

Driving Innovation With A Fintech Gateway

Ann Kuelzow, Global Head of
Financial Services – InterSystems

Jon Payne, Manager of Sales
Engineering and Education –
InterSystems



Data Challenges are #1 Issue for 81% of Fintechs*




- 81% of fintechs report data challenges as their biggest technical challenge
- 41% struggle to leverage data for analytics, AI, and ML
- 40% are unable to connect to customers' applications, data, and legacy systems
- 52% say improving scalability and reliability is their top technical priority, 48% say increasing agility, 34% say improving security, and 27% say improving customer experience
- Top planned investments include cloud (51%), data management (48%), AI/ML (45%), data fabric (42%)

*Global survey of 500 senior executives from fintech organisations

...And For Their Banking Partners



To what extent are the following challenges when partnering with fintechs?



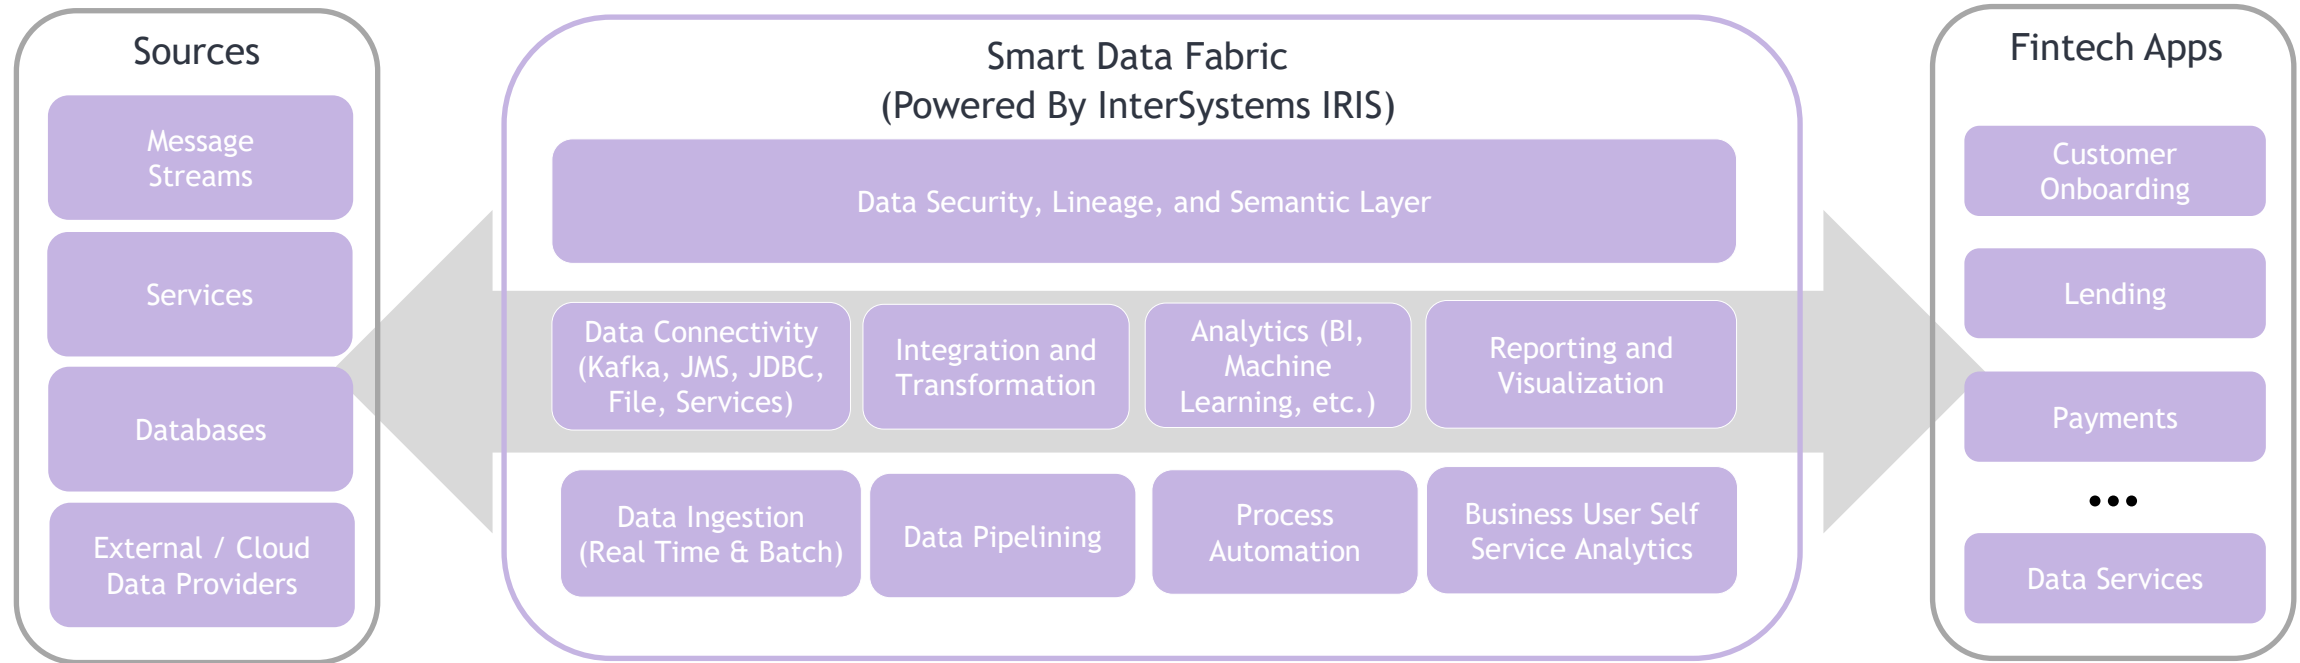
Fintech Gateway Solution



InterSystems Fintech Gateway Solution provides a real-time, bidirectional gateway between cloud-based fintech applications and customers' production applications and data assets

It enables financial services organisations to accelerate the implementation and production deployment of cloud-based fintech applications, delivering greater agility and better business outcomes.

Cloud Fintech Gateway Leverages The Smart Data Fabric



Fintech Gateway Solution




- ✓ **Increases Speed to Value.** Improves speed to value by enabling financial services organisations to more quickly and easily integrate cloud-based fintech solutions and technology into their existing infrastructures
- ✓ **Enhances Agility.** Enables financial services organisations to react to new opportunities and changes in their environments by making it faster and easier to incorporate the wide range of fintech applications and technologies available in the marketplace
- ✓ **Increases Operational Efficiencies.** Eliminates the complexity and inefficiencies of manual integrations and other legacy approaches to integration by leveraging a modern smart data fabric approach
- ✓ **Supports Real-Time Scenarios.** Leverages InterSystems IRIS data platform, which is ideal for handling mission-critical, real-time, transactional-analytic use cases at scale

InterSystems IRIS Data Platform



A Complete Software Platform for Developing Real-Time, Enterprise Data- and Analytics-Intensive Applications
For Financial Services Organisations



Multi-Model

Process any data in any format from any source
=> Superior flexibility

Multi-Workload

Ingest and simultaneously analyse incoming transactions and events with superior performance
=> Intelligent, scalable real-time processes and analytics

Analytics

Execute analytic SQL, BI, ML, NLP, and business rules and embed within real-time applications
=> Intelligent processes and improved decision making

Interoperability

Gain real time, on-demand access to distributed data, build streamlined end to end processes and “data fabrics”
=> Accurate real-time visibility and analytics leveraging disparate enterprise data

Flexible Cloud | Hybrid | On Premises Deployment Options

InterSystems IRIS Data Platform



InterSystems IRIS is real-time, high-performance data management software built for **performance at scale** with **enterprise-class security** and **mission-critical reliability**.

It simplifies architectures and reduces infrastructure costs by combining multiple technologies in a **single product** built from the ground up on a **common architecture**.

- Highly optimised persistent DBMS with in-memory performance
- Built-in distributed elastic cache with guaranteed consistency
- Built-in analytics platform (IntegratedML, BI, NLP, ...)
- Built in integration platform (API management, process orchestration, ...)
- Developer friendly development environment

Supports deployments on premises and in any cloud with a single API

Highly Optimised Persistent DBMS with In-Memory Performance



Feature	Benefit
Supports a wide range of different data models and access methods (relational, object, document, ...) with one physical representation	Eliminates the need for multiple engines; no data duplication
All data automatically and consistently indexed	Enables consistent and fast access to all data
High ingestion and transaction speed	Captures incoming data at high speed for high throughput real time use cases; supports low latency high performance use cases requiring ACID transactions; 100's of millions of transactions/sec sustained
High performance transactional-analytic processing	Enables real time use cases requiring analytics on incoming transactions with high performance at scale
Automatic schema-on-read and schema-on-write	Does not require creating schema upfront, subsequently access the data in many ways; flexibility to handle new data, formats, and requirements as they arise
Highly efficient use of storage (sparse arrays with automatic encoding)	Highly resource efficient (disk, I/O); no replication of data needed

Built-in Distributed Elastic Cache with Guaranteed Consistency



Feature	Benefit
Built in, consistent, high performance distributed cache	Guaranteed consistent, real-time access to any data. Provides the performance of “in memory only” solutions with built in durability, lower TCO and no restart delays. No need to implement and manage multiple technologies; replaces multi-tier infrastructures with a single technology
Single, consistent representation of data on disk, in memory, and on the wire	Performant, efficient, and self-managing. Simplifies development and maintenance
Scales easily, linearly, vertically and horizontally	Easier to operate; predictable capacity planning
Ability to scale data, ingestion workloads and analytic workloads independently	Maintain superior performance at scale with lower resource requirements and lower TCO
Add and remove query and ingest capacity without service interruptions	High reliability, greater flexibility, lower cost

Built-in Comprehensive Analytics Capabilities



Feature	Benefit
All analytics capabilities are fully embedded in the product; execute directly on the data	Moves the analytics to the data for highest performance; simpler development, deployment and maintenance
IntegratedML	Easily build and tune machine learning models, and operationalise models in processes and applications, all from within the platform. Use familiar SQL syntax to integrate machine learning into applications
Business Intelligence	Guaranteed consistent; updated via real time events, transactions, triggers
Business User Self Service	“Adaptive Analytics” enables data stewards and business users to easily visualise, analyse, and interrogate live data
Analytic SQL	Supports analytic SQL on large distributed data sets, supporting ad hoc SQL queries with high performance and superior reliability. Rich SQL vocabulary enables pushing computations close to the data.
Natural Language Processing	Out-of-the-box domain-independent content navigation and search capabilities
PMML Runtime Engine	Import and execute predictive models built in the lab using mainstream data science libraries for high-performance execution in production applications

Built-in Interoperability Capabilities



Feature	Benefit
Built-in data and application integration capabilities	Easily integrate with data and applications inside and outside the firewall; simplifies software requirements
Single consistent architecture spanning interoperability, database management and analytics	Simpler to build and maintain; eliminate duplication of data and effort
Highly scalable message-based integration	Handle high-volume message flows with low and predictable latency
Messages persisted and fully traceable by default	Enables transparency, forensics, maintainability, and analytics
Graphical & live Business Process and Business Rule editor	Enables business stakeholders to define and manage workflows and rules without requiring developer assistance or business interruption
Full lifecycle API management capabilities	Build API-first, execute, deploy, monetize, monitor, and support APIs and microservices based applications
Dynamic gateways to and from Java, .Net, Python, node.js	Integrate existing applications and components using the language of your choice

Developer Friendly



Feature	Benefit
Supports all major programming environments including Java, .NET, Python, Node.js, R	Enables application developers to be productive quickly with their existing skills, both for client- and server-side development
Near-data computing - embedded language runs right next to the data	Get extreme performance for data-intensive applications
Native multimodel platform extensible with your own “designer data structures”	Work with data in the best model for the task, avoiding the time and complexity of converting between models or duplicating data
VSCoDe-based modern tooling	Familiar, powerful development environment maximizes developer productivity
Production Extensions (PeX) - call in or out of applications and libraries in top languages	Tap into a huge ecosystem of prebuilt functions and libraries to build applications easily; integrate easily with existing code
InterSystems Developer Community, Global Masters, and OpenExchange	Provides expertise, components, and connections from a large and vibrant community

Flexible Deployment



Feature	Benefit
Runs in all major public clouds with the same API	Supports the cloud environment of your choice; seamlessly change cloud providers without making changes to the applications
Available in the major cloud marketplaces	Easy to try, easy to buy, flexible pricing schemes including pay-as-you-go
Supports distributed and hybrid deployment configurations	Flexibility to support public clouds, private clouds, on premises, and hybrid with high performance supporting distributed users
Containerised; supports Kubernetes	Supports many different workloads efficiently while maintaining isolation and security; fits within existing operations and DevOps processes

Fintech Case Study: Smart Data Fabric



Real time smart data fabric supporting dozens of applications and hundreds of clients

● Challenge:

- Needed an enterprise data fabric that consolidates and normalises data across business lines and asset classes; provides ability to perform complex calculations and analytics
- Improve data quality, consistency, usability, and availability; reduce latency
- Simplify architecture

● Results:

- Dynamic, enterprise data fabric integrates data from various sources and with internal and client applications with sub-second SLAs for near-real time updates to their suite of applications
- Calculates real time positions and “on the fly” aggregations for postings and balance data
- Performs various middle and back-office processing functions, e.g., trade reconciliations, enrichment and bookkeeping
- Complies with all regulatory requirements
- Reduces number of products required; completely cloud native deployment

Fintech Case Study: Branded Credit Card Fintech



High performance, scalable, reliable payment processing

● Challenge:

- Provide branded credit cards and loyalty programs for 6000+ retailers and 38,000+ concurrent users
- Reduce system latencies, improve reliability

● Results:

- Processes more than 12 million transactions per year
- Reduced credit card transaction processing from 16-20 seconds to 3 seconds
- Provides real time updates to the customer, retailer, and credit card brand owner
- Eliminated system outages experienced with prior data platform
- Delivering new functionality including virtual credit cards running on users' mobiles that are provisioned and activated immediately.

"Our applications run with high performance even during the weekends. My cell phone does not ring! The reliability of InterSystems IRIS allows us to spend the weekend in peace." - Luis Oliveira, CEO.

Fintech Case Study: Next Generation Wealth Management



Next generation, cloud-based wealth management software for major FinTech

● Challenge:

- Grow revenue and competitive differentiation through rapid innovation
- Develop cloud native, API based wealth management application
- Ensure high performance, resource efficiency, reliability, and scalability
- Support 22+ million client accounts with 7 years of data retention

● Results:

- 9X performance improvement on 30% of the infrastructure
- <1s response time; able to scale 5X on demand; supports 16K concurrent users; reduced RTO from 15 minutes to 4 seconds
- Replaced multiple data management software products (AWS Aurora, AWS Redshift, AWS DynamoDB, RedisLabs, Elasticsearch) with InterSystems IRIS
- Developed and deployed faster than projected

Any Questions?