



THE DEEPFAKE SUMMIT

March 2, 2026 | Houston, Texas

Resilient Trust™ in the Era of AI-Driven Fraud

Event Intelligence Report
Insights. Analysis. Findings

This report provides an analysis of the inaugural Deepfake Summit, held in Houston, Texas, on March 2, 2026. The summit was a half-day event that brought together fraud fighters, digital identity specialists, biometric technologists, AI security researchers, payments professionals, and governance experts for substantive, unfiltered dialogue.

This report synthesizes session discussions and participants' event and post-event commentary into a structured analysis of the event's key themes, findings, and recommendations.

www.thedepfakesummit.com



EVENT & REPORT SPONSORS

The Deepfake Summit and this Event Intelligence Report were made possible by the following sponsors:

Mitek

incode

 LexisNexis®
RISK SOLUTIONS



PARAVISION



persona

 1KOSMOS®

AWARE



Daon®



ENTRUST

moveris™



ROC

veriff 



IN THE ROOM WHERE IT HAPPENED

While The Deepfake Summit observed Chatham House rules, the following quotes—approved for publication—provide a direct window into the Summit’s energy, priorities, and concerns—and into what practitioners processed, responded to, and carried back to their organizations. These quotes reflect the opinions of the individuals and do not represent an endorsement by any organization

“The conversations weren’t just about the ‘hype’—they were about the practical, layered defenses needed to protect the future of digital identity. Key takeaways: deepfakes don’t live in a vacuum; they thrive in the gaps between identity and payments. Moving beyond single-point solutions to a mix of biometric, behavioral, and cryptographic signals. How we can build privacy-first frameworks without stifling the progress of agentic AI.”

Swati Satpathy · Senior Compliance Leader · Amazon

“This conference brought together an amazing group of people for this urgent and necessary conversation. Deepfakes are not theoretical. They’re here. What is your organization doing?”

Elizabeth Terry · CISSP · Secure Payments & Cybersecurity Expert

“It was energizing to spend time with like-minded peers, reconnect with old friends, and have candid conversations about where this industry needs to go next. If this first summit is any indication, the momentum is just getting started.”

Frances Zelazny · Principal · Identity Strategies

“What stands out about the Deepfake Summit is the focus on practical countermeasures and real dialogue—not vendor noise, not theoretical hand-wringing. We’re past awareness; this is about execution, accountability, and cross-disciplinary collaboration (fraud, identity, AI governance, and regulation all at the same table). If you’re responsible for protecting customers, citizens, or critical systems, this isn’t optional learning—it’s table stakes. Looking forward to the conversations this forum will unlock and the solutions that come out of it.”

Elizabeth Kiehner · Chief Strategy Officer · Nortal

“Insightful conversations regarding data protection, identity security in AI and especially Agentic AI concerns today at The Deepfake Summit”

Dikla Shabtey · Chief Executive Officer · AI · Risk · Cybersecurity

“The deep-dive content of this event was phenomenal bringing in various experts, solution providers, and practitioners together in an intimate size setting. It was great catching up with other fraud fighters and meeting new ones that all have unique niches in this extremely difficult risk to manage!”

Aaron Frye · Founder & CEO · Lucid Point Consulting



TABLE OF CONTENTS

- 1** INTRODUCTION
- 2** EXECUTIVE SUMMARY
- 4** SUMMIT CONTEXT & FROMAT
- 5** KEY THEMES
- 12** SUMMIT SESSION SUMMARIES
- 24** KEY FINDINGS
- 28** RECOMMENDATIONS
- 32** THE CHALLENGES AHEAD
- 33** ROAD TO THE NEXT DEEFAKE SUMMIT





INTRODUCTION



Welcome to The March 2026 Deepfake Summit Event Intelligence Report. This report introduces a new format for The Prism Project—a market education initiative created by Acuity Market Intelligence to bridge the gap between the identity technology intelligentsia and the enterprise professionals evaluating and deploying digital identity solutions to meet the challenges of digital transformation.

The Deepfake Summit is the first event to emerge from The Prism Project's research and reporting. Specifically, it brought the findings of the three 2025 Prism Project reports—Deepfake and Synthetic Identity Prism Report, Privacy and Compliance Prism Report, and Flagship Prism Report—into live, practitioner dialogue.

Where those reports diagnosed the threat environment and catalogued the ecosystem of solutions, the summit asked a harder question: what are practitioners on the front lines actually seeing, and what are they actually doing about it?

The answer, as you will read in the pages that follow, is both more alarming and more encouraging than expected. More alarming because the data confirmed what The Prism Project has been tracking: the fraud problem has fully entered the AI era, and most organizations remain calibrated for yesterday's threat landscape. More encouraging because the room was full of practitioners who are clear-eyed about the scale of the problem, technically sophisticated in their understanding of countermeasures, and deeply committed to the kind of cross-institutional collaboration that The Prism Project has long argued is essential.

The Prism Project is grounded in a philosophy of identity built on four key pillars:

- Digital identity belongs to the person it describes.
- True identity empowerment relies on government systems of record.
- Identity must be consistently and continuously orchestrated across both physical and online channels to remain secure.
- Biometrics must be central to any sustainable, reliable, and secure digital ecosystem, with identity flowing freely between virtual and physical worlds.

Every conversation at the Deepfake Summit reinforced these pillars—and illuminated how far the industry still has to go to fully realize them. The core thesis of The Prism Project's Resilient Trust™ framework (defined in the 2025 Flagship Prism Report) emerged from Houston with even greater urgency: without foundational biometric identity linking a physical human to their digital data, you do not truly have digital identity at all. And without continuous, layered, ecosystem-wide verification of that link, you have an open invitation for the synthetic fraud that is already inside the gates.

As ever, my collaborators and I are evangelists of strong identity and believe that the only way to move forward in our time of digital transformation is to take the ethics of human identity seriously in both the physical and digital realms. As a community, we believe that the identity industry and its stakeholders are morally obligated to develop and implement powerful digital technologies for the good of humanity. By reading and sharing this report, you are participating in the positive change required to close identity gaps. Together, we can usher in an identity-safe future for all



Authentically yours,

Maxine Most
Chair, The Deepfake Summit
Founder & CEO, The Prism Project
Principal, Acuity Market Intelligence



1





EXECUTIVE SUMMARY

The inaugural Deepfake Summit convened a curated group of practitioners, technologists, and policy experts to examine the accelerating convergence of AI-driven impersonation fraud, synthetic identity, and the evolving architecture of digital trust.

Across keynotes, practitioner panels, and moderated discussions, a consistent conclusion emerged: **the fraud problem has entered the AI era while most institutional defenses remain anchored in obsolete threat models.** Participants shared a broad recognition that the scale, speed, and accessibility of AI-enabled impersonation tools are fundamentally altering the operating environment for digital identity and financial services.

These were not theoretical warnings. They were expert practitioner assessments of the current state of affairs. Across sessions, participants described a threat landscape defined by three converging forces reshaping how organizations must think about identity assurance, fraud prevention, and digital trust infrastructure:

- The rapid industrialization of synthetic identity creation
- The broad accessibility of AI-powered impersonation tools
- The increasing automation of fraud operations




PARTICIPANT OBSERVATION

"You already have a large number of synthetic identities enrolled at your institutions today."



AI-DRIVEN FRAUD: FROM EMERGING RISK TO OPERATIONAL REALITY




8,000%

Deepfake-associated fraud has reached industrial scale. The World Economic Forum reported an 8,000% increase in deepfake use.




\$150-\$200

Fraud-as-a-Service is real, accessible, and rapidly expanding. KYC-validated bank accounts are openly for sale on dark web forums for \$150-\$200 each



>50%

Non-human interactions have crossed the majority threshold. More than 50% of all online activity is now non-human



15-20%

Synthetic identities are already inside enterprise systems. Estimates suggest 15-20% of some institutions' portfolios may already contain synthetic identities



\$16M+

The cost of inaction is not abstract. A single fraudulent transaction can result in losses of \$16 million or more



AI-Driven Fraud is Industrialized: scalable, automated, accessible, embedded, and highly profitable.



EXECUTIVE SUMMARY

These observations point to a structural shift in how digital trust must be managed. Traditional fraud defenses—largely designed to detect anomalies after the fact—are increasingly challenged by adversaries operating with automated tools, scalable synthetic identity generation, and rapidly evolving attack techniques.

Within this context, a central theme of the summit was the need to strengthen what The Prism Project describes as Resilient Trust™: the ability of digital ecosystems to maintain reliable identity assurance, transactional integrity, and privacy-centric governance accountability even as adversarial technologies evolve.

Participants repeatedly emphasized that the technologies required to build more resilient systems—biometrics, behavioral analytics, cryptographic credentials, device intelligence, and advanced fraud detection—already exist. The more difficult challenge lies in deploying these capabilities coherently across institutions, sectors, and governance frameworks.

The discussions at The Deepfake Summit, therefore, did not focus solely on emerging threats. They also highlighted the practical steps organizations can take to strengthen identity infrastructure, improve fraud resilience, and prepare for a future in which AI-enabled impersonation will remain a persistent feature of the digital landscape.

THE SUMMIT PRODUCED A CLEAR STRATEGIC AGENDA

The industry must move from reactive fraud detection to proactive identity assurance, from siloed defenses to ecosystem-wide collaboration, and from static onboarding to continuous identity verification. The tools exist. The standards are maturing.



What is missing is the urgency, coordination, and will to change.



PARTICIPANT OBSERVATION

“We are not just worried about AI. We are worried about the deployment of AI where governance, regulation, and understanding of liability are severely lacking.”





SUMMIT CONTEXT & FORMAT

The Deepfake Summit was conceived as a practitioner-focused forum designed to facilitate substantive dialogue across the communities most directly responsible for protecting digital identity and financial systems. Produced by C. Maxine Most of Acuity Market Intelligence and The Prism Project, the summit builds on broader research examining biometric-centric digital identity, fraud prevention, and the evolving role of identity infrastructure in the digital economy.

Unlike large industry conferences that often emphasize vendor announcements and pitches, or high-level commentary, the Deepfake Summit was intentionally structured to encourage direct operational exchange among practitioners confronting AI-enabled fraud in their daily work. The format was deliberately intimate: a curated event with a small group of hand-selected experts. As multiple participants noted, this created conditions for candid, practical conversation that larger conferences rarely permit.

Within this setting, speakers and attendees examined issues ranging from synthetic identity fraud and injection attacks on identity verification systems to the implications of agentic AI and the broader evolution of identity infrastructure required to support secure digital interactions. The summit functioned not only as an educational event but also as an opportunity for practitioners to compare operational experiences, share emerging threat intelligence, and explore collaborative approaches to strengthening fraud defenses across the ecosystem.

PARTICIPANT OBSERVATION

“The deep-dive content of this event was phenomenal, bringing various experts, solution providers, and practitioners together in an intimate setting.”

PARTICIPANT OBSERVATION

“The discussions were substantive, the format was genuinely interactive, and the commitment in the room made it clear this could be the start of something meaningful as we collectively confront the challenges ahead.”

The Summit included:

- Consumer banking, lending, healthcare, and energy fraud professionals
- Payments and fintech risk specialists
- Operational and cybersecurity experts
- Digital identity and biometrics technologists
- AI security researchers and deepfake detection experts
- Governance and regulatory policy experts
- Representatives from identity standards bodies and government programs





KEY THEMES

Six interconnected themes ran through every session at The Deepfake Summit.

Taken together, they define the strategic challenge facing every organization that must distinguish real humans from synthetic ones in digital interactions—which is to say, every organization that does business online.





THEME 1: The Scale Has Already Changed—and Most Organizations Haven't Caught Up

The single most repeated observation across the summit was the mismatch between how fast the threat is moving and how slowly institutional defenses are adapting. This was not theoretical concern—it was grounded in live data from practitioners with daily exposure to AI-enabled fraud.

- Panelists cited the World Economic Forum's reported 8,000% increase in deepfake use as a market-defining data point.
- LexisNexis data showed deepfake-associated fraud rising from 20% to 40% to nearly 70% of relevant transactions across 2023, 2024, and 2025.
- Fraud no longer occurs in occasional bursts. Attackers now use AI to test-and-learn at scale—running small probing attacks to identify control gaps, then striking hard and fast once a vulnerability is confirmed.
- Conversations with banks, governments, and fraud teams make the threat "viscerally clear"—yet most organizations remain focused on defending against yesterday's attacks.

What makes this particularly difficult to address is the nature of the measurement challenge. As one panelist summarized: you can only manage what you can measure, and you cannot measure attacks you do not know exist. This is not simply a matter of deploying new tools. It is a structural reckoning with the limitations of current fraud-detection systems based on an obsolete audit-based approach.



Key Challenge

You can only manage the things you can measure, and you cannot measure attacks that you do not know exist.

Critical Insight

The threat is not accelerating toward institutions—it has already arrived. The gap between the speed of attack evolution and the speed of institutional adaptation is widening every quarter. Catching up requires treating this as an operational emergency, not a technology roadmap item.

PARTICIPANT OBSERVATION

"The consistency is test and learning. The attacks are very small to find the gap. And then they become very big very quickly once they understand where the control gap is."





THEME 2: Deepfakes Are Not Visible In Existing Data—and That is a Core Problem

One of the summit's most important findings was what practitioners called measurement blindness: the systematic inability to see and quantify deepfake fraud in existing data. This is not a failure of effort—it is a structural limitation of current fraud detection systems that creates a dangerous and self-reinforcing dynamic.

Even expert practitioners cannot reliably distinguish AI-generated identities from real ones when shown a set of images. Organizations face the same limitation in their own transaction data. The consequence is that investment in deepfake defenses is systematically delayed—organizations cannot justify spending on a problem they cannot demonstrate exists. This is not irrational behavior; it is the predictable outcome of a measurement failure that must be diagnosed and corrected before meaningful remediation can begin.

- Fraud that appears to be a software error, a system glitch, or a processing anomaly may, in fact, be an injection attack or an identity probe.
- Synthetic identities with established histories and clean transaction records appear identical to legitimate customers in standard fraud-monitoring systems.
- The lag between attack evolution and defensive response is widening precisely because the attacks are designed to be invisible.

Key Challenge

Deepfake detection defenses investment is systematically delayed because organizations cannot justify spending on a problem they cannot demonstrate exists. This is not irrational—it is the predictable outcome of a measurement failure that must be corrected before any other remediation can begin.



Critical Insight

The absence of visible deepfake fraud in existing data is not evidence that it isn't there. It is evidence that current detection systems were not designed to surface it. Measurement blindness and actual absence look identical from inside a system that cannot tell the difference.

PARTICIPANT OBSERVATION

"It's hard to be proactive and spend money on a problem you can't really understand exists."





THEME 3: Synthetic Identity is a Long Game—and Organizations are Already Losing it

The Summit provided in-depth insight into how synthetic identity fraud actually operates in the wild—insight that should fundamentally change how institutions approach portfolio risk. Synthetic identity fraud is not a smash-and-grab operation. It is a long game, played with patience and precision, designed to exploit the structural incentives and measurement blind spots of IT and financial systems.

The following example of the synthetic identity lifecycle long game was presented:

- Synthetic personas apply for small loans from fintech startups, repay them reliably, and build credit profiles over months and years. Average credit scores settle in the 650–700 range.
- Fintech startups measured by account volume have a reduced incentive to scrutinize paying customers who may be synthetic. The incentive structure itself enables fraud.
- Once a strong credit profile is established, the synthetic identity moves up the chain to larger institutions capable of issuing larger loans—and that is when payments stop.
- Conservative estimates from multiple panelists: 15–20% of some institutions' portfolios may already contain dormant synthetic identities. They are already inside. The question is what to do tomorrow.
- The entire prior transaction history—every on-time payment, every modest credit line—builds the infrastructure for a single large-scale extraction.



Key Challenge

Standard fraud detection systems flag anomalies. Synthetic identities are engineered to produce none. The portfolios most at risk are those where the fraud looks most like success—on-time payments, growing credit lines, no alerts triggered.

Critical Insight

Don't think of fraud as having to result in a loss at your specific institution. You may simply be the steppingstone in a much longer play..

PARTICIPANT OBSERVATION

“About 15 to 20% of someone's portfolio could be synthetic. They're good. They look great. Nobody's doing the behavioral analytics on those accounts. They're already in.”





THEME 4: Injection Attacks are The Most Underrated Threat Vector


Multiple practitioners and researchers emphasized that injection attacks—where adversarial content is inserted directly into identity verification and authentication pipelines or backend systems—are significantly underappreciated relative to the threat they pose. Unlike presentation attacks (e.g., holding a photo up to a camera), injection attacks leverage vulnerabilities in capture devices or bypass them entirely, inserting synthetic media directly into the data stream at multiple points between the capture and the final verification. The distinction matters enormously for countermeasure design.

Participants highlighted that injection attacks were “often underrated” relative to the attention given to deepfake video—a gap that represents a significant tactical blind spot. Effective defense against injection requires a distinct set of controls:

- Detecting whether a device is stationary (rack-mounted devices are a red flag)
- Requiring that device motion during image capture corresponds to what an actual human holding that device would produce
- Detecting and blocking jailbroken or rooted devices, which are completely open to injection attacks
- Shifting all identity capture away from Windows machines and webcams to secured mobile devices
- Verifying that captured images actually come from the camera corresponding to the device’s IMEI

Key Challenge

The industry’s countermeasure investment is concentrated on presentation attacks and deepfake video. Injection attacks—which bypass the camera entirely—require a completely different set of controls that most institutions have not yet deployed, and do not yet recognize as a distinct threat category.



Critical Insight

Defending against injection is not an extension of existing liveness and deepfake detection. It is a separate discipline requiring device integrity controls, mobile-first capture architecture, and behavioral motion analysis. Organizations that conflate the two threat models will be underprotected against both.

PARTICIPANT OBSERVATION

“The biggest vulnerability is injection attacks—and it is a very difficult problem to solve. You cannot allow jailbroken phones. You cannot allow Windows machines and webcams anymore.”





THEME 5: Collaboration is The Missing Layer—Both Within and Between Organizations

Collaboration emerged as the summit's most resonant theme—the point at which virtually every speaker, panelist, and attendee converged. When one practitioner asked what single change her session counterpart would make with a magic wand, the answer came without hesitation: not more tools, not more alerts, but better coordination across teams, companies, and sectors. That answer crystallizes a foundational pillar of The Prism Project: the technology gap is narrower than the collaboration gap.

The collaboration gap operates at three distinct levels:

- Within institutions, fraud teams, cybersecurity, compliance, identity management, and AML teams frequently operate in separate silos with no shared data or common risk language—a dynamic that creates structural vulnerabilities no individual team can see or address on its own.
- Between institutions: a synthetic identity can build a legitimate profile at one institution and then exploit a second. Cross-institutional intelligence sharing is essential—but existing consortia are themselves siloed, sharing data only with their own members, not with each other.
- Across industries: law enforcement, government, regulators, and technology companies must all be part of the response. Fraudsters simply move between jurisdictions when collaboration is absent.

A critical structural point raised in discussion: consortium-based intelligence sharing is growing, but the consortia themselves don't share with each other. The result is still siloed information—just with slightly larger silos. True ecosystem intelligence requires standardizing risk signals across platforms and a willingness to compete on products and services rather than on fraud data.



Key Challenge

Consortium-based intelligence sharing is growing—but the consortia don't share with each other. The result is slightly larger silos, not ecosystem intelligence. True coordination requires standardized risk-signal taxonomies and a willingness to treat fraud data as pre-competitive infrastructure rather than as a proprietary advantage.

Critical Insight

The magic wand answer from the room was not more tools or more alerts—it was better coordination. The technology gap is narrower than the collaboration gap. Solving the coordination problem would unlock the full value of capabilities that already exist and are already deployed.

PARTICIPANT OBSERVATION

“Fraud signals have always existed and they will continue to evolve. Those signals must be shared across the organization and across the ecosystem. Silos create advantage for attackers.”





THEME 6: Agentic AI Demands Immediate Governance

The agentic AI sessions produced the summit's most forward-looking—and most sobering—discussions. Agentic AI systems do not merely generate content; they perceive, reason, plan, and take action autonomously. They are already performing everyday financial tasks—purchases, transfers, account management—on behalf of individuals, and the governance infrastructure required to verify that those agents are genuinely authorized by real humans does not yet exist.

This is not a future risk. It is a present-day gap being actively exploited.

- More than 50% of all online activity is already non-human. Organizations frequently cannot determine whether they are interacting with a verified human, an authorized agent, or a malicious bot.
- AI agents are already performing everyday financial tasks—purchases, transfers, and account management. Ensuring that every agent is backed by a verified person or organization is critical to consumer safety.
- Live coding now allows AI agents to autonomously spawn other AI agents—a capability that dramatically expands the attack surface for adversarial use.
- Agent workflows are commercially available on platforms like Etsy for under \$2. The barrier to deploying sophisticated automated fraud operations is, in practical terms, negligible.

The summit reached clear consensus: agentic AI is not inherently a threat or a solution—it is both. The same capabilities enabling legitimate automation can be weaponized against every stage of the identity lifecycle. The critical need is for governance frameworks that include monitoring, audit logging, traceability, and standards for delegated identity authority, so that every agent action can be traced to a verified human principal. Organizations that treat this as a future problem to be addressed after further development do so at their own risk.

Key Challenge

No adequate governance framework exists for verifying that an AI agent is genuinely authorized by a real human, for scoping the actions it may take on a principal's behalf, or for establishing liability when agent actions cause harm. Fraudsters are exploiting this gap now, while responsible actors debate how to close it.



Critical Insight

Agentic AI collapses the assumption that a human is making decisions. Every identity verification, authorization, and access control designed around human behavior must now account for agents that can act faster, at greater scale, and with fewer of the behavioral signals that fraud detection relies on.

PARTICIPANT OBSERVATION

“Non-human interactions are now greater than 50% of all activities online. Do we really know if we are engaging with other humans or unverified bots and agents?”





SUMMIT SESSION SUMMARIES

The sessions at The Deepfake Summit were intentionally designed and sequenced to create a logical progression of content, with each session building upon the previous ones.

The agenda served as a roadmap for Summit’s discussion of the current state, evolution, and future of AI-generated impersonation fraud.



AGENDA

WELCOME & INTRODUCTORY KEYNOTE -Tackling AI-Driven Impersonation Fraud: Putting the “Resilience” in Resilient Trust

Fireside Chat—The **Evolving Threat Landscape**—How Deepfakes, Synthetic Identity, and Agentic AI Intensify Fraud

Panel—State of Play: Real-World **Countermeasures**

Panel—It’s the **Ecosystem**, Stupid: Identity, Payments & the New Trust Infrastructure

KEYNOTE—How **Agentic AI** Is Changing the Game... or Not

Panel—Is Agentic-AI an Identity Problem or a Solution?

Panel: **Future Fast-Forward**: Where Do We Go From Here?

Wrap Up—**Key Insights**



OPENING KEYNOTE: Tackling AI-Driven Impersonation Fraud Putting the 'Resilience' in Resilient Trust™

CORE THESIS

Without foundational biometric identity linking a physical human to their digital data, you do not truly have digital identity at all.

KEY FRAMING POINTS

- Identity must evolve from a verification tool to foundational digital infrastructure—as essential as payment rails or network protocols. Organizations that treat it as a compliance function rather than a strategic asset are structurally underinvesting.
- Privacy and compliance must be treated as enablers of trust and strategic competitive assets, not as friction. Trust is a scarce and highly valued commodity, and it can be a genuine differentiator.
- AI has democratized fraud at scale. Fraud-as-a-service now makes synthetic identity generation accessible to anyone with an internet connection. Agentic AI pours fuel on that fire: attacks that once required human labor can now be automated, parallelized, and run continuously.
- Privacy and security are not competing priorities. The Resilient Trust™ framework unifies them: privacy-preserving identity technologies are essential to building the trusted ecosystems that make fraud prevention possible.
- Resilient trust is trench warfare—an ongoing, evolving discipline, not a one-time deployment. There is no end state; there is only continuous defense and continuous improvement.

The Summit opened by presenting The Prism Project's Resilient Trust™ framework as the intellectual architecture underpinning the event. The central argument rests on two critical points:

1. The identity systems we rely on for financial services, payments, government services, and all digital transactions must be reimaged.

These systems were never designed to withstand adversarial AI. The industry must redesign trust infrastructure around identity itself—treating identity not as a verification tool but as foundational digital infrastructure, as essential as payment rails or network protocols.

2. Privacy and security are not competing priorities.

Privacy-preserving identity technologies—verifiable credentials, biometric binding, decentralized identity architectures—are not obstacles to fraud prevention; they are essential to the trusted ecosystems that make fraud prevention possible.

The Deepfake Summit was convened to test that argument against the lived experience of practitioners, and it held.

The keynote established the data baseline: 300–3,000% reported increases in deepfake activity, with some sources reporting up to 8,000%. The implication is that fraud has shifted from a problem individual practitioners manage—one suspicious transaction at a time—to an automated, machine-scale pressure applied continuously and systematically against digital infrastructure. This is the operational environment that every attendee in the room was navigating.



FIRESIDE CHAT: The Evolving Threat Landscape

How Deepfakes, Synthetic Identity, and Agentic AI Intensify Fraud

CORE THESIS

The fraud industry's signal expertise is not a legacy to be replaced by AI — it is the foundation for AI-era defense. Every deepfake, every synthetic identity, every injection attack still produces signals. The discipline is in synthesizing them faster, more holistically, and in real time across every channel simultaneously.

Two practitioners conducted what multiple attendees described as the event's most grounded session—an unfiltered conversation between an identity and trust thought leader and a fraud professional who deals with AI-enabled fraud every day.

The session's central insight reframes how the industry should think about its existing capabilities: fraud signals have always existed and always will. The fraud industry's existing expertise in signal and pattern analysis is not a legacy to be discarded in favor of AI-native tools—it is the foundation for AI-era defense. What AI changes is the speed at which signals must be synthesized and acted upon, and the breadth of signals that must be considered simultaneously.

The magic wand answer—not more tools, not more alerts, but better coordination—crystallized the session's conclusion and set the tone for every discussion that followed. The industry does not have a technology problem. It has a coordination problem. Solving the coordination problem will unlock the full potential of the technology that already exists.



PANELIST OBSERVATION

"if we're not good at detecting an AI-generated identity document yet—there are still foundational signals attached to that image. We need to analyze those signals holistically and in real time."

KEY FRAMING POINTS

- Fraud signal expertise is the foundation for AI-era defense—not a legacy to be replaced. What AI changes is the speed and breadth of signal synthesis required.
- Deepfakes do not eliminate signals—they demand holistic, real-time analysis of all signals simultaneously, including those attached to AI-generated artifacts.
- New technology solutions stay ahead of attackers for roughly four months before being broken. The investment cycle is continuous; permanent solutions do not exist, and expecting one is itself a vulnerability.
- Injection attacks can masquerade as technology errors or system performance issues. Without real-time detection built in, organizations may spend weeks diagnosing—and inadvertently maintain—an active attack.
- The reactive posture—spending after losses rather than investing in prevention—is more expensive in every case. A single fraudulent transaction can cost \$16 million. The investment argument is straightforward once losses are quantified.
- Regulatory pressure is already arriving. The OCC and other bodies are scrutinizing fraud controls with new intensity, and cease-and-desist findings for inadequate frameworks are no longer theoretical.
- The coordination gap is wider than the technology gap. The single most impactful change practitioners could make is better coordination—across teams, companies, sectors, and with law enforcement and regulators. Fraud, identity, data risk, product, and compliance teams must operate from a shared picture—not sequential handoffs.





PANEL: State of Play—Real-World Countermeasures

CORE THESIS

There is no silver bullet, and there never has been. Effective defense against AI-era fraud requires layered, persistent, continuously updated controls — and it begins with measurement. You cannot fix what you cannot see, and you cannot see what your systems were never designed to surface.

This panel brought together fraud practitioners with live data on the evolving threat landscape. The discussion moved quickly from the scale of the problem to the specific technical and organizational requirements for effective defense. Panelists were consistent in their diagnosis: there is no silver bullet, there never has been, and the organizations that believe otherwise are the most exposed.

The panel outlined the full synthetic identity lifecycle in detail—small fintech loans, credit-building, dormancy, large-institution strike—and underscored the statistic that crystallized the room: 15–20% of some portfolios may already be synthetic, quietly building credit profiles indistinguishable from those of legitimate customers. The question at that point is no longer how to prevent enrollment but how to identify identities that have already passed verification.

Panelists agreed that the inability to clearly identify deepfake fraud in existing data is the root cause of underinvestment. Even expert practitioners cannot reliably distinguish AI-generated identities from real ones by visual inspection. This is not a training failure—it is a structural property of current-generation synthetic media that demands a structural response: deploying detection capabilities specifically designed to identify synthetic media, rather than relying on human judgment applied to data streams that were never designed to surface the relevant signals.

On investment, a panelist cited \$16 million lost in a single fraudulent transaction, then asked why banks would not invest even a fraction of that in prevention. The honest answer—banks cannot quantify what they cannot see—underscored the problem of measurement blindness with uncomfortable precision. Fix measurement first, and the investment case becomes self-evident.

Panelists drew an explicit connection to anti-money laundering: the fraud signal problem and the AML signal problem are not separate disciplines—they are fed by the same synthetic identity infrastructure. AML teams and fraud teams applying the same principles to the same synthetic identity vectors are duplicating effort. The cross-functional alignment argument applies here with full force.

On countermeasures, the panel was unambiguous—there is no silver bullet:

- Detection requires layered defenses: behavioral biometrics, device integrity, camera verification, liveness detection, and ongoing transaction monitoring.
- For injection attacks specifically: move all identity capture to secured mobile devices, block jailbroken phones, and verify camera-IMEI correspondence.
- Ongoing monitoring must replace one-and-done KYC. Once fraudsters are inside, every move must be tracked—the question shifts from ‘keep them out’ to ‘what does their movement look like?’



PANEL: State of Play—Real-World Countermeasures

KEY FRAMING POINTS

- There is no silver bullet—there never has been. Organizations that believe otherwise are the most exposed.
- Attackers test at low volume to identify control gaps, then strike hard and fast once the vulnerability is confirmed. The pattern is consistent and well-documented.
- Measurement blindness is the root cause of underinvestment. Even expert practitioners cannot visually distinguish AI-generated identities from real ones—which means the absence of visible fraud in existing data is not evidence of safety.
- Synthetic identity fraud is a long game. Estimates suggest 15–20% of some portfolios may already contain dormant synthetic identities with clean credit histories. The question has shifted from how to prevent enrollment to how to find those already inside.
- Layered defense is the only viable architecture: behavioral biometrics, device integrity, camera-IMEI verification, liveness detection, and continuous transaction monitoring must work in combination.
- Consortium-based intelligence sharing across institutions is essential—but fragmented. Organizations compete on services, not on fraud. Shared signals, shared data, and real-time intelligence exchange should be treated as pre-competitive infrastructure.

PANELIST OBSERVATION

“If you shut the front door and they find the back door, you have to look at the window, the second floor, and everything else. You always have to have a layered approach—not just for today but for tomorrow and every day after.”

”





PANEL: Identity, Payments, and the New Trust Infrastructure

CORE THESIS

The industry must shift from reactive fraud detection to proactive identity verification. Cryptographically verifiable credentials — anchored to a real human, issued by a trusted government system of record — are the strongest available authenticity signal. Deepfakes cannot yet spoof them. The work is adoption, not invention.

This session examined the structural gaps in the broader digital trust ecosystem, drawing on practitioners working at the intersection of identity standards, verifiable credentials, government identity infrastructure, and payments security. The conversation surfaced a fundamental architectural error at the heart of current financial services identity practice: KYC verification at onboarding is treated as sufficient protection for all subsequent transactions.

This is wrong—and it was wrong before the AI era. The gap between identity proofing and payment authorization—between the moment of enrollment and every subsequent moment—is where synthetic identities live and operate. Closing that gap requires continuous identity assurance, not a single checkpoint at the front door.

A long-standing pioneer of user-centric digital identity standards articulated the user-sovereignty dimension: identity and behavioral data must belong to the individual who generates it, not to the intermediaries who collect it. The tools for user-centric identity—digital wallets, verifiable credentials, decentralized identity—now exist and are maturing. The obstacle is not technical. It is the industry's institutional inertia, having invested heavily in existing infrastructure and being reluctant to acknowledge its fundamental inadequacy.

This exchange surfaced one of the summit's most important conceptual contributions—the distinction between fraud signals and authenticity signals:

Fraud Signals vs. Authenticity Signals

Fraud signals are reactive—they detect suspicious behavior after it has occurred. Authenticity signals are proactive—they verify that an identity is real before it acts. Digital IDs (mDLs, verifiable credentials) are authenticity signals: cryptographic, government-issued, deterministic rather than probabilistic.

The core insight is compelling: deepfakes can spoof voice, photo, and video with increasing fidelity. They cannot yet spoof a cryptographically verifiable government credential that requires physical presence to obtain. This makes government-issued digital IDs—an mDL, a digital passport, or a verifiable credential issued against a government system of record—potentially the strongest available proof of identity for online and in-person transactions.

The limitations of this approach were acknowledged with equal candor:

- Not everyone has a government digital ID. US financial institutions must serve people who don't have them, so credentials alone cannot be the gate.
- Gaining access through one door (credential verification) does not close the back door (system flooding, injection, behavioral manipulation).
- Credentials must be accompanied by behavioral authentication—verifying that the credential holder is behaving as expected, not just that they possess the credential.



PANEL: Identity, Payments, and the New Trust Infrastructure

KEY FRAMING POINTS

- KYC at onboarding is a structural error, not a calibration problem. The gap between identity proofing and payment authorization is where synthetic identities live and operate.
- Fraud signals are reactive. Authenticity signals are proactive. The industry must shift its orientation from detecting fraud after it occurs to verifying identity before action is taken.
- Cryptographically verifiable government credentials—mDLs, digital passports, verifiable credentials—are the strongest available proof of identity for online transactions. Deepfakes cannot yet spoof them. But adoption remains a critical unsolved challenge.
- User sovereignty over identity and behavioral data is both a privacy principle and a fraud defense. Until individuals can grant permission to access their own data—rather than ceding it to intermediaries by default—the ecosystem remains structurally exploitable.
- The governance gap for digital credentials is underappreciated: as issuers multiply beyond state DMVs, there is no registry of legitimate issuers and no standardized method for relying parties to verify credential provenance. Solving this is as important as issuing the credentials themselves.
- The incentive structure is misaligned: fraud and compliance teams are rewarded for blocking, not for accuracy. False positives are a real and invisible harm. Automation must include escalation paths for edge cases.
- Credentials alone are not sufficient. Behavioral authentication must accompany credential verification to confirm that the credential holder is acting consistently with their established identity.

PANELIST OBSERVATION

“Most in the IDV industry still talk about shared fraud signals. Maybe we should talk more about shared authenticity signals granted with user consent. That’s what digital IDs are. Fraud signals are reactive. Authenticity signals are proactive. Offense vs. defense.”





KEYNOTE: How Agentic AI Is Changing the Game, or Not

CORE THESIS

Agentic AI is categorically different from generative AI — it does not wait to be prompted. It perceives, reasons, plans, and acts. The governance infrastructure required to deploy it safely does not exist at the speed of deployment, and the gap between the two is being actively exploited by adversaries who face no governance constraints.

This session moved from fraud defense to the rapidly evolving architecture of AI systems themselves, reorienting the room's understanding of the threat environment in important ways. The speaker outlined what agentic AI actually means in practice—perception, reasoning, planning, action—and why it is categorically different from generative AI tools that simply produce content. An agentic AI system does not wait to be prompted. It acts.

The infrastructure enabling this is already widely deployed: model gateways, model hubs, tool marketplaces, agent-to-agent protocols, and agent workflow marketplaces that make sophisticated automated capabilities accessible to anyone with a few dollars and an internet connection. Agent workflows are available on commercial platforms for under \$2. The barrier to deploying sophisticated automated fraud operations is, in practical terms, negligible.

The governance requirements for responsible deployment—monitoring, observability, audit logging, traceability, guardrails for automated decision-making—are not optional features. They are prerequisites for safe deployment that are largely absent from current practice. The failure cascade was laid out plainly: when domain experts and agent builders cannot communicate, explainability breaks down, controls break down, feedback loops disappear, and the organization accumulates governance risk it cannot see.

KEY FRAMING POINTS

- Agentic AI is categorically different from generative AI. It does not wait to be prompted—it perceives, reasons, plans, and acts. The threat model must be updated accordingly.
- The infrastructure enabling adversarial agentic use is already widely deployed: model gateways, tool marketplaces, agent-to-agent protocols, and workflow platforms accessible for under \$2. The barrier to deploying sophisticated automated fraud operations is, in practical terms, negligible.
- Agent collusion and token drain are attack vectors that do not appear in traditional fraud frameworks. Multi-agent environments require governance that explicitly addresses how agents interact with each other, not just with humans.
- When domain experts and agent builders cannot communicate—when the tacit knowledge of identity, fraud, or compliance professionals does not translate to the people building agents—governance breaks down in a compounding cascade: poor explainability leads to inadequate controls, which leads to missing feedback loops, which produces systemic risk invisible to the organization.
- Governance requirements are not optional add-ons: monitoring, observability, audit logging, traceability, & guardrails for automated decision-making are prerequisites for safe agentic deployment, not afterthoughts.
- Total cost of agentic deployment is routinely underestimated. Governance, observability, and monitoring infrastructure are expensive. Cutting corners creates risks that are also expensive—just harder to see.
- The accessibility gap between legitimate users and bad actors has collapsed. Sophisticated automated operations are within reach of virtually any adversary with an internet connection and a few dollars.



PANEL: Is Agentic AI an Identity Problem or a Solution?

CORE THESIS

Agentic AI is simultaneously the most significant threat to identity infrastructure and the most promising tool for its defense. The balance tips entirely on governance. The technology is neutral. Will responsible actors build the frameworks — authorization standards, agent identity protocols, liability structures — before adversarial actors finish exploiting their absence?

This panel tackled one of the summit's most contested questions—and arrived at an honest answer: agentic AI is both a problem and a solution simultaneously. How the balance tips depends entirely on governance. The same generative AI capabilities driving deepfake attacks are also driving improvements in deepfake detection. The same agentic infrastructure enabling fraud automation is also enabling legitimate fraud prevention automation. The technology is neutral. The governance is not.

From the biometric technology perspective, the technical arms race is real, and detection must keep pace. A panelist offered the summit's sharpest framing of the identity challenge in an agentic world: a deepfake cannot walk into a bar—physical presence still anchors identity in the physical world. But in the digital world, deepfakes can exist everywhere, and the distinction between physical and digital presence is collapsing.

This led to discussion about the delegated authority challenge for agentic systems: verifying the human-to-agent relationship, scoping authorization to specific permitted actions, and establishing out-of-band verification that conveys both the access request and its purpose. Identity systems must now verify not just who someone is, but whether the entity acting is genuinely authorized to take the specific action it is attempting—a requirement that current infrastructure is not designed to meet.

The panel flagged an emerging concept gaining traction in standards forums: the idea of a birth certificate for an agent—a minimum identity standard that establishes who made the agent, what it is authorized to do, how mature it is, and what permissions it holds. Just as you would not let a 14-year-old manage a 401(k), you should not let a 14-second-old agent make high-value financial transactions. Scoping agent authority by maturity and context is a governance primitive that the industry does not yet have.

A hiring fraud example was cited, with one panelist reporting that three of the last four candidates in a recent hiring process were fake, and AI-generated identity representations appearing on video calls. This is not a financial fraud scenario—it is an employment integrity scenario. Agentic AI is breaching trust boundaries across every domain where identity matters, not just payments.

An equity issue was raised that cuts across every theme of the summit: the capabilities available to large institutions—advanced fraud detection, real-time signal analysis, sophisticated biometrics—are not accessible to community banks and credit unions. Fraud actors exploit the weakest points in the ecosystem. The entire ecosystem is only as resilient as its least-protected participant.



PANEL: Is Agentic AI an Identity Problem or a Solution?

KEY FRAMING POINTS

- Agentic AI is both a problem and a solution—simultaneously. The balance tips entirely on governance. The technology is neutral; the governance is not.
- Live coding enables AI agents to autonomously create other AI agents. This is not a future scenario. It is happening now, and it expands the attack surface well beyond what any single system's perimeter controls can address.
- The circle of problems is shifting, not shrinking. With agentic AI, existing threats will be retired as new threats emerge.
- A deepfake cannot walk into a bar—physical presence still anchors identity in the physical world. In the digital world, that anchor does not yet exist. An agentic AI can impersonate a person, manipulate a human agent on the other end of a call, and respond faster than any human can.
- The agent identity problem is broader than financial fraud. Three of four final candidates in one recent hiring process were fake AI-generated identities presenting on video. Trust boundaries are being breached across every domain where identity matters.
- Identity systems must now verify not just who someone is, but also whether the entity acting is genuinely authorized to take the specific action it is attempting—and whether the delegation chain from human to agent is intact and correctly scoped.

- Agents need identity, too. Emerging standards concepts—a 'birth certificate' for agents that establishes who built them, what they are authorized to do, and their maturity level—represent a governance framework the ecosystem does not yet have.
- Liability is unresolved and urgent. Accountability has not been legally defined. When an AI agent causes financial harm, who is responsible? The developer, deployer, or human principal? This ambiguity is being actively exploited.
- The equity gap is a systemic risk. Fraud actors exploit the weakest points in the ecosystem, and the entire ecosystem is only as resilient as its least-protected participant.

PANELIST OBSERVATION

“What keeps me up at night is making sure we can democratize technologies and resources across institution sizes. The resources JPMorgan Chase has are wonderful—but they're not helping the collective industry when a small community bank can't leverage the same data.”





PANEL: Future Fast Forward—Where do we go from here?

CORE THESIS

The gap between knowing and acting is the summit's defining challenge. The tools exist, the data is clear, and the practitioners in the room understand the threat. What the industry lacks is the prioritization to close the distance between what is possible and what is deployed — anchored by a universal standard: build systems safe enough for children, and they will be resilient enough for everyone.

The closing panel brought together the day's threads and asked a direct question: What must change for this to evolve? The moderator framed it as a question about the gap between knowing and acting, which had emerged as one of the summit's recurring themes. The question was what Summit participants would do differently now.

One of the opening contributions was a call for ecosystem mapping—a comprehensive, structured taxonomy of all the players in the identity and fraud ecosystem, where they overlap, and who is positioned to address which threat vectors. The AI agent landscape diagram in the keynote showed a model: the industry needs the same structured clarity across the full technology, business, and regulatory ecosystem. Without it, well-intentioned actors invest in adjacent solutions that leave gaps precisely at the seams.

The delegated and derived identity framework provided the session's conceptual anchor. A driver's license, a passport, a credit card are all derived identities—representations of an underlying human being. The ecosystem must be anchored to the human, not to the derived artifact. And the question of card-present versus card-not-present—a distinction designed for a pre-digital world—should be replaced by a more meaningful question: is a verified human being present and authorizing this transaction? That reframing, simple to state, requires substantial infrastructure to implement.

Independent testing emerged as an actionable recommendation. DHS and other bodies have developed independent testing standards for deepfake detection and biometrics—but procurement processes at financial institutions rarely reference them. Buyers must put independent testing standards into their RFPs. Vendors without third-party validation should not be considered comparable to those with it.

Consensus was clear on continuous verification as a non-negotiable foundation for identity solutions. One-and-done eKYC is insufficient. The passwordless identity vision must become an operational reality. As one panelist observed, the industry has understood the risks of passwords for years and yet remains largely dependent on them—a gap that encapsulates the distance between knowing and acting. The technology for passwordless, continuous, biometric-anchored identity exists. The adoption does not. Closing that gap is not a technology problem; it is a prioritization problem.

The panel concluded with a principle that The Prism Project's Resilient Trust™ framework has long advocated and that the summit's practitioners endorsed with conviction: if we build systems that are safe enough for children to operate in—for whom identity manipulation is potentially most harmful, and where the stakes of getting this wrong are highest—they will be resilient enough for everyone. Identity safety is not a specialized requirement for a subset of users. It is a universal standard that, when achieved, protects the entire ecosystem.





PANEL: Future Fast Forward—Where do we go from here?

KEY FRAMING POINTS

- The concept of card-present versus card-not-present is structurally obsolete. The question that actually matters is whether a verified human being is present—authorizing the transaction, not just possessing a credential number. Rebuilding identity infrastructure around that question is the work ahead.
- Derived identity is the right conceptual frame: a driver's license, a passport, a credit card are all derived identities—representations of an underlying human identity. The ecosystem must be anchored to that human, not to the derived artifact.
- Transaction data enrichment—tagging payments as human-initiated or agent-initiated, the vertical, and the transaction type—would give financial institutions a real-time signal to detect and respond to agentic fraud at the network level.
- Independent testing standards exist. DHS and others have developed them. Buyers should require third-party validated testing in their RFPs. Procurement that does not distinguish between validated and unvalidated solutions is not managing risk—it is deferring it.
- The ecosystem needs a comprehensive threat landscape map—covering all players, their overlaps, and who is positioned to address which threat vectors. Without that structured clarity, investment concentrates in well-understood areas, leaving gaps precisely at the seams.
- Regulatory mandate is the realistic mechanism for change at scale. Voluntary adoption of new identity standards does not move at the speed the threat requires. Industry-led frameworks serve as important precursors, but mandates are what drive universal adoption.

- Regulatory mandate is the realistic mechanism for change at scale. Voluntary adoption of new identity standards does not move at the speed the threat requires. Industry-led frameworks serve as important precursors, but mandates are what drive universal adoption. Continuous verification and biometric-anchored identity are the non-negotiable foundation. One-and-done eKYC is structurally inadequate. The technology for passwordless, continuously verified identity exists; the adoption does not. Closing that gap is a prioritization problem, not an invention problem.
- If we build systems safe enough for children to operate in confidently, they will be resilient enough for the entire ecosystem. Identity safety is a universal standard, not a specialized requirement.

PANELIST OBSERVATION

“If we build systems that are safe enough for kids to operate in confidently, they will be resilient and safe enough for all of us.”

”



KEY FINDINGS: Top Ten Takeaways

Practitioners, technologists, and policy experts reached a consistent conclusion: the fraud problem has entered the AI era while most institutional defenses remain anchored in obsolete threat models. AI-enabled impersonation is not an emerging risk to monitor — it is the dominant operating condition. What follows are the summit's ten most consequential insights.

Scale of Threat

Defense & Strategy

AI & the Road Ahead

1

Deepfake-Associated Fraud Has Already Reached Industrial Scale

2

Most Institutions Cannot See the Threat in Their Own Data

3

Synthetic Identities Are Already Inside—the Question Is How Many

4

One-and-Done eKYC Is a Structural Error, Not a Calibration Problem

5

Injection Attacks Are the Most Underdefended Vulnerability

6

The Coordination Gap Is Wider Than the Technology Gap

7

Agentic AI Has No Adequate Governance Framework—and Is Being Exploited Now

8

Fraud Resilience Is Distributed Unequally—and That Is Everyone's Problem

9

Fraud Signals Are Reactive. Authenticity Signals Are Proactive. Choose Offense.

10

Resilient Trust Is Trench Warfare—Continuous, Evolving, Never Finished



KEY FINDINGS: **Scale of Threat**

1

Deepfake-Associated Fraud Has Already Reached Industrial Scale

Data presented at the summit shows a consistent and accelerating upward trajectory: 20% of relevant fraud transactions in 2023 were associated with digital manipulation of identity; 40% in 2024; nearly 70% in 2025. The World Economic Forum reported an 8,000% increase in deepfake use in the two weeks prior to the summit. This is no longer an emerging threat to be monitored from a safe distance. It is the dominant fraud vector operating at scale in the institutions represented in the room and across the broader digital ecosystem. Organizations that are still treating deepfake fraud as a horizon risk are already behind.

2

Most Institutions Cannot See the Threat in Their Own Data

Measurement blindness—the systematic inability to identify deepfake fraud in existing transaction and identity data—is the primary barrier to investment and remediation. Even expert practitioners cannot reliably distinguish synthetic identities from real ones through visual inspection of the data their systems produce. This creates a self-reinforcing dynamic: organizations cannot justify the investment to solve a problem they cannot demonstrate exists, and the problem deepens while the justification is being assembled. Breaking this cycle requires dedicated measurement effort before any other remediation investment can be designed, sized, or defended.

3

Synthetic Identities Are Already Inside—the Question Is How Many

Conservative estimates suggest 15–20% of some institutions' portfolios may contain dormant synthetic identities. These identities are not flagged by standard fraud detection—they have clean transaction histories, strong credit profiles averaging 650–700, and no behavioral anomalies that current monitoring systems surface. They are invisible precisely because they were designed to be. The question facing enterprises is no longer how to prevent synthetic identity enrollment but how to identify the synthetic identities that have already passed verification and are currently building the profiles they will eventually exploit.



KEY FINDINGS: **Defense and Strategy**

4

One-and-Done eKYC Is a Structural Error, Not a Calibration Problem

The summit reached strong consensus that single-point identity verification at onboarding provides inadequate protection against synthetic identity fraud. This is not a failure of specific tools or processes or calibration problem—it is a structural inadequacy built into how the industry has conceived of identity verification. Continuous monitoring—behavioral biometrics, transaction pattern analysis, device signals, periodic re-verification for high-risk actions—must accompany all meaningful identity interactions throughout the full account lifecycle. The operative question is not ‘is this person who they say they are at enrollment?’ Is the entity acting on this account consistent with the verified identity over time?’

5

Injection Attacks Are the Most Underdefended Vulnerability

Injection attacks—inserting adversarial content directly into identity verification pipelines, bypassing cameras entirely—represent a significant and systematically under-addressed attack vector. The industry’s attention and countermeasure investment have concentrated on presentation attacks and deepfake video content; injection attacks have not received proportionate attention relative to the threat they pose. Combating injection requires a specific set of controls—secured mobile devices, jailbreak detection, camera-IMEI verification, motion analysis to detect rack-mounted static devices—that most institutions have not yet deployed. There is no single solution; these controls must operate in combination.

6

The Coordination Gap Is Wider Than the Technology Gap

The tools to fight deepfake fraud exist. Biometrics, behavioral analytics, verifiable credentials, device intelligence, signal orchestration—these are available, maturing technologies, and many of the leading vendors building them were represented in the room. What the industry lacks is the coordination infrastructure to deploy them effectively: cross-institutional intelligence sharing, standardized risk signal taxonomies, internal cross-functional alignment, and governance frameworks for AI deployment. Solving the technology problem without solving the coordination problem produces better-equipped silos. It does not produce resilience.



KEY FINDINGS: **AI and The Road Ahead**

7

Agentic AI Has No Adequate Governance Framework—and Is Being Exploited Now

More than 50% of online activity is already non-human. Agentic AI systems are already performing financial and other tasks on behalf of individuals. Yet no adequate governance framework exists for verifying that an agent is genuinely authorized by a real human, for scoping the actions an agent may take on a principal's behalf, or for establishing liability when agent actions cause harm. This gap is not theoretical; it is being actively exploited. The frameworks that responsible actors are debating while fraudsters exploit their absence must be developed and deployed with urgency—not treated as a long-term governance project.

8

Fraud Resilience Is Distributed Unequally—and That Is Everyone's Problem

Large financial institutions have meaningful resources to deploy advanced fraud detection, real-time signal analysis, and sophisticated behavioral biometrics. Community banks, credit unions, smaller fintechs, and emerging market participants do not. This equity gap is not simply a fairness concern; it is a structural systemic risk. Fraud actors exploit the weakest points in the ecosystem, and the weakest points are consistently the institutions with the least capacity to defend themselves. The entire ecosystem is only as resilient as its least-protected participant. Democratizing access to advanced fraud detection capabilities is therefore not an aspirational goal—it is a prerequisite for ecosystem-level resilience.

9

Fraud Signals Are Reactive. Authenticity Signals Are Proactive. Choose Offense.

The industry's default orientation is reactive: detect fraud after it occurs, flag anomalies after they appear, investigate losses after they are sustained. The next evolution of identity infrastructure is proactive: verify authenticity before action is taken. Cryptographically verifiable credentials—mobile driver's licenses, digital passports, verifiable credentials issued against government systems of record—are authenticity signals. They are deterministic rather than probabilistic, and deepfakes cannot yet spoof them. Shifting investment and organizational orientation from fraud detection to authenticity verification is the strategic reframe that the current threat environment demands.

10

Resilient Trust Is Trench Warfare—Continuous, Evolving, Never Finished

New detection and verification solutions stay ahead of adversarial capabilities for roughly four months before being broken or circumvented. There is no end state in which the problem is solved—only organizations that invest continuously and those that fall behind. Voluntary adoption of new identity standards has not moved at the speed the threat requires and will not; regulatory mandate is the only mechanism that has historically driven universal adoption across an ecosystem as distributed as financial services. Building identity systems that are safe enough for the most vulnerable users to operate in confidently is not a specialized goal—it is the right design standard for infrastructure that protects everyone.



RECOMMENDATIONS

These recommendations reflect the collective insights of summit practitioners, synthesized from session discussions, Q&A exchanges, and post-event analysis. They are addressed to enterprises, technology providers and identity industry players, and regulators and governance bodies. They are presented as the starting points the summit's practitioners identified as most urgent—not as the full scope of what is needed, as the challenges documented in this report are systemic and will require sustained effort across the ecosystem.

FOR ENTERPRISES



FOR TECHNOLOGY PROVIDERS AND THE IDENTITY INDUSTRY



FOR REGULATORS AND GOVERNANCE BODIES





RECOMMENDATIONS: Enterprises

FOR
ENTERPRISES



1

Replace One-and-Done eKYC with Continuous Identity Assurance

Move from single-point identity verification at onboarding to continuous monitoring across the full account lifecycle. Deploy behavioral biometrics, device signal analysis, transaction pattern monitoring, and periodic re-verification for high-risk actions. Design for the question that matters: is the entity acting on this account consistent with the verified identity over time—at enrollment, access authentication, transaction initiation, and account recovery?

2

Audit Existing Portfolios for Synthetic Identity Penetration

Given evidence that 15–20% of some portfolios may already contain synthetic identities, institutions should treat this as an urgent risk-management priority, not a future planning item. Retroactive behavioral analysis of existing accounts—looking for the specific pattern of credit-building without corresponding life-stage signals—is the starting point. Institutions that have never conducted this analysis should assume they have a problem until proven otherwise.

3

Harden Against Injection Attacks Specifically

Implement the technical controls required to defend against injection attacks. Shift all identity capture to secured mobile devices. Block jailbroken and rooted phones. Verify camera-IMEI correspondence. Implement device motion analysis to detect static, rack-mounted devices. Deploy liveness detection at PAD Level 3 where possible. These controls specifically address injection attack vectors that most current identity capture infrastructure leaves open.

4

Break Down Internal Silos Between Fraud, Identity, AML, and Cyber Teams

Address the internal coordination gap as a priority. Fraud teams, identity management, AML, cybersecurity, compliance, and audit must operate from shared data and aligned risk frameworks. This is an organizational design challenge, not a technology challenge. In the AI era, enterprise-wide risk management—not departmental optimization—is the required operating model for threats. Regulators are already moving in this direction; institutions that get ahead of it will be better positioned when the mandate arrives.

5

Quantify the Measurement Blindness Problem Before Investing in Detection

The primary barrier to effective deepfake defense is not the absence of detection technology—it is the inability to see the problem in existing data. Before investing in new detection technologies, establish a baseline: what proportion of your fraud data is potentially deepfake-associated? Work with detection specialists to apply current capabilities retroactively to existing transaction and identity data. The investment case for everything else follows from this measurement. Fix measurement first.



RECOMMENDATIONS: Technology Providers & the Identity Industry

FOR TECHNOLOGY PROVIDERS AND THE IDENTITY INDUSTRY



1

Shift the Narrative from Fraud Signals to Authenticity Signals

The industry's default orientation is reactive: detect fraud after it occurs. The next evolution of the identity ecosystem is proactive: verify authenticity before action is taken. Cryptographically verifiable credentials—mobile driver's licenses, digital identity wallets, verifiable credentials issued against government systems of record—are authenticity signals. They are deterministic, not probabilistic. The industry should invest in the infrastructure that makes these signals universally accessible, deployable, and interoperable.

2

Break Consortium Silos—Build True Ecosystem Intelligence

Intelligence-sharing consortia are necessary but not sufficient. Multiple competing consortia that do not share information reproduce the siloed architecture they were meant to replace at a slightly larger scale. The industry needs standardized risk signal taxonomies that enable cross-institutional intelligence sharing, with privacy-preserving protocols that protect individuals while allowing fraud signals to be shared across the ecosystem.

3

Democratize Fraud Detection Capabilities

Advanced fraud detection capabilities must be accessible to community banks, credit unions, and smaller organizations—not just institutions with the largest technology budgets. Technology providers should develop tiered access models, shared infrastructure, and API-based capabilities that bring enterprise-grade detection within reach of the full ecosystem. The equity gap in fraud resilience is a market failure with systemic consequences, and the identity industry has both the capability and the responsibility to address it.



RECOMMENDATIONS: Regulators & Governance Bodies

FOR REGULATORS AND GOVERNANCE BODIES



1

Establish Governance Frameworks for Agentic AI Deployment

AI agents are already performing financial transactions, managing accounts, and taking consequential actions on behalf of individuals without adequate governance frameworks for any of it. Governance frameworks must establish: standards for verifying human-to-agent authorization; requirements for audit logging and traceability; liability frameworks for agent actions; and minimum standards for agent identity verification. This gap is being actively exploited while frameworks are debated. Organizations waiting for these frameworks to arrive before planning for them are already behind.

2

Accelerate Digital Identity Infrastructure—including mDL Adoption

Government-issued digital credentials—mDLs and their equivalents—represent the strongest available proof of identity for online transactions. Regulatory bodies should accelerate adoption programs, work with financial institutions to establish acceptance infrastructure, and address the barriers preventing wider uptake. The TSA lists 21 states with mDL programs; adoption remains low relative to the opportunity. This must change, and regulators have a direct role in driving that change.

3

Mandate Cross-Functional Fraud Governance Standards

Regulators (OCC and equivalent) are already moving toward greater accountability for fraud program frameworks. This direction should be accelerated: require institutions to demonstrate cross-functional fraud governance, shared data environments across fraud and identity teams, and continuous monitoring programs—not just onboarding controls. Institutions that have already built these capabilities should be recognized and rewarded; institutions that have not should be given clear timelines and clear consequences.



THE CHALLENGE AHEAD

What emerged from the discussions in Houston was not simply heightened concern about deepfakes, synthetic identities, and the way AI accelerates and magnifies impersonation fraud. It was a shared recognition that this is no longer a theoretical risk but an operational challenge reshaping how organizations must think about digital identity and fraud resilience today.

While the threat landscape continues to evolve, The Deepfake Summit demonstrated that the industry is not starting from zero. In many respects, one of the most striking insights from the summit was that the technologies required to respond already exist. Biometrics, behavioral analytics, device intelligence, cryptographic credentials, and collaborative fraud-intelligence platforms are widely available and continue to mature.

The greater challenge lies in accelerating adoption and strengthening coordination within and across institutions and market sectors. Silos must be broken:

- Institutions must treat identity verification, authentication, and transaction monitoring as components of organization-wide, coordinated identity solutions.
- Industry sectors must improve collective intelligence sharing to thwart fraud rather than engage in whack-a-mole tactics.
- Governance frameworks must evolve to address the growing presence of automated agents and AI-enabled identity manipulation.

In this context, Resilient Trust™—the ability of digital systems to maintain secure, reliable, and privacy-centric identity assurance and transactional integrity in the face of adversarial innovation—will remain central to the future of digital identity.



“It was energizing to spend time with like-minded peers, reconnect with old friends, and have candid conversations about where this industry needs to go next.

Frances Zelazny • Principal • Identity Strategies



The work of building resilient, trusted digital ecosystems is only beginning. The inaugural Deepfake Summit demonstrated the value of bringing together fraud practitioners, digital identity technologists, payments leaders, and policy experts in an intimate setting designed for candid operational dialogue. The conversations that began in Houston are just the beginning.

The central question facing the ecosystem is no longer whether identity systems must evolve; it is how quickly organizations can deploy the architecture required to sustain Resilient Trust™ in an AI-driven world.

CORE THESIS



Identity is critical infrastructure — treat it accordingly

Identity infrastructure built on numbers, passwords, and tokens was never designed to verify individual human beings or protect PII. Privacy and security are two sides of the same coin. The Resilient Trust™ framework: organizations that demonstrably protect their users' identity occupy a differentiated, defensible competitive position.



SUMMIT CONSENSUS



We don't have a technology problem. We have a coordination problem.

The tools to fight deepfake fraud are available, maturing, and were represented in the room. Cross-institutional intelligence sharing, internal cross-functional alignment, standardized risk signal taxonomies, and AI governance frameworks are not. Solving coordination unlocks the full value of the technology that already exists.





THE ROAD TO THE NEXT DEEPFAKE SUMMIT

The inaugural Deepfake Summit exceeded the expectations of its participants and demonstrated a clear appetite for exactly this kind of practitioner-focused, intimate, substantive dialogue. Multiple attendees explicitly called for a second summit, and the conversations that began in Houston are continuing on LinkedIn, through direct collaboration, and in ongoing Prism Project research.

The summit also revealed how much work remains. The threat is accelerating. The tools to respond exist, but are not deployed at scale. Governance frameworks for agentic AI are absent from the institutions that need them most. The intelligence-sharing infrastructure is fragmented across consortia that do not share information with one another. And the equity gap between large institutions and the rest of the ecosystem represents an unresolved systemic vulnerability that fraudsters are already exploiting.

These are not problems that will be solved by the next generation of technology tools, however capable those tools may be. They are problems of institutional will, organizational design, and industry coordination—and they require the kind of frank, practitioner-driven dialogue that the Deepfake Summit was designed to enable.

For the next summit, the following were identified as priority topics:

- Progress report: Have institutions begun acting on the synthetic identity penetration problem?
- Agentic AI governance: What frameworks have emerged, and how are institutions deploying them?
- Authenticity signals in practice: What has mDL adoption looked like, and how are verifiable credentials being integrated into fraud prevention?
- Regulatory development: What guidance has emerged from the OCC and equivalent bodies?
- Cross-institutional intelligence: Have any genuine consortium-to-consortium sharing models emerged?

“The discussions were substantive, the format was genuinely interactive, and the commitment in the room made it clear this could be the start of something meaningful as we collectively confront the challenges ahead.”

Frances Zelazny • Principal • Identity Strategies








THE ROAD TO THE NEXT DEEPFAKE SUMMIT

Prepare for the next Deepfake Summit by reading The Prism Project reports



Making sense of the digital identity ecosystem so you don't have to.



Continue the conversation at



THE DEEPFAKE SUMMIT

September 1, 2026 | Washington, DC



Anchoring Trusted Identity:
Alignment. Coordination. Governance.

**REGISTRATION
IS OPEN**

**SPONSORSHIPS
AVAILABLE**

www.thedepfakesummit.com

Brought to you by

