Deloitte.

Together makes progress



Fraud in the age of synthetic trust:

How frontier ai and quantum computing are reshaping digital trust

FRAUD, AI & QUANTUM


Together, Frontier AI and quantum computing challenge how organizations establish, verify, and protect trust.

Trust is becoming a fraud attack surface

Fraud is becoming faster, more personalized, and harder to distinguish from legitimate activity. Frontier AI refers to increasingly capable AI systems that can work across multiple modalities, adapt to changing conditions, and perform increasingly autonomous tasks across connected environments. Unlike Generative AI used primarily to create content, Frontier AI can easily orchestrate activity, make decisions, interact with systems, and adjust its behavior based on feedback. These capabilities enable criminals to scale social engineering, synthetic identities, voice cloning, deepfakes, and adaptive evasion—often across multiple channels at once.

At the same time, quantum advances may create longer-term risks to encryption, digital signatures, authentication, payment integrity, and sensitive investigative data, which are widely acknowledged in addition to enabling other new, unexpected fraud actor behavior. The "harvest now, decrypt later" threat means information stolen today could become usable in the future.

Together, Frontier AI and quantum computing challenge how organizations establish, verify, and protect trust. Organizations need an enterprise response that connects fraud prevention, intelligence, cybersecurity, investigations, technology governance, and practical readiness.

Fraud defense has to address two horizons

Organizations should consider two related but distinct horizons:

1. **AI-enabled fraud today**—where Frontier AI increases the speed, scale, personalization, and adaptability of fraud, and
2. **Quantum-related trust risks ahead**—where advances in quantum computing could challenge the cryptographic foundations that support identity, authentication, transactions, and secure information as well as provide new avenues for committing fraud.

By taking a dual-horizon approach, organizations can strengthen current fraud defenses and prepare for longer-term shifts in digital trust.



How Frontier AI is rewiring fraud

The next wave of fraud is being defined less by better phishing, deepfakes, and synthetic identities—which grab headlines—and more by how AI changes the way fraud actors organize and execute fraud. The more consequential shift is from AI-generated deception to AI-operated fraud: systems that can identify targets, create synthetic actors, coordinate activity across environments, learn from responses, and adapt as controls react.

The deeper shift is operational. Frontier AI can automate routine work across the fraud lifecycle—from reconnaissance and target selection to persona creation, engagement, execution, and concealment. Human operators can then focus on exceptions, high-value decisions, and other activities that require judgment. This allows fraud actors to run more attempts simultaneously, test more variations, learn quickly from failed interactions, personalize the approaches most likely to succeed, and deliver maximum gain to the fraud actor. Fraud becomes not only more scalable, but more continuously optimized.

As fraud becomes increasingly AI-operated, organizations should use AI not simply to generate more alerts, but to connect signals, challenge manufactured trust, and respond at the speed of the threat.

AI can compress the fraud lifecycle

In the past, a fraud scheme had people gathering target information, creating convincing communications, managing interactions, operating accounts, and responding when a target or control reacts. Frontier AI can increasingly support or automate portions of each activity:

- **Targeting**: Assemble fragmented public, stolen, or brokered information into profiles that reveal likely points of influence.
- **Persona creation**: Generate identities, communications, documents, voices, images, and supporting narratives tailored to a specific target or business context.
- **Engagement**: Maintain conversations, answer objections, change tone, and preserve a consistent backstory over time.
- **Coordination**: Support distinct roles across reconnaissance, impersonation, account activity, payments, and concealment.
- **Adaptation**: Learn from engagement and control responses, then adjust messages, channels, timing, counterparties, or transaction behavior.

This is how fraud velocity and effectiveness increase together. AI reduces the effort required to initiate and sustain an attempt while improving the actor's ability to select targets, adapt in real time, and replicate successful patterns across channels.

The advantage is not that every attempt succeeds. It is that perpetrators can fail more cheaply, learn faster, and scale the approaches that work.



Familiar fraud schemes can become AI-operated

Traditional fraud schemes remain relevant, but Frontier AI can change their speed, scale, and effectiveness:

- **Synthetic identities may become synthetic enterprises**. Rather than creating one false identity, perpetrators could coordinate apparently independent customers, employees, vendors, accounts, devices, and digital histories to create the appearance of a legitimate business relationship.
- **Business email compromise may become authority hijacking**. As organizations delegate approvals, procurement, payments, and service decisions to AI agents, perpetrators may target the data, instructions, or context those agents rely on—rather than compromising a person directly.
- **Social engineering may become adaptive**. AI can support rapid testing of messages, channels, timing, and authority cues to identify what is most likely to prompt disclosure, approval, or payment.
- **Collusion may operate at machine speed**. Multiple AI-enabled processes could support reconnaissance, impersonation, communications, account activity, and concealment while appearing decentralized to the organization being targeted.
- **Control evasion may become continuous**. Rather than execute a fixed scheme, perpetrators could learn from control responses and adjust behavior, timing, counterparties, and transaction patterns in real time.
- **The detection model may become a target**. Attackers could manipulate the data or context used by fraud models so malicious activity appears ordinary—or legitimate activity appears ambiguous.
- **Evidence may be manufactured at scale**. AI can generate consistent records, explanations, communications, and activity histories that make a false narrative appear independently supported.

The central shift is from **isolated deception to AI-operated fraud**—more coordinated, adaptive, persistent, and difficult to distinguish from legitimate activity.

Organizations therefore need to defend not only against suspicious events, but against the broader systems of identity, authority, evidence, and technology that make those events appear trustworthy.

How organizations can fight AI-delivered fraud with AI

Frontier AI can also strengthen fraud prevention and investigations when deployed with appropriate governance and human oversight. Organizations can use advanced analytics, behavioral monitoring, network analysis, and unsupervised modeling to identify relationships and patterns that traditional rules may miss. They can enumerate their fraud risk surface in ways that a human assessment could not. For example, Frontier AI can string together a number of low-risk issues into one high risk issue, which a traditional fraud risk assessment would not find.

The opportunity for defenders is to apply AI at the same points in the lifecycle that fraud actors are seeking to automate:

- **Targeting**: Link identity, transaction, behavioral, access, and network signals to identify coordinated activity earlier.
- **Persona creation**: Test whether identities, accounts, communications, and supporting records are behaviorally consistent over time.
- **Engagement**: Monitor interactions across channels for changes in language, timing, escalation patterns, and unusual persistence.
- **Coordination**: Use network analysis to connect seemingly unrelated accounts, devices, cases, and events.
- **Adaptation**: Continuously test controls against changing behavior, model evasion, and emerging fraud hypotheses.
- **Evidence**: Preserve provenance, corroborate records across systems, and generate explainable evidence packages for human review.

The objective is not to replace investigators or automate decisions without oversight. It is to help teams move from isolated alerts to coordinated pattern detection, faster investigation, and more defensible action at higher speed.



Govern AI-enabled fraud prevention

AI-enabled fraud prevention should be secure, privacy-preserving, transparent, explainable, fair, accountable, and reliable. Models and agents should support human decision-makers rather than operate as unsupervised authorities.

Organizations should establish:

- Clear permissions and approval thresholds for AI agents
- Human review for high-impact fraud and investigative decisions
- Monitoring for prompt injection, model manipulation, deepfakes, and adversarial behavior
- Continuous oversight of model performance, drift, explainability, and unintended outcomes
- Evidence retention and escalation paths that support defensible action
- Clear accountability for model owners, investigators, risk leaders, and control functions

Governance is therefore not a final check after deployment. It should be built into how organizations design, test, operate, and monitor AI-enabled fraud controls.

Quantum challenges and defense for foundations of digital trust

AI is changing how trust can be impersonated today. Quantum computing represents an emerging threat to digital trust yet also opens new capabilities for broad defense.

A quantum computer of sufficient scale will be able to disrupt security infrastructure; it also provides new computational capabilities that adversaries can leverage to gain organizational intelligence. While the exact day of this threat is unknown, the combination of Frontier AI with quantum computers is bringing the risk closer and not farther[1].

Quantum computing could affect fraud prevention and financial-crime controls by challenging:

- Encryption protecting long-lived sensitive data
- Digital signatures and certificates
- Identity and authentication mechanisms
- Payment integrity and transaction validation
- Investigative records and retained evidence
- Find vulnerabilities and loopholes organizational controls

Beyond cryptographic disruption, organizations should consider how advanced computational capabilities could affect fraud intelligence itself. Potential areas include target prioritization, relationship mapping, attack-path optimization, vulnerability discovery, and analysis of complex fraud networks.

These possibilities remain forward-looking and should be treated as scenarios for structured analysis—not as established fraud techniques. The practical question is how organizations can preserve trust in identities, records, decisions, and investigative evidence as the underlying methods of computation evolve.

Quantum computing and quantum-inspired methods can yield new defense mechanisms to prevent against fraud by detecting subtle signals in data and hidden structure in fraud or financial-crime networks.

Quantum-related planning should follow the same disciplined principles: validated risk assessments, documented dependencies, approved use cases, and clear ownership. Organizations should distinguish practical near-term actions:

1. Create post quantum resiliency—such as creating a cryptographic inventory and improving crypto-agility
2. Explore defensive capabilities with quantum-inspired methods—such as enhancing models with quantum indicators and quantum enhanced network analytics

The objective is not to replace investigators or automate decisions without oversight.

[1] New Advances Bring the Era of Quantum Computers Closer Than Ever | Quanta Magazine



What fraud leaders should do differently

Frontier AI and quantum computing are changing the nature of fraud risk—not simply adding new tools to an existing fraud environment. Frontier AI enables more adaptive, coordinated, and autonomous fraud activity. Quantum computing introduces a longer-term challenge to the cryptographic foundations that support identity, authentication, transactions, and secure information.

As a result, fraud leaders should reconsider several long-standing assumptions.

1. **Fraud is no longer limited to the transaction**
   A transaction may be only the visible endpoint of a broader attack on the trust stack. Fraud leaders should consider the full chain of trust supporting a decision or payment—including identities, credentials, devices, AI agents, approval paths, data, digital signatures, and investigative records.

   The critical question is no longer only, "Was this transaction anomalous?" It is also, "Which elements made this transaction appear trustworthy, and could any of them have been manipulated?"

2. **The adversary is becoming adaptive**
   Traditional fraud controls often assume that a perpetrator executes a relatively stable scheme and that the organization can improve detection rules over time. Frontier AI challenges that assumption. AI-enabled actors may be able to observe control responses, alter their behavior, and coordinate activity across channels.

   Fraud leaders should therefore shift from static control effectiveness to adaptive resilience—the ability to detect changing behaviors, test controls against new scenarios, and adjust before losses become systemic.

3. **AI agents may be manipulated without compromising accounts**
   As organizations deploy AI agents in customer service, procurement, finance, payments, and investigations, the potential attack surface expands. A perpetrator may not need to compromise an account directly if they can manipulate the data, instructions, permissions, or context that an AI-enabled decision-maker relies on.

   Fraud leaders should treat AI agents as participants in the control environment, with defined authority, traceable actions, escalation thresholds, and human review for high-impact decisions.

4. **Fraud intelligence must become connected**
   The next generation of fraud may be distributed across identities, accounts, devices, communications, transactions, and systems. No single alert may be decisive. The risk may emerge only when seemingly unrelated signals are viewed as part of a coordinated network.

   Fraud organizations should move beyond isolated case management and toward connected intelligence that brings together fraud prevention, cybersecurity, identity, investigations, technology, and risk.

5. **Quantum risk should be treated as a trust-transition issue**
   Quantum computing should not be framed as an immediate fraud technique or as a reason to pursue speculative technology programs. The more practical concern is that sufficiently capable quantum systems could eventually weaken selected cryptographic mechanisms supporting digital signatures, certificates, authentication, payment validation, and long-lived sensitive information.

   Fraud leaders should view post-quantum readiness as part of preserving trust over time. The relevant question is not when quantum computing will arrive, but which trust relationships would be most difficult to replace if cryptographic assumptions change.

6. **Success should be measured by resilience**
   Losses recovered and alerts closed will remain important, but they cannot fully measure readiness for AI-operated fraud. Leaders should also consider:

   - How quickly emerging threats are detected
   - Whether coordinated activity is identified across systems
   - How many high-impact decisions receive independent verification
   - Whether AI-agent actions are explainable and traceable
   - How effectively controls adapt to new behavior
   - Whether critical cryptographic dependencies are understood
   - Whether investigations can produce defensible evidence

The imperative: Make trust resilient

Frontier AI and quantum computing are not separate technology issues. They expose the same underlying dependency: the ability of an organization to establish, verify, and preserve trust across identities, decisions, transactions, evidence, systems, and relationships.

Frontier AI is changing fraud today by making deception more adaptive, coordinated, and autonomous. Quantum computing may create a longer-term challenge to the cryptographic mechanisms that support digital trust. The timing and form of these risks will continue to evolve, but the strategic implication is already clear: organizations cannot treat trust as a fixed foundation.

Readiness should not be viewed as a one-time assessment, a technology implementation, or a prediction of the next fraud scheme. It is the ability to understand where trust is created, recognize how it can be manipulated, adapt controls as threats change, and preserve defensible decision-making as the underlying technology evolves.

Organizations that build this capability should be better positioned to:

- Detect coordinated and adaptive fraud before it becomes systemic
- Govern AI models and agents as part of the control environment
- Connect fraud, cybersecurity, identity, investigations, technology, and risk intelligence
- Prepare for potential changes to cryptographic trust mechanisms
- Preserve evidence and decision quality when traditional signals become less reliable

The organizations best positioned for what comes next will not necessarily be those that deploy the most advanced technology. They will be those that understand where trust is created, recognize how it can be manipulated, and adapt faster than the threat evolves.



Contact us



Mike Weil
Managing Director
Deloitte Financial Advisory Services LLP
miweil@deloitte.com



Shuba Balasubramanian
Principal
Deloitte Financial Advisory Services LLP
subalasubramanian@deloitte.com



Mekena McGrew
Senior Manager
Deloitte Consulting LLP
mmcgrew@deloitte.com

Deloitte.

This document contains general information only and Deloitte is not, by means of this document, rendering accounting, business, financial, investment, legal, tax, or other professional advice or services. This document is not a substitute for such professional advice or services, nor should it be used as a basis for any decision or action that may affect your business. Before making any decision or taking any action that may affect your business, you should consult a qualified professional advisor. Deloitte shall not be responsible for any loss sustained by any person who relies on this document.

As used in this document, "Deloitte" means Deloitte Financial Advisory Services LLP, a subsidiary of Deloitte LLP. Please see www.deloitte.com/us/about for a detailed description of our legal structure. Certain services may not be available to attest clients under the rules and regulations of public accounting.

Copyright © 2026 Deloitte Development LLC. All rights reserved.