Omdia by informa techtarget

APRIL 2026

# Fraud Detection and Credit Risk Assessment with Cloudera AI, Accelerated by NVIDIA on Dell Technologies

Stephen Catanzano, Senior Analyst

Abstract: Financial institutions face mounting pressure to detect fraud in real time and assess credit risk with greater accuracy, all while navigating increasingly complex regulatory requirements. These efforts are driven by the need to improve the customer experience, grow revenue, and protect both the firm's and its customers' assets. Traditional machine learning (ML) models deliver fast scoring but lack the contextual reasoning and explainability that regulators and investigators demand. A new approach combines traditional ML with generative and agentic AI to move from simple binary decisions to intelligent, explainable outcomes. Cloudera AI, built on NVIDIA-accelerated computing and AI software, deployed on Dell Technologies infrastructure, provides the trusted data foundation and high-performance AI capabilities financial institutions need to enhance their fraud detection and credit risk assessment.

## The AI imperative in financial services

**78% of financial services organizations said AI agents are their top priority or a high priority compared to other AI initiatives.**

Financial services institutions are rapidly prioritizing AI agents as a strategic imperative. The promise of automating complex workflows, improving decision accuracy, and reducing operational costs has elevated AI from experimental to a board-level priority. Within financial services specifically, organizations are targeting AI agents for their most consequential processes. Anti-money laundering detection, credit scoring and underwriting, risk management, and fraud detection rank among the top use cases being deployed or actively considered, as shown in Figure 1.[1] These are not peripheral applications; they represent the core operational and compliance functions that define competitive advantage and regulatory standing.

The shift reflects a fundamental change in what financial institutions expect from AI. Where earlier implementations focused on isolated predictions, today's requirements demand systems that can reason across multiple data sources, explain their decisions, and take coordinated action. This evolution from simple scoring to intelligent reasoning is reshaping how institutions approach their most critical challenges.

[1] Enterprise Strategy Group (now Omdia) Research Report, *AI Agents: The Game-changing Generative AI Use Case*, August 2025. All Enterprise Strategy Group research references and charts in this Showcase are from this report.

**Figure 1.** Top financial services use cases for AI agents

**Which of the following use cases for AI agents are currently deployed or under consideration within your organization? (Percent of respondents, N=65, multiple responses accepted)**


*Source: Omdia*

## Why traditional approaches fall short

Despite significant investment in ML over the past decade, financial institutions continue to struggle with the limitations of traditional approaches. In fraud detection, rules-based engines and ML scoring models generate high volumes of false positives, overwhelming investigation teams and frustrating legitimate customers. Point solutions for fraud detection often rely on transaction data alone, missing important contextual clues from other data sources, such as customer behavior patterns, device information, and external signals. When a transaction is flagged, investigators must manually reconstruct the context by reviewing transaction history, customer interactions, and related signals to determine whether fraud occurred. In many cases, the ML model provides a score but no explanation for why the transaction was suspicious.

**Credit risk assessment faces parallel challenges. 43% of financial services organizations cited managing risk and AI-driven decision-making in financial transactions as one of their most significant challenges in adopting AI agents.**

Credit risk assessment faces parallel challenges. Traditional approaches rely heavily on credit bureau scores and structured payment history, creating thin-file rejections for viable customers who lack conventional credit records. For commercial lending, analysts spend more time gathering and organizing data than actually assessing risk. Reviews of loan applications remain slow and inconsistent, with critical signals buried in unstructured documents like earnings reports, news articles, and market analyses going unexamined.

Both domains share a common problem: Existing systems produce decisions without reasoning. Financial institutions need AI that cannot only score but also investigate, explain, and act.

## The convergence of traditional ML, generative AI, and agentic AI

The most sophisticated financial institutions are now deploying a layered AI architecture that combines the strengths of traditional ML, generative AI (GenAI), and agentic AI.

Traditional ML continues to provide the foundation, with high-speed scoring models that can evaluate transactions in under 50 milliseconds, identifying anomalies and calculating risk probabilities at scale. These models excel at pattern recognition across structured data and deliver the real-time performance that fraud detection and credit decisioning require.

GenAI adds the capability for synthesis and explanation. Large language models can consume thousands of pages of financial documents, extract key risk factors, and generate narrative summaries that explain complex situations in natural language. For fraud cases, GenAI can produce investigation briefs that contextualize why a transaction was flagged. For credit decisions, it can draft preliminary credit memos that synthesize diverse data sources.

Agentic AI orchestrates these capabilities into coordinated workflows. Rather than requiring human analysts to manually gather information and connect systems, AI agents can autonomously retrieve relevant data, invoke appropriate models, and assemble comprehensive assessments. The agent performs the first level of investigation, significantly reducing the manual workload while maintaining human oversight for final decisions.

This convergence represents a shift from credit scoring to credit reasoning and from fraud flagging to fraud investigation. Financial institutions gain both faster decisions and better-explained, more defensible outcomes.

## Use case: Real-time fraud detection

Traditional fraud detection forces a tradeoff: Cast a wide net and overwhelm investigators with false positives or narrow the criteria and miss sophisticated fraud schemes. The combination of Cloudera, NVIDIA, and Dell Technologies eliminates this tradeoff.

**The solution begins with Cloudera Data in Motion. 31% of financial services organizations said ensuring AI-driven fraud detection aligns with compliance requirements was a key challenge in adopting AI agents.**

The solution begins with Cloudera Data in Motion, capturing real-time transaction streams alongside customer interaction history, device metadata, and behavioral signals. Traditional ML models, accelerated by NVIDIA GPUs and AI software, score each transaction in real time. When a transaction is flagged, the system shifts from scoring to investigation. Cloudera AI orchestrates an agentic workflow that retrieves relevant context such as recent transaction patterns, customer communication history, geographic signals, and related account activity. Large language models accelerated by NVIDIA then generate a case file with a natural-language summary of the risk factors and evidence. Investigators receive not just a fraud score but a preliminary investigation brief explaining why the transaction warrants review. The time from flag to informed decision drops dramatically while compliance requirements for explainability and audit trails are satisfied. These advancements lead to tangible business outcomes, including reduced fraud, fewer false positives, improved customer experience, and lower operational costs, as shown in Figure 2.

**Figure 2.** Cloudera AI – Accelerating world-class AI apps and agents for financial institutions

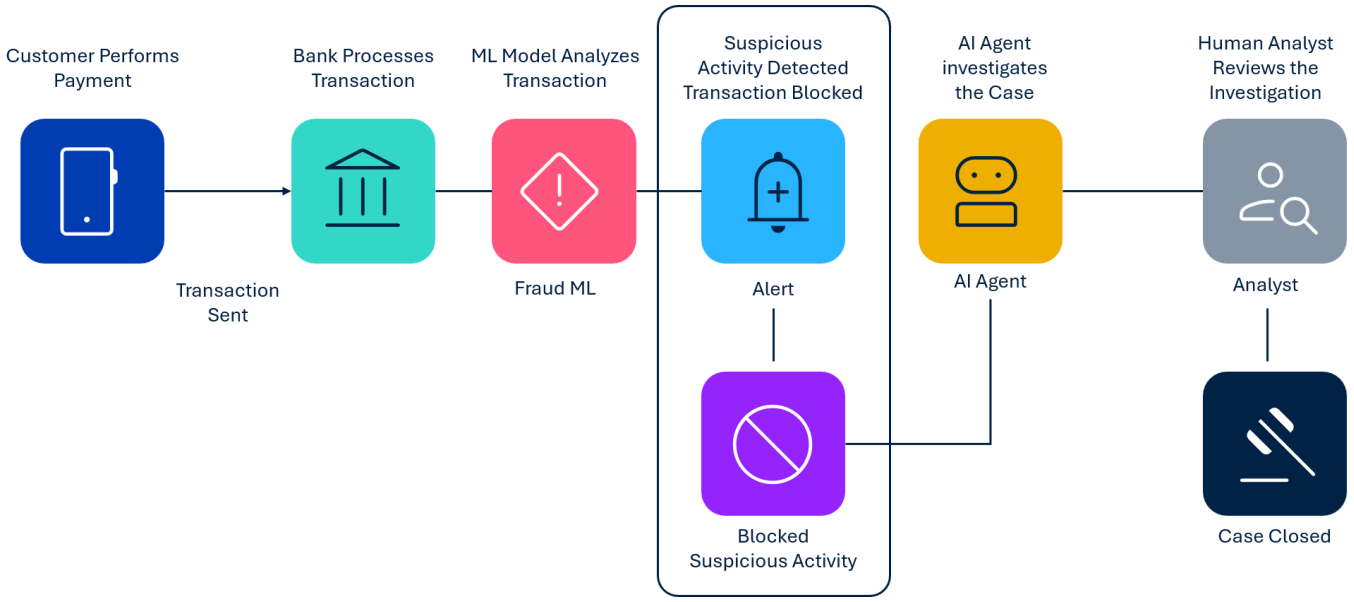


*Source: Cloudera and Omdia*

## Use case: Intelligent credit risk assessment

**45% of financial services organizations said they were concerned about the inaccuracy of AI agent decision-making.**

Credit risk assessment has long been constrained by the data it could practically analyze. Bureau scores and payment history provide a narrow view of creditworthiness, while the unstructured information that could enrich decisions, financial statements, market analyses, and news coverage remained too labor intensive to incorporate systematically.

Cloudera AI transforms this process by unifying structured and unstructured data in a governed lakehouse environment. For consumer lending, alternative data sources can be incorporated alongside traditional bureau data to serve thin-file customers who would otherwise be rejected. For commercial lending, the platform can ingest and analyze thousands of pages of financial documents, earnings reports, and market intelligence.

NVIDIA-accelerated computing enables two critical capabilities. First, complex risk simulations run dramatically faster, enabling more sophisticated risk modeling without extending decision timelines. Second, large language models can synthesize complex document sets and generate preliminary credit memos that highlight key risk factors for analyst review.

The result is an AI-powered credit analysis assistant that augments human judgment rather than replaces it. Loan officers receive applications with risk factors already extracted, relevant comparables identified, and credit memo drafts prepared. The platform ensures that AI-generated insights are traceable and accurate, which is critical for regulatory compliance. These capabilities translate into significant business benefits: an improved customer experience, greater revenue opportunities, reduced risk, and access to a larger total addressable market, as firms can confidently expand their credit offerings.

## Cloudera AI with NVIDIA and Dell Technologies: The platform for trusted financial services AI

Deploying AI across fraud detection and credit risk assessment demands more than model capabilities; it requires a platform that ensures data governance, regulatory compliance, and operational performance. Together, Cloudera, NVIDIA, and Dell Technologies provide a complete solution that is purpose-built for the requirements of financial services (see Figure 3).

**Figure 3.** The combined solution for financial services

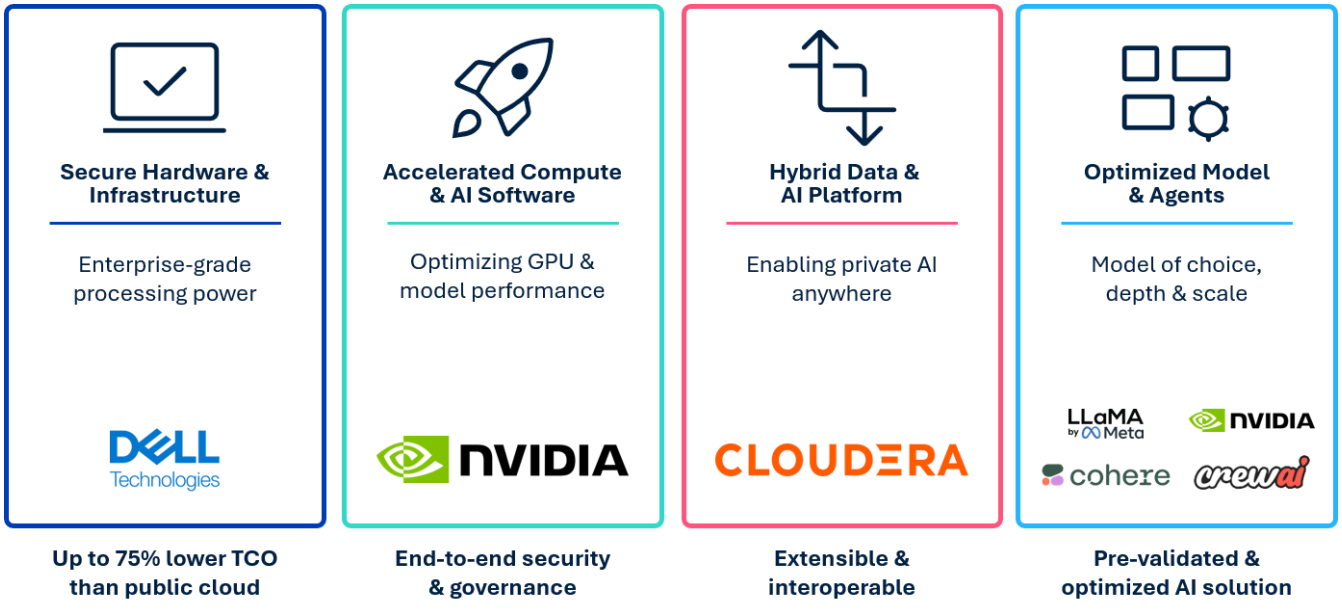


*Source: Cloudera and Omdia*

Cloudera provides the trusted data foundation. The Cloudera data platform unifies diverse data types in an open data lakehouse architecture, ensuring sensitive personally identifiable information never leaves its jurisdiction and providing fine-grained access controls and complete data lineage. Cloudera Data in Motion captures real-time transaction streams and integrates historical data across hybrid environments, avoiding unnecessary data movement. Cloudera's Shared Data Experience delivers consistent governance and security whether data resides on premises, in private clouds, or across public cloud services.

**51% of financial services organizations said they were concerned about data privacy, handling sensitive financial data, and consent when using AI agents.**

NVIDIA accelerates the AI pipeline. NVIDIA Blackwell GPUs dramatically reduce training time for traditional ML models and enable large-scale inference for neural networks and large language models. For credit risk, NVIDIA acceleration transforms Monte Carlo simulations for probability-of-default calculations from overnight batch jobs into near-real-time analyses. For fraud detection, GPU-accelerated inference ensures that even sophisticated models meet the sub-50 millisecond latency requirements of real-time transaction processing. Cloudera AI Inference, powered by NVIDIA technology, enables organizations to deploy and scale AI models on premises, including NVIDIA Nemotron for LLMs, fraud detection, computer vision, and more. Leveraging the NVIDIA AI stack, Blackwell GPUs and NVIDIA Dynamo-Triton Inference Server ensure secure, governed AI deployment at enterprise scale with superior performance and cost efficiency.

Dell Technologies provides cost-effective, enterprise-grade infrastructure optimized for AI workloads. Dell PowerEdge servers with NVIDIA Blackwell GPUs deliver the compute density financial institutions need while keeping sensitive data on premises where privacy and regulatory requirements demand. The Cloudera Data and AI platform, powered by NVIDIA and deployed on Dell infrastructure, creates an integrated environment for developing, deploying, and operating AI at scale.

The key differentiator of this combined platform is the ability to bring AI processing to the data rather than moving sensitive financial data to external AI services. This architecture minimizes data movement, maintains regulatory compliance, and leverages NVIDIA acceleration to reduce inference latency, critical requirements for real-time financial services applications.

## Conclusion

Financial services institutions can no longer afford the limitations of traditional approaches to fraud detection and credit risk assessment. The convergence of traditional ML, generative AI, and agentic AI offers a path forward—but only when built on a foundation of trusted, governed data and accelerated by infrastructure designed for AI at scale.

Cloudera AI, leveraging NVIDIA accelerated compute and AI software and deployed on Dell Technologies infrastructure, delivers this foundation. The platform enables financial institutions to move from simple scoring to intelligent reasoning: fraud detection that investigates as well as flags and credit assessment that synthesizes and explains as well as scores. Unified governance ensures that AI decisions remain explainable, auditable, and compliant with regulatory requirements.

For financial services organizations ready to transform their approach to fraud detection and credit risk, Cloudera, NVIDIA, and Dell Technologies provide the trusted platform to make it possible. Omdia recommends engaging with Cloudera to explore how this solution can accelerate your AI initiatives while maintaining the governance and compliance posture your business demands.

To learn more, visit the Cloudera Financial Services page.



**Get in touch:** www.omdia.com askananalyst@omdia.com