



THE DEEPPFAKE SUMMIT

September 1, 2026 | Washington, DC

Building Trust Infrastructure for the AI Economy Event Intelligence Report

Insights. Analysis. Findings.

The Washington, DC Deepfake Summit, held on September 1, 2026, brought together fraud fighters, digital identity specialists, biometric technologists, AI security researchers, governance experts, and professionals from banking, payments, and the public sector for substantive, unfiltered dialogue.

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

www.thedeepfakesummit.com

Brought to you by The Prism Project
©2026 Acuity Market Intelligence



EVENT & REPORT SPONSORS

This Event Intelligence Report was made possible by the following Deepfake Summit sponsors:

SIGNATURE SPONSOR



PLATINUM SPONSORS



GOLD SPONSORS





IN THE ROOM WHERE IT HAPPENED

While The Deepfake Summit followed the Chatham House Rule, the following public quotes and commentary provide a direct window into the Summit experience, conversations, and insights—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 future trust gap is not primarily a technology problem. It’s an ownership, process, data, and people problem. One call to action may have been the most practical of all: call ahead, build relationships before you need them—whether it’s emergency management, payments, fraud, or life in general, trust is built before the crisis arrives. Innovation matters. Execution matters more. The adversaries see the entire lifecycle. We can’t continue defending in organizational silos.”

Andrea Valentin · Chief Risk Officer · Maxwell State Bank

“Had a great time and learned a great deal at The Deepfake Summit. Packed with great people in the industry and deep discussions on where the future is heading.”

Jim Pham · Fraud Operations Manager · Amivero

“As agentic AI becomes more prevalent: AI agents should be verified and tied to the individuals or entities that authorize them; identity and accountability should be built into AI-enabled activity, not added later; and as agent-to-agent interactions grow, actions should remain traceable to the people or organizations that initiated them. We have the tools; we need meaningful action. Let’s do the work!”

Amit Sharma · Managing Director, Digital Identity · Fragomen

“The deepfake conversation is no longer about whether AI can create convincing fake content. It’s about whether organizations can still trust the digital environment in which they operate. AI governance has to mature from governing models to governing outcomes, actions, and trust. The organizations that win with AI won’t simply be the ones that adopt it fastest—they’ll be the ones that build the trust infrastructure to scale it responsibly.”

Michaela Hatcher Blackwell, CISA, APRP, CDPSE, CIA · Senior Director, Information Technology Audit · Truist

“I really enjoyed the Summit and hope you continue to bring it to the industry. It was the perfect size, venue, and industry mix. It’s one of the most truly unique and productive identity conferences I’ve ever attended. I can’t tell you a time I’ve ever stayed at a conference to the end when I was literally getting bounced out of the club.”

Jeremy Haynes · Head of Sales · Trinsic

“A synthetic identity can reportedly be bought online for less than ten dollars. What surprised me most was hearing that the technology to fix this largely already exists—the holdup is organizational. A point that hit harder than expected: kids today are handing over academic records, family financial details, and biometric data through routine things like college applications, without ever agreeing to it. Someone posed a simple test for anyone building these systems: would you want this happening to your own kid’s data? Identity is becoming boardroom infrastructure, not a back-office IT problem.”


Dustin J. Eaton · VP, Fraud Intelligence & Financial Crimes · Vouched



IN THE ROOM WHERE IT HAPPENED

WHO ATTENDED THE DEEPPAKE SUMMIT?

By Sector



Other: AI, Big Tech, NGO, Privacy, HR, Healthcare

By Title






TABLE OF CONTENTS

- 1 INTRODUCTION
- 2 EXECUTIVE SUMMARY
- 4 SUMMIT CONTEXT & FORMAT
- 5 KEY THEMES
- 14 SUMMIT SESSION SUMMARIES
- 36 KEY FINDINGS
- 40 RECOMMENDATIONS
- 44 THE CHALLENGE AHEAD
- 45 THE ROAD AHEAD

The technologies undermining trust are evolving faster than the institutions responsible for PROTECTING it.



AI-driven impersonation is reshaping how organizations establish identity, verify authenticity, govern AI, and preserve trust.

The Deepfake Summit is where leaders from government, industry, and technology come together to build the TRUST INFRASTRUCTURE required for the AI economy.



INTRODUCTION

Welcome to the September 2026 Deepfake Summit Event Intelligence Report. This report captures and synthesizes session discussions and commentary from the Summit that took place on September 1, 2026, in Washington, DC.

The Deepfake Summit brings together senior executives, policymakers, researchers, and innovators at the forefront of AI, digital identity, cybersecurity, fraud prevention, compliance, privacy, and digital trust for a day of discussion, debate, and collaboration designed to strengthen identity assurance, counter AI-enabled fraud, and build resilient trust infrastructure for the AI economy.

The Summit is highly interactive and unscripted by design, creating dynamic opportunities for speakers and participants alike to offer bold ideas, challenge assumptions, exchange practical experience, and build the cross-sector and cross-silo relationships that will drive the future of trust.

The inaugural Summit in Houston in March 2026 established a clear diagnosis: the impersonation fraud problem has fully entered the AI era, while most institutional defenses remain calibrated for yesterday's threat landscape. The DC Summit asked the follow-on question: if the threat is here and the tools exist, what will it actually take to build the trust infrastructure that the AI economy requires—and what is standing in the way?

Over the course of the day, four realities came into focus:

- Deepfakes are no longer just a fraud instrument—they are a weapon aimed at human judgment itself.
- The attack surface has expanded from onboarding to the entire identity lifecycle.
- Agentic AI is transacting at scale while its governance is still being debated.
- The gap between the speed of AI innovation and the speed of enterprise transformation—the Trust Velocity Gap—is the defining strategic challenge for every public- and private-sector enterprise doing business online and in person.

The discussions were candid and the debates intense but respectful. The connections made and relationships fostered were equally important to the challenge at hand: how to collaborate across disparate sectors and areas of expertise to coalesce around a shared vision, shared values, and practical priorities for a future built on trust infrastructure.

To comply with the Chatham House Rule, this document captures and reports content without revealing the identity or affiliation of any speaker or participant.

My deepest gratitude to our sponsors, speakers, and engaged attendees who made the Summit a powerful experience. I hope this report provides valuable insight into what transpired and encourages you to join us for the next Deepfake Summit in 2027.



Authentically yours,

C. Maxine Most
Founder & Chair, The Deepfake Summit
Founder & Lead Analyst, The Prism Project
Founding Principal, Acuity Market Intelligence

The Deepfake Summit is an initiative of Acuity Market Intelligence (www.acuitymi.com) and The Prism Project (www.the-prism-project.com).





EXECUTIVE SUMMARY

The second Deepfake Summit convened a curated group of practitioners, technologists, intelligence and security veterans, and policy experts in Washington, DC, on September 1, 2026. Six months after the inaugural Summit in Houston established that AI-driven impersonation fraud, synthetic identity, injection attacks, and agentic automation had become operational realities, the DC Summit pushed past diagnosis to a harder question: What will it actually take to build trust infrastructure capable of keeping pace with AI-era threats?


The DC Summit’s answer began with an architectural reframe. Digital systems were built for connectivity and access, with security layered on top to control that access. AI is now attacking the assumption beneath that model: possessing the right credential, password, token, device, face, or voice no longer proves that the identity, transaction, content, or interaction behind it should be trusted. The required progression is therefore ACCESS → IDENTITY → TRUST. Identity is foundational to trust, but the two are not the same: most of the industry’s tools remain calibrated to identity, while the threat has moved to trust.

The opening keynote extended that argument from systems to human judgment. Drawing on the human-intelligence recruitment cycle—spot, assess, develop, recruit—it revealed how deepfakes and synthetic personas can compress months of trust-building into minutes. The warning captured one of the day’s defining insights: “Deepfakes are the disguise. Your judgment is the target.” Technical defenses remain essential, but human decision-making itself has become part of the attack surface.

The sessions that followed translated that insight into institutional requirements:

- Identity assurance must persist across the entire lifecycle, particularly recovery and re-entry.
- Autonomous AI agents require governed delegation, not merely authentication. This includes an identifiable agent, a verified principal, a defined mandate, bounded authority, continuous auditability, and human escalation for consequential actions.
- Data and organizational silos must be bridged so defenders can see the same lifecycle that adversaries exploit.
- Governance must move beyond static policy toward executable decision logic.
- Privacy must be designed into trust architecture rather than treated as a competing priority.
- Testing must reflect real-world performance, integration effects, adversarial conditions, and long-tail failures—not simply vendor claims or compliance status.

AI FRAUD: FROM OPERATIONAL REALITY TO BUILDING TRUST INFRASTRUCTURE



Sources & Notes: ^{1,2} Deepfake Summit DC Fireside Conversation: Overcoming Obsolete Threat Models, citing 2025 estimates. ³ Deepfake Summit DC Dialogue, participant citing an onboarding-fraud study. ⁴ Participant estimates, Deepfake Summit DC. ⁵ Surfshark Research, July 2026. ⁶ Deepfake Summit DC Leadership Keynote.





EXECUTIVE SUMMARY (Continued)

Across these discussions, one conclusion emerged repeatedly: the largest gap is increasingly institutional, not technological. Biometrics, behavioral analytics, cryptographic and verifiable credentials, deepfake detection, privacy-enhancing technologies, and orchestration capabilities are advancing rapidly. Procurement, integration, governance, testing, budgeting, and organizational change are not. This Trust Velocity Gap—innovation at AI speed, transformation at enterprise speed—is where risk accumulates.

Privacy emerged as particularly important. Rather than assuming stronger identity requires ever-larger centralized data stores, participants pointed toward minimization, selective disclosure, reusable credentials, consent-aware access, and privacy-enhancing technologies. The objective is not less identity assurance, but stronger assurance with less unnecessary exposure.

Over the course of the day, the threat landscape was defined by three converging forces:

- Agents that execute transactions and take consequential actions autonomously—in one documented case, carrying out a multi-stage deception with only limited human intervention.
- The identity lifecycle itself—onboarding, recovery, offboarding, re-entry—becoming the primary attack surface, not just the front-door credential check.
- Deepfakes targeting human judgment directly, compressing the human-intelligence recruitment cycle from months to a single video call.

The closing Leadership Forum tied these strands together into an ecosystem model spanning identity, continuous assurance, provenance, delegated authorization, privacy, testing, and governance—underpinned by relationships among government, industry, and practitioners that must exist before a crisis, not be improvised during one.

Beneath all of this sat a point nearly every speaker made in different words: **you can verify exactly who someone is—their credential, device, and organization—and still have no defensible basis for trusting the specific transaction, behavior, or AI agent now acting in that person's name.**

THE SUMMIT PRODUCED A CLEAR STRATEGIC AGENDA

The tools exist. The standards are maturing. What is missing is enterprise velocity—and the governance to direct it.



- Shift the security model from access and identity to TRUST.
- Move from one-time verification to continuous lifecycle monitoring.
- Govern AI agents with explicit mandates before they act.
- Prove performance with honest, independent testing.
- Restructure enterprises to move at the speed the threat demands.

EXECUTIVE CONCLUSION

Trust infrastructure is not a product category; it is an operating capability that continuously establishes, challenges, and re-establishes confidence in identity, authority, authenticity, and interaction integrity as conditions change.



SUMMIT CONTEXT & FORMAT

The Deepfake Summit is a practitioner-focused forum designed to facilitate substantive dialogue among the communities responsible for protecting digital identity, financial systems, public-sector programs, and critical infrastructure.

Produced and hosted by Maxine Most of [Acuity Market Intelligence](#) and [The Prism Project](#), the Summit builds on decades of research and analysis examining biometric-centric digital identity, fraud prevention, and the evolving role of identity and trust infrastructure in the digital economy.

The DC Summit retained the intimate, highly interactive format established in the inaugural Houston event in March 2026, while expanding both the scope and duration of the conversation. The full-day program brought together leaders from government, financial services, technology, digital identity, fraud prevention, cybersecurity, risk, privacy, and standards—adding intelligence-community tradecraft, government benefits delivery, federal testing programs, and critical-infrastructure perspectives to the banking, payments, healthcare, and identity expertise represented in Houston.

The agenda extended the discussion, analysis, and findings from the inaugural event, and, like the Houston agenda, was deliberately cumulative rather than a sequence of disconnected presentations. Each Keynote, Fireside Conversation, and Summit Dialogue built on the one before it, reflecting the premise that AI-driven deception crosses identity, fraud, cybersecurity, privacy, governance, payments, data, compliance, and organizational boundaries. The Chatham House Rule supported candid operational exchange, while the keynotes and the chair's opening and closing framing provided a shared strategic architecture.


While the Houston Summit (see the [Houston Event Intelligence Report](#) for details) focused primarily on diagnosing the emerging threat environment, the DC Summit moved toward the institutional response.

The DC Summit progressed from the changing trust landscape and obsolete threat models through lifecycle assurance, agentic AI, governance, identity transformation, privacy and data control, and ultimately the collaborative infrastructure required to build and sustain trust.

A Different Kind of Summit

Rather than a traditional conference built around a series of disconnected presentations and one-way communication, The Deepfake Summit is a daylong EXECUTIVE CONVERSATION.

Each Keynote, Fireside Conversation, and Summit Dialogue builds on the discussions that came before it. The interactive format includes both speakers and participants in candid discussion, informed debate, and practical collaboration. Everyone contributes, challenges assumptions, advances ideas, strengthens relationships across sectors and silos, and helps shape the future of trust.






KEY THEMES

Eight interconnected themes ran through the sessions at The Deepfake Summit.

Taken together, these themes describe the strategic challenges every organization faces in distinguishing real humans from synthetic ones.
Today, that means every organization doing business online and, increasingly, in person.





THEME 1: Authenticity Is a Decision Problem—Human Judgment Is Part of Trust Infrastructure

Deepfakes are most dangerous not because they can fool a detector, but because they can borrow authentic relationships and manipulate consequential decisions. The classic intelligence recruitment cycle—spot, assess, develop, recruit—has not disappeared; AI compresses it from months into minutes and increasingly automates parts of the operation.

That distinction makes authenticity a decision problem, not only a media-detection problem. Human sight and sound are unreliable, and even perfect detection of one modality cannot establish whether the request, authority, context, or purpose should be trusted.

Trained human verification behavior must be part of trust infrastructure. Because attackers deliberately exploit urgency, secrecy, authority, fear, and flattery, individuals must learn to pause, question what they see, verify information, and confirm high-stakes requests through an independent channel.

KEY POINTS

- Deepfakes borrow existing trust relationships rather than build new ones from scratch.
- AI agents automate targeting, impersonation, and follow-through, compressing attacks from months to minutes.
- Neither human senses nor media detection alone can establish whether a request should be trusted.
- Consequential decisions should require corroboration through a channel controlled by neither the claimant nor the suspicious interaction.
- Urgency, secrecy, authority, fear, and flattery are attack signals that should trigger escalation.
- Verification is a practiced skill: employees need rehearsal, not just policy, to slow the tempo under pressure.



Key Challenge

Most security programs are built to validate credentials, media, or access paths. Adversaries increasingly target the decision-maker's perception, urgency, and confidence—human factors that are rarely trained or measured with the same discipline as technology.

Critical Insight

Human judgment is part of trust infrastructure. Technology can detect, authenticate, and verify, but consequential decisions still require trained discipline: pause, corroborate through an independent channel, and treat urgency as a warning signal.



SUMMIT OBSERVATION

"Deepfakes are the disguise. Your judgment is the target."



THEME 2: Identity Is Foundational—But Assurance Must Persist Across the Full Lifecycle

The Summit reinforced a widely accepted view that one-time identity verification is no longer adequate; identity must be verified across its full lifecycle. The attack surface includes onboarding, credential binding, authentication, transactions, device and privilege changes, leave and return, recovery, and offboarding. Strong onboarding cannot compensate for weaknesses at any other vulnerable point in the identity lifecycle.

Account recovery is an especially high-value target because many organizations silently downgrade assurance when a device is lost, a help desk is called, or an exception is triggered. The MGM breach¹ illustrated the point: attackers bypassed technical controls through a human reset process.

The operational implication is twofold: return to foundational identity when confidence falls, and combine biometrics, cryptography, device, behavior, context, transaction, provenance, and human review so assurance persists rather than restarting at each touchpoint.

KEY POINTS

- Onboarding and recovery should be treated as high-risk trust events, not administrative workflows.
- Identity must persist; each touchpoint should not start fresh or fall back to unrelated weak factors.
- Recovery should restore assurance by returning to foundational identity rather than bypassing it.
- Multi-level signals—biometrics, device, behavior, context, transaction, provenance—should work together.
- Human review remains necessary for red flags and high-consequence exceptions.
- Continuous monitoring should evaluate whether activity remains consistent with the verified identity over time.

Key Challenge

Many organizations have strong controls at enrollment but weaker evidence requirements at recovery, re-entry, help-desk reset, and other exception points. Assurance is effectively discarded exactly where adversaries concentrate.

Critical Insight

Identity is not a login, and authentication is not the end of the identity process. Resilient identity is a closed lifecycle that can return to the foundational human whenever the system's confidence is challenged.



SUMMIT OBSERVATION

“The adversary sees the whole picture. Defenders only see their silo.”

Sources & Notes: ¹ Bloomberg, “MGM Resorts Hackers Broke in After Tricking IT Service Desk,” Sept. 16, 2023.



THEME 3: Agentic AI Collapses the Traditional Accountability Chain


Agentic AI breaks the assumption that one person, one account, and one action form a clear accountability chain. Agents can invoke tools, communicate with other agents, execute transactions, and repeat mistakes at machine speed. AI performed most of the operational work in one cited espionage campaign, demonstrating that this is already more than a theoretical risk.

For AI agents, both identity and authority must be established. Knowing which agent is present is necessary but insufficient; organizations must know who or what authorized it, what objective it is pursuing, what data and systems it may access, what actions it may take, for how long, and under what escalation rules.

The emerging model is governed delegation: bind the agent to an accountable principal, make the mandate explicit and time-bounded, apply role and transaction limits, continuously re-verify and log behavior, shadow-test agents before expanding authority, and retain human control over high-consequence outcomes.

KEY POINTS

- Agents should have their own identifiable operating identity while remaining cryptographically or operationally bound to an accountable principal.
- Agent mandates should be explicit, scoped, time-bounded, and task-specific. Permission structures should prevent an agent from taking an action a corresponding human would not be authorized to take.
- Shadow testing should compare agent behavior with human decisioning before production authority expands.
- High-consequence actions require human verification or escalation.
- Logs must preserve the chain from principal to delegation to agent action to outcome.
- Agent-to-agent interactions and agent marketplaces create new systemic attack surfaces that current IAM models do not fully address.



Key Challenge

Most current authorization systems were designed around human users, service accounts, or static APIs. They do not express dynamic delegated mandates, changing context, agent-to-agent action, or liability for autonomous behavior.

Critical Insight

The agentic enterprise needs a new trust layer: identity + delegation + authorization + attestation + audit. Authentication answers who or what is present; governed delegation determines what that actor is allowed to do now.



SUMMIT OBSERVATION

“There’s always an identity underneath ...—even if it’s an agent that’s been hijacked, or one that’s been built to commit fraud. The best-built systems need a permission structure that doesn’t allow the actions that shouldn’t be enabled, and continuous re-verification of who is actually taking that action.”



THEME 4: The Coordination Gap Is Wider Than the Technology Gap

The Summit returned repeatedly to the idea that what most often holds back identity defense is coordination, ownership, and culture, not missing technology. Fraud, cybersecurity, IAM, KYC, privacy, compliance, risk, and business teams buy differently, use different taxonomies, and see only part of the identity lifecycle, while adversaries see all of it.

The answer is not indiscriminate centralization. Shared taxonomies, authoritative sources, identity or knowledge graphs, and cross-functional scorecards can create shared context while preserving least privilege, segmentation, selective disclosure, and privacy boundaries.

The unresolved tension: some silos should be broken down because they hide common risk; others should be strengthened because they contain breaches and protect privacy. The goal is not fewer boundaries at any cost, but interoperable trust signals, clear accountability, and privacy protections across the boundaries that remain.

KEY POINTS

- Culture and governance are larger barriers than technology availability.
- Identity strategy is fragmented across cyber, fraud, IAM, KYC, privacy, compliance, risk, and business ownership.
- Board attention follows incentives: liability, litigation, insurance, regulation, resilience, customer loss, and economic impact.
- Not every silo should fall: break down those that hide common risk, strengthen those that contain breaches and protect privacy, and build interoperable trust signals and clear accountability across the boundaries that remain.
- Fraud and identity telemetry should be normalized across teams using shared taxonomies, authoritative sources, identity graphs, and cross-functional scorecards rather than raw-data pooling.



Key Challenge

Better technology without better coordination gives each silo better tools but does not make the institution more resilient. Participants argued both for breaking down silos and for strengthening them, but agreed the current setup is not working and that data storage, sharing, ownership, and control must be rethought.

Critical Insight

Shared context matters more than raw data volume. Trust signals must be interoperable, attributable, and governed across the lifecycle, so teams can act on them while protecting data and preserving privacy.



SUMMIT OBSERVATION

"It is a people problem, not a technology problem."



THEME 5: Governance Must Shift From Static Control to Executable Trust

Governance was reframed as decision infrastructure. When organizations must continuously decide whether to trust a person, credential, agent, transaction, or interaction, policies must be operationalized and enforced by the systems that make those decisions, not just written down in policy documents.

Governance spans corporate decision rights, technical trust frameworks, model governance, regulation, standards, and culture. The practical requirement is to connect those layers so consent rules, retention requirements, agent mandates, transaction limits, and escalation thresholds become enforceable controls and measurable behavior.

The deeper challenge is organizational literacy: executives cannot govern risks they do not understand, and fragmented ownership produces policies without execution. Governance becomes trust infrastructure when it specifies who or what may act, under which evidence and conditions, with what accountability and recourse.

KEY POINTS


- Define governance as the mechanisms that connect evidence, authority, policy, accountability, and recourse.
- Align governance with customer trust, operational performance, resilience, and growth—not only loss prevention.
- Create common baselines that can be mapped into sector-specific and jurisdiction-specific control frameworks.
- Use cryptographic provenance and trusted digital credentials to move governance toward positive authenticity signals.
- Ensure executives understand identity, authentication, fraud, and agent risks deeply enough to govern them.
- Recognize that public- and private-sector friction tolerances differ; governance must reflect mission and access obligations.
- Do not wait for a complete regulatory framework before implementing practical trust controls.

Key Challenge

Governance is often distributed among teams that optimize for different outcomes and use different risk languages. Policies exist, but ownership and execution are fragmented.

Critical Insight

Governance becomes trust infrastructure when it is operationalized and moves from documentation to decisioning: who or what may act, under which conditions, using which evidence, with what accountability, and how that decision changes when risk changes.



SUMMIT OBSERVATIONS

“You cannot provide meaningful governance if the organization does not understand the subject deeply enough to govern it.”

“AI governance has to mature from governing models to governing outcomes, actions, and trust.”



THEME 6: The Trust Velocity Gap Is the Defining Institutional Risk

The Trust Velocity Gap describes a two-speed market: threats and defensive capabilities evolve in days or weeks, while budgeting, procurement, legal review, governance, integration, testing, and change management can take months or years. The technology gap is narrowing; the adoption gap is not.

Legacy structures and practices are the bottleneck: identity ownership is fragmented, vendor onboarding can take 18–36 months, and testing and governance still run on older cycles, even as AI compresses product development from months to days. The result is a growing gap between the technologies and solutions available to deploy and what enterprises can actually integrate. That gap is where the risk concentrates.

The objective is not speed for its own sake, but faster, safe integration: standardized evidence, independent testing, modular architectures, repeatable procurement pathways, continuous logging, and governance that can evolve without abandoning assurance.

KEY POINTS

- Threat evolution and defensive innovation occur at AI speed; institutional adoption and integration occur at enterprise speed.
- Procurement, integration, governance, and testing are now part of the security perimeter.
- Operational testing must scale with AI-based product development, not trail it by multiple release cycles.
- Standardized evidence and independent testing can reduce repetitive buyer diligence and accelerate safer adoption.
- Modular, interoperable architectures reduce the cost of replacing controls as threats evolve.
- Continuous logging and post-deployment evaluation are necessary because model and environment behavior changes over time.

Key Challenge

Institutions cannot adopt needed innovations at the necessary speed using procurement, governance, and testing processes designed for slower technology cycles. But bypassing those processes creates a different form of risk and is not an option in highly regulated industries.

Critical Insight

The solution is not to choose between speed and assurance. It is to create a flexible translation layer—standards, testing, procurement, integration, and governance—so institutions can adopt and integrate trustworthy innovation faster.



SUMMIT OBSERVATION

“The technology required to build resilient, trusted digital identity ecosystems already exists. The gap is no longer invention. The gap is institutional adoption.”



THEME 7: Privacy and Security Must Be Co-Designed Around Data Minimization and User Control

The assumption that privacy and fraud prevention are competing objectives was rejected. Strong identity does not require repeatedly exposing or indefinitely centralizing the underlying data; privacy-enabled, consent-driven credentials can preserve assurance while reducing attack surface.

The discussion started from an uncomfortable baseline: organizations did not fully control or even understand their identity data before AI, and AI makes that accumulated data easier to discover, combine, infer from, and repurpose. Hybrid architectures—trusted systems of record plus wallets, verifiable credentials, selective disclosure, purpose limitation, privacy-enhancing computation, and consent-aware authorization—offer a more resilient alternative.

The goal is not decentralization as ideology. It is to minimize unnecessary exposure while preserving trusted evidence for recovery, accountability, and high-assurance decisions. Privacy is a fraud-resilience measure when systems prove more while disclosing less.

KEY POINTS

- Collect less data where the same assurance can be achieved through proofs or selective disclosure.
- Separate proof of an attribute from disclosure of underlying data whenever possible.
- Make purpose-of-use explicit by binding consent to specific data, tasks, and relying parties.
- Use privacy-enhancing technologies and disaggregation to limit centralized exposure.
- Treat retention, deletion, and third-party data flows as part of the identity threat model.
- Use trusted credentials and wallets to reduce repeated exposure of the same identity data across organizations.
- Position demonstrable privacy protection as a customer trust and competitive differentiator, not only a compliance requirement.



Key Challenge

Foundational identity needs durable evidence, but centralizing personal data creates a persistent high-value target. AI makes both the value of the data and the consequences of uncontrolled reuse larger.

Critical Insight

The privacy paradox is an architecture problem. Strong identity and strong privacy can reinforce each other when systems minimize disclosure, constrain purpose, preserve provenance, and return control to the verified individual where practical.



SUMMIT OBSERVATIONS

“Privacy and security should be two sides of the same coin, not a forced choice.”

“I should be able to be asked who I am by multiple verifiers and not have to re-expose identity constantly in an open web architecture.”



THEME 8: Trust Infrastructure Is an Ecosystem Capability—Technology and Social Capital Must Work Together

Trust infrastructure must be broadened beyond enterprise architecture. Fraudsters collaborate across technologies, institutions, sectors, and borders; defenders rarely do. Standards bodies, government programs, vendors, consortia, and testing and certification organization relationships each provide part of the response, but none is sufficient alone.

Trust is therefore relational as well as technical. In high-consequence environments, response depends on organizations that already know how to share signals, interpret evidence, escalate quickly, and believe one another before a crisis begins. Few participants had cross-sector emergency relationships in place, which made the gap tangible.

Transparent testing and shared market language are equally important. Buyers need comparable evidence rather than cherry-picked claims, and identity works as a system of complementary technologies—biometrics, cryptography, liveness, secure hardware, provenance, authorization, and risk signals—rather than a single product.

KEY POINTS

- Build cross-sector relationships and response mechanisms before crisis, not during it.
- Develop shared definitions and signal taxonomies so intelligence can move across institutions.
- Expand objective, operational testing and publish comparable evidence buyers can use.
- Use common standards as a baseline while allowing sector-specific control mappings.
- Treat fraud intelligence as pre-competitive infrastructure where legally and operationally possible.
- Engage government, standards and certification bodies, technology providers, academia, and regulated industries in the same trust architecture.
- Invest in market education so non-specialist decision-makers understand the controls, trade-offs, and limitations.



Key Challenge

The organizations that collectively possess the necessary signals, technology, authority, and expertise do not share a single operating model, vocabulary, incentive structure, or response mechanism.

Critical Insight

Trust infrastructure is socio-technical infrastructure. Identity, cryptography, analytics, testing, and governance provide the mechanisms; standards, relationships, shared language, and social capital make those mechanisms work across an ecosystem.



SUMMIT OBSERVATIONS

“No single layer proves a claim true. Together, the layers strengthen the decision.”

“Innovation matters. Execution matters more.”




SUMMIT SESSION SUMMARIES

Nine sessions, framed by the chair's opening Welcome and closing Reflections, made up the day's agenda. The program was deliberately sequenced to move from the evolving nature of the threat and the limitations of existing threat models to lifecycle assurance, the agentic enterprise, governance, privacy, and the broader institutional transformation required for the AI economy.

Together, the sessions provide a roadmap for the Summit's exploration of forged identity, persistent assurance, delegated agency, and the infrastructure of trust. Each summary provides a narrative of the conversation then distills the Core Thesis and Key Framing Points from each discussion. In accordance with the Chatham House Rule, session discussions are summarized without attributing comments to individual participants.

AGENDA ROADMAP



CONTENT AGENDA

8:45 AM — Welcome and Framing What is The Deepfake Summit? Why are we here?

9:00 AM — Opening Keynote Forged—Deepfakes and the Oldest Trick in Espionage

9:30 AM — Fireside Conversation Overcoming Obsolete Threat Models: Can We REALLY Build Trusted and Resilient Identity Ecosystems?

10:00 AM — Deepfake Summit Dialogue Proactive Fraud Detection Throughout the Identity Lifecycle: Breaking Down Organizational Silos to Achieve Enterprise-Wide Identity Assurance

11:30 AM — Deepfake Summit Dialogue The Agentic Enterprise—Building Trust Among Humans, Machines & AI Agents: Leveraging Rapidly Evolving Innovation to Build Trusted Identity-Centric Systems

12:15 PM — Deepfake Summit Dialogue Turning Controls on Their Head: Reimagining Governance as a Driver of Trust, Security, and Growth

2:30 PM — Leadership Keynote Moving at the Speed of Trust: Accelerating Identity Modernization Before Crisis—or Regulation—Demands It

3:00 PM — Fireside Conversation Identity-First Transformation: Cautionary Tales From Real-World Deployments

3:30 PM — Deepfake Summit Dialogue Privacy, Data Ownership & User Control: Can Privacy Become a Competitive Advantage?

5:00 PM — Leadership Forum Building the Trust Infrastructure for the AI Era

5:45 PM — Closing Reflections Where Do We Go From Here?



WELCOME & FRAMING

CENTRAL SUMMIT QUESTION: *How do we preserve trust when authenticity itself can no longer be taken for granted?*

The Welcome set the stage and expectations for the day, and the Framing provided context for the market's current state of play. AI-driven impersonation was framed as an architectural problem, not simply a new fraud category. Forty years ago, computer systems were designed to resolve two critical problems: connectivity and information accessibility. Networking and relational databases have been remarkably successful in addressing these challenges, creating a globally interconnected digital ecosystem now straining under its own success.

AI has exposed this ecosystem's limitations. Layered access-control security is fundamentally incompatible with, and therefore ineffective in, a system designed to make information easy to network, search, and share. The credentials, documents, biometrics, and now agents introduced to facilitate security can all be used to carry out fraudulent interactions. This shift was compressed into a simple progression: CONNECTIVITY → ACCESSIBILITY → ACCESS CONTROL → AI ATTACKS AUTHENTICITY.


The Trust Velocity Gap model was introduced to illustrate how threats and defensive technologies move at AI speed, while procurement, governance, testing, integration, policy, and organizational change move at enterprise speed. Closing that gap requires faster adoption, cross-functional ownership of trust by enterprises, and an interface capability that bridges the speed gap.

The Prism Project's Identity Hierarchy model was leveraged to show how credentials and context link to real humans through trusted foundational identity. Then the latest model from The Prism Project's [2026 Financial Services Prism Report](#)—the Resilient Trust™ Pyramid—was referenced to demonstrate that even decentralized systems that continuously assess trust through device, behavior, environment, content, transaction, and agent signals must still be grounded in foundational identity.

Finally, the question was raised whether 20th-century hierarchical enterprise structures are capable of the level of responsiveness required to address fraud evolving at AI speed.

THE TRUST VELOCITY GAP


Innovation at AI speed. Transformation at enterprise speed.



THE IDENTITY HIERARCHY

Identity is layered and built on a trusted foundation

Source: The Prism Project






WELCOME & FRAMING (Continued)

CORE THESIS: Modern digital infrastructure was built for connectivity, information accessibility, and access control. AI now attacks authenticity, exposing the limits of credential-centric security. The required shift is from access to identity to trust—and ultimately to trust infrastructure capable of evolving at the speed of AI.

KEY FRAMING POINTS


- **AI did not create the underlying weakness in digital infrastructure;** it accelerates the exploitation of systems designed for connectivity, accessibility, and access.
- **Access control addresses permission.** It does not establish that the identity, content, transaction, or interaction is authentic.
- **The security paradigm must evolve from ACCESS → IDENTITY → TRUST,** with trust treated as a dynamic condition rather than a one-time decision.
- **Identity is foundational to trust, but it is not synonymous with trust.** Device, behavior, environment, content, transaction context, and delegated agent activity all influence whether trust should persist.
- **Identity is layered** and must ultimately resolve to a real human and trusted foundational identity. Authenticating a false identity more effectively does not make it real.
- **The Trust Velocity Gap**—innovation at AI speed versus transformation at enterprise speed—is where institutional risk accumulates.
- **Closing the Trust Velocity Gap requires a two-pronged approach:**
 - Faster procurement, testing, governance, integration, standards translation, and cross-functional decision-making, not simply another fraud or authentication tool.
 - An interface layer that bridges the pace of technological change required to respond to threats evolving at AI speed with enterprises' ability to integrate technology advancements.
- **Trust infrastructure** is therefore both a technology architecture and an operating model spanning identity, fraud, cybersecurity, privacy, data, governance, risk, policy, and organizational design.



THE RESILIENT TRUST PYRAMID™

The third dimension: foundational identity

Source: The Prism Project, 2026 Biometric Digital Identity Financial Services Prism Report





KEYNOTE: Forged—Deepfakes and the Oldest Trick in Espionage

STRATEGIC QUESTION: *Why has trust become the defining strategic challenge of the AI era?*

Human-intelligence tradecraft provided the keynote's lens for AI-driven deception. The keynote opened with a video of the speaker and later returned to it with a question: how would the room know it was real, and not a synthetic file built from years of public footprint?

Drawing on a career spent countering the world's most capable intelligence services, the speaker argued that human intelligence is a business of people, not lies. An officer finds someone with access, learns their motivations, and earns trust over months or years: a cycle called spot, assess, develop, recruit. AI has not changed the cycle's purpose, only its speed and economics. Criminal operations are now using technology to exploit and amplify the cycle.

The speaker then ran the cycle live on the audience under the code name Operation Unmask, suggesting that anyone who had told a friendly stranger at breakfast where they worked had already been spotted. Deepfakes have changed development the most: months of trust-building can now happen in a single call.

This mechanism was illustrated with a recent case: a finance director joined a video call with convincing synthetic versions of their CEO and senior colleagues, then wired \$499,000. The attackers never breached the network; they borrowed years of trust embedded in familiar faces and voices. A human still directed that operation, but that constraint is changing.

Deepfakes, the speaker noted, make disguise instant, cheap, and scalable; agentic AI increasingly changes who conducts the operation. The keynote pointed to a state-linked campaign in which AI carried out 80–90% of the operational work, with humans stepping in at only a few decision points. The reveal that the opening video was a deepfake made the point: appearance and voice can no longer establish identity.

The strategic consequence, the speaker argued, is broader than fraud loss. Synthetic media creates a "liar's dividend," allowing genuine evidence to be dismissed as fake and making doubt itself a weapon. Detection technology alone cannot carry the burden. The operational target is judgment. Urgency, authority, familiarity, secrecy, fear, and flattery are the levers used to force action before trust is independently established.

A deception only has to survive until the funds land, the speaker observed, but the doubt it leaves behind lasts far longer. What is at stake is cognitive liberty, the ability to direct attention, form judgments, and act on one's own values.

The countermeasure was framed as Mind Sovereignty™: reclaim attention, challenge reality, and verify consequential information. A cloned-voice call to a Ferrari executive failed when he asked a question only the real CEO could answer. The SPY Check (Source, Purpose, You) makes that discipline repeatable, adding a human check to detection, authentication, and provenance at the moment of consequence.



SUMMIT OBSERVATIONS

"The mind is the one perimeter that must never be outsourced. Thinking for yourself, in a world designed to think for you, is how we defend that perimeter."

"For centuries, a forged passport represented the height of a spy's craft—an artisan's labor, carried by one officer at real risk over a border. This morning, a portrait of me at a summit in Washington was assembled in minutes from the traces of an ordinary public life."



KEYNOTE: Forged—Deepfakes and the Oldest Trick in Espionage

(Continued)

CORE THESIS: Deepfakes change the economics and speed of deception, but the real target is human judgment. Trust infrastructure must combine strong technical safeguards with trained decision discipline, corroborated channels, and the ability to establish identity independently of appearance.

KEY FRAMING POINTS

- **Deepfakes are a disguise mechanism;** manipulation of judgment is the operational objective.
- **AI compresses the recruitment cycle:** spot, assess, develop, and recruit now happen in minutes rather than months, at machine scale.
- **Appearance is no longer a reliable identity signal;** trusted systems and independent channels must establish authenticity.
- **A controlled test embedded in the keynote**—a deepfake of the speaker shown at the start of the talk—illustrated that even the trained identity and fraud experts in the room struggled to detect a well-made deepfake by sight and sound alone.
- **The liar's dividend** means deepfakes also weaken trust in genuine evidence by making denial plausible.
- **AI agents are beginning to execute social-engineering operations autonomously**—not merely serve as better disguises for human operators. Evidence was cited that AI performed 80–90% of the operational work in one espionage case.
- **Verification should be normalized** as standard operating procedure, not treated as insubordination toward a legitimate, urgent-seeming request.
- **Mind Sovereignty™**—reclaim attention, challenge reality, verify information—forms a trainable human defense layer that completes, rather than competes with, technical detection. The SPY Check (Source, Purpose, You) makes it a repeatable practice.

MIND SOVEREIGNTY™: THE SPY CHECK

Three questions before acting on a consequential request

S

SOURCE

Who is really behind the request?

P

PURPOSE

Why is it being made, and who benefits?

Y

YOU

Why are you being targeted or pressured?

If any answer is unclear, pause and verify through an independent channel.



FIRESIDE CONVERSATION: Overcoming Obsolete Threat Models

Can We *REALLY* Build Trusted and Resilient Identity Ecosystems?

STRATEGIC QUESTION: *Why are organizations still preparing for yesterday's attacks?*

The Fireside Conversation opened with the moderator naming the problem: AI has permanently changed the threat landscape, so what is stopping organizations from adopting the technologies and processes it now demands? Put differently, why do institutions still rely on evidence requirements, buying centers, and risk models built for an earlier era?

One panelist framed it as an evidence problem, not a tooling problem. Identity processes still rely on evidence, such as a printed utility bill, that generative tools can now fake convincingly, and AI-based detection tools hide their blind spots until someone exploits them. The discussion therefore moved quickly away from whether better technology exists and toward whether institutions are organized to use it.

A second panelist rejected the technology frame and named three obstacles. The first is culture: cyber, fraud, KYC, and resilience teams each buy identity solutions separately, and a Global Entry or passport background check yields nothing reusable for opening a bank account. In other words, the panelist said, the problem is not a lack of high-assurance evidence, but that trust established in one context rarely carries over to another. The second is a false choice between risk mitigation and innovation, when privacy-enabled, consent-driven credentials already exist. The third is agentic AI: always-on agents produce fraud that looks legitimate, and a synthetic identity sells for \$6–\$8. The economics, the panelist argued, have shifted from crime to commerce.

Asked who should own the problem, the panel placed identity at board level. Though one panelist went further: no single organization owns identity end to end, and the driver's license, for example, became the default U.S. credential by historical accident. Naming that gap, the panelist argued, must come before fixing it.

What would get boards to act before a catastrophe? Aligning incentives so privacy-by-design supports growth, one panelist said. An audience member was blunter, arguing that boards resist change until a crisis forces it, but that liability is already shifting as courts allow a company breached with stolen credentials to sue the company that lost them. Watching that liability chain play out in court, the audience member argued, will move boards faster than most voluntary governance efforts.

One panelist expected incremental change at best, requiring many organizations to move together. The urgency came from one figure: bots and agents already made up 55–60% of online traffic in 2025. The test offered to every practitioner was personal: would you deploy this system for your own children?



SUMMIT OBSERVATION

"We've got the technology—so when is this all going to change? A friend of mine said: well, we haven't had enough blood in the streets yet. Boards and management teams don't like to change as long as what they're doing is keeping their shareholders happy right now."



FIRESIDE CONVERSATION: Overcoming Obsolete Threat Models

Can We *REALLY* Build Trusted and Resilient Identity Ecosystems? (Continued)

Audience members noted that privacy-enabled, reusable credentials already make portable assurance feasible; the barrier, they agreed, is governance, liability, and incentives.

The sharpest turn as another audience member said that perhaps identity needs more silos, not fewer, since separation limits what one compromised credential can reach. The panel did not dismiss it, but distinguished organizational fragmentation from useful technical segmentation.

The panel's answer was to break down ownership silos while keeping least privilege and selective disclosure, and to return identity to the individual. As one panelist put it, today a bank tells the customer who they are; the customer should prove it, and the bank should simply verify the proof.

Asked why the U.S. has no national ID, the panel pointed to distrust of federal government and federal-state fragmentation. The session closed on an open question that led into the next panel: if everything reduces to trust, what mechanism can extend it down to the individual?

SUMMIT OBSERVATION

"The challenges are not technology. Culture, with a big C and a small c—organizationally, in government as well as corporate. If I'm a buyer in a bank for security solutions on the cyber team, my buying decisions are different from the anti-fraud team, from the KYC onboarding team, from the enterprise resilience team. What's the through line? Identity."

CORE THESIS: The principal obstacle to trusted and resilient identity ecosystems is no longer the absence of technology. It is the institutional model: siloed ownership, culture, incentives, liability, governance, and political will prevent organizations from using available capabilities as one coherent trust system.

KEY FRAMING POINTS

- **Identity assumptions from an earlier threat era need re-examination.** Fabrication, impersonation, and the growing volume of bot and agent activity are undermining evidence and policy models that were never designed for today's conditions.
- **The primary barriers to identity modernization are institutional, not technological.** Culture, fragmented ownership, siloed functions, and misaligned incentives matter more than the availability of solutions; no public- or private-sector organization currently owns identity coherently end-to-end.
- **Identity infrastructure remains fragmented across contexts.** High-assurance government evidence is rarely reusable elsewhere—forcing repeated verification and disclosure rather than portable credentials and proofs designed around the individual.
- **Breaking down silos does not require eliminating useful boundaries.** Functions that share an identity thread need common context and accountability, while technical segmentation, selective disclosure, and privacy protections should remain intact.
- **Incentives determine whether privacy and resilience become operational priorities.** Data monetization can work directly against minimization and user control, while liability, litigation, insurance, regulation, resilience, customer loss, and economic impact increasingly provide the forcing functions that get board attention.
- **The agentic future is already part of the operating environment.** With bots and agents already representing a substantial share of online activity, organizations must modernize identity architecture and governance for a world in which the actor may increasingly be non-human.



SUMMIT DIALOGUE: Proactive Fraud Detection Throughout the Identity Lifecycle

Breaking Down Organizational Silos to Achieve Enterprise-Wide Identity Assurance

STRATEGIC QUESTION: *How do organizations move beyond one-time identity verification to enterprise-wide identity assurance?*

The dialogue translated the Summit's architecture into an operational requirement: identity assurance must persist across the full lifecycle, not stop at onboarding. The panel opened by naming the structural problem: onboarding has long been treated as a one-and-done event, yet generative AI now breaks nearly every point of identity verification across the lifecycle. Even ordinary moments, such as parental leave, rarely trigger a review of what an employee should still be able to access. The lifecycle, the panel agreed, is full of old vulnerabilities that AI now finds and exploits systematically.

Participants noted that organizations often make recovery easier precisely because initial onboarding was difficult, effectively discarding assurance where adversaries concentrate. Leave, PTO, device changes, and re-entry can create similar gaps when identity evidence and access rights are not reassessed.

One panelist reframed the problem from the attacker's side: adversaries do not care which standard an organization follows; they look for the path of least resistance. For decades, a single onboarding event has created a credential the system then trusts on sight. But a credential being real, the panelist argued, does not mean it should be trusted.

Account recovery, the panelist said, is today's biggest weak link. The 2023 MGM Resorts breach¹ was reported as an authentication failure, but attackers did not defeat cryptography; they simply called the help desk, impersonated an employee, and talked their way in. Because onboarding is deliberately hard, organizations assume resets can be easy, and with federation and single sign-on, one compromised recovery flow can reach many organizations at once.

Participants described a connected risk surface that starts with verifying candidates, employees, and customers and continues through access and transactions, device or privilege changes, leave and return, credential reset and account recovery, re-entry, consent, and

offboarding. Generative AI and agents make it easier to find and exploit the discontinuities between these verification events.

A second panelist, drawing on experience at massive transaction scale, turned to data governance. Every lifecycle stage generates its own data, and when that data is siloed, no one sees the full picture. Large institutions may already possess billions of useful signals, but onboarding, transactions, fraud, cybersecurity, consent, recovery, and data-governance teams often see different slices of the same identity.

Securing the data is step one; the harder work is standardizing it well enough to power AI-driven detection. The panelist asked how many in the room could say their telemetry was standardized enterprise-wide. No hands were raised.

A participant called for normalized definitions, documented ownership, authoritative sources, and identity or knowledge graphs that connect systems, business processes, data, risk controls, and incident management. Shared scorecards can create cross-domain visibility without indiscriminately centralizing sensitive raw data, and vendors should return intelligence rather than functioning as one-way consumers of institutional telemetry.



SUMMIT OBSERVATION

"Human in the loop is important. Even if we are doing automation using AI, there has to be someone catching the red flags—and those individuals need to be empowered to push the button, whatever it takes to stop that fraud."

Sources & Notes: ¹ Bloomberg, "MGM Resorts Hackers Broke in After Tricking IT Service Desk," Sept. 16, 2023.



SUMMIT DIALOGUE: Proactive Fraud Detection Throughout the Identity Lifecycle (Continued)

Breaking Down Organizational Silos to Achieve Enterprise-Wide Identity Assurance

Testing was treated as part of lifecycle assurance. The group called for representative, production-like datasets, population-specific metrics, and scenarios capable of exposing slow-moving identity poisoning or swarm attacks that can look normal to conventional models. Human review remains necessary for red flags and high-consequence exceptions—and reviewers must have real authority to stop activity. The broader governance goal is to translate legal, consent, vendor, and policy obligations into executable controls tied to monitoring and incident response.

The moderator offered a counterpoint: onboarding itself is the highest-risk moment, because once an adversary is in, everything downstream gets easier. Asked who used personal AI agents, nearly every hand went up; almost none of those agents had their own identity.

Another panelist pushed back on the idea that this is new: service accounts and robotic automation have borrowed human identities for years. The real mistake, repeated for decades, is failing to design a distinct trust relationship for automated actors.

The final panelist turned to data as an asset: share patterns across functions, make multi-level signal verification the baseline, and keep a human in the loop with authority to stop a transaction, not just flag it. Data that should have been destroyed often sits undefended; when you are done with it, get rid of it.

SUMMIT OBSERVATION

“The credential might be real, but that doesn’t mean I should inherently trust it. If you’re asking me where the biggest weak link is today, it is account recovery mechanisms—any reset mechanism—because we’ve made it too easy, on the theory that since onboarding was hard, the back end must somehow be safe to make easier.”

CORE THESIS: One-and-done identity verification and KYC at onboarding were never designed to withstand AI-era fraud. The identity lifecycle—onboarding, leave, return, recovery, offboarding—is now the attack surface. Organizations often see only their own silo while the adversary sees the whole picture.

KEY FRAMING POINTS

- **The full lifecycle is one connected risk surface, not an onboarding event.** Onboarding, access, transactions, leave, return, recovery, re-entry, and offboarding all create trust transitions that adversaries can exploit.
- **Recovery and exception paths are first-class identity controls.** Help desks, account recovery, PTO or leave-related access changes, and re-entry often fall back to weaker evidence; the MGM breach¹ illustrated how a human-process failure can bypass otherwise strong technical controls.
- **Adversaries operate across the lifecycle; defenses cannot remain siloed.** Fraud, identity, cybersecurity, compliance, and operations need shared ownership, normalized telemetry, common definitions, authoritative sources, and measurable cross-functional outcomes so signals can be correlated across functions.
- **Shared intelligence does not require centralizing raw data.** Identity and knowledge graphs, shared scorecards, and cross-domain signal exchange can improve visibility while preserving appropriate data boundaries.
- **Automation does not eliminate the need for human authority.** Multi-level signals should drive automated detection, but human reviewers must retain the ability to stop high-risk transactions and resolve consequential exceptions.
- **Controls must be tested and governed for real-world conditions.** Use representative datasets, population-specific metrics, and slow-moving or adversarial scenarios, and translate governance requirements into executable controls connected to incident management.

Sources & Notes: ¹ Bloomberg, “MGM Resorts Hackers Broke in After Tricking IT Service Desk,” Sept. 16, 2023.



SUMMIT DIALOGUE: The Agentic Enterprise—Building Trust Among Humans, Machines & AI Agents

Leveraging Rapidly Evolving Innovation to Build Trusted Identity-Centric Systems

STRATEGIC QUESTION: *How do organizations establish and govern trust among humans, machines, and autonomous AI agents?*

The Agentic Enterprise session opened with the moderator naming what identity has always offered: “a throat to choke, someone to blame,” a clear ownership chain when something goes wrong. Agentic AI is challenging the traditional accountability model: agents are already moving from experimentation into production, invoking tools, communicating with other systems, and taking actions without a human at every decision point. Historically, there was always a person, account owner, or organizational role to hold responsible. Autonomous agents weaken that chain unless identity, authority, and accountability are designed in from the start.

One panelist set the tone with a simple rule: agents cannot simply be let loose. They are tested against previously decided cases and are not allowed to talk to each other, so a failure stays isolated. In one wallet-tracing test, an agent misread a transaction receipt as a person and repeatedly labeled it an actor linked to child sexual abuse material (CSAM). Like a sixth grader’s algebra, its output has to show its work.

Operational risk was a major concern. Agents can make “confident mistakes,” repeating a misunderstanding at machine speed rather than stopping to ask for clarification.

Panelists drew sharp distinctions between identifying an agent and defining what it is authorized to do. Another panelist described scoping every agent to an explicit mandate, time-bound and limited to specific merchant categories, and always separating the human principal from the agent acting for them. That produces a three-part custody chain: what the agent did, whether it met its fraud criteria, and whether the merchant met the same criteria. Card networks are now writing liability rules around that chain.

A third panelist gave a careful yes for some agent tasks and a hard no for others. Regulators once required an anti-money laundering (AML) automation built to hibernate low-risk alerts to

run in shadow mode for two years to prove it would not have missed a single suspicious activity report (SAR). AI agents are already well suited to supporting fraud decisions and gathering context for investigations. But if an agent invents details that end up in a filed SAR, the institution faces regulatory and legal risk.

Another offered a simple test: would the action be authorized if a human in the same role proposed it? If a person cannot approve a \$4 million transaction without going through procurement, neither should an agent.

The moderator then shared a personal example. Waiting for a rideshare, the moderator told an unguarded personal agent “I want to go home,” and it booked a \$3,743 trip to Toronto instead of a hotel a few blocks from the venue. How, the moderator asked, should guardrails work in practice?

Panelists answered with a checklist: correctness, cost per conversation, whether it is safe for an entry-level employee, and whether it is customer-facing or internal. Onboarding rarely asks whether a customer’s agent will move funds, and crypto has no chargeback. One panelist added that removing the human from the loop does not remove human accountability. Another said vendors cannot invent a client’s risk tolerance: clients must define what an agent may touch, who can invoke it, and its ceiling.

The session closed on one point: whether an agent acts for a human, a business, or a fraudster, there is always an identity underneath it. Fraud has always been partly automated; one panelist recalled watching fraud rings take lunch breaks and naps. Agentic AI gives them a tireless partner.



SUMMIT DIALOGUE: The Agentic Enterprise—Building Trust Among Humans, Machines & AI Agents (Continued)

Leveraging Rapidly Evolving Innovation to Build Trusted Identity-Centric Systems

CORE THESIS: Autonomous agents must be treated as governed actors: identified, bound to an accountable principal, scoped, tested, logged, and overseen by humans for consequential outcomes.

KEY FRAMING POINTS

- **Every agent action must resolve to an accountable identity.** Agents need explicit identities and mandates—not borrowed human credentials with invisible scope—and every delegated action should remain bound to a verified human, business, or other responsible principal.
- **Delegation must be explicit, bounded, and auditable.** Role- and workflow-specific permissions, transaction ceilings, time limits, category constraints, and prohibited-action rules should define what an agent may and may not do, regardless of how confidently it proposes an action.
- **Authority must be continuously re-verified, not granted once.** Logging, monitoring, and ongoing verification are necessary to detect hijacking, context drift, or changes in who or what is actually operating the agent.
- **Provenance matters more as agents scale action.** Agents amplify both trustworthy and fraudulent evidence, increasing the importance of source credentials, document provenance, and verifiable chains of authority.
- **Agent controls must reflect the organization's real workflows.** Generic guardrails can over- or under-protect; agents should be tested in context, including shadow testing against human decisions before production authority expands.
- **Humans must remain accountable for consequential outcomes.** Legal, financial, regulatory, and safety-critical decisions still require clear human escalation and responsibility.

THE AGENTIC TRUST LAYER

Each action must clear all five controls

- 1 IDENTITY** Which agent is acting?
- 2 DELEGATION** On whose behalf, bound to which accountable principal?
- 3 AUTHORIZATION** What may it do, and within what limits?
- 4 ATTESTATION** Can it prove its state and the provenance of its inputs?
- 5 AUDIT** Is every action traceable from principal to outcome?



SUMMIT OBSERVATIONS

“An agent should never be able to take an action that the equivalent human role would not be authorized to take.”

“When AI or an agent makes a mistake, it makes a very confident mistake. It’s not going to come back to you and say, I didn’t understand—can you repeat that? Like a human would.”



SUMMIT DIALOGUE: Turning Controls on Their Head

Reimagining Governance as a Driver of Trust, Security, and Growth

STRATEGIC QUESTION: *How must governance evolve when trust—not fraud detection—becomes the primary objective?*

The session opened with a terminology problem: compliance, AI governance, model risk management, and model governance mean different things, yet practitioners use them interchangeably. The discussion converged on a useful definition—governance connects evidence, roles, authority, acceptable behavior, accountability, and recourse. If the objective shifts from detecting bad activity to continuously establishing trust, governance must enable high-confidence action rather than function only as a brake on innovation.

The framing question then sharpened: if trust and culture change are the real objectives, what should governance stop optimizing for, and what should replace it?

One panelist noted that publishing a framework and achieving adoption are different accomplishments; the first output of the Linux Foundation’s Trust Over IP effort was a governance framework that has not seen wide adoption. Pressed for a definition, the panelist pointed to the passport, arguably the world’s most widely accepted digital identity. Its trust rests on ICAO (International Civil Aviation Organization), whose 193 member states jointly set evidence-of-identity standards. Real governance at global scale is slow, multilateral, and durable.

Another panelist argued that banks compete with digital challengers on infrastructure built for branches, and that much of what looks like governance failure is an education failure. Executives cannot govern authentication, fraud prevention, or agentic risk that no one has explained to them in plain terms. A unified customer identity experience across a bank’s own systems is foundational to keeping customers. Governance therefore must translate shared principles into operating models that fit different missions and levels of maturity.

The moderator added an example from the day before: a founder building an AI platform asked what “AI governance” meant, then answered the question mid-sentence while describing how the platform controlled access at different levels.

A third panelist noted that most banks and credit unions are small community institutions without the budget or staff to build identity platforms. Any solution that does not integrate natively with core banking providers such as Fiserv and Jack Henry will stall, leaving institutions that serve small businesses without protections larger banks can afford.



SUMMIT OBSERVATIONS

“You cannot provide meaningful governance if the organization does not understand the subject deeply enough to govern it.”

“Fraud actors are attacking the federal and state governments in a horizontal fashion, and we’re responding in a vertical fashion. Culturally, in government, we’re not even equipped to have these conversations, because people think an identity layer means you’re trying to deny a legitimate person their benefit—not protect the program that serves them.”



SUMMIT DIALOGUE: Turning Controls on Their Head (continued)

Reimagining Governance as a Driver of Trust, Security, and Growth

A fourth panelist said federal and state agencies trail financial services on deepfakes, in part because most government program leaders never hear conversations like this one. Fraud actors attack programs horizontally, across agencies and channels, while government responds vertically, agency by agency. Asked how it confirms a caller is really a veteran, one agency answered, in effect, “they’re calling because they want a benefit.” People in government want to serve, veterans especially, and attackers exploit that instinct: a convincing impersonation of a real beneficiary is easy to build and aimed at people disinclined to doubt. The same culture makes identity checks politically hard, since they are sincerely seen as friction that blocks entitled benefits rather than as protection for a program’s ability to serve.

The panel also noted a difference in obligation: a bank can decline a transaction in circumstances where a government benefits program may still have an obligation to serve.

Two participants closed with concrete news.

- The Department of Homeland Security soft-launched an extension of its Remote Identity Validation Rally (RIVR) to include deepfake detection. RIVR is an objective benchmark that tests vendors against thousands of real and fake documents and deepfake video attacks.
- The American Association of Motor Vehicle Administrators reported mobile driver’s licenses live in 23 jurisdictions and a forthcoming standard adding a cryptographic signature to physical license barcodes, so any relying party can verify the data against the state record.



SUMMIT OBSERVATION

“[An] identity credential from the DMV that’s done that identity proofing is fundamental to a person proving who they say they are—at onboarding, and potentially at every step throughout your process.”

CORE THESIS: Governance must evolve from static oversight and blocking logic into an adaptive trust framework that connects culture, identity, authorization, standards, data policy, and business objectives—enabling high-confidence action rather than merely preventing bad action.

KEY FRAMING POINTS

- **Governance must enable trusted action, not merely constrain risk.** Once the objective shifts from blocking fraud to continuously establishing trust, governance must connect evidence, authority, policy, accountability, and recourse to customer experience, resilience, performance, and growth.
- **Common trust baselines can support different operating environments.** Governance should establish shared principles and evidence requirements that can be mapped into sector-, jurisdiction-, and mission-specific controls; public agencies and private enterprises cannot always tolerate the same level of friction or denial.
- **Trusted credentials show that governance can create positive authenticity signals.** Passports, mobile driver’s licenses, and emerging cryptographic signature standards work best when layered with authoritative record matching and biometric verification, rather than relying on any single control.
- **Governance requires informed leadership and action before regulation is complete.** Executives need sufficient understanding of identity, authentication, fraud, and agentic risk to make sound decisions, and organizations should implement practical trust controls rather than waiting for a comprehensive regulatory framework.
- **Institutional modernization is as important as standards development.** Legacy organizations—particularly banks built around branch-era processes—must adapt governance, identity architecture, and operating models to meet the same trust expectations as digitally native competitors.



LEADERSHIP KEYNOTE: Moving at the Speed of Trust

Accelerating Identity Modernization Before Crisis—or Regulation—Demands It

STRATEGIC QUESTION: *How should enterprises adapt to identity modernization, strategically and operationally?*

The keynote speaker opened by describing a forthcoming book that grew out of frustration from running a startup. Buyers, the speaker said, rarely disputed the case for stronger identity practices; they simply could not get it approved inside their own organizations, and investors often treated compliance as the finish line rather than the floor.

The speaker then noted that machine identities now outnumber human identities inside the enterprise, but said the agentic moment is not really an “agent problem.” Agents are simply pursuing objectives across infrastructure whose identity, data, recovery, and accountability mechanisms were already incomplete. The task now, the speaker said, is to fix those fundamentals while the window is still open.

The speaker described two problems driving the urgency. The first is data: as long as personal data sits in a centralized honeypot, it will eventually be breached. Quoting a former co-founder with an offensive cyber background, the speaker said attackers need only one way in and have unlimited patience. Encryption, multi-factor authentication, and perimeter controls help, but the speaker called for privacy-preserving storage and computation, tighter retention rules, and auditable access.

The second problem, the speaker said, is human: deepfake-enabled fraud, scams, and sextortion have become an international emergency. The speaker described this as a “scamdemic”.

To explain why these problems persist, the speaker quoted the founder of an early cybercrime marketplace, now an FBI consultant: “A Band-Aid is cheap and is the path of least resistance for everyone with a stake in the system.”

The speaker walked through six factors that keep these problems in place.

- Authentication has mature standards, but data protection, data sharing, and fraud nomenclature do not.
- Insurers absorb losses, weakening the incentive to fix root causes.
- Security budgets are shrinking as spending shifts to cloud and AI.
- Onboarding a vendor at a large bank can take 18–36 months.
- Intelligence sharing breaks down without shared terms.
- Compliance lags technology.

Next, the speaker introduced two identity frameworks: The Prism Project’s Identity Hierarchy, which places the real human at the center, and the Circle of Identity², which ties onboarding, credential binding, authentication, and recovery to one foundational identity. As the speaker noted, flawless onboarding means little if recovery falls back to one-time passcodes, the gap behind the MGM Resorts breach¹. The speaker acknowledged that foundational identity needs stored data, which recreates the honeypot.

The speaker then presented the “Crossing Boundaries” model for agents: a verified human principal authorizes a specific agent for a scoped task, with a signature that resolves back through a trust registry. The speaker warned of scope creep, context drift, accountability gaps, and data governance failure.

The speaker closed with a “Monday Morning” action list for four audiences. **Regulators** should set principles and outcomes rather than technology mandates, and build the conditions for trustworthy identity. **Consumers** should demand better protections and use their spending power to reward security over convenience. **Identity and technology providers** should hold themselves to being “in the trust business” and be honest about their limits. **Enterprises** should be honest about where their identity risks sit, stop equating compliance with security, and take a first concrete step.

Sources & Notes: ¹ Bloomberg, “MGM Resorts Hackers Broke in After Tricking IT Service Desk,” Sept. 16, 2023. ² The Circle of Identity framework was developed by Frances Zelazny.





LEADERSHIP KEYNOTE: Moving at the Speed of Trust (continued)

Accelerating Identity Modernization Before Crisis—or Regulation—Demands It

CORE THESIS: The technology required to build resilient, trusted digital identity ecosystems already exists. The gap is no longer invention—it is institutional adoption. Closing that gap requires executive leadership and organizational alignment, not another point solution.

KEY FRAMING POINTS

- **The identity crisis is not primarily an agent problem.** It is the consequence of unresolved weaknesses in data architecture, foundational identity, lifecycle continuity, standards, incentives, governance, procurement, intelligence sharing, and compliance. The window to fix those fundamentals is still open—but closing.
- **Identity modernization must remain anchored to the foundational human across the full lifecycle.** The Identity Hierarchy¹ and Circle of Identity² keep credentials, devices, and behaviors—and every stage from onboarding to recovery—tied to that same identity. Biometrics remain foundational within a layered system that includes liveness, injection protection, and contextual signals.
- **Centralized personal data is itself a structural vulnerability.** Better containment is necessary but insufficient; resilient identity requires changing how sensitive data is stored, shared, minimized, and controlled rather than simply building stronger walls around ever-larger repositories.
- **Agentic AI extends identity into delegated authority.** The keynote’s “**Crossing Boundaries**” model binds an attested agent to a verified human principal, scoped authority, and a trust registry. That chain must persist as agents act, because scope creep, context drift, accountability gaps, and data-governance failures can erode trust even when the original authorization was valid.
- **Compliance is not security.** Adversaries do not care whether required controls were documented or boxes were checked. Identity assurance must be evaluated against whether controls actually preserve identity, authorization, and trust under real operating conditions.
- **The technology gap has largely become an adoption gap.** Procurement cycles, budgets, incentives, standards, shared vocabulary, intelligence sharing, and organizational alignment now determine how quickly institutions can modernize.



SUMMIT OBSERVATIONS

“Trust is designed. It is not assumed. The technologies will evolve. The threats will evolve with them. Our responsibility will not.”

“The good guys are allowing it to happen because, let’s be honest, do they really want to stop, or do they want to just put a Band-Aid on it and collect some money?”

Sources & Notes: ¹ The Identity Hierarchy was developed by The Prism Project. ² The Circle of Identity framework was developed by Frances Zelazny.





FIRESIDE CONVERSATION: Identity-First Transformation

Cautionary Tales From Real-World Deployments

STRATEGIC QUESTION: *What separates successful identity transformation from another stalled technology initiative?*

This session moved from strategy to implementation. The moderator opened with an informal poll: Who in the room was in the middle of an active identity deployment, and who had completed one successfully? Many attendees were mid-deployment but almost none had completed one, and, the moderator added, he was confident that everybody in the room had a horror story. The panel's diagnosis: identity transformation is a people and operating-model problem as much as a technology one. Programs stall when ownership is unresolved, success is defined differently, procurement focuses on product, or change management comes last.

Asked what separates success from another stalled initiative, the first panelist answered without hesitation: failing to establish ownership at the start. Everyone agrees identity is a problem; almost no one agrees who owns fixing it, and onboarding, fraud, and technology teams each see only part of it. Successful programs set decision rights before technology selection and carry them through rollout and escalation; otherwise, each function optimizes locally while the transformation fragments.

The second panelist named testing as another dividing line between success and stalled initiatives. They noted that governance and rollout plans are designed in controlled, "sane" environments, but real-world use challenges those assumptions. They recommended four critical factors to drive successful deployments:

- **Test before selection:** use organizations experienced with the technology, such as the Maryland Test Facility for face recognition, rather than choosing a vendor on faith.
- **Monitor continuously:** assumptions shift once the system is live.
- **Measure the tails, not the average:** if a system works 92% of the time, ask what happens in the other 8%, to whom, and under what conditions.
- **Treat change management as ongoing:** even a correct installation will need new versions or parameters, and handling that without a stalled rollback decides whether it survives.

The second panelist also called this "change management plus": design for a model, parameter set, or entire technology to be replaced without institutional paralysis, especially as AI compresses development cycles faster than testing and governance can keep pace.

The moderator stated that procurement takes 18–36 months depending on the institution, and a panelist argued that most procurement processes focus on the wrong question. A bake-off picks whoever tests best, but the decision should turn on what happens in the institution's own operations at go-live. The best liveness system is still the wrong choice if it stops 20% of legitimate customers at the door. The moderator's summary: don't pick the vendor who takes you to the most steak dinners.

Asked where the first cracks typically show, the first panelist said they appear when talk turns from the problem to who will do the work. Smaller institutions need a partner that understands threats such as deepfakes and injection attacks and can explain them to non-technical or identity experts. A 120-page report can be an excellent reference (the panelist cited the Prism reports approvingly), but short insights with a link are what travel into boardroom conversations.



SUMMIT OBSERVATION

"The damaging thing is not working out ownership at the very beginning of your project—everybody will agree there's an identity problem, but owning it is a different story."





FIRESIDE CONVERSATION: Identity-First Transformation (Continued)

Cautionary Tales From Real-World Deployments

An audience member asked how smaller institutions can afford rigorous testing and partner education. The first panelist was direct: don't build identity infrastructure in-house; use an orchestration platform for anything that is not a differentiator.

The second panelist offered two structural fixes: shared standards for newer categories such as deepfake detection, and pooled testing, following small law-enforcement agencies that aggregate face-recognition testing. The moderator connected this to trade associations, naming EPCOR, which represents about 1,700 small financial institutions.

An attendee asked which metrics matter once a system is live. The second panelist said each layer needs its own metrics, since one combined success rate hides which layer is failing; the business measure is how often legitimate users are sent to secondary checks.

The panel closed on two principles: share success metrics across fraud, product, technology, and operations so customer experience does not work against fraud control and keep regulated identity data within deliberate boundaries rather than in generic external AI services.



SUMMIT OBSERVATION

"It is a people problem, not a technology problem."

CORE THESIS: Identity transformation rarely fails because the technology doesn't work. It stalls because organizations underestimate the complexity of integrating identity across business units, legacy systems, and competing institutional priorities—and because no one resolved ownership at the start.

KEY FRAMING POINTS

- **Identity transformation is fundamentally an operating-model challenge, not a technology project.** Clear ownership, decision rights, and shared success measures across fraud, risk, product, technology, compliance, marketing, and operations should be established before technology is selected.
- **Fragmented teams create fragmented visibility.** Onboarding, fraud, identity, and technology functions often see different pieces of the same problem; transformation requires shared context and accountability across those boundaries.
- **A layered identity-defense system needs per-layer performance metrics—presentation-attack detection, injection-attack detection, deepfake detection, biometric matching—rather than a single blended pass rate that obscures which specific layer is actually underperforming.**
- **Independent testing should precede deployment—and continue afterward.** Organizations should validate technologies with realistic, independent testing rather than vendor demonstrations alone, including integration effects, population differences, and long-tail failure modes.
- **Testing and governance must move at AI speed.** Deployment inevitably challenges assumptions made during design, so organizations need continuous logging, evaluation, and post-deployment testing as models, inputs, workflows, and environments change. Shared testing infrastructure can also reduce the burden on smaller institutions.



SUMMIT DIALOGUE: Privacy, Data Ownership & User Control

Can Privacy Become a Competitive Advantage?

STRATEGIC QUESTION: *How can evolving models of identity ownership, user control, and data governance strengthen security, improve customer experience, and create competitive advantage?*

The privacy dialogue rejected the familiar framing that stronger security requires weaker privacy. The moderator's first question was whether organizations really controlled identity ownership and privacy even before AI. The answers came back almost uniformly "no," each for a different reason.

The first panelist described data scattered across warehouses and odd corners of the organization, with leadership unsure how it drove the business.

The second panelist was the bluntest: organizations had, and still have, incentives not to control data responsibly, because "the purpose of a system is what it does." Calling the result an accident understates how intentional those incentives are.

The third panelist offered a hedge: smaller organizations often control their own data reasonably well, but government is rarely small, and control gets harder once data is shared across offices, agencies, and other governments.

The fourth panelist reframed the question: companies have been very effective at monetizing personal data, but that is not stewardship. No one chooses their birth date, name, or country of birth, they argued, so the issuing government authority, not the individual, owns identity provenance. Whoever owns it, collecting less is better privacy practice.

Asked whether AI has made this technical debt better or worse, the panel answered at once: worse. The third panelist said data now flows to places and uses that even the organizations that collected it cannot fully see or predict. The first panelist added that AI is reaching data already forgotten or siloed, and that organizations were often sold a picture of their data and AI tools that did not match what those systems could do responsibly.

The discussion separated ideas that are often collapsed together: custody is not control, access is not authorized purpose, and nominal consent is not meaningful consent. Voice recordings, website terms accepted by continuing a transaction, and automated license-plate collection were cited as weak cases; biometric exit programs, with published privacy analysis, as clearer trust relationships.

Participants therefore focused on purpose-bound consent and minimization. A relying party may be authorized to access data but still lack permission to use it for another purpose. Wallets, verifiable credentials, selective disclosure, and cryptographic proofs let people prove what is necessary without repeatedly handing over complete identity records, while hybrid architectures preserve institutional systems of record and shift more control over disclosure to the individual.

The EU Digital Identity Wallet was discussed as a working example: any relying party must register with the issuing authority before using it, creating an auditable relationship rather than open collection. In the more fragmented U.S., wallets and verifiable credentials are maturing technically but lack a unifying regulatory anchor, so industry standards carry more of the weight for now.



SUMMIT OBSERVATION

"Organizations had incentives not to control identity data responsibly. They still have incentives not to. The purpose of a system is what it does—and the system was designed the way it was for a reason."



SUMMIT DIALOGUE: Privacy, Data Ownership & User Control

Can Privacy Become a Competitive Advantage? (Continued)

An audience member who has built AI systems for more than a decade asked how many people had used an AI tool to manage their own consent settings. Almost no hands went up. In their view, years of talk about user control have not yet reached the people it is meant to serve, which was framed as a starting point for building it, not a reason to give up.

The conversation also looked beyond conventional consent screens. Personal AI and user-managed access models could eventually help individuals interpret terms, express preferences, and enforce policy across services, making consent operational rather than ceremonial. Global digital identity deployments show that reusable, more decentralized identity is already emerging, although national approaches differ in governance and public trust.

Biometric and foundational identity data were treated as especially high-consequence assets: a payment card can be reissued; a face or voice cannot be replaced after compromise. That raises the standard for cryptography, retention, access governance, liveness and injection protection, and lifecycle controls. The session's commercial conclusion: organizations that establish trust with less data, and make those practices visible and understandable, can turn privacy into a fraud-resilience measure and a competitive differentiator.



SUMMIT OBSERVATION

"This opens up a lot of conversations about wallets and credentials—is privacy versus security really a choice we should be building our identity and AI systems to force people to make, or are they two sides of the same coin?"

"Those who can sell trust with less data will be in a stronger position than those who continue to accumulate exposure."

CORE THESIS: Privacy is not merely something trust infrastructure has to accommodate; well-designed privacy is itself part of the trust and fraud-resilience architecture.

KEY FRAMING POINTS

- **Privacy and security are complementary, not competing.** Collecting less, proving more, and limiting purpose-of-use reduce risk and build trust.
- **Start with the data you already have.** Inventory what identity data you hold, where it sits, who can access it, and why, before adding another policy layer.
- **Minimize disclosure without weakening assurance.** Selective disclosure, cryptographic proofs, and verifiable credentials let people prove only what is needed.
- **Consent must govern purpose, not just access.** Authorization to access data is not permission for every use; consent should bind both.
- **User control must become operational, not theoretical.** People need practical ways to understand, manage, and revoke permissions, not consent that exists only on paper.
- **Hybrid trust architectures can preserve authoritative evidence while reducing centralized exposure.** Institutional records can pair with user-controlled credentials, so not every relying party holds the same sensitive data.
- **Foundational identity and biometric data require a higher standard of protection.** Unlike a password or card number, a face cannot be reissued, so it demands strong cryptography, retention discipline, and tightly governed access.
- **Privacy can become a competitive advantage.** Organizations that build trust with less data, and show it, will be better positioned than those accumulating exposure.



LEADERSHIP FORUM: Building the Trust Infrastructure for the AI Era

STRATEGIC QUESTION: *If no single organization can solve this challenge alone, what must we build together?*

The closing Leadership Forum brought the day's threads together: trust infrastructure cannot be reduced to a single technology, product, or institution. Durable trust depends on layers (biometrics, cryptography, liveness, secure hardware, device and behavioral signals, provenance, and human context) in which multiple forms of evidence strengthen the decision.

The sharpest critique focused on evidence: cherry-picked claims and certifications do not show how systems perform especially under degraded or adversarial conditions. The first panelist called publishing results from cherry-picked videos indefensible ("that's insane") and said iBeta certifications say little about a 360p video dropping every fourth frame. The industry should not wait for its own disaster, as aviation did, to raise the bar.

The second panelist called the industry reactionary: decades and hundreds of billions of dollars spent on systems that prove an attack happened but prevent nothing, "like a vaccine, but the next iteration of the disease wipes out millions of people." AI compresses attack timelines while procurement, testing, and regulation run on slower cycles, so he called for a "complete brain transplant" toward proactive defense. Two dozen competing identity frameworks, none widely adopted, suggest the field needs fewer standards that are broadly understood, interoperable, independently tested, and enforced.

Asked by an audience member how to rebuild trust against swarm attacks, the third panelist made the Forum's most distinctive point: trust infrastructure includes social capital. Only one attendee had a relationship with a government security professional or a local emergency manager. Expertise and authority across sectors help only if relationships and escalation paths exist before a crisis; suggesting water infrastructure is an easy target.

The second panelist described identity as a system of symbiotic systems in which no layer works alone, and joked that unlike a biometrics salesman, a used-car salesman will tell you what is wrong with what he is selling.

The first panelist described declining a foreign private-equity request to put the company's technology into drones ("don't do evil" and "family comes first") and called for a shared baseline on biometric data after a recent exposure of millions of faces.

The fourth panelist described "AI vertigo": catching up on a technology only to find it has moved on, under policies dating to the late 1990s; closing the capability gap falls to industry. The third panelist added that rebuilding institutions outlasts any election cycle, describing quietly ghostwriting policy initiatives on both sides of the aisle.

On transparency, the first panelist said vendors who cannot stand behind their performance data should "get the hell out of our industry." Participants called for standardized testing, public benchmarks, and model-card reporting that shows where a technology fails, which requires separating defensible evidence from marketing, not exposing proprietary models. Publishing false accept rate (FAR), equal error rate (EER), and area under the curve (AUC) does not help adversaries, who already have more compute than most vendors.

Closing the gap requires continuous testing, adaptable architectures, evidence reusable across buyers, and faster paths from validation to production without sacrificing assurance. The second panelist closed with W.C. Fields: "you can fool some of the people some of the time, and that's enough to make a good living."



LEADERSHIP FORUM: Building the Trust Infrastructure for the AI Era (Continued)

CORE THESIS: Trust infrastructure is a layered ecosystem capability. Identity must be foundational, but durable trust also requires interoperable standards, real-world testing, transparent evidence, privacy, cross-sector relationships, shared intelligence, and social capital that allow institutions to coordinate before and during crisis.

KEY FRAMING POINTS

- **Trust infrastructure is layered, not product-specific.** Identity is foundational but not sufficient; biometrics, cryptography, liveness, secure hardware, device, provenance, behavior, authorization, and context must be complementary.
- **Operational evidence must replace marketing claims.** Buyers need standardized, comparable performance data under real-world conditions—not marketing-focused accuracy figures, favorable demonstrations, or certifications that reveal little about operational performance.
- **The ecosystem needs fewer, better-adopted standards.** Proliferation of non-interoperable identity and cryptographic frameworks creates unnecessary friction. Convergence around broadly adopted, interoperable, and independently tested standards is more valuable than adding more frameworks.
- **AI-speed threats require a more proactive operating model.** Reactive, detect-after-the-fact defenses are insufficient. Product development, testing, procurement, governance, and deployment cycles must become faster, more continuous, and more adaptable without sacrificing assurance.
- **Social capital is part of trust infrastructure.** Government, industry, standards bodies, academia, emergency management, critical infrastructure, and technology providers need working relationships, shared language, and escalation paths in place before a crisis occurs. Continuity of operations depends on people prepared to coordinate under uncertainty, not only systems that establish authenticity.
- **Building trust at scale requires sustained collaboration among all constituencies** and radical honesty about the capabilities and limitations of current technology and solutions. Waiting for regulation to drive honesty is unsustainable.



SUMMIT OBSERVATIONS

“Trust infrastructure is both systems that help establish what is authentic and people prepared to act wisely when uncertainty remains.”

“If we have a catastrophic event while we’re sitting here debating privacy and a massive something happens—who do you think is going to be in charge of response or continuity of operations for the next 24 hours? Do you have a point of contact here?”

FORUM CONCLUSION

The Forum closed on a pragmatic note. Building trust infrastructure does not require waiting for universal agreement, comprehensive regulation, or a single architecture embraced by the entire ecosystem. Progress can come through focused working groups, independent testing programs, shared performance metrics, interoperable standards, pre-competitive intelligence exchange, and cross-sector relationships built around specific problems.

The larger lesson from the day was that **no single organization can rebuild trust at the scale AI-era threats require. The task is collective—and the mechanisms for collaboration must be built with the same intentionality as the technology itself.**



CLOSING REFLECTIONS

Where Do We Go From Here?

STRATEGIC QUESTION: *How do we build trust infrastructure for the AI economy?*

The closing returned to the identity model introduced in the morning and refined it through the day's discussion. Credentials—licenses, passports, academic records, professional certificates—are representations of identity, not identity itself. Their trust ultimately depends on resolving to the same foundational human. Assurance must persist from that anchor across the lifecycle, while remaining dynamic enough to strengthen, weaken, or return to stronger controls as context and risk change.

The **Trust Velocity Gap** re-emerged as the institutional challenge. AI accelerates both threats and defensive innovation, while procurement, governance, testing, integration, policy, and organizational change move far more slowly. The unresolved question is whether 20th-century functional silos can manage a trust problem that increasingly behaves as a fluid, cross-functional network.

The Summit's final message was therefore an operating mandate, not a product prescription: **close the identity lifecycle, govern delegated agency, reduce the Trust Velocity Gap, test under real-world conditions, redesign data practices, and turn the relationships built at the DC Summit into shared implementation.** Trust cannot simply be assumed. It must be designed, continuously evaluated, challenged, and re-established as conditions change.

THE IDENTITY HIERARCHY

Identity is layered and built on a trusted foundation; every level depends on the integrity of the layers beneath it. Identity is not a login—authentication is not the end of the identity process.

THE RESILIENT TRUST PYRAMID

Adds a third dimension to the decentralized Triangle of Trust by creating a biometric-anchored link between a physical human and their digital representations—underpinning every layer of assurance above it.

IDENTITY ALONE IS NOT TRUST

The challenge expands from verifying identity to continuously assessing the trustworthiness of the interaction, including the person, device, and increasingly complex environmental and risk signals.

SHIFT THE SECURITY PARADIGM

From ACCESS to IDENTITY to TRUST. Another fraud tool or authentication layer will not close the gap. The security model itself must evolve.

THE TRUST VELOCITY GAP

Threats and defenses evolve at AI speed; institutional adoption moves at enterprise speed. The mandate is twofold: increase enterprise speed AND bridge the gap at the interface—procurement, budgets, governance, and shared intelligence.

THE 21ST-CENTURY ENTERPRISE

Does this transformation require a new enterprise structure? The hierarchical 20th-century model may be the bottleneck. The alternative: fluid, cross-functional networks built for AI-era tempo.

CORE THESIS: Building trust infrastructure is not a destination. The work is a continuous cycle of implementation, verification, governance, collaboration, and learning. The DC Summit's value lies in converting a day of shared insight into concrete institutional and ecosystem action.

KEY FRAMING POINTS

- **Build for trust, not simply access.** Identity must remain foundational and continuous, but trust depends on the integrity of the entire interaction—not just the credential.
- **Close the Trust Velocity Gap.** Institutions must increase the speed at which they can safely evaluate, adopt, govern, and operationalize new capabilities.
- **Make governance executable.** Policies and principles have to become operating controls, decision rights, monitoring, and measurable outcomes.
- **Treat privacy, testing, standards, and collaboration as core infrastructure.** They are not supporting functions; they are part of how resilient trust is designed and sustained.
- **Turn Summit relationships into action.** Carry the conversations forward through focused working groups, shared testing, interoperable standards, intelligence exchange, and implementation.





KEY FINDINGS: Top Ten Takeaways

Practitioners, technologists, and policy experts reached a consistent conclusion at the DC Summit: the primary constraint on digital trust is increasingly not the availability of detection technology, but the lifecycle coverage, governance, and enterprise velocity required to deploy these capabilities effectively. What follows are the Summit's ten most consequential insights.

1

The Attack Surface Is the Full Identity Lifecycle—Account Recovery Is the Weakest Link

2

Deepfakes Target Judgment—and the Liar's Dividend Erodes Trust Itself

3

Identity Alone Is Not Trust—Shift From Access to Identity to Trust

4

Foundational Identity Must Anchor a Closed Circle of Identity¹

5

Data and Organizational Silos Are Security Defects—Telemetry Without Governance Is Noise

6

Privacy Is a Fraud-Resilience Measure, Not Merely a Compliance Obligation

7

Real-World Testing Beats Lab Claims—Model Cards, Bake-Offs & Independent Rallies

8

Agents Need Mandates, Not Just Identities

9

The Trust Velocity Gap Is the Defining Enterprise Risk

10

Trust Infrastructure Is Built Before the Crisis—Relationships, Standards & Education

Where Trust Fails

How Trust Must Be Built

What Comes Next





KEY FINDINGS: **Where Trust Fails**

1

The Attack Surface Is the Full Identity Lifecycle—Account Recovery Is the Weakest Link

Generative AI breaks verification assumptions across the full identity lifecycle—from onboarding and access through recovery, re-entry, privilege changes, and offboarding. Recovery is often the weakest point, with help desks, device changes, call centers, and fallback processes relying on weaker evidence than initial enrollment. Synthetic identities can also remain dormant for months before becoming active threats. Identity assurance must therefore persist across the lifecycle, with recovery treated as a first-class, high-risk identity event and stronger foundational verification triggered whenever risk rises.

2

Deepfakes Target Judgment—and the Liar's Dividend Erodes Trust Itself

Synthetic media is the modern disguise in a classic deception operation whose real target is human judgment. Recruitment cycles that took months are compressed to moments; an estimated 55–60% of online traffic is already non-human¹; synthetic identities sell for \$6–\$8². The second-order damage is worse: once anything can be faked, genuine evidence can be dismissed—the liar's dividend—degrading decision-making across markets, newsrooms, and leadership tables. Defending Mind Sovereignty™ requires trained verification habits alongside technical controls.

3

Identity Alone Is Not Trust—Shift From Access to Identity to Trust

Knowing exactly who someone is does not establish that the device, behavior, environment, content, or delegated AI agent should be trusted. The security paradigm must evolve from access (permission), to identity (who or what), to trust (should this interaction proceed—and should that trust persist?). Point solutions cannot close this gap; the model itself must change.

Sources & Notes: ^{1,2} Deepfake Summit DC Fireside Conversation: Overcoming Obsolete Threat Models, citing 2025 estimates.



KEY FINDINGS: How Trust Must Be Built

4

Foundational Identity Must Anchor a Closed Circle of Identity¹

Onboarding, credential binding, authentication, and account recovery must operate as one closed loop anchored to a foundational, biometric-linked human identity—with fallback to that foundation whenever risk signals spike. When these run as separate cycles stitched together with OTPs, SMS, and passwords, attackers exploit the seams. The paradox must be managed, not wished away: foundational identity implies storage, and storage creates risk—even government-issued credentials sit atop repositories.

5

Data and Organizational Silos Are Security Defects—Telemetry Without Governance Is Noise

Even the largest platforms hold billions of records they cannot fully use: signals are siloed by team, inconsistent in format, and ungoverned in ownership. Lifecycle-spanning attacks—swarm attacks, identity poisoning—are invisible to fragmented monitoring. The path forward is an identity/knowledge graph linking data sources, ownership, processes, risks, and controls, with taxonomies and policy-as-code turning legal and consent constraints into enforceable, automated guardrails.

6

Privacy Is a Fraud-Resilience Measure, Not Merely a Compliance Obligation

Privacy and security are mutually reinforcing. Excessive collection, centralized storage, and repeated disclosure of identity data increase both attack surface and the consequences of compromise. Data minimization, selective disclosure, reusable credentials, and stronger user control can reduce exposure while preserving strong identity assurance—making privacy an integral part of fraud resilience and trust infrastructure.

7

Real-World Testing Beats Lab Claims—Model Cards, Bake-Offs & Independent Rallies

Cherry-picked accuracy claims corrode the credibility of the entire detection industry. The corrective agenda: operational testing under low-resolution, frame-drop, adversarial, and swarm conditions; buyer-run bake-offs; independent programs such as DHS's Remote Identity Validation Rally (RIVR), which was recently extended to include deepfake detection; and honest model cards with standardized, comparable metrics—equal error rate (EER), area under the curve (AUC), and presentation-attack and injection performance. Independent testing requirements belong in every RFP.



KEY FINDINGS: **What Comes Next**

8

Agents Need Mandates, Not Just Identities

AI agents are executing transactions across payments, crypto, identity, and customer channels—and can create other agents on the fly. Verifying an agent's identity is necessary but insufficient: every agent needs an explicit, inspectable mandate—time-bound, merchant- and category-limited, formally separated from its human principal—with permissioning mapped to roles and workflow actions, shadow-testing before production, and humans retained on legally consequential outcomes. Agents hallucinate confidently; in compliance contexts, tolerance is near zero.

9

The Trust Velocity Gap Is the Defining Enterprise Risk

Threats and defensive technology evolve at AI speed; institutions adopt at enterprise speed. Procurement cycles run 18–36 months¹, security budgets are diluted by cloud and AI spend, insurance economics blunt the incentive to fix root causes, and compliance box-checking substitutes for security. Closing the gap requires increasing enterprise speed—procurement, budgeting, governance—and may require the enterprise structure itself to evolve from hierarchical silos to fluid, cross-functional networks.

10

Trust Infrastructure Is Built Before the Crisis—Relationships, Standards & Education

Trust is social capital as much as technology: cross-sector relationships—private industry, government, emergency management, standards bodies—must exist before the crisis that tests them, particularly where high-consequence infrastructure (water, power, medical systems) is in scope. Deployable standards beat theoretical completeness (~25 competing identity frameworks were cited), digital credentials are scaling (mDLs in 23 U.S. jurisdictions), and market education—plain-language definitions, honest reporting—is itself part of the infrastructure.


Sources & Notes: ¹ Deepfake Summit DC Leadership Keynote.



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; to technology providers and the identity industry; and to regulators and governance bodies. They are presented as the starting points the Summit’s practitioners identified as most urgent—not the full scope of what is needed. The challenges documented in this report are systemic and will require sustained effort across the ecosystem.

FOR
ENTERPRISES



FOR TECHNOLOGY &
IDENTITY INDUSTRY
PROVIDERS




FOR REGULATORS &
GOVERNANCE BODIES





RECOMMENDATIONS: Enterprises

FOR
ENTERPRISES



1

Harden Onboarding and Account Recovery as First-Class Risk Controls

Attackers go where entry is cheapest: onboarding and recovery. Apply multi-level verification and diverse signals rather than single credential checks; strengthen help-desk and call-center identity verification; keep humans in the loop on red flags; and never let recovery flows demand less assurance than the original enrollment.

2

Extend Verification Across the Entire Identity Lifecycle—Agents Included

Instrument candidate screening, onboarding, access, leave and return, recovery, and offboarding as one continuous risk surface. Give agents their own identities, bound to accountable human owners, and monitor for agent-driven anomalies, swarm attacks, and slow-moving identity poisoning.

3

Close the Circle of Identity¹ with a Foundational, Biometric-Anchored Fallback

Anchor onboarding, binding, authentication, and recovery to one foundational human identity, and fall back to it whenever risk spikes—new device, new location, anomalous behavior. Retire the OTP/SMS/password seams that break the circle.

4

Normalize, Own, and Govern Your Telemetry

Establish authoritative sources, ownership, and shared taxonomies for identity signals across teams; link them in an identity/knowledge graph; and convert legal, consent, and vendor constraints into policy-as-code with automated enforcement connected to incident management.

5

Train the Human Layer—and Rehearse the Deepfake Plan

Institutionalize the pause: train employees on the Mind Sovereignty™ SPY Test to assess Source, Purpose, and You—who is behind the request, why it is being made, and why they are being targeted—then independently verify consequential requests through a trusted second channel. Make verification standard operating procedure, not insubordination, and rehearse high-pressure scenarios before the urgent 2 a.m. call arrives.



RECOMMENDATIONS: Technology & Identity Industry Providers

FOR TECHNOLOGY & IDENTITY INDUSTRY PROVIDERS



1

Publish Honest Model Cards with Standardized Operational Metrics

Report comparable figures—equal error rate (EER), area under the curve (AUC), and presentation-attack and injection-attack performance across datasets—including weaknesses. Cherry-picked accuracy claims corrode buyer trust in the entire category; standardized disclosure raises the floor for everyone.

2

Submit to Independent, Real-World Testing

Participate in DHS's Remote Identity Validation Rally and equivalent programs; support buyer bake-offs and shared testing labs; and validate under operational conditions—low-resolution capture, frame drops, adversarial and swarm attacks—not just lab benchmarks.

3

Build for Document Provenance and Verifiable Credentials

Document fraud is the underrated blind spot. Support cryptographic document provenance—including the emerging cryptographic signature standard for driver's license barcodes—DMV-backed record matching, and mDL acceptance as adoption expands past 23 U.S. jurisdictions.

4

Democratize Capability for Smaller Institutions

Community banks, credit unions, SMBs, and local public agencies cannot run 36-month integrations. Deliver orchestration-ready, API-based capabilities; support consortium purchasing and aggregated testing; and price shared labs within reach. The ecosystem is only as resilient as its least-protected participant.

5

Converge on Deployable Standards

Roughly 25 competing cryptographically secured identity frameworks represent fragmentation, not choice. Prioritize pragmatic, deployable standards—and interoperate—over proprietary completeness.



RECOMMENDATIONS: Regulators & Governance Bodies

FOR REGULATORS & GOVERNANCE BODIES



1

Establish Baseline Governance for AI Agents

Require explicit agent mandates (time-bound, scoped, principal-separated), permissioning mapped to authorized human actions, shadow-testing before production for consequential decisioning, comprehensive audit logging, and human accountability for legally consequential outcomes—including AML and SAR filings.

2

Reduce Regulatory Fragmentation with Cross-Industry Baselines

Federal-state patchwork and shifting administrations leave institutions without a consistent floor. Advance common baseline frameworks that each sector can adapt—drawing lessons from the EU's faster, more concrete rulemaking—while industry standards fill the near-term gap.

3

Fund and Expand Government-Anchored Identity and Testing Infrastructure

Accelerate mDL adoption beyond 23 jurisdictions, support cryptographic document-provenance standards, and scale independent validation programs such as the DHS Remote Identity Validation Rally so that buyers—public and private—can procure on evidence.

4

Reframe Verification Friction as Beneficiary Protection

Especially in public benefits delivery, identity verification is miscast as denying access. Set expectations—and fund customer education—that verification protects the legitimate beneficiary from organized criminal networks, and design friction to be proportionate, personalized, and explained.



THE CHALLENGE AHEAD


The Gap Is Adoption and Coordination, Not Invention

The Summit did not reveal a shortage of possible solutions. It revealed the difficulty of turning available capabilities into institutional capacity. The unresolved work is increasingly organizational, architectural, and systemic.

- **Architectural legacy.** Digital systems were built for connectivity, accessibility, and convenience, with security layered on afterward. Moving from access control to trust requires changes that extend well beyond the IAM stack.
- **Institutional speed.** Procurement, budgeting, governance, testing, legal review, integration, and change management still operate on cycles far slower than AI-driven threats and innovation. Institutions must accelerate safe adoption without weakening assurance.
- **Ownership.** Identity spans fraud, cyber, IAM, privacy, compliance, data, product, operations, and leadership. Organizations need shared outcomes and clear decision rights while preserving appropriate segmentation and privacy boundaries.
- **Delegated agency.** The ecosystem lacks mature, widely adopted models for identifying who authorized an AI agent, defining its mandate and limits, governing agent-to-agent actions, and assigning accountability when something goes wrong.
- **Evidence.** Buyers need credible, comparable information about how identity, biometric, deepfake, and AI controls perform under real-world conditions. Independent testing, common metrics, transparent limitations, and model-card-style reporting must replace reliance on marketing claims and isolated benchmarks.
- **Privacy.** Strong foundational identity must coexist with reduced data exposure. That requires trusted systems of record combined with minimization, selective disclosure, cryptographic proofs, privacy-enhancing technologies, bounded purpose, and meaningful user control.
- **Collective action.** Trust failures cross institutions and sectors, while no organization sees or controls the entire chain. Resilience therefore depends on shared language, interoperable evidence, intelligence pathways, standards, and relationships strong enough to support coordinated action.

Identity verification has improved substantially, but verification is no longer the finish line. The harder question is whether a specific interaction, behavior, transaction, or autonomous agent should be trusted—and whether that trust should persist as conditions change. Many of the required capabilities are already emerging: independently tested identity technologies, cryptographically verifiable credentials, delegated-authority models for agents, and trainable human defenses against AI-enabled deception. The remaining gap is not primarily invention. It is adoption, coordination, and institutional execution.

The central challenge is therefore whether the ecosystem can build and operationalize trust infrastructure—architecturally, organizationally, and collectively—at the speed the AI economy now requires.





THE ROAD AHEAD

The DC Summit confirmed what Houston revealed: there is a real and growing appetite for practitioner-focused, intimate, substantive dialogue on AI-driven impersonation fraud and the architecture of digital trust. Washington widened the circle—bringing intelligence tradecraft, public-sector program integrity, standards bodies, and federal testing programs into the conversation—and moved the agenda from diagnosing the threat to building the infrastructure required to address it.

It also made clear how much work remains. The identity lifecycle has become an end-to-end attack surface while organizations are still determining how to govern AI agents already operating in production. Existing consent frameworks are being strained by new forms of data use and automated decision-making, and procurement, budgeting, governance, and organizational change still move at enterprise speed. The challenge is increasingly institutional: whether organizations can adapt their operating models, decision structures, and cross-sector relationships quickly enough to keep pace with AI-driven change.

The next phase must therefore be measured by implementation, not awareness. The question for every stakeholder is no longer whether the trust landscape is changing, but which part of the Trust Velocity Gap it can reduce now.

Trust infrastructure is not a static architecture or a one-time deployment. It is a continuous institutional capability that must be designed, tested, governed, challenged, and strengthened as technology, threats, and operating environments evolve. The work ahead is to turn the principles explored at the Summit into repeatable practice—and to build the organizational and ecosystem capacity required to sustain trust at AI speed.

PRIORITIES FOR FUTURE WORK

When The Deepfake Summit reconvenes in 2027, it will take stock of where the ecosystem stands in each of these areas.

- **Lifecycle assurance:** Are organizations hardening recovery, re-entry, and other exception paths to the same standard as onboarding?
- **Agent governance:** Are enterprises binding agents to verified principals, explicit mandates, scoped permissions, continuous logs, and human escalation?
- **Operational testing:** Are shared datasets, model cards, independent testing, and benchmarks buyers can use becoming standard practice?
- **Data architecture:** Are institutions reducing unnecessary centralization, improving retention discipline, and adopting selective disclosure and privacy-enhancing approaches?
- **Trust Velocity Gap:** Are procurement, integration, governance, and change-management cycles becoming faster without sacrificing assurance?
- **Standards and intelligence:** Are organizations converging on shared definitions, provenance mechanisms, and pre-competitive signal exchange?
- **Social capital:** Are government, industry, standards bodies, emergency management, and technology leaders building durable relationships before the next high-consequence event?

**In the age of agentic AI, identity is necessary but not sufficient.
Trust must be engineered, verified, and continually re-earned.
The work of building that infrastructure has only just begun.**








CONTINUE THE RESEARCH

For more insight into the biometric-centric digital identity ecosystem, read The Prism Project reports



THE PRISM
— PROJECT —

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



Continue the conversation at



THE DEEPFAKE SUMMIT

September 2027 | Washington, DC

www.thedepfakesummit.com

Brought to you by The Prism Project
©2026 Acuity Market Intelligence

