JP Morgan Chase Saves $50M By Transforming Fraud Detection with Real-Time Graph Analytics

THE CHALLENGE →

Existing fraud detection methods missed complex patterns, costing hundreds of millions annually. Legacy systems struggled with accuracy and couldn't scale to handle 50M+ daily transactions under 80ms response time requirements.

THE SOLUTION →

Implemented TigerGraph's high-speed graph database to analyze relationships and patterns in real-time. Integrated ML algorithms with graph features to detect synthetic fraud through complex queries and centrality clustering of suspicious accounts.

THE RESULTS ✢

Achieved $50M in operational savings while protecting 60M households. Enhanced fraud detection accuracy with 30TB+ data processing capacity and sub-80ms response times for real-time transaction analysis.

Introduction:

JP Morgan Chase's Consumer and Commercial Banking division stands as the largest credit card issuer in the United States, serving over 140 million accounts and processing more than 50 million transactions daily. With such massive scale and responsibility for protecting millions of customers' financial assets, the bank required cutting-edge technology solutions to combat increasingly sophisticated fraud schemes while maintaining seamless customer experiences.

The Challenge:

JP Morgan Chase faced a critical challenge in fraud detection despite having systems that were considered industry-standard. Undetected fraudulent transactions and applications were costing the bank hundreds of millions of dollars annually, highlighting significant gaps in their existing approach. The bank's legacy fraud detection methods struggled with accuracy and were unable to scale effectively to meet the demands of their massive transaction volume.

The core issue was the need to identify transactional and relationship patterns far more complex than what traditional data table analysis could reveal. These patterns needed to be discovered and delivered in under 80 milliseconds to support real-time decision-making for their machine learning inference pipeline. With over 50 million daily transactions from a large, varied customer base generating data points at massive scale, the existing systems were simply inadequate. The legacy approaches were not only difficult to adapt to changing regulations but also fundamentally unable to handle the sophisticated fraud schemes that were evolving in the market.

The Solution:

JP Morgan Chase implemented TigerGraph's advanced graph database technology to revolutionize their fraud detection capabilities through real-time relationship analysis. The solution enabled the bank to aggregate signals and uncover complex patterns by identifying suspicious relationships through sophisticated graph queries that went far beyond traditional database capabilities. The system integrated machine learning algorithms specifically designed to detect synthetic fraud, utilizing centrality analysis and clustering techniques to identify suspicious account networks.

The architecture combined transaction data, account holder information, and merchant data into a comprehensive graph structure that could be updated in bulk on a regular basis without interrupting ongoing query operations. At the heart of the system was a real-time streaming pipeline that processed requests for graph patterns and algorithms, characterizing the risk level of each transaction and its associated parties. Results were delivered as numeric vectors for seamless integration into existing machine learning models, which were retrained frequently to adapt to new and evolving fraud patterns. To ensure high availability and fault tolerance across JP Morgan Chase's geographic range, TigerGraph replicas were deployed across multiple regions, providing the scalability and reliability required for such critical operations.

The Results:

The implementation delivered transformative results that exceeded expectations across multiple dimensions of performance and value creation. JP Morgan Chase achieved $50 million in operational savings while simultaneously protecting 60 million households from fraudulent activities. The system successfully processes over 30TB of raw data capacity, demonstrating its ability to handle massive scale operations without compromising performance.

The architecture now features high-availability TigerGraph database replicas, regularly scheduled machine learning model training that combines graph features with traditional data features, and real-time streaming ML inference pipelines for fraud prediction. Multiple high-speed, event-driven graph queries run for each credit card transaction, providing comprehensive risk assessment in real-time. The system maintains sub-80-millisecond response times while processing the bank's enormous daily transaction volume, proving that advanced graph analytics can operate at enterprise scale without sacrificing speed or accuracy.

Impact Summary:

VALUE	SPEED	REDUCED RISK
$50M in operational savings achieved through improved fraud detection accuracy and reduced false positives	60M households protected from fraudulent transactions and identity theft through enhanced pattern recognition	30TB+ raw data processing capacity enabling comprehensive analysis of transaction relationships and patterns

About TigerGraph

TigerGraph, the enterprise AI infrastructure and graph database leader, delivers massively parallel storage and computation that scales independently and without size limits, to meet the changing workloads and growing data volumes required for crucial business needs and AI adoption within companies. By providing visibility into the multidimensional data connections and relationships, TigerGraph has become a trusted partner to leading companies including JPMC, Mastercard, Microsoft, and Unilever successfully solving fraud detection, entity resolution, customer 360, supply chain management, and other problems. TigerGraph is headquartered in Silicon Valley, California.

www.tigergraph.com