

AI-ready banking starts with trusted, resilient data

Why the data foundation determines whether banking AI moves from pilot to production

How banks operationalize AI is evolving. From real-time fraud detection to online payments, AI enables banks to make smarter decisions faster. Modernizing infrastructure is the first step, but infrastructure performance is just part of the equation. Banks must also ensure their data is trusted and available in real time to support AI at scale.



Executive summary

Why pilots stall

More than half of AI proof of concepts never reach production, and 94%¹ of organizations say data quality is critical to AI success. Those numbers illustrate a reality many banks are experiencing firsthand. **The challenge is no longer proving AI can work.** It's building an operational foundation that allows AI to perform reliably in production.

What production requires

As AI becomes embedded in fraud detection, real-time payments, and risk management, success depends far more than the model itself. **Banks need reliable data, clear governance, operational resilience, and the ability to deliver the right information at the right time.** Without that foundation, even promising AI pilots struggle to scale into production.

Why bank data is the obstacle

Banks already have enormous volumes of information. The challenge is understanding what exists, what it means, who owns it, and whether it can safely participate in AI workflows. As AI initiatives expand, those questions become increasingly important, and increasingly difficult, to answer.

What the bank needs to build

Building an AI-ready bank requires more than infrastructure modernization. **Banks need to understand and govern the data AI depends on, while ensuring that data remains available and resilient as workloads scale.** Understanding those requirements is the first step toward moving AI from pilot to production.

In this ebook, we'll discuss...

- Why most banking pilots never reach production
- What separates production AI from proof of concept
- How leading banks are building AI-ready data foundations

56%¹
of AI proof of concepts
fail to reach production

94%¹
of organizations identify data
quality as important
or extremely important
to AI success

Why production AI is different from pilot AI

Developing production AI requires a dynamic understanding of enterprise data. That includes continuously discovering and classifying information, understanding semantic relationships between datasets, maintaining a business-aware view of enterprise information, ensuring context evolves as data changes, and tracking the provenance of AI-driven decisions and actions.

Without that context, banks risk relying on incomplete, inaccurate, or duplicative information, exposing sensitive data, or using data that is not appropriate for the intended AI task.

AI and data governance

Before data enters an AI workflow, banks should be able to answer four questions:



Where does the data reside?



Who can access it?



Which jurisdiction(s) governs it?



Is it approved for AI use?



Production AI changes banking risk

In banking, AI increasingly supports high-volume, time-sensitive workflows where automation and human oversight must work together. As decisions happen faster, the need for clear review and intervention becomes even more important, especially when data is incomplete, inaccurate, or lacks the context needed to interpret it correctly.

Production AI depends on trusted, real-time data that remains available across critical banking systems. Infrastructure disruptions, poor data quality, or governance gaps interrupt AI-driven workflows and affect payments, fraud prevention, and other business-critical operations.

When data is not properly prepared for AI, it can have serious ramifications:



FRAUD EXPOSURE

AI evaluates transactions in real time. Incomplete or inaccurate data increases the odds that fraud may go undetected.

Potential for financial losses, regulatory exposure



FALSE POSITIVES

AI makes decisions based on the provided data. When incomplete or inaccurate data is provided, it can lead to unwarranted transaction delays.

Potential for customer frustration, operational costs, attrition



COMPLIANCE ISSUES

AI uses data from multiple sources. Without embedded oversight, banks may not be able to demonstrate compliance.

Potential for regulatory penalties, audit findings, reputational damage

Why banking AI pilot efforts stall

Banks often prove an AI use case in a pilot, but operating it in production introduces a different set of requirements. AI systems must continuously access validated data, operate within policy requirements, remain observable, recover from failures, and scale across the enterprise. Without those operational capabilities, even successful pilots struggle to become business-critical applications. Issues may include:

The real test is not whether a model works in a pilot. It is whether the bank can trust it in production.

Inadequate data infrastructure.

Production AI requires a resilient data foundation that can deliver trusted information in real time and remain available as workloads scale.

Missing context.

Data access alone is not enough. Banks need the business and policy context that allows AI to interpret enterprise information correctly and determine whether it is appropriate for a given use.

Disconnected governance.

Governance cannot be added at the end of the AI workflow. Banks need policy, ownership, and oversight embedded throughout the data lifecycle so they can maintain control and demonstrate how data is being used.



Building the AI-ready data foundation

Production AI is built incrementally. Banks that establish a trusted data foundation before expanding AI initiatives are better positioned to scale securely and deliver consistent business value.

1

Identify one production use case

Start with a specific banking workflow and define the desired outcome.

2

Identify the required data

Determine what data the use case needs, where it resides, who owns it, and how current it is.

3

Understand and govern the data

Identify sensitivity, lineage, business context, policy requirements, and approved use.

4

Address data and infrastructure gaps

Resolve silos, access issues, availability constraints, and other barriers that could prevent the workload from operating reliably.

5

Validate production readiness

Test the use case under realistic conditions, including scale, performance, observability, disruption, and recovery.

6

Improve and expand

Close remaining gaps, then apply the model to additional workloads and data sources.

AI readiness checklist

AI readiness starts before production. Use these five questions to identify gaps across data, governance, security, infrastructure, and AI operations.

Ask your Chief Data Officer:

Do we know what data is required for the use case, where it resides, who owns it, and how current it is?

What good looks like:

Relevant data is identified across systems, with clear ownership, quality, source, and availability.

Watch for:

Unknown data sources, duplicate copies, unclear ownership, or data that cannot be reliably located.

Ask your Head of Data Governance:

Can we explain where the data came from, what it means, and whether it is approved for the intended AI use?

What good looks like:

The data has clear lineage, classification, business context, and documented rules for how it can be used.

Watch for:

Missing lineage, unclear business meaning, or governance applied only after the data enters the AI workflow.

Ask your CISO:

Do we know who can access the data, where it resides, and which privacy, security, and jurisdictional requirements apply?

What good looks like:

Sensitive data is identified, access is controlled, and location and policy requirements are understood before the data is used.

Watch for:

Unclassified sensitive information, unclear access rights, or data crossing environments (or jurisdictions) without appropriate controls.

Ask your CIO:

Can our infrastructure deliver the performance, availability, observability, and recovery required for the use case?

What good looks like:

The workload has been tested under realistic production conditions, including scale, disruption, monitoring, and recovery.

Watch for:

Manual processes, inconsistent performance, limited monitoring, or recovery plans that have not been tested.

Ask your Chief AI Officer:

Do we understand how the AI system will use the data, how its outputs will be evaluated, and where human oversight is required?

What good looks like:

The use case has clear success criteria, approved data, evaluation standards, escalation paths, and accountable owners.

Watch for:

Unclear outcomes, undefined evaluation criteria, weak oversight, or no owner for AI-driven decisions.

What to do next

If any answer is unclear, treat it as a production-readiness gap. Assign an owner, define what must be validated, and resolve the issue before expanding the use case. If you have any questions, contact Everpure.

Introducing Everpure Data Intelligence

Production AI depends on more than access to data. Banks need to know what data exists, what it means, how it is related, how sensitive it is, and whether it is appropriate for a specific AI use. Everpure Data Intelligence provides that context across hybrid environments.

By continuously discovering, classifying, and contextualizing enterprise data, Data Intelligence helps data, governance, and technology teams understand what they have, where it resides, and how it can be used safely in AI and compliance workflows.

Discover and map enterprise data

with complete, continuous visibility into sensitive data estates across all environments.

Understand data in context

Classify and enrich enterprise data with semantic, sensitivity, lineage, relationship, and policy context so AI systems can use it appropriately.

Enforce governed AI access

by providing a trusted semantic catalog that agents, applications, and workflows can use to find relevant data while respecting policy and sensitivity context.

Prepare data for AI and retrieval-augmented generation

by identifying sensitive, stale, copied, or low-trust data before it is used in AI pipelines, helping reduce risk, improve relevance, and optimize token usage.

What Everpure Data Intelligence delivers

Data accuracy

Attain 98.6%² precision mapping and classification of all data across all environments.

Cost reduction

Reduce scanning costs by 99%².

Compliance

Know your sensitive data is identified, protected, and compliant with the latest regulations, from GDPR to CCPA and beyond.

AI readiness

Find relevant, trusted, and policy-approved data before it enters AI pipelines or agent workflows.

How Everpure helps banks operationalize AI

Moving AI from pilot to production in banking requires infrastructure that can deliver data reliably at scale, support AI applications across hybrid environments, simplify operations, and govern how enterprise data is used.

Deliver and protect the data

Provide resilient, high-performance access to the structured and unstructured data production AI depends on.

Run and scale AI applications

Provide data services for cloud-native AI applications across Kubernetes and hybrid environments.

Operate the environment

Simplify infrastructure operations through management, automation, and observability.

Understand and govern the data

Discover, classify, and contextualize enterprise data so teams can determine what data exists and how it can be used appropriately.

Together, these capabilities help banks operate the data foundation required to run AI reliably in production.

AI infrastructure



Real world evidence and proof

Before banks can safely use enterprise data in AI, they need to know where sensitive information exists, how it is classified, and whether it is appropriate for use. This example shows how automated discovery and classification can help address that foundational challenge at scale.

Challenge

A national payment processor lacked visibility into tens of millions of misclassified credit card records, increasing liability risks.

They also faced challenges meeting regulatory compliance due to fragmented data systems.

Solution

The company implemented Data Intelligence (formerly 1touch.io) for automated sensitive data discovery and protection across complex systems.

Outcomes

The payment processor was able to identify and remediate risks in 91 million sensitive records, eliminating significant liabilities.

It also strengthened its compliance posture, fostering trust among both internal and external stakeholders.

91 million
sensitive credit card records identified and remediated.

Helped reduce compliance risk
through automated sensitive data discovery.

Strengthened security and compliance posture.

Enabled ongoing automated discovery and classification
of sensitive data for proactive risk management.

Your data foundation determines AI success or failure

Successful AI initiatives are built long before models reach production. Banks that move beyond pilots establish a data foundation that delivers trusted information in real time, strengthens policy guardrails without slowing innovation, and provides the operational resilience required for business-critical AI workloads.

Everpure combines resilient data infrastructure, modern application data services, operational management, and data intelligence to help banks build a stronger foundation for production AI.

To learn more, visit [Everpure Banking](#)

1. Building AI-Ready Data Storage Infrastructure: Insights for Enterprise Success, IDC, March 2026
2. Unlock Enterprise AI at Scale with Everpure Data Intelligence, Everpure, 2026



©2026 Everpure, the Everpure P Logo, Pure Storage, Portworx, and the marks in the Everpure Trademark List are trademarks or registered trademarks of Everpure, Inc. or its licensed subsidiaries in the U.S. and/or other countries. The Trademark List can be found at everpuredata.com/trademarks. Other names may be trademarks of their respective owners.

Contact