customers
Product	InterSystems IRIS enterprise data platform
Website	www.intersystems.com

Source: InterSystems

As an enterprise data platform and data management provider, InterSystems' main products include InterSystems IRIS® data platform, InterSystems® Data Fabric Studio™, and a variety of industry-specific applications that are variations of these core products. InterSystems product portfolio also includes a variety of solution-focused cloud services built on InterSystems IRIS. The platform has a specialized emphasis on supporting data fabrics, a nondisruptive architectural approach that allows organizations to easily access, transform, and harmonize data from diverse and distributed sources across their entire business. Moreover, InterSystems distinguishes its solutions as providing a "Smart Data Fabric" architecture, which connects to live data from multiple sources and embeds generative AI (GenAI), business intelligence, natural language processing, and machine learning capabilities directly within the fabric. All these capabilities are provided via a self-service, managed cloud service.

In the first quarter of 2025, InterSystems briefed Celent on its financial services industry strategy and InterSystems IRIS and InterSystems Data Fabric Studio solutions. Celent spoke with:

- Joe Lichtenberg, Global Head of Product and Industry Marketing, Data Management
- Irene Galperin, Senior Advisor, Financial Services
- Kate Kennedy, Senior Manager, Data Platform Marketing

Financial Service Context

This Celent Solution Brief reviews the InterSystems IRIS and InterSystems Data Fabric Studio solutions in the context of financial institutions and their need to improve the availability and quality of data to support more complex analytics scenarios—notably for AI and decision intelligence. Smart data fabrics are applicable across the financial services industry in a wide variety of local and enterprise-level, mission-critical initiatives. Examples include:
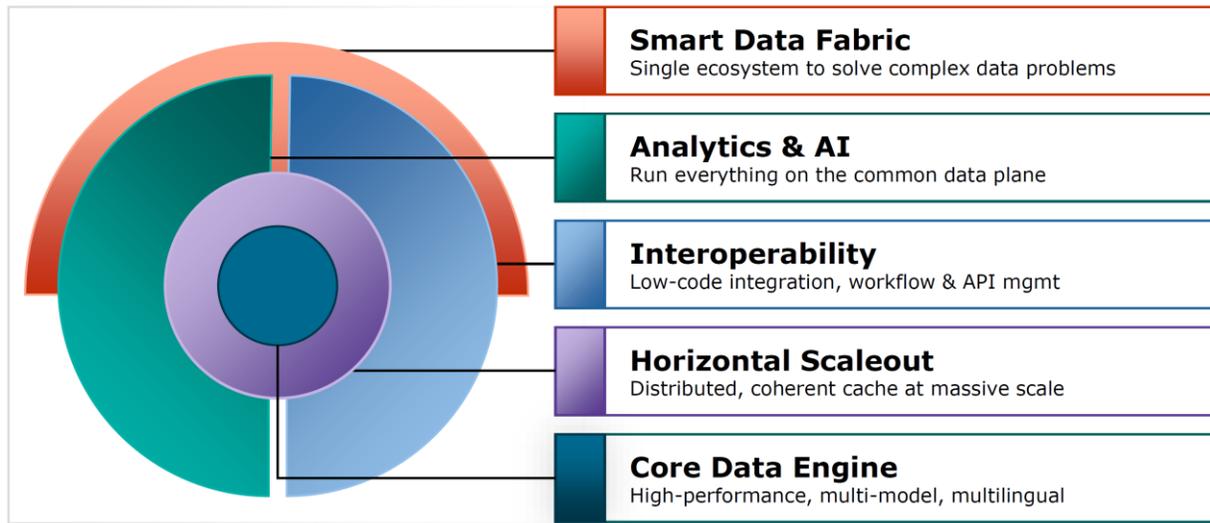
- enterprise risk and liquidity management
- intelligent decisioning for credit, fraud prevention, and customer intelligence
- GenAI-enabled advisors and copilots
- scenario planning and decision support
- business management reporting, analytics, and quantitative models
- real time monitoring, visibility, and alerting of critical processes

As a multi-industry and enterprise data platform provider, InterSystems supplies solutions to retail and commercial banks, wealth managers, and asset managers worldwide, as well as fintech and data providers. Example clients include Northern Trust, UBS, Lloyds Bank, NatWest, Harris | Oakmark, Fiserv, and FIS.

Solution Overview

InterSystems IRIS is a comprehensive data platform designed to facilitate the development and deployment of data-intensive applications. It integrates a range of capabilities including database management, integration and interoperability, analytics, GenAI, and machine learning, all within a unified framework. This cloud-first data platform makes it easier to build and deploy high-performance, AI-enabled applications and models that connect data across application silos, and it provides data analytics capabilities to analysts, quants, data science professionals, and developers. Figure 2 spotlights some of the key capabilities and characteristics of the InterSystems IRIS Smart Data Fabric platform—all of which are important criteria for banks as they navigate their way to enhanced data processing and data analytics capabilities.

Figure 2: InterSystems IRIS Architecture

Smart Data Fabric
Single ecosystem to solve complex data problems

Analytics & AI
Run everything on the common data plane

Interoperability
Low-code integration, workflow & API mgmt

Horizontal Scaleout
Distributed, coherent cache at massive scale

Core Data Engine
High-performance, multi-model, multilingual

Source: InterSystems

Key Product Briefing Insights

Banks should demand the highest levels of quality and execution of transactions and operational processes. Although legacy platforms persist and may compromise reliability, a data transformation program must be based on a foundation of quality, end-to-end execution, and resilience. It is through this lens that Celent highlights key briefing insights.

InterSystems Common Data Plane

A unique feature of the InterSystems IRIS Smart Data Fabric platform is the Common Data Plane. Within the core data engine, this data fabric architecture weaves among a bank's data repositories, regardless of underlying data infrastructure. It persists information via highly optimized, multidimensional arrays— which means it can connect and distribute data to and from any data model and data type, inside or outside a bank. This could be streaming market data feeds, loan origination systems, payment systems, enterprise data warehouses, and risk models. The Common Data Plane supports data access by any data consumer, using any format or access method, without the need to copy or replicate data.

Figure 3: InterSystems Common Data Plane
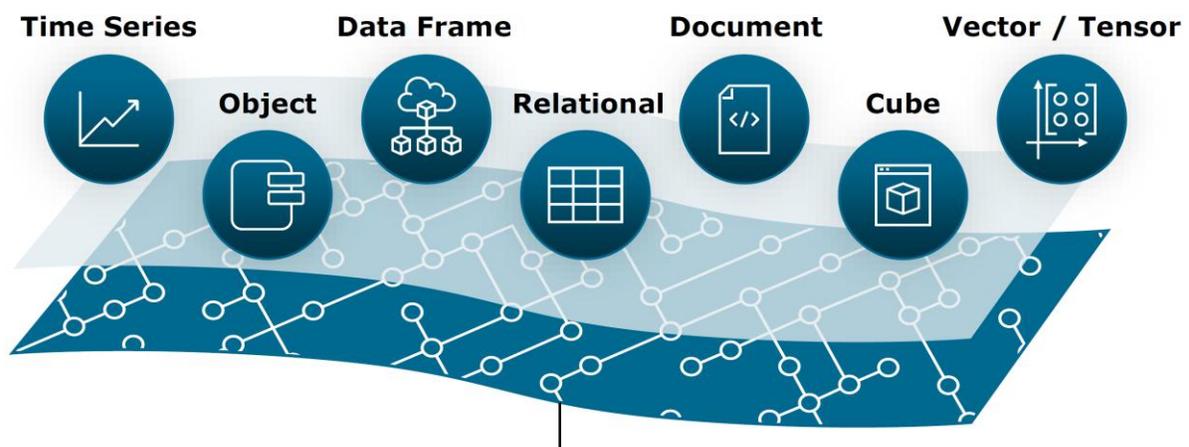
Source: InterSystems

InterSystems IRIS Data Platform—Significant Characteristics

Speed and Scale

A key requirement of any mission-critical data platform is low-latency performance and the ability of the platform to scale up or down according to application and analytics/model workloads. This may take several forms, such as:

- Increased transactional or low-latency workloads—such as those found in payment processing and trade execution. These workloads can be episodic throughout the day (such as at start of day and before cutoff times). At peak load, they may stress analytics models used in fraud detection or sanctions screening.

- High-intensity computational models requiring bursts of compute power and data processing. Examples include non-transactional models for pricing analysis, portfolio evaluation, risk simulations, and credit risk profile calculation.

Celent research finds that the ability to deliver high-quality data, enriched as needed and without contaminating other data views or flows, is a primary driver to improving critical data processes and AI models in financial services.[2] The InterSystems IRIS Smart Data Fabric platform addresses these challenges with the following capabilities:

- low-latency processing and delivery of data (as evidenced by capital markets implementations of InterSystems)

- data persistence across the InterSystems IRIS data fabric for high availability and resilience

- exponential data processing scalability across the platform (versus through an architecture of multiple data platforms and services)

Integration and Interoperability

Any data management or data fabric platform requires an intrinsic ability to integrate and interoperate seamlessly with other data platforms, applications, and processes. It is its reason for being.

The InterSystems IRIS Smart Data Fabric platform weaves together data from applications, APIs, data stores, data warehouses / data lakes / lakehouses, quantitative analytics, and artificial intelligence models into a unified data platform. The platform includes built-in support for native APIs that enable solutions based on InterSystems IRIS to integrate with applications. Support for popular data formats, including JSON and XML, cuts down time to connect to other systems, and a comprehensive library of adapters provides out-of-the-box connectivity and data transformations for packaged applications, databases, industry standards, protocols, and technologies. The platform is designed to connect and deliver data to the point of need at the right time. Importantly, the platform complements existing data warehouses and data lakehouses rather than replacing them.

Connectivity is important, but so, too, is the ability to orchestrate and monitor activity across the platform. InterSystems IRIS provides full lifecycle API management capabilities, low-code graphical business process orchestration capabilities, a business rules engine, and a workflow engine that enables

---

[2] Celent Decision Intelligence and Data Fabric Survey, Celent, March 2025.

banks to automate enterprise-wide business processes and create new composite applications that integrate with their *existing* infrastructure. This includes next-generation microservices and digital transformation initiatives—two hallmarks of modern banking technology environments.

InterSystems IRIS can seamlessly store and analyze messages as they flow through the system for business activity monitoring, alerting, real time business intelligence and analytics, and event processing. The platform offers a management portal that provides information about bank solutions, including all messages, process state, and metadata captured from connected systems. Debugging, when necessary, is facilitated by the "Visual Trace" feature, which gives administrators the power to trace messages and examine their contents as they move through the system.

Security and Reliability

Because the business of banking is rooted in data, financial services institutions demand the highest levels of security around data and operational processes. Reliability of systems is also of utmost importance. Too many banks have unwittingly been the center of news stories when core banking or payment systems suffer an outage. The key is balancing security and resilience with speed of execution. It is a fine line to walk.

Security
The security model for InterSystems IRIS is based on four core principles:

1. **Authentication—**InterSystems IRIS supports a variety of common authentication mechanisms (LDAP, Kerberos, direct passwords, OpenAM, and OpenID), and includes support for multi-factor authentication as needed.

2. **Authorization—**Management of InterSystems IRIS privileges (including role-based and application-based privileges) are accomplished through APIs and interactive applications. Row- and column-level security is also supported.

3. **Encryption—**The InterSystems IRIS platform provides mechanisms for encrypting both data at rest and data in motion. Data-at-rest encryption encrypts the entire database, including indexes. In addition, data-element encryption is available to encrypt highly sensitive information.

4. **Audit—**All InterSystems IRIS platform events are recorded in a tamper-proof append-only log. Any SQL-based BI or query tool can be used to review and analyze audit logs.

Reliability
Operational resilience has become a strategic focus for regulators since the turn of the decade. Although this topic is perennially important for banks, new regulations such as DORA[3] in Europe and similar requirements in North America and Australia have placed this at the top of bank boardroom discussions. While speed of recovery and reducing the impact of disruption remains important, the emphasis has extended to preventing disruption through investment in systems and processes as necessary by proactively looking for threats and potential weaknesses.

The InterSystems IRIS platform makes extensive use of database mirroring, sharding, clustering, and virtualization to provide banks with a variety of high availability and recoverability options.

---

[3] European Insurance and Occupational Pensions Authority. *Digital Operational Resilience Act (DORA)*. Available at: https://www.eiopa.europa.eu/digital-operational-resilience-act-dora_en (Accessed: 1 June 2025).

- **Database Mirroring**—This is a logical grouping of two InterSystems IRIS systems. Upon startup, the mirror automatically designates one of these two physically independent systems as the primary system; the other one automatically becomes the backup system. Application recovery time is typically reduced to seconds and also enables minimal downtime upgrades.

- **Database Sharding**—A database mirror is set up for each database shard, thereby eliminating any single point of failure. Such solutions can be deployed on-premise and in cloud environments.

- **Clustering and Virtualization**—In cases where a longer recovery time is acceptable, clustered or virtualized instances of InterSystems IRIS may be deployed. These instances are dependent on shared disk access but with only one InterSystems IRIS system instance active at a time. If the active system fails, InterSystems IRIS is automatically started on another server that takes over the processing responsibilities.

Generative AI

As a relatively recent product innovation, InterSystems IRIS supports native vector embeddings with the data. It enables semantic search and GenAI capabilities that allow banks to build "smart" applications that combine GenAI with traditional AI and analytics atop any dataset—including low-latency streaming data. This is delivered through the combination of native vector search and RAG (retrieval-augmented generation) without the need for a separate vector database.

Vector Search	Traditional search is based on keyword matching. Vector search is a technique that uses numerical representations (vectors) of words, entities, documents, images, or videos. Vector search captures the *semantic relationship*s between elements to enable more effective processing by machine learning models and artificial intelligence applications.
Retrieval-Augmented Generation (RAG)	Retrieval-augmented generation is a technique for enhancing the accuracy and reliability of GenAI models with information from specific and relevant data sources. RAG modifies interactions with an LLM with integration to additional data sources or domain-specific information. The objective is to improve GenAI response accuracy to a higher level than is possible with trained LLMs.

The InterSystems IRIS platform provides embedded vector search capabilities that allow users to search unstructured and semi-structured data. Data is converted to vectors and then stored and indexed in InterSystems IRIS for semantic search, RAG, text analysis, recommendation engines, and other use cases. The RAG architecture combines vector search and large language models (LLMs) to overcome some LLM accuracy limitations (including hallucinations) and generate more accurate responses using fresh, authoritative information. This can be based on proprietary datasets. Although production implementations of GenAI are still in their infancy in financial services, Celent does see high levels of activity as banks pursue proofs of concepts and pilots as stepping stones to production deployments.

Implications for Financial Institutions

Data fabric use cases in financial services are diverse. They include supply and distribution of data for managing fraud risk, payments, credit risk, customer intelligence, and loan/account origination.

Common characteristics of banking use cases include: the need for real time decisioning and customer engagement; large volumes of data; and dynamic models that process data and that must be monitored regularly for drift and decay and updated accordingly. Celent's research[4] shows that data availability, quality, and latency are among the top challenges that banks face when building better decision intelligence solutions. Retail loan underwriting is a classic and broadly understandable use case that has the characteristics described above and that data fabric architecture can help automate.

Retail Loan Underwriting
Processing, underwriting, and approving a loan application requires data collection and analysis of identity, fraud, credit, collateral, and customer data to process, underwrite, approve, and fund a loan. This data requires collection from external and internal sources, normalization, and accessible storage. It must also be collected in real time, integrated into customer-facing and middle-office processes, and inputted into decisioning models for analysis.

Figure 4 depicts an IT architecture for loan underwriting. Technologies are segmented into front, middle, and back office. High-level processes within each IT stack layer are shown on the left, and the benefits of data fabric are shown on the right. Conceptually, the data fabric architecture is woven throughout the credit decisioning IT stack, although in practical terms this occurs with integrations across datasets, applications, models, and data science workbenches.

---

[4] Celent Decision Intelligence and Data Survey, Celent, 2025.

Figure 4: Retail Loan Underwriting Supported by Data Fabric Architecture
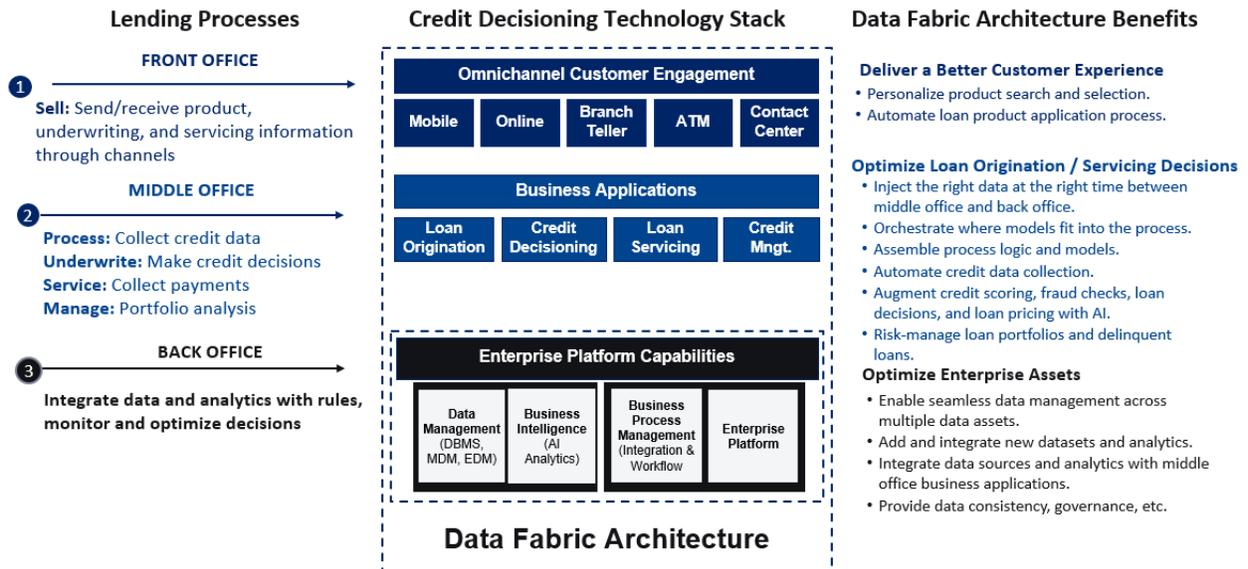
Source: Celent

Because the data fabric provides a common data connectivity and access layer, every process, application, and decision system can access the same data "truth" in the context it needs—without needing to understand the complex data models and views managed in a traditional data lake. The data fabric acts as an abstraction layer that insulates both the applications and data science tools from the underlying data sources and dynamically enriches or transforms data as required.

Lending has numerous decision points across the many stages of the lending business, making it a prime target for data integration and decision automation. For example, the multiple steps of a loan origination process might be connected via a business orchestration layer. At each point of the journey, whether feeding credit reference data from third party providers, building personalized risk profiles, or providing insights to the relationship manager, the data fabric connects with each process and presents it with the data it needs in the right format, at the right time. Underpinning that is an integrated data management toolset to ensure the data supply and distribution adheres to data governance standards. An abstract example of this is shown in Figure 5.

Figure 5: Lending and Decisioning Process Overview (Abstract)
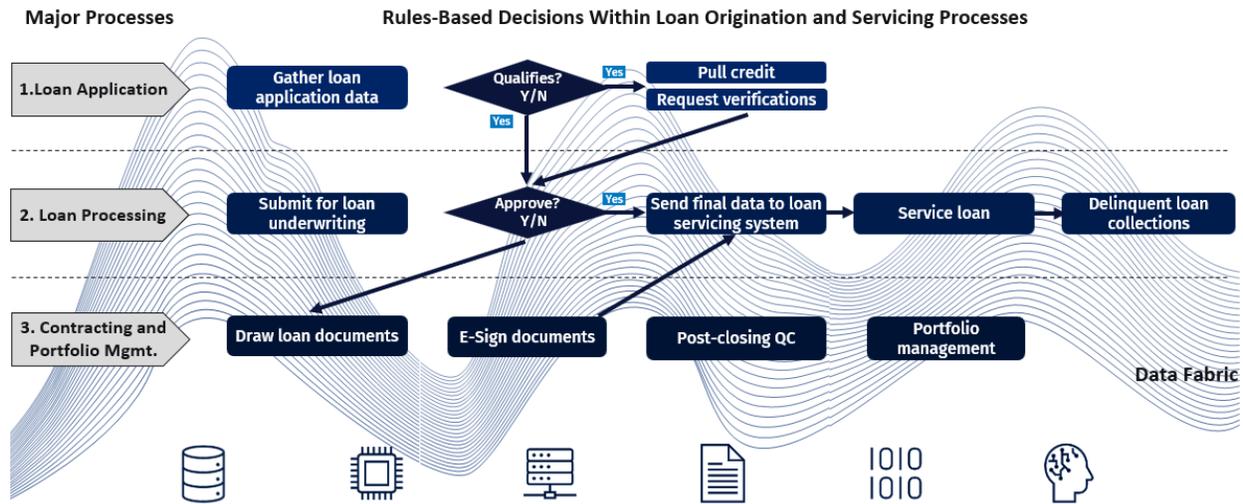
Source: Celent

Banks look to improve analytics and decision intelligence throughout their business operations and have a high dependence on accurate and timely data. Traditional data architecture approaches, such as data warehouses, data lakes, and data lakehouses have all added value, yet data challenges remain. A data fabric architecture has the potential to augment and complement all these prior investments, and also address the root cause of model performance challenges. When a data fabric delivers more accurate, complete, timely, and unified data to the point of need banks have the opportunity to innovate and compete more effectively.

Leveraging Celent's Expertise

If you found this report valuable, you might consider engaging with Celent for custom analysis and research. Our collective experience and the knowledge we gained while working on this report can help you streamline the creation, refinement, or execution of your strategies.

Support for Financial Institutions

Typical projects we support include:

Vendor short listing and selection. We perform discovery specific to you and your business to better understand your unique needs. We then create and administer a custom RFI to selected vendors to assist you in making rapid and accurate vendor choices.

Business practice evaluations. We spend time evaluating your business processes and requirements. Based on our knowledge of the market, we identify potential process or technology constraints and provide clear insights that will help you implement industry best practices.

IT and business strategy creation. We collect perspectives from your executive team, your front line business and IT staff, and your customers. We then analyze your current position, institutional capabilities, and technology against your goals. If necessary, we help you reformulate your technology and business plans to address short-term and long-term needs.

Support for Vendors

We provide services that help you refine your product and service offerings. Examples include:

Product and service strategy evaluation. We help you assess your market position in terms of functionality, technology, and services. Our strategy workshops will help you target the right customers and map your offerings to their needs.

Market messaging and collateral review. Based on our extensive experience with your potential clients, we assess your marketing and sales materials—including your website and any collateral.

Related Celent Research

It's All About the Data! Improving Decision Intelligence Models in Banking
May 2025

Next-Generation Intelligent Decisioning Platforms: Shifting to the Enterprise and Cloud to Modernize Front and Back Office Retail Loan Decisioning
January 2025

Digital Sovereignty: The Impact on Data, AI, and Next Gen Banking Solutions
October 2024

In Search of Shangri-La – the Expensive Pursuit of Data Management in Banking
September 2024

Detangling Data: The Art and Science of Managing Banking Data
June 2023

Improving Operational Excellence While Migrating to a Hybrid Cloud, Multi-Cloud World: What to Move to The Cloud, When, and How
December 2023

Navigating to Value with Corporate Banking Data: Recommendations for Managing Data Strategies
September 2022

Oceans of Data: Recognizing the Perils in Corporate Banking Data Transformation, The Challenges of Developing and Implementing a Data Strategy
August 2022

Financial Institutions Get Smart About Data (Fabric): FIs embrace Smart Fabric approaches to better leverage enterprise-wide data for advanced insight capabilities
April 2021