

Agentic AI in Fintech Commercial Blueprint



Solving the Unstructured
Data Crisis with A.I. vs. A.I.

By Leon Liao

Foreword: Execution over Creativity

Execution over Creativity — The True Inflection Point for Financial AI in 2026

In 2026, the focus of financial AI is shifting from content generation to controlled execution, cross-checking and compliance accountability.

From “Can Generate” to “Can Execute”

If 2023–2024 was AI’s “romantic era,” when the market was captivated by LLMs’ generative power, then by 2026 the fintech narrative has turned the page.

In finance, where regulation is strict and tolerance for error is low, the real need is not poetry writing, but execution, verification and compliance. Agentic AI can plan, invoke tools, cross-verify and complete complex business loops.

This whitepaper further proposes an “A.I. vs. A.I.” mechanism: different agents review one another, using traceable evidence to reduce the risk of single-model errors.

The Market Has Entered Controlled Deployment

52% of surveyed financial services firms have piloted Agentic AI or reached a more mature stage; only **23%** are at the scale-up or transformation stage.

Official disclosures from BNY, Citi and Lloyds show that multi-agent “digital employees,” Agentic AI workspaces and real-time anti-fraud systems are already being deployed.

These cases show controlled deployment is expanding, but should not be read as universal enterprise maturity or as evidence of a common industry-wide ROI.

KEY TAKEAWAY

The next commercial inflection point for financial AI is not a more eloquent model, but an executable, reviewable and accountable workflow.

Source: CCAF 2026; BNY 2025; Citi; Lloyds 2026. See original whitepaper references [1], [10]–[12].

Contents

Contents

00	Foreword: Execution over Creativity	03	Scaling the Paradigm: Agentic AI in Institutional Fintech
01	Cognitive Reshaping: Why Finance Must Embrace Agentic AI	04	Building a Responsible Financial AI Ecosystem
02	Anatomy of a Use Case: Practical Insights from Checkie-AI	05	Conclusion & Call to Action
		R	References

This whitepaper consists of five parts: problem, evidence, product practice, institutional mapping and governance recommendations.

The structure moves from market inflection to operating model, then to governance and action.

01

Chapter 01: Cognitive Reshaping

Why the Financial Industry Needs Agentic AI

Cognitive Reshaping: Why the Financial Industry Needs Agentic AI

Chapter Focus

Starting from the limits of traditional LLMs, this chapter explains why Agentic AI better fits financial workflows, unstructured data and high-compliance environments.

1.1 The Achilles' Heel of Traditional LLMs in Finance

Why foundation models alone do not satisfy financial operating requirements

In recent years, many financial institutions tried to plug foundation LLMs directly into business flows, but quickly hit a ceiling. Traditional generative AI shows three critical limits in financial use cases.

01 Hallucinations

When handling precise financial terms, models may generate wrong information with high confidence. For compliance, transactions and customer decisions, this is not merely a content flaw, but an unacceptable risk source.

02 Lack of Multi-step Reasoning

Financial review often requires reading across pages, tracking dependencies and making multiple logic jumps. A single prompt cannot reliably support this depth of reasoning and cross-checking.

03 Closed Ecosystem / Inability to Invoke Tools

Without real-time access to external databases such as credit interfaces, sanctions lists or transaction systems, model information quickly becomes stale and cannot close the end-to-end business loop.

OPERATING IMPLICATION

Finance does not need a model that simply generates better content; it needs a controlled workflow that can plan, invoke tools, verify evidence and accept review.

1.2 Agentic AI: From “Brain” to “Digital Employee”

Four capabilities turn a model into an executable operating system

Agentic AI is not a single model, but a new agent architecture. It embeds model capability into executable processes and supports commercialization through four core capabilities.

01 Planning

Breaks complex tasks into executable subtasks, such as decomposing trade finance review into document recognition, rule validation, sanctions screening and discrepancy handling.

02 Tool Use

Uses APIs to access OCR, enterprise databases, compliance registries and approved knowledge sources, connecting analysis to real business environments.

03 Memory

Combines short-term context with long-term business memory so key information is not lost across multi-round processing.

04 Reflection

Reviews and corrects errors before output, echoing the financial sector’s Maker–Checker principle.

Together, these capabilities move AI from “answering questions” to “completing tasks,” and extend governance from the model to the entire workflow.

1.3 The Core Pain Point: Unstructured Data Meets Zero Tolerance


Unstructured data meets a zero-tolerance operating environment

Financial operations rely heavily on complex unstructured data such as contracts, documents, financial statements and emails. The real barrier is not format alone, but the combined pressure of data quality, governance boundaries and audit requirements.

Industry Evidence

Up to about 60% of corporate and investment banking customer service work still relies on email and manual document processing. CCAF finds that data availability and quality hinder AI adoption for 49% of traditional financial institutions. In a KPMG survey, 72% of banking respondents cited data quality as a concern.

Exhibit 1: Data availability and quality remain major barriers to AI adoption



Note: Respondent groups use different denominators; percentages should not be added.

Source: Cambridge CCAF (2026); fieldwork Oct 2025–Jan 2026.

Traditional RPA fits fixed formats; RAG and Agentic systems need traceable sources, controlled data flows and human review to support true auditability.

1.4 Technology Evolution: Why Agentic AI Wins

A design framework for comparing automation and AI operating models

Exhibit 2: Capability evolution in financial workflows

Capability Dimension	Traditional RPA	Traditional LLM / RAG	Agentic AI
Unstructured Data Processing Unstructured data	WEAK Requires fixed templates	GOOD Prone to hallucinations	STRONG Cross-modal perception + cross-verification
Multi-step Logical Reasoning Multi-step reasoning	NONE None	LIMITED Limited by a single conversation	STRONG Goal decomposition and autonomous planning
Compliance Explainability Auditability	HIGH Based on fixed rules	LOW Algorithmic black box	HIGH Evidence-based decision traceability
External Tool / API Invocation Tool & API use	HARD-CODED Requires hard-coding	LIMITED Limited	DYNAMIC Autonomous invocation by task
Boundary Exception Handling Exception handling	FAILS Fails when exceptions occur	UNPREDICTABLE Unpredictable	ADAPTIVE Self-reflection and routing correction

Note This table should be read as a design framework, not a universal capability ranking. Explainability, tool use and exception handling all depend on deployment, controls and workflow design.

Source: D2 Intelligence analysis; comparison framework from the original whitepaper draft.

02

Chapter 02: Anatomy of a Use Case

Practical Insights from the Consumer Risk-Avoidance Tool Checkie-AI

Anatomy of a Use Case: Practical Insights from Checkie-AI

Chapter Focus

The consumer-side environment is extremely unstructured and noisy, making it a stress test for multi-agent architecture. As a technical sandbox and MVP, Checkie-AI accumulates boundary cases such as blurry screenshots, multilingual contracts and scam scripts, forming a data flywheel: consumer-side edge data feeds enterprise risk-control capabilities.

2.1 Scenario Entry Point: Reducing Information Disadvantage and Cognitive Overload

Case 1 — translating a complex contract into a decision signal


Case 1 | Spotting a Hidden Unfair Clause Instantly

A student newly arrived in London was about to sign a long English rental contract. Unfamiliar with local legal wording and worried about missing key terms, the student uploaded the contract to Checkie-AI for an additional risk assessment.

The system extracted the contract text and found a deeply buried early-termination clause: even if the tenant moved out early, they might still owe rent for the full year. The relevant content was clearly marked red / high risk.

The product’s value is not asking users to learn more legal terminology, but translating complex clauses into verifiable risk prompts that help users pause before signing.

Case Exhibit A: The product interface maps a high-risk score to a human confirmation action



A high-risk score is not the endpoint; it should trigger evidence verification and human decision-making.

Simulated case; the interface screenshot shows how the product presents risk.

Source: Checkie-AI product interface; case narrative from the original whitepaper draft.

Case 2 | A Mirror for Multi-step Scams


Case 2 — interpreting contract terms and conversational pressure together

A freelance illustrator received a high-paying outsourcing contract said to come from an overseas client. The opportunity first looked normal, but before work began, the other party suddenly asked for an upfront “cross-border handling fee.”

The illustrator uploaded both the contract and chat history. Beyond contract terms, the system detected high-pressure language and the abnormal request to “pay first before cooperation begins,” marking the overall case red / high risk.

By understanding documents and chat context together, the product turns an attractive-looking opportunity into a clear risk picture, prompting the user to re-check before transferring money.

Case Exhibit B: The full product journey from upload and text extraction to scanning and risk results



Source: Checkie-AI product interface; simulated scenario from the original whitepaper draft.

2.2 Deconstructing the Multi-agent Architecture

A digital expert team behind a deliberately simple interface

Checkie-AI is not a simple chat box. Behind it is a precisely coordinated “digital expert team.” The front end shows only the necessary processing status; the true complexity is handled by the multi-agent workflow.


Perception Layer

The cross-modal OCR Agent handles blurry screenshots, multilingual PDFs and non-standard layouts, first converting input into verifiable structured information.

Expert Layer

- Legal & Compliance Agent: compares rules and regulations to identify non-compliant clauses.
- Hidden-fee Extraction Agent: identifies late fees and hidden management charges.
- Scam / Emotional Manipulation Warning Agent: identifies fraud patterns and high-pressure tactics.

Case Exhibit C: The processing status interface turns multi-agent operation into an understandable progress signal




The UI only says “scanning for risks”; routing, tool use and expert division of labor remain inside the system.

Source: Checkie-AI product interface and workflow description, original whitepaper draft.

Checkie-AI Multi-agent Architecture

A flat, evidence-oriented view of the multi-agent workflow

Exhibit 3: Multi-agent flow from user input to actionable risk recommendations



Source: D2 Intelligence / Checkie-AI product architecture, reproduced from the original whitepaper draft.

2.3 UI/UX as Risk Control: Let AI Move Behind the Scenes

The interface translates model output into an auditable action signal

Checkie-AI translates complex analysis into red, yellow and green risk signals. This is not decorative product display, but a mapping from model judgment to clear review and decision actions.

Exhibit 4: How the traffic-light interface turns risk scores into actionable decisions



Source: Checkie-AI product interface; annotations map the original risk language to business actions.

The best risk-control interface does not show more AI; it shortens the distance from risk signal to decision action.

03

Chapter 03: Scaling the Paradigm

Mapping Agentic AI to Institutional Fintech

Scaling the Paradigm: Mapping Agentic AI to Institutional Fintech

Chapter Focus

If reviewing a consumer rental contract is “street fighting,” reviewing a multimillion-dollar letter of credit and trade document set is a “large-scale battle.” The scale differs, but both rely on cross-document checks, anomaly detection, external verification and traceable decisions.

3.1 A Unified Logic: From Rental Contracts to Trade Finance

Cross-verification and discrepancy checking are the common operating core

Whether for consumer contracts or complex international trade finance, the core business action is cross-document comparison and anomaly detection. This makes it possible to transfer and scale the Checkie-AI workflow to institutional scenarios.


Transaction Complexity

A single transaction may involve over 20 entities, 10–20 paper documents and around 5,000 data-field exchanges. ICC DSI separately identifies 36 key trade document types across the end-to-end supply chain; this uses a different denominator and should not be added to per-transaction document counts.

Exhibit 5: Structural complexity of a single trade finance transaction

Anatomy of a document-heavy trade transaction

Structural complexity is visible across participants, documents and data exchanges



Interpretation rule: participants, documents per transaction, field exchanges and end-to-end document types use different denominators and should not be added.

Source: BCG / SWIFT / ICC; ICC Digital Standards Initiative KTDDE.

Industry Data and Competitor Trends: Controlled Deployment by Global Banks

Verified public cases show targeted production and staged expansion

Exhibit 6: Public Agentic AI cases at global banks and evidence boundaries

Public Case	Verified Stage	What the Case Supports	Evidence Boundary
BNY Eliza	In production	160 enterprise AI solutions are in production, with 134 multi-agent “digital employees” working alongside human colleagues.	No unified accuracy, time-saving or ROI figures disclosed.[^10]
Citi Stylus Workspaces	Early access for thousands of employees; staged expansion	Connects internal data and enterprise tools to execute multi-stage research and analysis workflows.	Early access does not mean full company-wide production.[^11]
Lloyds Multi-agent Anti-fraud	Deployed	Multiple agents perform identity, transaction and fraud-risk checks in parallel; employees retain final accountability and override rights.	Not a trade finance Maker–Checker case; existing anti-fraud results should not be attributed to the new deployment.[^12]
NatWest Cora and Trials	Early customer launch and internal trials	Agentic AI extends to customer assistance, fraud support and engineering scenarios.	10x productivity comes only from a specific engineering trial and cannot be generalized across the bank.[^13]

Cross-case reading: banks are moving along two paths—controlled workflow deployment and staged expansion—but the cases do not yet provide a generalizable ROI.

Source: BNY, Citi, Lloyds and NatWest official disclosures; original whitepaper refs [10]–[13].

Trade Finance Digitization: Electronic Originals, Shared Data and APIs

Production cases are moving beyond document scanning

Exhibit 7: Public trade digitization cases and supported conclusions

Public Case	Verified Stage	What the Case Supports	Evidence Boundary
J.P. Morgan Electronic Bill of Exchange	Live in the UK; expanded to the U.S.	Digital trade is moving beyond document scanning toward legally effective electronic transferable instruments.	No per-transaction time or cost improvement disclosed. ^[14]
HSBC Bahrain Electronic Bill of Lading	Live transaction completed	One import transaction was completed within hours; moving paper documents usually takes days or weeks.	Only one transaction; should not be presented as an industry average. ^[15]
Citi Consolidate	Live in the U.S. and Canada	A shared structured environment automates matching, approvals and discrepancies.	No unified cycle-time reduction or accuracy figure disclosed. ^[16]
Standard Chartered ICC-Swift API Guarantee	Live transaction completed; expansion planned	Industry API standards and automated instruction receipt are entering real trade finance workflows.	No issuance time or cost reduction disclosed. ^[17]

Evidence Boundary

The cases support that digitization is entering real workflows, but do not prove that banks broadly use dual-agent review, nor allow different metrics to be combined into a unified ROI. UCP 600's five banking days are the maximum rule-based period, not average manual checking time.

Source: Official disclosures; UCP 600; original whitepaper refs [6], [14]-[17].

3.2 Agent Design for Intelligent Trade Finance and Regulatory Alignment

A controlled architecture aligned with sandbox principles—not a claim of approval

Document Consistency Review Agent

Internalizes UCP 600 rules and automatically cross-checks commercial invoices, bills of lading and certificates of origin to identify mismatched amounts, backdated documents and other anomalies.

External Compliance and AML Agent

Uses controlled APIs to access external data such as global sanctions lists and vessel tracks, checking whether counterparties or routes trigger AML red flags.

Regulatory Environment and Public Policy Signals

The HKMA GenA.I. Sandbox provides authorized institutions with a risk-controlled environment to develop, test and pilot innovative AI solutions in real banking scenarios.

The Sandbox++ joint circular mentions using an “A.I. vs. A.I.” strategy for automated quality checks on generated outputs while maintaining human oversight for key decisions. This supports responsible innovation, but does not mean HKMA has endorsed a specific architecture or that all MAS / EU requirements are met.

“The launch of the GenA.I. Sandbox++ marks a significant milestone... These partnerships are unlocking the potential of agentic AI to create new opportunities.”

— Eddie Yue, HKMA Chief Executive (2026)

Source: HKMA GenA.I. Sandbox and Sandbox++ materials; original whitepaper refs [7]-[8].

3.3 The Breakthrough: “A.I. vs. A.I.” Quality Control

Separate creation, challenge and escalation responsibilities

In multimillion-dollar transactions, the key question is not whether a model can be made error-free, but how to make errors more detectable, ensure disagreements are escalated and preserve final human judgment.

Exhibit 8: Maker–Checker–Supervisor separation of responsibilities

01	MAKER AGENT Maker Agent	Extracts information from unstructured documents and produces structured data and an initial review report.
02	CHECKER AGENT Checker Agent	Does not generate; only challenges, pressure-tests, falsifies logic and verifies evidence.
03	SUPERVISOR AGENT Conflict Resolution Mechanism	Intervenes when confidence is low or views diverge; supplements approved knowledge sources and, when needed, triggers yellow status and human judgment.

The governance value of A.I. vs. A.I. lies in responsibility separation and conflict escalation—not mistaking agreement between two models for absolute correctness.

Source: D2 Intelligence design framework, original whitepaper draft.

Quality Control: Evidence and Applicability Boundaries

What the design can support—and what it must not claim

Design Evidence and Applicability Limits

Maker–Checker can use a second AI verification layer to challenge or validate outputs and escalate disagreements to human decision-making. It can improve evidence traceability, error detectability and auditability, but cannot guarantee the elimination of model errors or constitute universal regulatory approval. HKMA also notes that prompt engineering is a mitigation, not a guarantee.^[8]

Adjacent Industry Evidence

Lloyds’ anti-fraud system shows that multiple AI agents can run different risk checks in parallel while employees retain final accountability and override rights. This is a relevant production case, but does not prove that dual-agent document review is widespread in trade finance.^[12]

External Messaging Principles

Emphasize / SUPPORTED

Second-layer verification and evidence challenge
Disagreement escalation and final human judgment
Traceable sources, processing records and auditability

Avoid / NOT SUPPORTED


Complete elimination of hallucinations
Broad industry deployment
General regulatory certification

Source: HKMA and Lloyds disclosures; original whitepaper refs [8], [12].

Trade Finance A.I. vs. A.I. Review Flow

A flat operating model from source documents to an accountable decision

Exhibit 9: Extraction, verification, escalation and decision flow for trade documents



Source: D2 Intelligence / Checkie-AI operating model, original whitepaper draft.

04

Chapter 04: Building a Responsible Financial AI Ecosystem

Responsible A.I. Framework for Commercialization

Governance is the precondition for scaling financial AI

Chapter Focus

In fintech, even advanced technology remains experimental if it cannot meet regulatory and internal-control requirements. Commercializing Agentic AI must be built on a Responsible A.I. framework of explainability, auditability, human-AI collaboration and enterprise-grade security.

4.1 Explainability: Breaking the Algorithmic Black Box

Explainability requires an evidence-based audit trail

Financial decisions cannot rely on a “black box.” The architecture should embed decision traceability: every risk warning or discrepancy must be based on verifiable evidence.

The system highlights specific clauses from source documents in the UI and retains source, transformation and human-review records. AI should not only give judgments, but also show the basis for them; this supports explainability and auditability, but should not be described as having passed an HKMA XAI audit.

Governance Evidence: The BCBS 2023 progress report shows that in the 2022 assessment of 31 global systemically important banks, only two fully complied with all BCBS 239 principles.

Exhibit 10: BCBS 239 compliance remains the exception



Governance implication: retain data lineage, transformation records, evidence sources and human decision logs.

Source: Basel Committee on Banking Supervision (2023 progress report; assessment period 2022).

4.2 Human-in-the-Loop: A New Paradigm for Human-AI Collaboration

AI acts as a Super Copilot; humans retain final accountability

Agentic AI is not designed to fully replace humans, but to serve as a “super copilot” for experts. AI can help handle standardized checks under controlled conditions; for high-risk and ambiguous exceptions, final approval and accountability remain with human experts.

Governance Case: Lloyds clearly states that employees using its multi-agent anti-fraud system remain fully responsible for outcomes and can override AI recommendations. This supports human decision rights, agent-behavior monitoring and auditable escalation mechanisms.

Exhibit 11: Human-AI decision responsibilities mapped to traffic-light signals

GREEN

Standard compliant cases

Can be automated under controlled conditions; retain traceable notices.

YELLOW

Ambiguous or low-confidence

Submit for human review; add evidence or re-verify.

RED

High-risk exception

Final decision by a human, with complete records retained.

Source: D2 Intelligence decision logic; Lloyds disclosure, original whitepaper ref [12].

4.3 Data Privacy and Enterprise-grade Security Compliance

InfoSec and data residency claims must be deployment-specific and verifiable

Given the strict data-security requirements of financial institutions, the architecture should assess data classification, access control, encryption, logging, retention, vendor risk and data residency against actual regulatory and security conditions.

Exhibit 12: Four core control domains for enterprise deployment

01

Data Classification and Access Control

Define data types, permission boundaries, least-privilege rules and access logs.

02

Encryption and Logging

Protect data in transit and at rest, record key processing steps and support later audits.

03

Vendor Risk

Assess the security and compliance risks of external tools, APIs, models and data supply chains.

04

Data Residency

Design hybrid-cloud, private-deployment or regional processing options according to client and regulatory requirements.

Claim Boundary

Claims such as SOC 2 Type II, ISO 27001, zero trust, hybrid cloud or private deployment should only be used externally when they apply to the actual product and deployment and have been internally verified.

Source: D2 Intelligence control framework; original whitepaper draft.

4.4 A Zero-risk Enterprise Adoption Path

A progressive path from historical validation to controlled production

To reduce concerns about “replacing existing systems,” adoption should follow a progressive path that is reversible, comparable and monitorable.

Exhibit 13: Three-stage enterprise adoption path



Source: D2 Intelligence adoption framework, original whitepaper draft.

CHAPTER 05

05

Chapter 05: Conclusion and Commercial Call to Action

Conclusion & Call to Action

In 2026, we are witnessing a deep industry transformation. Agentic AI and the derived “A.I. vs. A.I.” mechanism are moving financial services away from labor-intensive processes reliant on headcount and toward agent-supported intelligence-intensive operating models.

From protecting ordinary consumers’ rights to safeguarding funds at multinational financial institutions, the underlying capabilities consistently center on unstructured-data understanding, cross-document verification, evidence traceability and human accountability.

Evidence Supplement: Existing evidence supports that pilots, staged launches and targeted production deployments are expanding, and shows that regulators emphasize controlled innovation, validation and human oversight.

However, current cases do not prove that all financial institutions have reached mature deployment, nor that “A.I. vs. A.I.” trade-document review has become an industry norm.^{[1][7][10][12][14][17]}

Agentic AI is not a promise of “full automation,” but a realistic path that recombines execution, evidence and accountability.

NEXT STEP

Ready for the Agentic Revolution in 2026?

Contact our expert team for an Agentic AI adoption assessment for your institution's trade finance / risk-control workflow, or book a Checkie-AI Enterprise Demo.

01

CONNECT WITH THE AUTHOR

Leon Liao | [linkedin.com/in/leon-l-5322b32b/](https://www.linkedin.com/in/leon-l-5322b32b/)

02

BUSINESS INQUIRIES

leon.liao@d2-ai.com | **leonliao4@gmail.com**

03

WEBSITE

checkie-ai.com

EXECUTION / EVIDENCE / ACCOUNTABILITY

References

References [1]–[9]

-
- [1] Cambridge Centre for Alternative Finance (CCAF), The 2026 Global AI in Financial Services Report (2026): sample n=352; 52% at pilot or more mature stages; 23% at scale-up or transformation stages.
-
- [2] McKinsey, corporate and investment banking customer service study (2023): up to about 60% of service work still relies on email and manual documents.
-
- [3] KPMG, Intelligent Banking (2025): banking respondent group n=183; 72% expressed concern about data quality.
-
- [4] BCG, SWIFT and ICC, Digital Ecosystems in Trade Finance (2019): over 20 entities, 10–20 paper documents and about 5,000 data-field exchanges.
-
- [5] ICC Digital Standards Initiative, Key Trade Documents and Data Elements: 36 key trade document types are an end-to-end classification.
-
- [6] ICC, UCP 600, Article 14(b): banks have a maximum of five banking days to determine whether a presentation appears compliant on its face.
-
- [7] HKMA, Joint Circular on the Expansion of Generative Artificial Intelligence Sandbox (5 March 2026).
-
- [8] HKMA, Responsible Innovation with GenA.I. in the Banking Industry — Practical Insights from the GenA.I. Sandbox (31 October 2025), especially p.28 and pp.38–39.
-
- [9] Basel Committee on Banking Supervision, Progress in adopting the Principles for effective risk data aggregation and risk reporting (2023); assessment based on 2022 supervisory findings.
-

References

References [10]–[17]

[10]

BNY, Annual Report 2025 (February 2026): 160 enterprise AI solutions in production, 134 multi-agent “digital employees,” and nearly 50% of employees involved in building AI agents. Official source: <https://www.bny.com/corporate/global/en/investor-relations/annual-report-2025.html>

[11]

Citi, Citi Unveils Citi Stylus Workspaces with Agentic AI (22 September 2025): early access first offered to thousands of employees, with staged expansion. Official source: <https://www.citigroup.com/global/news/press-release/2025/citi-unveils-citi-stylus-workspaces-agentic-ai-turbocharging-productivity>

[12]

Lloyds Banking Group, deploys agentic AI to strengthen real-time fraud protection (8 June 2026): parallel identity, transaction and fraud-risk checks, with employees retaining final accountability and override rights.

[13]

NatWest Group, Being a trusted partner for tomorrow’s banking through technology, data and AI (13 February 2026): Cora customer plans and specific Agentic engineering trials.

[14]

J.P. Morgan, Brings Electronic Bill of Exchange Solution to the U.S. (20 May 2026): expanded to the U.S. after live UK deployment.

[15]

HSBC Bahrain, HSBC Leads Digital Transformation in Trade Finance (23 June 2024): one electronic bill of lading and ePresentation transaction was completed within hours.

[16]

Citi, Citi Is Now Live with a Trade Digitization Solution (28 July 2026): a shared structured data environment automates matching, approvals and discrepancies.

[17]

Standard Chartered, first ICC–Swift API standards digital bank guarantee transaction through Komgo (15 January 2026): live API-based digital bank guarantee transaction.

ABOUT THE AUTHOR



Leon Liao

Founder of Checkie-AI
Head of AI Product at D2 Intelligence

✉ leon.liao@d2-ai.com

🌐 checkie-ai.com

