zensar

Building a Governed Intelligence Layer for Scalable AI in Financial Services

Whitepaper

An RPG Company

zensar

Table of Contents

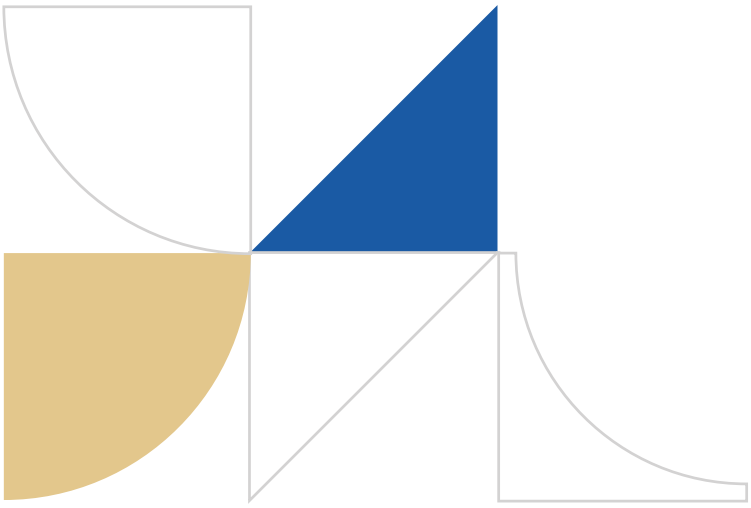

Executive summary

Financial services organizations are rapidly adopting AI, yet most initiatives do not progress beyond pilot phases. The core issue is not the strength of the models, but rather fragmented intelligence foundations incapable of supporting trust, governance, and reuse. This white paper contends that lasting AI advantage is achieved through a governed intelligence layer — data products with built-in lineage, freshness, and regulatory accountability. Databricks supplies the unified data and AI platform building blocks, while Zensar and ZenseAI.data operationalize these for regulated financial environments. By combining these capabilities, organizations can make faster, more explainable decisions without elevating operational or regulatory risk.

The intelligence layer imperative

Why financial services AI breaks at scale—and how governed intelligence changes the outcome

Financial services executives are consistently promised that artificial intelligence will accelerate operations, reduce fraud, personalize engagement, streamline onboarding, improve credit decisions, and automate processes that were never designed for human speed. However, internal conversations often reveal a different reality: pilots appear impressive but struggle to survive production environments.

The main reason for AI initiative failures in financial services is not weak models, but the lack of an intelligence foundation designed to be shared, governed, fresh, and reusable at scale. The core takeaway is that in financial services, the competitive edge lies not in the model itself but in the governed intelligence layer — data products with embedded lineage, freshness, and the capacity for reuse. Databricks delivers the governed platform essentials, while Zensar and ZenseAI.Data enable safe, scalable deployment and operation of these components.

This white paper explains why this is significant, how it addresses executive concerns, and why Databricks and Zensar are increasingly aligned on the future of data and AI in regulated industries.

The industry problem no one can ignore anymore

Financial services have invested heavily in data over the decades, through core banking systems, payment platforms, risk engines, CRM tools, digital channels, acquisitions, and regulatory systems — all of which generate large volumes of data. The challenge is not a lack of data, but rather its fragmentation.

Customers are often present in multiple systems, each with slight variations in their definitions. Transactions may process at different speeds depending on the platform. Valuable institutional knowledge is trapped in documents, statements, contracts, call transcripts, and emails, which typically fall outside analytics workflows. Governance tends to function as a late-stage checkpoint instead of an integral design principle.

This fragmentation was manageable when reports were run overnight, and decisions occurred the next day, with humans mediating between data and action. However, AI disrupts this balance. AI systems do not reconcile definitions or intuit context; they learn from existing patterns and act rapidly. Any ambiguity in the data is magnified in the resulting decisions.

The central challenge is transforming decades of fragmented data into a single, governed intelligence layer that can reliably enable real-time, automated decision-making. Without this layer, AI initiatives are unsustainable and built on unstable foundations.

Why this keeps the c-suite awake

For CEOs, the primary concern is strategic: competitors are reducing decision times, and customers now compare experiences across industries, not just within banking. If processes take days instead of minutes, market share erodes gradually but steadily. The risk is not in adopting AI, but in doing so either too slowly or without sufficient caution.

CTOs and CIOs grapple with technical debt as architectures are stretched beyond their intended scope—pipelines originally built for small teams become critical to the entire enterprise, and essential transformations reside in legacy code. Issues in these systems can have unpredictable, widespread impacts, raising questions about the business's dependence on outdated designs.

Chief Data Officers experience the erosion of trust most acutely. When insights are delivered, they are often interrogated for their source, definitions, and reconciliation. If trust must be rebuilt with every insight, organizational speed is lost.

Chief AI Officers are concerned with scale. While a flawed dashboard can be ignored, a flawed autonomous system cannot — if an AI workflow makes thousands of decisions overnight, the organization must be able to explain, justify, and correct those actions by the next day.

These are not hypothetical worries; they are operational realities in an environment where regulators demand explainability, boards expect resilience, and customers require instant responses. The fundamental question is how to move faster without introducing systemic risk.

How these problems show up in the real world

Common scenarios in retail banking illustrate these challenges. A fraud model may flag a transaction, another system may approve it based on customer history, and yet another may block the account pending review. Each system acts correctly based on its data, but from the customer's perspective, the experience is broken, and the bank lacks confidence to automate decision-making.

In credit decisioning, different teams may build strong models using different data sources, but in production, inconsistencies in customer identity and data freshness cause unexpected edge cases, slowing approvals as exceptions accumulate.

For regulatory reporting, data may exist across structured tables, PDFs, scanned documents, and historical systems, making responses possible but slow and error-prone. These examples highlight how innovation often stalls not because of poor ideas, but because the foundational infrastructure cannot support them.

Where Databricks changes the equation

Databricks addresses these foundational issues by focusing on the system that generates intelligence, not just individual use cases. The core concept is that data engineering, analytics, and AI should all run on a shared, governed foundation rather than isolated platforms.

Delta Lake provides necessary reliability with transactional consistency, time travel, and support for both batch and streaming workloads, eliminating the need for parallel architectures and ensuring intelligence remains current. Lakeflow and Delta Live Tables transform pipelines from custom code to reusable, observable infrastructure with embedded quality checks and lineage.

Unity Catalog centralizes governance, offering unified controls and lineage tracking across all data and AI assets, ensuring trust can scale without manual intervention. MLflow bridges experimentation and production, allowing models and agents to be versioned, evaluated, promoted, monitored, and rolled back as needed. It also supports continuous improvement via evaluation and human feedback loops.

Vector Search extends governance to unstructured content, making documents, transcripts, and contracts governed and traceable rather than shadow data sources. Together, these features form the building blocks of a robust intelligence layer.

How this aligns with Databricks Industry Data Intelligence outcomes

Databricks' Industry Data Intelligence outcomes focus on delivering results that matter to executives: enabling faster decisions without sacrificing trust, creating reusable data products instead of fragile pipelines, supporting real-time intelligence anchored in governed data, and deploying AI systems that withstand audits, outages, and regulatory reviews. The intelligence layer is the mechanism that makes these outcomes achievable.

Core use cases such as fraud detection, credit underwriting, personalization, compliance, and finance become different ways of leveraging the same governed intelligence products. Speed improves not by cutting corners, but by eliminating unnecessary duplication.

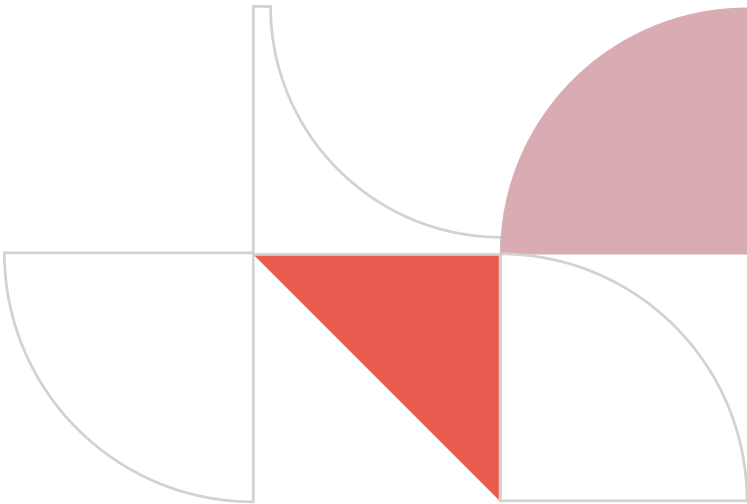

Where Zensar fits — and why that matters

Platforms alone cannot deliver operational outcomes — operating models are essential. Zensar brings practical financial services experience to the intelligence layer, defining data products around business meaning, aligning governance with regulatory expectations, and phasing adoption from proof of concept to production without introducing new risk.

Zensar recognizes that assets like Customer 360 and fraud signals must be governed, reusable products that multiple teams can rely on. This operational discipline transforms platform capabilities into executive confidence.

How ZenseAI.Data accelerates what usually takes years

Even with a strong platform and operating model, execution can be slow. ZenseAI.Data addresses this by automating processes that typically slow AI initiatives. Legacy data ingestion becomes repeatable, pipelines are generated using native Databricks services, and metadata enrichment, classification, and enforcement occur by default.

Observability expands to cover data quality, freshness, drift, and behavior, enabling business users to access governed intelligence directly. Agentic workflows are implemented with guardrails and continuously refined through MLflow evaluations and human feedback, ensuring that autonomy grows alongside trust. This approach allows organizations to transition from impressive demos to dependable, production-ready systems.

Databricks and Zensar together: Why the combination matters

Databricks provides the foundational platform, Zensar brings domain execution and operational rigor, and ZenseAI.Data accelerates delivery and reuse. Together, they cover the full intelligence lifecycle — from ingestion and governance to analytics and AI, from experimentation to production.

This collaboration enables financial institutions to reduce data latency, quickly build reusable pipelines, and scale AI adoption without increasing risk. The outcome is not just faster projects, but a more resilient and adaptable enterprise.

The message this white paper delivers

AI success in financial services is not about selecting the best model. It is about creating a system where intelligence is trustworthy, reusable, and actionable at scale. That system is the intelligence layer.

Why Databricks. Why Zensar.

Databricks solves the platform challenge at enterprise scale, while Zensar addresses execution in regulated environments. Together, they resolve the confidence issues that prevent the adoption of operational AI. After the proof-of-concept stage, platforms like Databricks evolve from project tools into shared enterprise services. Success at this stage depends not just on platform features, but on the ability to operationalize them consistently. Zensar provides this operational layer, transforming Databricks into a reliable, governed intelligence service for the business.

Databricks supplies the core data and AI tools, and Zensar defines its enterprise-scale use. With ZenseAI.Data, processes like ingestion, pipeline management, governance, and observability are automated and standardized, replacing custom engineering with reusable, production-grade assets. Intelligence is created once, governed by design, and safely used across teams, use cases, and regulatory boundaries.

This ensures that pilots are sustainable within the enterprise, that data products persist beyond individual projects, and that governance is efficiently scaled. The result is faster delivery without extra risk — turning experimentation into a durable organizational capability rather than technical debt.

The outcomes customers are actually targeting

The objectives are concrete and measurable across the financial services sector. In asset and wealth management, the emphasis is on trusted intelligence to enhance portfolio decisions, risk monitoring, and regulatory reporting, with faster access to insights, explainable risk assessments, and the ability to extract intelligence from diverse sources without sacrificing governance. In banking, leaders seek personalized engagement that respects privacy, straight-through processing with human oversight, and automated compliance responses.

All these outcomes depend on intelligence that is timely, explainable, governed, and reusable. A governed intelligence layer transforms isolated improvements into scalable, enterprise-wide capabilities.

The final question

Financial services leaders no longer debate whether AI will play a role in their future — it already does. The real question is whether their AI will be based on raw data or on intelligence by design. Databricks provides the foundational platform.

What kind of intelligence layer are you building?

Key takeaways and conclusion

AI success in financial services hinges on intelligence by design, not mere experimentation. Organizations relying on fragmented data struggle with trust, scalability, and regulatory scrutiny. A governed intelligence layer transforms data into reusable, auditable products that enable real-time decision-making across fraud, credit, compliance, and personalization. Databricks delivers the platform foundation, while Zensar and ZenseAI.Data add the operational rigor necessary to move from pilot projects to production systems. The result is resilient AI that executives, regulators, and customers can trust at scale.

Author:

Daniel Gomez
Head - Partner Sales,
Data and Analytics

zensar

An RPG Company

At Zensar, we're 'experience-led everything.' We are committed to conceptualizing, designing, engineering, marketing, and managing digital solutions and experiences for over 145 leading enterprises. Using our 3Es of experience, engineering, and engagement, we harness the power of technology, creativity, and insight to deliver impact.

Part of the $4.8 billion RPG Group, we are headquartered in Pune, India. Our 10,000+ employees work across 30+ locations worldwide, including Milpitas, Seattle, Princeton, Cape Town, London, Zurich, Singapore, and Mexico City.

For more information, please contact: **info@zensar.com | www.zensar.com**