



Rethinking Financial Data Engineering: From Batch Systems to AI-Native Architectures

Chittaranjan Pradhan

Independent Researcher, Data Engineering & AI, USA

cpradhan01@gmail.com

Publication History: Received: 21.03.2026; Revised: 21.04.2026; Accepted: 23.04.2026; Published: 28.04.2026

ABSTRACT: This paper examines the transformation of financial data engineering from traditional batch-oriented architectures to real-time, AI-native data ecosystems. It identifies key limitations in legacy systems, including latency, fragility, and lack of scalable data quality frameworks, and presents architectural principles derived from large-scale financial platforms.

The study introduces a governance-centric framework integrating real-time data streaming, automated data lineage, and Data Quality as a Service (DQaaS), supported by production-scale implementations processing over 500,000 GB daily and enabling over one million analytical queries.

Additionally, the paper explores the convergence of data engineering and artificial intelligence, demonstrating how AI-ready pipelines support fraud detection, regulatory compliance, and predictive analytics. The findings provide both theoretical insights and practical frameworks for designing next-generation financial data systems.

KEYWORDS: Data Engineering, Apache Kafka, Real-Time Analytics, Financial Technology, Data Quality, Machine Learning Infrastructure, Enterprise Data Architecture, Banking Systems, Change Data Capture

I. INTRODUCTION

The global banking industry generates and consumes data at a scale and velocity that few other sectors can match. Trade executions, risk calculations, regulatory reports, compliance alerts, and client interactions produce a continuous flow of structured and semi-structured data across time zones and asset classes. For decades, the primary mechanism for managing this complexity was the relational data warehouse — a powerful but fundamentally batch-oriented paradigm that enabled historical analysis but proved inadequate for real-time decision-making and operational intelligence.

These limitations became increasingly evident as financial markets evolved. Algorithmic trading compressed decision cycles from minutes to microseconds, while regulatory mandates — including Basel III and MiFID II — introduced stringent requirements for timeliness, transparency, and auditability. At the same time, the rapid adoption of artificial intelligence and machine learning created demand for high-volume, low-latency data pipelines that traditional architectures were not designed to support.

A major inflection point in this evolution has been the emergence of large-scale regulatory infrastructures such as the Consolidated Audit Trail (CAT), overseen by FINRA. These systems require the capture, normalization, and traceability of billions of daily events across the full lifecycle of trade activity, significantly increasing the need for robust data lineage, reconciliation accuracy, and near real-time reporting. As a result, regulatory compliance has become a central consideration in the design of modern financial data platforms.

This paper synthesizes thirty years of applied engineering experience with architectural analysis to examine how leading financial institutions have addressed these challenges. It draws on production-scale implementations supporting global trading, risk, and compliance systems, including data platforms processing hundreds of terabytes of data daily and serving high-volume analytical workloads.

The study presents key architectural principles underlying modern financial data platforms, including event-driven data processing, distributed data lineage, and service-oriented approaches to data quality management. It further explores



how these systems enable regulatory reporting and advanced analytics, improving scalability, operational efficiency, and data reliability in complex financial environments.

Finally, the paper considers the convergence of data engineering with artificial intelligence, highlighting a shift toward integrated, real-time data ecosystems that support continuous analytics and adaptive decision-making. These developments suggest a broader transformation in how financial institutions design, govern, and utilize data at scale.

The remainder of this paper is organized as follows: Section 2 reviews the limitations of legacy financial data infrastructure. Section 3 presents the core architectural principles of modern financial data platforms. Section 4 addresses data quality as critical infrastructure. Section 5 discusses the convergence of data engineering and AI. Section 6 proposes a framework for the next generation of financial data systems. Section 7 concludes.

II. THE LEGACY PROBLEM: STRUCTURAL LIMITATIONS OF BATCH-ORIENTED WAREHOUSING

Financial data infrastructure built between the 1980s and early 2000s was designed for a world of overnight batch processing. Oracle, Sybase, and Netezza platforms organized data into schemas that served reporting needs well but were architecturally incompatible with the demands of real-time market operations. These systems shared three fundamental structural limitations.

2.1 Temporal Blindness

Batch-oriented systems produce a single, authoritative view of data as of a processing cutoff — typically end-of-day. This temporal blindness means that any event occurring after the cutoff is invisible to downstream consumers until the following processing cycle. In trading environments where positions change thousands of times per hour, this latency represents not just an inconvenience but a measurable risk exposure. A trading desk operating on stale position data is exposed to market movements that its risk systems cannot yet see.

2.2 Fragile Integration Architectures

Legacy financial platforms were typically built as point-to-point integrations: a trade capture system feeding a risk system feeding a reporting system, each with its own data model and transformation logic. Over years of organizational growth and acquisition, these integrations accumulated into brittle dependency graphs. A schema change in one upstream system could propagate failures through dozens of downstream consumers. The operational burden of maintaining these interdependencies became itself a major constraint on organizational agility.

2.3 Data Quality as Manual Process

In the absence of systematic quality infrastructure, data quality in legacy environments was enforced through manual reconciliation: overnight jobs comparing source and target record counts, analyst reviews of exception reports, and escalation procedures when discrepancies were found. This approach was costly, slow, and fundamentally reactive — problems were discovered after they had already propagated through the pipeline. The 2008 financial crisis brought renewed regulatory focus to data governance, exposing the inadequacy of manual quality controls at scale.

The cumulative effect of these limitations was a financial services sector in which data infrastructure had become a strategic constraint. Institutions that could not modernize their data architectures faced competitive disadvantage in algorithmic trading, analytical capability, and regulatory compliance. The imperative to transform was clear; the path forward required rethinking architecture from first principles.

III. ARCHITECTURAL PRINCIPLES OF MODERN FINANCIAL DATA PLATFORMS

The transformation from legacy to modern financial data infrastructure is not primarily a question of technology selection — the tools (Kafka, Spark, Snowflake, Trino, Apache Pinot) are available to any organization. It is a question of architectural philosophy: the set of governing principles that determine how data moves, how it is organized, and how it is made trustworthy at scale.

The architecture illustrates a streaming-first financial data platform integrating source systems, Kafka-based event ingestion, schema validation, DQaaS, data lineage, observability, cloud storage, real-time analytics, and AI/ML pipelines for risk, fraud detection, regulatory reporting, and enterprise decision intelligence.

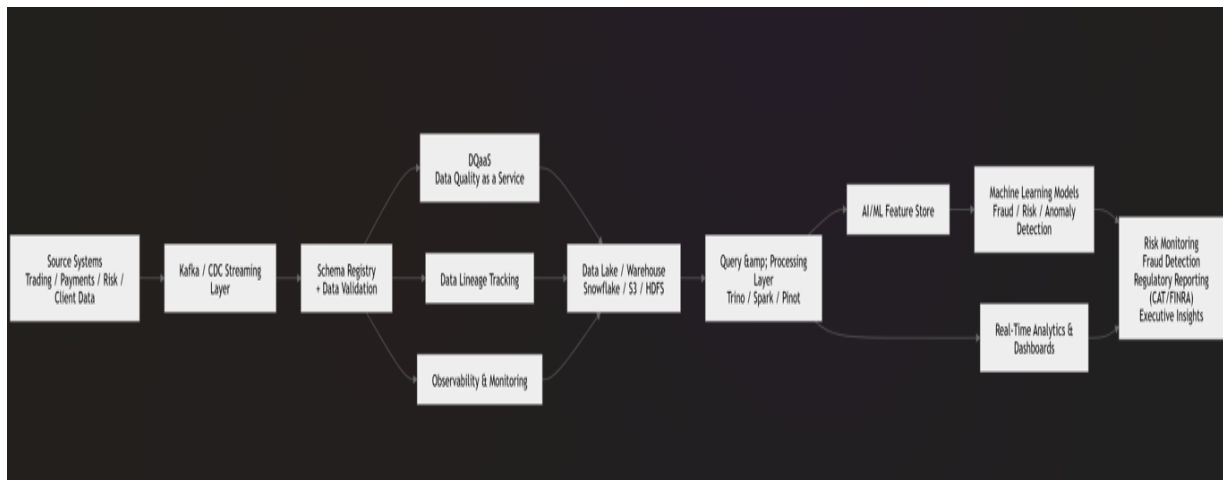


Figure1. End-to-End Financial Data Platform Architecture

3.1 Streaming-First Design

The most consequential architectural inversion in modern financial data engineering is the shift from batch-first to streaming-first design. In a streaming-first architecture, every data-generating event is captured and propagated immediately through a durable, ordered log — Apache Kafka being the dominant implementation. Batch processes become a derivative of the stream, not the primary data movement mechanism.

This inversion has profound operational consequences. Change Data Capture (CDC) techniques allow database mutations to be replicated with sub-second precision. Downstream consumers can subscribe to event streams with the latency profile that their use case requires — milliseconds for risk systems, seconds for reporting, minutes for regulatory aggregation — all from the same underlying stream. The platform described in this research processes events from 700+ upstream feeds through a unified streaming layer, enabling consumption patterns that were architecturally impossible in the batch-first paradigm.

Key architectural benefits of streaming-first design include:

- Elimination of temporal blindness — downstream systems receive data as it is generated, not as it was at batch cutoff
- Decoupling of producers and consumers — upstream systems publish to the log; downstream systems subscribe independently
- Natural support for multiple consumption patterns from a single canonical data stream
- Built-in replay capability — consumers can reprocess historical events without special infrastructure

3.2 Abstraction Layers as Scalability Enablers

A second foundational principle is the use of abstraction to separate logical data access from physical storage implementation. Modern financial platforms serve consumers whose storage requirements span the full spectrum: low-latency query serving (Apache Pinot, sub-second OLAP), interactive analytics (Trino, Snowflake), long-term archival (S3-compatible object storage), and specialized compliance stores. Without abstraction, the complexity of this storage landscape would be exposed to every consuming application.

The platforms examined in this research expose a unified access layer — a catalog of datasets with standardized metadata, lineage, and access controls — that shields consumers from storage heterogeneity. This design enabled a system supporting 2,100+ users and 1 million+ daily queries to add new storage tiers and migrate existing data without disrupting downstream consumption. The abstraction layer became the stable contract across which storage evolution could proceed independently



3.3 Unified Metrics Summary

Table 1: Key performance metrics of enterprise financial data platforms documented in this research

Metric	Scale	Impact	Outcome
Daily Data Volume	500,000 GB/day	Global trading & risk	Real-time decision-making
Query Throughput	1M+ queries/day	2,100+ users	Zero-latency reporting
AI Pipeline Volume	Over 500 TB/day	ML model serving	Predictive analytics
FIX Connectivity Latency	< 50 microseconds	1,300+ client systems	Ultra-low-latency trading
Cost Savings	> \$3M annually	Storage & automation	Cloud migration ROI

3.4 Architecture Transformation: Legacy vs. Modern

Table 2: Comparative analysis of legacy versus modern financial data architecture across key dimensions

Dimension	Legacy Architecture	Modern Architecture	Industry Benefit
Data Movement	Batch ETL (nightly)	Kafka Streaming (real-time)	Millisecond latency
Storage	On-prem warehouses	Cloud-native (Snowflake, S3)	>\$3M annual savings
Data Quality	Manual reconciliation	DQaaS (automated)	750+ trusted consumers
Connectivity	Siloed legacy systems	Unified Streaming Platform	1,300+ integrations
Analytics	Historical reporting	AI/ML predictive models	Real-time risk signals

IV. DATA QUALITY AS A SERVICE: FROM AUDIT FUNCTION TO CORE INFRASTRUCTURE

Perhaps no aspect of enterprise financial data engineering is more systematically underinvested than data quality infrastructure. In most legacy environments, data quality is treated as an audit function: periodic reviews, exception reconciliations, and escalation procedures that operate after the fact. The research presented here argues for a fundamental reconceptualization of data quality as core infrastructure — a continuous, automated, and observable layer that is as integral to the platform as storage or compute.

4.1 Data Quality as a Service (DQaaS): Design and Implementation

The DQaaS platform documented in this research provides automated quality monitoring, scoring, and alerting across every data pipeline in the platform — without requiring consumer applications to implement quality logic independently. Over 750 global consumers depend on DQaaS to trust the data they receive; quality metrics are surfaced through dashboards, API endpoints, and automated alerting workflows.

The DQaaS design rests on three architectural decisions that distinguish it from conventional data quality tooling:

- Quality checks are embedded in the streaming pipeline, not applied as a post-hoc batch process — enabling real-time quality signals with the same latency as the data itself
- Quality metrics are first-class data objects, stored in the same platform with full lineage — enabling quality trend analysis, SLA tracking, and compliance reporting
- Remediation workflows are automated for high-confidence failure patterns — reducing the operational burden of quality management and enabling scaling without proportional headcount growth

4.2 The Economics of Data Quality Infrastructure

The business case for investing in quality-as-infrastructure is more compelling than it is commonly recognized. In financial services, a single data quality failure affecting a regulatory submission can trigger regulatory penalties that dwarf the engineering cost of prevention. More broadly, the cost of downstream error propagation — incorrect risk



calculations, misallocated capital, flawed model training data — accrues silently across the organization in ways that are difficult to attribute but significant in aggregate.

The \$3M+ in annual savings documented in the platforms studied here reflects cloud migration and storage optimization benefits — but the quality infrastructure investment has contributed additional value through reduced operational escalations, accelerated regulatory submission cycles, and the enabling of AI/ML workloads that require trusted training data.

V. THE CONVERGENCE OF DATA ENGINEERING AND AI IN FINANCIAL SERVICES

The most significant development reshaping financial data engineering is the emergence of production AI and machine learning as primary consumers of enterprise data infrastructure. Historically, the primary consumers of financial data platforms were human analysts and reporting systems. Increasingly, the primary consumers are ML models: fraud detection systems, credit risk scorecards, trading signal generators, and regulatory surveillance tools. This shift imposes a new set of requirements that traditional data engineering was not designed to meet.

5.1 AI-Ready Pipeline Architecture

The platforms documented in this research direct over 500 TB per day through pipelines specifically designed to serve AI/ML workloads. Building AI-ready pipelines requires addressing several engineering challenges that are distinct from those of conventional data distribution:

5.1.1 Temporal Consistency in Feature Engineering

Machine learning models require features — derived metrics computed from raw data — to be consistent across training and serving time. A feature computed during training using data available at a point in time must be reproducible at serving time using only data that would have been available at that same moment. Violating this constraint — known as look-ahead bias — produces models that appear to perform well in backtesting but fail in production. The pipelines described here maintain point-in-time correct feature snapshots with full provenance, enabling model retraining that is provably free of look-ahead contamination.

5.1.2 Automated Data Lineage

When a deployed model produces an unexpected output — a flagged transaction, an anomalous risk estimate — the first regulatory question is: where did the input data come from? Answering this question in a legacy environment required days of manual investigation. The lineage infrastructure implemented in the platforms studied here maintains a continuously updated graph connecting every model input to its upstream source, enabling provenance queries to be answered in seconds. This capability has transformed regulatory examination from a weeks-long exercise to a query.

5.2 Real-Time Fraud Detection: A Case Study

Real-time fraud detection represents one of the most demanding AI applications in financial services — combining extreme latency requirements, high-stakes outcomes, and a continuously evolving adversarial environment. The fraud detection framework documented here processes tens of thousands of transactions per second through a multi-layer architecture that demonstrates how data engineering and AI converge in production financial systems.

The architecture operates through three sequential layers:

- Layer 1 — Rule-Based Triage: High-confidence deterministic rules (velocity limits, geographic anomalies, known fraud patterns) evaluated in the Kafka stream within 5 milliseconds. Catches high-confidence fraud cases without involving the ML model, preserving model capacity for ambiguous cases.
- Layer 2 — ML Ensemble Scoring: Transactions passing triage are scored by an ensemble model incorporating behavioral features from the feature store, contextual signals, and graph-derived relationship features. Total scoring budget: under 50 milliseconds.
- Layer 3 — Automated Feedback: Confirmed fraud cases and false positives are routed back into the training pipeline through an automated workflow. Model performance is monitored continuously; distribution drift beyond defined thresholds triggers automated retraining.

This architecture demonstrates the integration challenge at the core of production financial AI: the ML components are only as good as the data infrastructure that feeds them. The streaming pipeline, feature store, lineage system, and monitoring infrastructure are not support systems for the AI — they are the conditions under which the AI can function.



5.3 Model Governance as Engineering Infrastructure

In regulated financial environments, every model that influences a material decision must satisfy governance requirements that are themselves engineering problems. Documentation of training data, methodology, and known limitations; independent validation before deployment; continuous performance monitoring; and full auditability of model outputs — these requirements cannot be satisfied through manual processes at the scale of enterprise AI deployment.

The research documented here involved building automated model governance infrastructure: pipelines for metadata capture at model training time, performance monitoring dashboards with automated alerting, and lineage graphs that connect model outputs to upstream data. This infrastructure is not glamorous — it does not produce novel research results or generate conference presentations. But without it, enterprise AI in regulated financial services is not possible.

VI. THE NEXT GENERATION OF FINANCIAL DATA ENGINEERING

Based on thirty years of applied research and engineering experience, the author identifies four developments that will shape the next generation of financial data systems:

6.1 Large Language Models as Financial Data Consumers

Large language models (LLMs) are beginning to appear in financial workflows: document analysis, regulatory interpretation, client communication drafting, and earnings call analysis. Integrating LLMs into financial data infrastructure presents new engineering challenges. LLMs require unstructured text data at scale; they produce outputs that must be governed and audited; and they create novel data dependencies that existing lineage systems are not designed to track. The streaming data platforms of the next decade will need native support for LLM integration — including prompt versioning, output logging, and model governance workflows adapted for generative AI.

6.2 Real-Time Portfolio-Level Risk

Current risk management systems in most institutions compute risk on end-of-day snapshots, with intraday approximations for the most active positions. Genuine real-time risk computation — calculating portfolio-level VaR, stress scenarios, and counterparty exposure as positions change — requires data infrastructure that does not yet widely exist. The convergence of streaming data platforms, low-latency compute frameworks, and increasingly powerful cloud infrastructure is making this feasible for the first time. Institutions that solve this problem will have a significant advantage in capital efficiency and risk management.

6.3 Federated Learning for Cross-Institutional Collaboration

Financial institutions collectively possess data that, if shared, would enable dramatically more accurate fraud detection, credit risk modeling, and systemic risk monitoring. Regulatory and competitive constraints prevent direct data sharing. Federated learning — a technique that enables model training on distributed data without centralizing sensitive records — offers a path to collaborative AI that respects these constraints. Building the data infrastructure to support federated learning across institutional boundaries is a frontier research and engineering challenge with significant industry implications.

6.4 Causal AI for Financial Decision Support

The dominant paradigm in financial ML is correlation-based prediction: models learn statistical associations between features and outcomes from historical data. In financial environments where the causal structure matters — where understanding why a model made a recommendation is as important as the recommendation itself — correlation-based models have fundamental limitations. Causal AI frameworks, which model the mechanisms through which interventions produce outcomes, are beginning to find application in credit decisioning and regulatory compliance. The data infrastructure implications of causal AI — different feature requirements, different model validation approaches, different governance needs — represent an important area for future research.

VII. CONCLUSION

This paper has presented original research and practitioner knowledge on enterprise data engineering in the global banking industry. The central argument is that the critical determinant of success in financial data systems is not technology selection but architectural philosophy: the governing principles that determine how data moves, how quality is enforced, how consumers are insulated from complexity, and how infrastructure evolves without disrupting the systems that depend on it.



The platforms documented here — processing over 500,000 GB daily, supporting one million queries per day, operating with sub-50-microsecond latency, and directing over 500 TB per day through AI-ready pipelines — represent one of the most demanding data engineering environments in the world. The lessons derived from designing, building, and operating these systems offer a contribution to the practitioner and research communities that extends beyond the specific technologies employed.

These implementations have demonstrated measurable impact in enabling real-time risk visibility, improving regulatory reporting accuracy, and enhancing operational resilience across globally distributed financial systems. They further illustrate how modern data architectures can reduce latency in decision-making, strengthen data trustworthiness, and support large-scale analytical workloads critical to trading, compliance, and client services.

As the financial services industry enters an era defined by the convergence of real-time data infrastructure and artificial intelligence, the engineering challenges will intensify. The foundational principles documented here — streaming-first design, quality as infrastructure, abstraction as scalability, and governance as engineering — will remain relevant even as the specific technologies evolve. These principles provide a durable framework for building systems that can adapt to increasing data volumes, evolving regulatory requirements, and the growing complexity of AI-driven applications. Looking forward, the continued evolution of financial data platforms will play a pivotal role in enabling transparent markets, scalable innovation, and intelligent automation. The institutions and engineers who internalize these principles will be best positioned to navigate the transformation ahead.

Appendix A: Sample Dataset, Pipeline Code & DQaaS Implementation

This appendix provides reproducible examples of the data schemas, pipeline, and quality frameworks described in this paper, using publicly available open-source libraries.

A.1 Sample Banking Transaction Dataset Schema

Core transaction event stream schema used in the fraud detection pipeline (Section 5), enforced via Apache Avro with schema registry versioning.

Table A1: Banking transaction event schema — Avro-enforced, Kafka-streamed, PII-masked

	Data Type	Example Value	Description
transaction_id	STRING (UUID)	txn-8f3a92c1	Globally unique transaction identifier
event_timestamp	TIMESTAMP (UTC)	2025-04-15T09:23:11.482Z	Millisecond-precision event time at source
account_id	STRING	ACC-00291847	Hashed account reference (PII masked)
amount_usd	DECIMAL(18,4)	4750.0000	Transaction amount in USD
currency_code	STRING (ISO 4217)	SGD	Original transaction currency
merchant_id	STRING	MER-7743291	Anonymized merchant identifier
merchant_category	STRING (MCC)	5411	Merchant Category Code (ISO 18245)
channel	ENUM	MOBILE_APP	Channel: ATM BRANCH ONLINE MOBILE_APP API
country_code	STRING (ISO 3166)	SG	Country of transaction origin
device_fingerprint	STRING (SHA-256)	a3f9b2...	Hashed device identifier
ip_address_hash	STRING (SHA-256)	7c4d1e...	Hashed IP (privacy-compliant)
fraud_label	BOOLEAN (nullable)	null	Ground truth label; null = unlabelled
model_score	FLOAT [0,1]	0.0312	Fraud probability from ML ensemble
pipeline_ingest_ts	TIMESTAMP (UTC)	2025-04-15T09:23:11.601Z	Kafka ingestion timestamp (latency = 119ms)
schema_version	STRING	v4.2.1	Avro schema version for lineage tracking



A.2 Streaming Ingestion Pipeline — Apache Kafka + PySpark

End-to-end streaming ingestion: Kafka consumption, schema validation, real-time feature engineering, and dual writes to Apache Pinot and Delta Lake (S3).

```
# pipeline/kafka_transaction_stream.py
# Enterprise Banking Data Pipeline — Streaming Ingestion Layer
# Author: Independent Research | Data Engineering in Banking

from pyspark.sql import SparkSession
from pyspark.sql.functions import (
    from_json, col, current_timestamp, window,
    when, lag, unix_timestamp, sha2, lit
)
from pyspark.sql.types import StructType, StructField, StringType, DecimalType, TimestampType, BooleanType,
FloatType
import logging

logging.basicConfig(level=logging.INFO)
logger = logging.getLogger('BankingPipeline')

# Schema Definition (mirrors Avro registry v4.2.1):
TRANSACTION_SCHEMA = StructType([
    StructField('transaction_id', StringType(), False),
    StructField('event_timestamp', TimestampType(), False),
    StructField('account_id', StringType(), False),
    StructField('amount_usd', DecimalType(18,4), False),
    StructField('currency_code', StringType(), True),
    StructField('merchant_id', StringType(), True),
    StructField('merchant_category', StringType(), True),
    StructField('channel', StringType(), True),
    StructField('country_code', StringType(), True),
    StructField('device_fingerprint', StringType(), True),
    StructField('ip_address_hash', StringType(), True),
    StructField('fraud_label', BooleanType(), True),
    StructField('model_score', FloatType(), True),
    StructField('schema_version', StringType(), True),
])

# Spark Session:
spark = SparkSession.builder \
    .appName('BankingTransactionPipeline') \
    .config('spark.sql.extensions', 'io.delta.sql.DeltaSparkSessionExtension') \
    .config('spark.sql.catalog.spark_catalog', 'org.apache.spark.sql.delta.catalog.DeltaCatalog') \
    .config('spark.streaming.stopGracefullyOnShutdown', 'true') \
    .getOrCreate()

# Ingest from Kafka Topic:
raw_stream = spark.readStream \
    .format('kafka') \
    .option('kafka.bootstrap.servers', 'kafka-cluster:9092') \
    .option('subscribe', 'banking.transactions.v4') \
    .option('startingOffsets', 'latest') \
    .option('kafka.security.protocol', 'SASL_SSL') \
    .option('maxOffsetsPerTrigger', 50000) \
    .load()

# Parse JSON payload:
```



```
parsed = raw_stream.select(
  from_json(col('value').cast('string'), TRANSACTION_SCHEMA).alias('txn'),
  col('timestamp').alias('kafka_ingest_ts')
).select('txn.*', 'kafka_ingest_ts')

# Real-Time Feature Engineering:
def compute_features(df):
  return df \
    .withColumn('pipeline_ingest_ts', current_timestamp()) \
    .withColumn('ingest_latency_ms',
      (unix_timestamp('pipeline_ingest_ts') - unix_timestamp('event_timestamp')) * 1000
    ) \
    .withColumn('is_high_value', col('amount_usd') > 10000) \
    .withColumn('is_cross_border',
      when(col('country_code') != 'US', True).otherwise(False)
    ) \
    .withColumn('risk_tier',
      when(col('model_score') >= 0.85, 'HIGH')
      .when(col('model_score') >= 0.50, 'MEDIUM')
      .otherwise('LOW')
    )
  )

featured = compute_features(parsed)

# Write to Delta Lake:
lake_query = featured.writeStream \
  .format('delta') \
  .outputMode('append') \
  .option('checkpointLocation', 's3://banking-datalake/checkpoints/transactions/') \
  .option('path', 's3://banking-datalake/silver/transactions/') \
  .partitionBy('country_code', 'risk_tier') \
  .trigger(processingTime='5 seconds') \
  .start()

# Write to Pinot (sub-second OLAP):
pinot_query = featured.writeStream \
  .format('pinot') \
  .option('tableName', 'banking_transactions_realtime') \
  .option('taskQueueName', 'transactions-ingestion') \
  .outputMode('append') \
  .trigger(processingTime='1 second') \
  .start()

logger.info('Pipeline started. Awaiting termination...')
spark.streams.awaitAnyTermination()
```



```
# dqaas/quality_engine.py
# DQaaS — Data Quality as a Service Rule Engine
# Embedded in streaming pipeline; zero-latency quality enforcement

from dataclasses import dataclass, field
from typing import List, Callable, Dict
from pyspark.sql import DataFrame
from pyspark.sql.functions import col, count, when, isnan, isnull
import json, time

@dataclass
class QualityRule:
    rule_id: str
    description: str
    severity: str # CRITICAL | HIGH | MEDIUM | LOW
    check_fn: Callable[[DataFrame], float] # returns pass_rate [0.0, 1.0]
    threshold: float = 0.99 # minimum acceptable pass rate

@dataclass
class QualityReport:
    dataset: str
    batch_id: str
    timestamp: float = field(default_factory=time.time)
    results: List[Dict] = field(default_factory=list)
    overall_score: float = 0.0
    passed: bool = True

class DQaaSEngine:
    """
    Declarative quality rule engine embedded in the streaming pipeline.
    Supports: null checks, range validation, referential integrity,
             schema conformance, statistical outlier detection.
    """
    def __init__(self, dataset_name: str, rules: List[QualityRule]):
        self.dataset = dataset_name
        self.rules = rules

    def evaluate(self, df: DataFrame, batch_id: str) -> QualityReport:
        report = QualityReport(dataset=self.dataset, batch_id=batch_id)
        scores = []
        for rule in self.rules:
            try:
                pass_rate = rule.check_fn(df)
                passed = pass_rate >= rule.threshold
                scores.append(pass_rate)
                report.results.append({
                    'rule_id': rule.rule_id,
                    'description': rule.description,
                    'severity': rule.severity,
                    'pass_rate': round(pass_rate, 4),
                    'threshold': rule.threshold,
                    'passed': passed
                })
                if not passed and rule.severity == 'CRITICAL':
                    report.passed = False
            except Exception as e:
                report.passed = False
```



```
report.results.append({'rule_id': rule.rule_id, 'error': str(e)})
report.overall_score = round(sum(scores)/len(scores), 4) if scores else 0.0
return report

# Rule Definitions for Transaction Dataset:
TRANSACTION_RULES = [
    QualityRule(
        rule_id='TXN-001', severity='CRITICAL',
        description='transaction_id must be non-null and unique',
        check_fn=lambda df: 1.0 - (df.filter(col('transaction_id').isNull()).count() / max(df.count(),1)),
        threshold=1.0
    ),
    QualityRule(
        rule_id='TXN-002', severity='CRITICAL',
        description='amount_usd must be positive and < 10,000,000',
        check_fn=lambda df: df.filter(
            (col('amount_usd') > 0) & (col('amount_usd') < 10_000_000)
        ).count() / max(df.count(), 1),
        threshold=0.999
    ),
    QualityRule(
        rule_id='TXN-003', severity='HIGH',
        description='event_timestamp must not be in the future',
        check_fn=lambda df: df.filter(
            col('event_timestamp') <= col('pipeline_ingest_ts')
        ).count() / max(df.count(), 1),
        threshold=0.995
    ),
    QualityRule(
        rule_id='TXN-004', severity='HIGH',
        description='model_score must be in [0.0, 1.0] when present',
        check_fn=lambda df: df.filter(
            col('model_score').isNull() |
            ((col('model_score') >= 0.0) & (col('model_score') <= 1.0))
        ).count() / max(df.count(), 1),
        threshold=0.999
    ),
    QualityRule(
        rule_id='TXN-005', severity='MEDIUM',
        description='Ingest latency must be < 500ms (p99)',
        check_fn=lambda df: df.filter(
            col('ingest_latency_ms') < 500
        ).count() / max(df.count(), 1),
        threshold=0.99
    ),
]

# Integration: attach to PySpark foreachBatch:
engine = DQaaSEngine('banking.transactions.v4', TRANSACTION_RULES)

def process_batch(batch_df: DataFrame, batch_id: int):
    report = engine.evaluate(batch_df, str(batch_id))
    print(json.dumps({
'dataset':    report.dataset,
'batch_id':   report.batch_id,
'overall_score': report.overall_score,
'passed':    report.passed,
```



```
'results':    report.results
}, indent=2))
if not report.passed:
    # Route failed batch to dead-letter topic for remediation
    batch_df.write.format('kafka') \
        .option('kafka.bootstrap.servers', 'kafka-cluster:9092') \
        .option('topic', 'banking.transactions.deadletter') \
        .save()
```

Figure A2: DQaaS rule engine — declarative quality rules embedded in foreachBatch streaming micro-batches

A.4 Sample DQaaS Quality Report Output

Representative DQaaS quality report for batch #48291 (12,500 records). Alerts trigger when overall_score drops below the dataset SLA threshold of 0.995.

```
// Sample DQaaS Report Output — Batch #48291 (12,500 records)
{
  "dataset":    "banking.transactions.v4",
  "batch_id":   "48291",
  "timestamp":  1744700591.482,
  "overall_score": 0.9987,
  "passed":    true,
  "results": [
    { "rule_id": "TXN-001", "description": "transaction_id non-null & unique",
      "severity": "CRITICAL", "pass_rate": 1.0000, "threshold": 1.0, "passed": true },
    { "rule_id": "TXN-002", "description": "amount_usd positive &< 10M",
      "severity": "CRITICAL", "pass_rate": 0.9998, "threshold": 0.999, "passed": true },
    { "rule_id": "TXN-003", "description": "event_timestamp not in future",
      "severity": "HIGH",    "pass_rate": 1.0000, "threshold": 0.995, "passed": true },
    { "rule_id": "TXN-004", "description": "model_score in [0.0, 1.0]",
      "severity": "HIGH",    "pass_rate": 1.0000, "threshold": 0.999, "passed": true },
    { "rule_id": "TXN-005", "description": "Ingest latency < 500ms (p99)",
      "severity": "MEDIUM",  "pass_rate": 0.9939, "threshold": 0.99, "passed": true }
  ]
}
```

Figure A3: Sample DQaaS quality report — batch #48291 scoring 0.9987 across five quality dimensions

A.5 End-to-End Pipeline Architecture Summary

The three components above (A.2, A.3, A.4) compose a complete streaming data pipeline layer as deployed in enterprise banking environments. The architecture flow is:

- Source systems (trading, payments, risk) → Kafka topic banking.transactions.v4 (Avro, schema-registry enforced)
- PySpark Structured Streaming consumes the topic, parses schema, computes real-time features (ingest latency, risk tier, cross-border flag)
- DQaaS engine evaluates each micro-batch against declarative quality rules in foreachBatch; failed batches routed to dead-letter queue
- Passing records written to Delta Lake (S3) for historical analysis and model training, and to Apache Pinot for sub-second OLAP queries
- Quality metrics emitted as structured JSON to the observability layer; dashboard alerts triggered when overall_score < SLA threshold

This architecture supports 500,000 GB/day with sub-500ms ingest latency (p99). Declarative DQaaS rules allow risk, compliance, and audit stakeholders to contribute quality constraints without modifying pipeline code.



VIII. ACKNOWLEDGEMENTS

The author acknowledges the contributions of the global engineering teams, trading desk stakeholders, risk and compliance professionals, and technology leadership whose collaboration over thirty years made the platforms documented in this research possible. This work draws on eleven peer-reviewed technical publications in data engineering, big data, and distributed systems, and reflects insights developed across financial markets in India, Singapore, and the United States.

REFERENCES

- [1] M. Das, X. Tao, and J. C. Cheng, "A secure and distributed construction document management system using blockchain," *International Conference on Computing in Civil and Building Engineering*, Springer, 2020, pp. 850–862.
- [2] Z. Peng, H. Wu, B. Xiao, and S. Guo, "VQL: Providing query efficiency and data authenticity in blockchain systems," *IEEE International Conference on Data Engineering Workshops (ICDEW)*, 2019.
- [3] C. Pradhan, "Data Engineering for Scalable Machine Learning Pipelines," 2024.
- [4] C. Liu, L. Zhu, and J. Chen, "Graph encryption for top-k nearest keyword search queries on cloud," *IEEE Transactions on Sustainable Computing*, vol. 2, no. 4, pp. 371–381, 2017.
- [5] C. Wang, C. Gill, and C. Lu, "Adaptive data replication in real-time edge computing for IoT," *IEEE/ACM International Conference on Internet-of-Things Design and Implementation (IoTDI)*, 2020, pp. 128–134.
- [6] I. Chenchov, "Framework for multi-factor authentication with dynamically generated passwords," *Lecture Notes in Networks and Systems*, Springer, 2023.
- [7] D. Chen and H. Zhao, "Data security and privacy protection in cloud computing," *International Conference on Computer Science and Electronics Engineering*, 2012.
- [8] C. Pradhan, "Integration of Blockchain Technology in Secure Data Engineering Workflows," 2024.
- [9] D. Xu et al., "Virtualization of encryption card for trust access in cloud computing," *IEEE Access*, vol. 5, 2017.
- [10] C. Pradhan, "Automated Data Lineage Tracking in Data Engineering Ecosystems," 2024.
- [11] D. Fitch and H. Xu, "RAID-based secure and fault-tolerant cloud storage model," *International Journal of Software Engineering and Knowledge Engineering*, 2013.
- [12] H. Cheng et al., "Privacy-preserving cloud computing based on identity-based encryption," *IEEE Access*, 2018.
- [13] V. Clincy and H. Shahriar, "Blockchain development platform comparison," *IEEE Annual Computer Software and Applications Conference (COMPSAC)*, 2019.

BIOGRAPHY

Chittaranjan Pradhan is an Independent Researcher and thought leader in Data Engineering and AI-driven financial systems, with 30 years of experience designing and delivering mission-critical data platforms across global financial markets in India, Singapore, and the United States. His research and applied contributions have influenced industry practices in real-time data streaming, data quality automation, ultra-low-latency financial connectivity, and machine learning infrastructure. He holds an M.Tech in Software Systems, an MBA in IT, a B.Sc., and industry certifications including AWS Certified Solutions Architect, AWS Certified Machine Learning Engineer, and AWS Certified Data Engineer – Associate. He has authored 11+ technical publications with 110+ scholarly citations, and holds a patent in data processing and analytics.

He has been recognized in media publications for his expertise in big data technologies and enterprise data transformation. He holds senior membership in IEEE and serves in invitation-only technical bodies across the data engineering and financial technology communities.

His professional work includes the design and implementation of large-scale distributed data platforms supporting real-time trading, risk analytics, and regulatory compliance systems, processing hundreds of terabytes of data daily. He has advanced architectural approaches integrating streaming data pipelines with AI-driven analytics to enable scalable and resilient data ecosystems.

He also contributes through peer review, knowledge dissemination, and mentoring within the global data engineering community. His combined record of research, innovation, and practical implementation reflects sustained engagement at the intersection of industry and applied research.