INDIGO CONSULTING

Get a Free Assessment

E-BOOK

Verified Trust: Navigating the AI Frontier in CIAM and Fraud Prevention

Author: Nicholas Seigneur
Published in March 2026

info@indigoconsulting.ca

INDIGO CONSULTING

Get a Free Assessment

Table of Contents

Introduction: The New Imperative for Digital Identity 1
Section 1: The AI-Accelerated Threat Landscape 3
1.1 A Force Multiplier for Fraud: The Democratization of Cybercrime 3
1.2 Anatomy of the Modern AI-Powered Attack 4
1.2.1 Phishing 2.0 & Hyper-Personalized Social Engineering 4
1.2.2 The Deepfake Deception: Weaponizing Sight and Sound 4
1.2.3 Synthetic Realities: The Rise of AI-Generated Identities and Documents 5
1.3 Quantifying the Impact: A System Under Siege 6
Section 2: The Defender's Edge: AI and Cloud-Scale 8
2.1 Beyond Human Capability: The Inevitability of AI in Defense 8
2.2 The Foundational Synergy: AI and the Cloud Platform 9
2.3 From Reactive to Predictive: The Power of Data Intelligence 10
Section 3: AI in Action: Practical Applications in Fraud Detection 11
3.1 Anomaly Detection: The Digital Immune System 11
3.2 Behavioral Biometrics: You Are How You Act 12
3.3 Adaptive Authentication and Intelligent Friction 13
Section 4: The Proactive Ecosystem: Unifying Defenses with the Shared Signals Framework 17
4.1 The Problem of Security Silos 17
4.2 Introducing the Shared Signals Framework (SSF) 18
4.3 Key Protocols in Action 18
4.4 Use Case in Practice: From Detection to Remediation in Milliseconds 19
Section 5: The Perfect IAM Program: Balancing Security and User Experience 23
5.1 The Core CIAM Dilemma: Security vs. Convenience 23
5.2 The Solution: "Invisible Security" and Frictionless Authentication 24
5.3 From Security Tax to Competitive Advantage: The Win-Win of AI-Driven CIAM 25

Section 6: The North Star: Future Trends and Strategic Vision 26
6.1 Lessons from Identiverse: Andre Durand's Vision for the Future of Identity..... 26
6.2 The Duality of Generative AI: Fraud Copilot vs. Defender's Assistant 27
6.3 The Road Ahead: Decentralization, Privacy, and the Future of Verification 29
Conclusion: Future-Proofing Your Identity Strategy..... 30
Actionable Recommendations..... 31
Works Cited 33

Introduction: The New Imperative for Digital Identity

The field of cybersecurity is at a critical inflection point. The rapid democratization and advancement of Artificial Intelligence (AI) have irrevocably altered the balance of power between malicious actors and enterprise defenders. [1]In this new paradigm, digital identity has transcended its traditional role as a component of security architecture to become the primary battleground where trust is contested, exploited, and defended. [3]The perimeter has dissolved, and identity is now the nexus of control for access to critical data and services.

This report advances a central thesis: in an era defined by AI-driven threats of unprecedented scale and sophistication, trust can no longer be a passive assumption; it must be an active, continuous, and contextually verified state. This concept of "Verified Trust," a vision articulated by industry leaders like Ping Identity CEO Andre Durand, necessitates a fundamental shif beyond traditional Customer Identity and Access Management (CIAM) strategies. [5]It demands the adoption of AI-powered defenses, the creation of interconnected security ecosystems, and a meticulously calibrated balance between robust security and a seamless user experience. The challenge is no longer merely to authenticate users at the point of entry but to maintain a high-assurance state of trust throughout every digital interaction, invisibly and in real-time.

To navigate this new frontier, this report will provide a comprehensive analysis structured to guide security, IT, and digital transformation professionals. It begins by dissecting the AI-accelerated threat landscape, detailing how attackers now operate at a speed and scale that overwhelm human-led defenses. It then pivots to the defender's necessary response, exploring how the synergy of AI and cloud-scale platforms provides the only viable countermeasure. The analysis will delve into the practical applications of defensive AI, from anomaly detection to behavioral biometrics, that are reshaping fraud prevention.

Subsequently, it will introduce the imperative for a proactive, unified ecosystem enabled by standards like the OpenID Shared Signals Framework (SSF), which breaks down the dangerous silos between security tools. Finally, the report will address the core business challenge of CIAM—balancing security and experience—and conclude with a strategic vision for the future of identity, incorporating key industry insights and the transformative potential of Generative AI. This comprehensive examination will equip organizations with the knowledge required to future-proof their identity strategies in the age of AI.

Section 1: The AI-Accelerated Threat Landscape

The current threat landscape is characterized by a fundamental asymmetry, where AI has become a force multiplier for cybercriminals. This has enabled a new generation of identity-based attacks that operate at a scale, speed, and level of sophistication that renders traditional, static defenses inadequate. Understanding the mechanics and impact of this shift is the first step toward building a resilient defense.

1.1 A Force Multiplier for Fraud: The Democratization of Cybercrime

Historically, the most sophisticated cyberattacks were the domain of well-funded nation-state actors. Today, AI has dramatically lowered the barrier to entry, effectively democratizing advanced cybercrime capabilities.[1] This shift is not merely incremental; it represents a paradigm change in the accessibility of powerful attack tools. The modern threat landscape is now defined by the rise of illicit "as-a-service" models, where complex attack methodologies are packaged and sold on the dark web. The emergence of Malware-as-a-Service (MaaS), Phishing-as-a-Service (PhaaS), and uncensored Large Language Models (LLMs) like WormGPT and FraudGPT means that even criminals with limited technical skills can now launch attacks that were once the exclusive purview of elite hacking groups.[1]

The core advantage these tools provide to attackers is the potent combination of automation and scale. AI algorithms can automate nearly every stage of the attack lifecycle, from initial reconnaissance and vulnerability discovery to the mass execution of campaigns.[6] This allows threat actors, such as the North Korea-linked group "Famous Chollima," to sustain an exceptionally high operational tempo, conducting hundreds of intrusions in a single year—a pace impossible to achieve through manual means.6 For defenders, this means that attacks now materialize far faster than human-led security teams can possibly track, analyze, or respond to, creating a significant and dangerous gap between detection and remediation.[12]

This new reality creates an asymmetric warfare scenario in cybersecurity. Previously, organizations with larger security budgets and more human analysts held a defensive advantage. Now, a single attacker leveraging inexpensive, powerful AI tools can generate an attack volume and level of sophistication that can overwhelm a large, traditionally-equipped Security Operations Center (SOC).[14] This technology directly undermines biometric authentication systems and liveness checks, which are increasingly used for identity verification.

The cost and effort for the attacker have plummeted, while the resources required for a defender using legacy methods to keep pace have skyrocketed. This fundamental shift in the economic and operational balance of power is what makes the AI-driven threat so formidable.

1.2 Anatomy of the Modern AI-Powered Attack

The democratization of AI has given rise to a new class of identity attacks that are more personalized, more believable, and more difficult to detect. These attacks specifically target the weakest link in the security chain: human trust and perception.

1.2.1 Phishing 2.0 & Hyper-Personalized Social Engineering

The era of poorly worded phishing emails with obvious grammatical errors is over. Generative AI and LLMs are now used to craft flawless, contextually aware, and hyper-personalized phishing campaigns delivered via email, SMS (smishing), and social media messages.[1] These AI-generated communications eliminate the classic red flags that users have been trained to spot, making them significantly more convincing and effective.[14]

Furthermore, AI automates the entire social engineering lifecycle. Malicious actors use AI tools to scrape publicly available data from social media profiles, professional networking sites, and personal blogs to build detailed profiles of their targets. This information is then fed into LLMs to generate highly tailored lures that exploit known psychological biases, reference specific personal or professional details, and mimic the communication style of trusted individuals or organizations. [10] This level of personalization at scale was previously unattainable, requiring immense manual effort.

1.2.2 The Deepfake Deception: Weaponizing Sight and Sound

Perhaps the most alarming development is the weaponization of deepfake technology. Powered by Generative Adversarial Networks (GANs), deepfakes allow for the creation of hyper-realistic video and audio clones from just a few seconds of source material.

Successful deepfake attacks have already been documented against financial institutions, where synthetic media was used to bypass security controls during loan applications.[23]

participating in a video conference call with what they believed were the company's senior officers. In reality, every participant on the call, except for the victim, was a deepfake recreation.[10] Similarly, AI-powered voice cloning is being used to bypass voice-based authentication and to conduct highly convincing vishing (voice phishing) attacks, a threat that has prompted specific warnings from law enforcement agencies like the FBI.[9]

This trend has a critical second-order effect: the systematic erosion of human intuition as a defense mechanism. Security awareness programs have traditionally relied on training employees to recognize anomalies—a suspicious link, a grammatical error, an unusual request. AI-powered attacks are engineered to eliminate these anomalies. Deepfakes fundamentally challenge our ability to trust what we see and hear.[5] The consequence is that organizations can no longer depend on their employees to be a reliable last line of defense against sophisticated social engineering. Security must become systemic, automated, and predicated on a Zero Trust principle that assumes any user can and will be deceived. Verification must become a technical control, not a human judgment call.

1.2.3 Synthetic Realities: The Rise of AI-Generated Identities and Documents

Beyond impersonating existing individuals, AI is also being used to create entirely new, fraudulent identities from scratch. Synthetic identity fraud, which combines real (often stolen) and fabricated personally identifiable information (PII) to create a "new" person, is one of the fastest-growing financial crimes.[25] AI and automation enable fraudsters to create and scale these synthetic identities across multiple platforms, with projected annual losses in the U.S. expected to reach $23 billion by 2030.[25]

A parallel and equally concerning trend is the explosion in AI-assisted digital document forgery. For the first time, digital forgeries have surpassed physical counterfeits, with a staggering 244% year-over-year increase.[7] AI tools can now generate highly convincing fake government IDs, utility bills, and financial statements that include sophisticated security features, making them difficult for both human reviewers and basic Optical Character Recognition (OCR) systems to detect.[7] This directly threatens the integrity of digital onboarding and Know Your Customer (KYC) processes across industries.

1.3 Quantifying the Impact: A System Under Siege

The financial and operational impact of this AI-accelerated threat landscape is immense. Projections from Deloitte indicate that losses from generative AI-driven fraud in the United States alone will reach $40 billion by 2027.1 On a global scale, the total cost of cybercrime is forecasted to hit an astonishing $10.5 trillion annually by 2025.[7]

Identity has become the primary vector for these attacks. Data from Microsoft's identity infrastructure reveals the sheer volume of this onslaught: the company observes over 600 million identity attacks daily, with password-based attacks comprising over 99% of this traffic. This translates to Microsoft blocking an average of 7,000 password attacks every second, a testament to the pervasive and automated nature of modern identity threats.[3] These figures underscore a critical reality: the digital identity ecosystem is under a constant, high-velocity siege, driven by the very AI technologies that promise to revolutionize our world.

The table below provides a comparative analysis of traditional versus AI-enhanced attack methods, illustrating the significant evolution in threat capabilities.

Attack Vector	Traditional Method (Pre-AI)	AI-Enhanced Method	Key AI Enabler	Impact
Phishing	Mass, generic emails; often contain grammatical errors and obvious tells.	Hyper-personalized, context-aware, and grammatically flawless emails generated and distributed at massive scale.	Large Language Models (LLMs)	Significantly higher success rate; bypasses traditional spam filters and user training.
Social Engineering	Manual research on targets; generic pretexts and scripts.	Automated profiling using public data; deepfake video/ voice calls; dynamically generated, convincing scripts.	Generative AI, LLMs	Drastically increases believability and psychological manipulation; enables high-stakes fraud.
Document Fraud	Physical counterfeiting; manual digital editing of document images.	Fully AI-generated digital forgeries with realistic security features (holograms, micro-text).	Generative Adversarial Networks (GANs)	Bypasses standard OCR and human review processes for onboarding and KYC.
Credential Attacks	Simple brute force and password spraying against common passwords.	AI-driven credential stuffing using breached password lists; ML-powered password cracking that predicts variations.	Machine Learning Models	Higher efficiency in compromising accounts; circumvents basic password policies.
Biometric Spoofing	Static photos or pre-recorded videos used in presentation attacks.	Real-time deepfake videos with generated liveness artifacts (blinking, head movement); realistic voice clones.	Generative Adversarial Networks (GANs)	Bypasses many liveness detection and voice authentication systems.

Section 2: The Defender's Edge: AI and Cloud-Scale Defense

In the face of an adversary armed with AI, automation, and scale, a reactive, human-centric defense is no longer viable. The only effective countermeasure is a defensive strategy built on the same foundational principles: leveraging artificial intelligence, embracing automation, and operating at the massive scale enabled by modern cloud infrastructure. This section outlines the strategic imperative for adopting an AI-powered, cloud-native defense to regain the advantage in the battle for digital identity.

2.1 Beyond Human Capability: The Inevitability of AI in Defense

The sheer volume, velocity, and sophistication of modern identity attacks have surpassed the limits of human cognitive and operational capacity. A security analyst, no matter how skilled, cannot manually sift through the petabytes of log data, network traffic, and user events required to detect a subtle anomaly in real-time.[1] Similarly, static, rule-based security systems are brittle and quickly become obsolete as AI-driven malware and attack patterns constantly mutate and evolve.[13] Relying on these legacy approaches in the current environment is akin to bringing a knife to a gunfight.

The adoption of AI in defensive security is therefore not a choice, but a necessity. AI-powered security platforms are designed to operate at machine speed, analyzing vast datasets to identify complex patterns and correlations that are invisible to human analysts. Empirical data supports this transition; organizations that make extensive use of AI and automation in their security operations see significant benefits, including a reduction in alert fatigue, earlier breach detection, and faster, more precise incident response. This translates into tangible financial advantages, with such organizations saving an average of USD 1.76 million per data breach compared to those without these capabilities.[31]

2.2 The Foundational Synergy: AI and the Cloud Platform

The effectiveness of any defensive AI system is contingent on two critical resources: massive amounts of data for training and powerful, scalable compute for real-time analysis. This is where the synergy with cloud platforms becomes essential.
Modern AI and machine learning (ML) models for fraud detection require access to diverse, high-volume datasets to learn the nuances of both legitimate and fraudulent behavior. Cloud platforms provide the ideal environment for this, offering elastically scalable storage for data lakes, high-performance computing resources (including specialized hardware like GPUs), and robust networking to handle real-time data streams.[32]

This combination is particularly critical for CIAM. A retail business, for example, may experience a 100-fold increase in login and transaction volume during a holiday sale. A cloud-native CIAM platform can dynamically scale its resources to handle this surge without any degradation in performance, ensuring a smooth customer experience during peak traffic. An on-premise solution with fixed capacity would likely fail under such a load, resulting in lost revenue and customer frustration.[32] This synergy enables a crucial operational shift from slow, batch-based analysis—where fraud might only be discovered hours or days after the fact—to real-time, sub-millisecond decisioning. This allows fraudulent transactions to be identified and blocked *before* they are completed, preventing financial loss and protecting the customer.[29]

The adoption of cloud infrastructure is, therefore, no longer simply a matter of IT strategy focused on cost or agility; it has become a strategic imperative for a viable security posture. An organization's ability to defend against sophisticated, AI-powered threats is now directly proportional to its ability to leverage the data aggregation and computational power of the cloud. In the modern era, a "cloud-first" strategy is synonymous with a "security-first" strategy.

2.3 From Reactive to Predictive: The Power of Data Intelligence

The most significant advantage of an AI-driven, cloud-scale defense is the ability to move beyond a purely reactive posture. Traditional security systems are designed to respond to known threats—they rely on signatures, rules, and blacklists of previously identified malicious activity. This approach is inherently backward-looking and ineffective against novel, zero-day attacks.

AI-powered defense, in contrast, enables a predictive and proactive stance. By ingesting and analyzing vast, unified datasets that combine structured transaction data with unstructured information like user behavior and device telemetry, ML models can identify non-linear patterns and subtle anomalies that signal emerging fraud schemes.[29] This allows the system to detect novel attack methods without having a pre-existing signature.

This capability extends to anticipating future threats. Advanced AI systems can model how existing fraud techniques might evolve and then use generative AI to create high-fidelity synthetic data representing these future attack vectors. This synthetic data can be used to train and test fraud detection models, effectively inoculating the system against new threats before they are ever encountered in the wild.[30] This shift from reactive detection to proactive, predictive protection is the ultimate goal of a modern, AI-driven security strategy.

This transformation also brings about a fundamental economic inversion in security operations. Historically, the effectiveness of a SOC was tied to the number of human analysts it employed—a linear and expensive scaling model plagued by high alert volumes and analyst burnout.[13] AI-powered defense platforms disrupt this model by automating up to 90% of routine security tasks, such as alert triage, event correlation, and initial investigation.[13] This automation dramatically reduces the manual workload and allows human analysts to be elevated from ticket-closers to strategic threat hunters, model trainers, and incident commanders.[32] Consequently, a smaller, more highly-skilled team augmented by AI can achieve a greater level of security effectiveness than a much larger team operating with legacy tools. The value proposition shifts from scaling human labor to scaling human expertise through AI, fundamentally changing the economics and efficacy of enterprise security.

Section 3: AI in Action: Practical Applications in Fraud Detection

Moving from strategic principles to tactical implementation, this section examines the specific AI and machine learning technologies that form the core of modern CIAM and fraud prevention systems. These technologies work in concert to create a multi-layered, intelligent defense that can distinguish between legitimate customers and sophisticated attackers in real-time.

3.1 Anomaly Detection: The Digital Immune System

At its core, anomaly detection serves as the foundational layer of AI-driven defense, functioning much like a digital immune system. The underlying principle is to first establish a dynamic, individualized baseline of "normal" behavior for each user, and then to identify and flag any statistically significant deviations from that baseline.[17] This approach is exceptionally powerful because it does not rely on pre-defined signatures of known attacks; instead, it focuses on identifying any activity that is out of the ordinary for a specific user, making it effective against novel and evolving threats.

Real-world applications of anomaly detection are widespread and integral to modern security:

- **Contextual and Global Anomalies:** Systems can identify contextual anomalies, such as a login attempt from Russia at 3:00 AM for a user who consistently logs in from New York during standard business hours.[42] They can also detect global anomalies, like a $10,000 credit card transaction for a customer whose average monthly spending is $2,000.[43] These deviations from established patterns immediately raise a red flag.
- **Real-Time Transaction Monitoring:** In the financial and e-commerce sectors, AI algorithms analyze multiple variables for every transaction in real-time, including the amount, geographic location, merchant category, frequency, and velocity (the rate of transactions over time). This data is used to calculate a risk score that determines whether the transaction should be approved, challenged, or blocked.[1]

- **Access Pattern Analysis:** Anomaly detection is also applied to monitor user access patterns within corporate or cloud environments. For instance, if a marketing employee whose role involves accessing customer relationship management (CRM) and analytics platforms suddenly attempts to access a source code repository or a financial database, the system would flag this as a severe anomaly. This is a critical control for detecting both external account takeover and malicious insider threats.[42]

3.2 Behavioral Biometrics: You Are How You Act

Behavioral biometrics represents a significant evolution in authentication, moving beyond static credentials. It complements traditional factors like "what you know" (passwords), "what you have" (tokens), and physiological biometrics like "what you are" (fingerprints), by introducing a dynamic, continuous factor based on "how you behave".[38] This technology creates a unique, living profile of a user's interaction patterns, which is exceptionally difficult for an attacker—even one with stolen credentials—to replicate.

The key modalities used to build these behavioral profiles include:

- **Keystroke Dynamics:** This involves analyzing the unique rhythm and cadence of a user's typing. The system measures metrics such as typing speed, the time between consecutive key presses (flight time), and the duration a key is held down (dwell time).[47]
- **Mouse and Touchscreen Dynamics:** The way a user interacts with a mouse or touchscreen is also highly individualized. Algorithms track patterns in mouse movement, speed, acceleration, click duration, and scrolling behavior. On mobile devices, they analyze swipe gestures, touch pressure, and the angle at which the device is held.[47]
- **Gait Analysis:** Using the accelerometers and gyroscopes present in modern smartphones, systems can even analyze a user's unique walking pattern, or gait, as a passive authentication factor.[46]

The primary application of this technology is to enable continuous authentication. Unlike traditional methods that verify a user only at the initial login, behavioral biometrics monitors these patterns passively and continuously throughout a user's entire session. If the system detects a significant deviation from the established profile—for example, if the fluid, learned mouse movements of a human user are suddenly replaced by the robotic, linear movements of a bot, or if the typing rhythm changes drastically—it can flag the session as high-risk in real-time. This allows for immediate intervention, such as forcing a step-up authentication or terminating the session, without ever interrupting a legitimate user whose behavior remains consistent.[46]

3.3 Adaptive Authentication and Intelligent Friction

Adaptive authentication is the mechanism that translates the risk intelligence gathered from anomaly detection and behavioral biometrics into concrete, real-time security actions. It represents the end of the one-size-fits-all approach to security, where every user is subjected to the same level of friction regardless of context. Instead, adaptive authentication applies security controls dynamically, tailoring the level of friction to the assessed risk of each specific action.[1]

The risk-based workflow operates as follows:

1. A user initiates an action, such as logging in, transferring funds, or accessing sensitive data.
2. In milliseconds, the AI risk engine analyzes hundreds of contextual signals. These can include the user's device fingerprint (is it a known, trusted device?), IP reputation and geolocation (is the connection coming from a suspicious location?), the time of day, the user's historical behavior, and the nature of the transaction itself.[52]
3. Based on this multi-faceted analysis, the engine calculates a real-time risk score.[36]

4. This score is then compared against pre-defined organizational policies to trigger an automated response:
    - Low Risk: The action is deemed safe. The user is granted access seamlessly, experiencing a completely frictionless process.
    - Medium Risk: The action carries some element of risk. The system introduces “intelligent friction” by prompting the user for a step-up authentication, such as a push notification via a mobile app or a biometric check.
    - High Risk: The action is flagged as highly suspicious. The user is blocked from completing the action, the session may be terminated, and a high-priority alert is sent to the security team for investigation.

This dynamic, risk-based approach is the key to resolving the central tension in CIAM between providing robust security and delivering a smooth, user-friendly experience. It ensures that security measures are proportional to the actual risk, making the system both more secure and less intrusive.

The table below summarizes these core AI/ML techniques and their roles in a modern fraud detection framework.

Technique	Description	Key Data Inputs	Primary Use Case	Benefit
Anomaly Detection	Establishes a baseline of normal user behavior and identifies statistically significant deviations.	Login times, IP address, geolocation, transaction data, device information.	Detecting account takeover, payment fraud, and insider threats.	Flags novel and suspicious events that pre-defined rules would miss.
Behavioral Biometrics	Continuously verifies a user's identity based on their unique patterns of interaction with a device.	Keystroke dynamics, mouse movements, touchscreen gestures, gait analysis.	Continuous authentication post-login, bot detection, session hijacking prevention.	Provides a powerful, invisible layer of security throughout a user's session.
Adaptive Authentication	Dynamically adjusts the required level of security verification based on a real-time risk score.	Calculated risk score, device trust status, user context, policy rules.	Balancing user experience with security.	Applies friction only when necessary, creating a seamless experience for legitimate users.
Network Analysis (Graph ML)	Uncovers hidden, non-obvious relationships between different entities (users, devices, accounts) to identify coordinated fraud rings.	User accounts, devices, transaction IDs, IP addresses, email domains.	Detecting organized crime, synthetic identity fraud, and money laundering networks.	Identifies large-scale, coordinated fraudulent activity that appears as isolated incidents.

The implementation of these technologies represents a profound evolution in the concept of security. The most significant shift is the move away from a static, moment-in-time security posture to a continuous, session-aware model. Traditional security architecture primarily focuses on verifying a user at the “front door” during the initial login and then largely trusts that session until it expires.[50] This creates a critical vulnerability to attacks like session hijacking or instances where a legitimate user leaves a workstation unlocked. By continuously monitoring behavioral patterns during the session, these AI technologies effectively extend the authentication event across the entire user journey.[46] This transforms security from a single, static checkpoint into a dynamic, ongoing process, dramatically shrinking the window of opportunity for any attacker who manages to compromise an active session.

Ultimately, these applications are creating a new, much richer definition of digital identity. An identity is no longer merely a username and password, a hardware token, or even a fingerprint. It is a complex, multi-faceted, and dynamic profile composed of a user’s habits, rhythms, devices, locations, and relationships over time. This “behavioral identity” is far more robust and orders of magnitude more difficult for an attacker to steal or spoof than any single, static credential. AI is not just verifying a credential; it is continuously verifying the consistency of a complex, living identity pattern.

Section 4: The Proactive Ecosystem: Unifying Defenses with the Shared Signals Framework

While the AI-powered tools discussed in the previous section represent a monumental leap forward in fraud detection and identity verification, their effectiveness can be severely limited if they operate in isolation. A modern enterprise security stack is a complex ecosystem of disparate tools, and without a mechanism for them to communicate and act in concert, critical threat intelligence remains trapped in silos. This section introduces the OpenID Shared Signals Framework (SSF) as the essential standard for creating a collaborative, proactive, and automated security ecosystem that can respond to threats at machine speed.

4.1 The Problem of Security Silos

A typical enterprise environment is protected by a multitude of best-of-breed security solutions. An Endpoint Detection and Response (EDR) agent monitors laptops for malware, an Identity Provider (IDP) manages authentication, a network security solution inspects traffic, and a cloud security platform monitors SaaS applications. Each of these tools generates valuable security signals, but they rarely speak the same language.[56]

This lack of interoperability creates a dangerous latency gap between threat detection and response. For example, if an EDR agent detects malware on a user's laptop, that critical piece of risk information is not automatically relayed to the IDP. The IDP, unaware of the device's compromised state, may continue to grant that user access to sensitive applications. This forces security analysts to engage in "swivel-chair integration"—manually correlating alerts from one system and then pivoting to another to take action. This manual process is slow, error-prone, and completely inadequate for countering automated, high-velocity attacks.[59]

4.2 Introducing the Shared Signals Framework (SSF)

The Shared Signals Framework (SSF), an open standard developed by the OpenID Foundation, is designed to solve this exact problem. It provides a standardized, secure, and privacy-preserving protocol for different systems and services to share security-related events in near real-time.60 In essence, SSF acts as a common language or a "central nervous system" for the entire security stack, allowing any tool to communicate a change in risk posture to every other tool that needs to know.[56]

The framework operates on a simple yet powerful publish/subscribe model with three core components:

- **Transmitter (Publisher):** Any entity that observes a security event and publishes a signal. This could be an EDR platform like CrowdStrike detecting a compromised device or an IDP like Ping Identity detecting a password spray attack.[56]

- **Receiver (Subscriber):** Any entity that subscribes to these signals and is configured to take an action based on them. This could be a critical SaaS application terminating a session or a Zero Trust Network Access (ZTNA) platform blocking a user's network access.[56]

- **Security Event Token (SET):** The signal itself. A SET is a specialized, cryptographically signed JSON Web Token (JWT) that contains standardized information about the security event, such as who it happened to, what happened, and when it happened.[56]

4.3 Key Protocols in Action

SSF is the underlying framework that enables the transmission of specific types of events, which are defined in profiles. For CIAM and fraud prevention, two of the most critical profiles are:

- **Continuous Access Evaluation Profile (CAEP):** This profile is focused on communicating real-time changes related to a user's session or the state of their device. Key CAEP events include session-revoked, token-claims-change (e.g., a user's role has changed), and device-compliance-change (e.g., a device is no longer compliant with security policy). CAEP is the mechanism that allows for the near-instantaneous termination of active sessions across all federated applications when a threat is detected anywhere in the ecosystem.[57]

- **Risk & Incident Sharing and Coordination (RISC):** This profile is designed for sharing information about high-risk events related to a user's account itself. Important RISC events include credential-compromise (a user's password has been found in a breach), account-disabled, and recovery-activated (a user has initiated a password reset). RISC is a powerful tool for proactively preventing account takeover and mitigating the impact of credential stuffing attacks.[57]

4.4 Use Case in Practice: From Detection to Remediation in Milliseconds

The true power of SSF is best illustrated through a practical, multi-vendor scenario. Consider the real-world interoperability demonstration conducted by AppOmni, Cisco Duo, and SGNL, which showcases the framework's ability to automate response across a distributed security environment.[59]

The automated workflow unfolds as follows:

- **Detection (Transmitter 1):** AppOmni, a SaaS Security Posture Management (SSPM) platform, is monitoring a user's activity within a critical cloud application. It detects anomalous behavior indicative of a session hijacking attack.
- **Signal Transmission:** AppOmni immediately generates and transmits a CAEP session-revoked SET to its configured receiver, Cisco Duo.

- **Initial Action (Receiver 1):** Duo, acting as the identity and authentication provider, receives the signal. It instantly validates the SET and takes action by revoking its own session token for that user. This effectively logs the user out and would require them to re-authenticate, potentially with a stronger, step-up method, to regain access.
- **Signal Propagation (Transmitter 2):** Recognizing the severity of the event, Duo then acts as a transmitter itself. It generates a RISC credential-compromise SET and broadcasts it to other subscribed services in the ecosystem, including SGNL.
- **Comprehensive Remediation (Receiver 2):** SGNL, a dynamic access management platform, receives the signal from Duo. It immediately enforces policies to terminate the user's active sessions across all other connected SaaS applications, such as Salesforce, Microsoft 365, and Workday.

This entire chain of detection, signaling, and remediation occurs automatically, in near-real-time, without any human intervention. It shrinks the attacker's window of opportunity from potentially hours (the time it would take for a human analyst to correlate alerts and manually revoke access) to mere seconds. This is the power of a proactive, interconnected ecosystem.[56]

The table below details some of the key event types within the SSF, making the standard more tangible.

Profile	Event Type URI	Description	Example Scenario
CAEP	.../secevent/caep/event-type/session-revoked	Informs receivers that a specific user session has been terminated by the transmitter.	An administrator manually revokes a user's active session from the IDP dashboard due to a suspicious helpdesk call.
CAEP	.../secevent/caep/event-type/device-compliance-change	Notifies receivers that a device's security posture has changed (e.g., from compliant to non-compliant).	An EDR tool detects malware on a user's laptop and updates its status, triggering a signal to the IDP to block access from that device.
RISC	.../secevent/risc/event-type/credential-compromise	Alerts receivers that a user's credential is known to be compromised and should not be trusted.	A threat intelligence feed informs the IDP that a user's corporate password has appeared in a newly discovered breach dump.
RISC	.../secevent/risc/event-type/account-disabled	Communicates that a user's account has been disabled by the transmitter, often due to security concerns.	An AI-powered fraud detection engine observes a high-velocity series of impossible travel logins and automatically locks the account.

The adoption of SSF is the architectural key to operationalizing the principles of a Zero Trust security model. Zero Trust is predicated on the idea of "never trust, always verify," which requires continuous, real-time verification of trust at every single access request.[4] This verification, in turn, requires real-time data about the user, their device, their location, and the broader context of their request. This data originates from the multiple, disparate security tools across the enterprise.[58] SSF provides the critical, standardized communication bus—the plumbing—that connects these various signal sources (Policy Information Points) to the decision-maker (Policy Enforcement Point). It is the mechanism that makes the theoretical ideal of "continuous verification" a practical reality at enterprise scale.

Furthermore, the widespread adoption of SSF is poised to drive a significant shift in the cybersecurity market. For years, large security vendors have benefited from creating proprietary, "walled garden" ecosystems, incentivizing customers to purchase their entire suite of products to achieve better integration.[57] SSF dismantles these walls by creating a universal standard for interoperability.[61] This empowers customers to build a truly best-of-breed security stack, selecting the top solution for each function (endpoint, identity, network) with the confidence that they will work together seamlessly.[58] This shifts the competitive landscape away from vendor lock-in and toward a model where vendors must compete on the quality of the signals their platforms generate and the intelligence of their response actions. This ultimately benefits the customer, who gains greater flexibility, a more effective integrated defense, and a higher return on their existing security investments.[58]

Building Your Proactive Ecosystem

Implementing the Shared Signals Framework (SSF) and eliminating security silos is the next frontier of Zero Trust. At **Indigo Consulting**, we help organizations transition from isolated security tools to a unified, AI-driven identity architecture. We specialize in designing the "connective tissue" that allows your stack to respond to threats in milliseconds.

Contact Indigo Consulting to discuss how we can help you unify your IAM strategy.

Section 5: The Perfect IAM Program: Balancing Security and User Experience

The central and most persistent challenge in Customer Identity and Access Management (CIAM) is navigating the inherent tension between implementing robust security measures and delivering a frictionless user experience. This section addresses this core dilemma, demonstrating how the AI-powered technologies previously discussed are not just security tools, but are the essential mechanisms for achieving this delicate balance and creating an IAM program that is both formidable to attackers and transparent to legitimate customers.

5.1 The Core CIAM Dilemma: Security vs. Convenience

Unlike traditional workforce IAM, which governs a captive audience of employees, CIAM is designed for external customers. For this demographic, the quality of the digital experience is not just a convenience—it is a primary driver of engagement, loyalty, and revenue. Research consistently shows that customers have a low tolerance for friction; a significant percentage will abandon a service or an online shopping cart after just a single bad experience.[54]

Excessive security friction—manifested as complex password creation rules, mandatory Multi-Factor Authentication (MFA) on every login, convoluted account recovery processes, and repetitive identity checks—directly and negatively impacts key business metrics. It frustrates users, leads to higher rates of cart abandonment, and can ultimately drive customers to competitors with a smoother experience.[52]

Conversely, an under-emphasis on security can have even more catastrophic consequences. Weak security controls lead to account takeover fraud, financial loss, and the theft of sensitive personal data. A security breach erodes the most valuable asset a brand possesses: customer trust. Once lost, this trust is incredibly difficult to regain, leading to long-term brand damage and customer churn.[54] This creates the fundamental tension that every CIAM program must resolve: how to implement security that is strong enough to protect customers and the business without creating an experience that is so cumbersome it drives them away.

5.2 The Solution: "Invisible Security" and Frictionless Authentication

The key to resolving this dilemma lies in a strategic shift towards "invisible security"—the practice of embedding powerful security controls that operate in the background, having little to no perceptible impact on legitimate users during normal, low-risk interactions.[68] This is made possible by the AI-driven technologies that enable a deep, real-time understanding of user behavior and context.

The practical implementation of invisible security is achieved through a suite of frictionless authentication methods:

- **Passwordless Authentication:** The single greatest source of both user friction and security risk is the traditional password. Passwords are hard to remember, frequently reused, and the primary target of phishing and credential stuffing attacks. Adopting passwordless methods, such as device-bound passkeys that leverage biometrics like Apple's Face ID or Windows Hello, simultaneously enhances security and dramatically improves the user experience. The login process becomes as simple and secure as unlocking a phone.[4]
- **Continuous Authentication:** As detailed previously, behavioral biometrics provide a powerful layer of invisible security that functions after the initial login. By passively monitoring a user's unique interaction patterns, the system can continuously verify their identity throughout a session without requiring any active input. This protects against session hijacking and other post-authentication threats in a completely frictionless manner.[46]
- **Risk-Based Authentication:** This is the core engine that powers the concept of intelligent friction. Instead of treating all actions equally, it uses AI to assess the risk of each interaction in real-time. This ensures that the vast majority of legitimate, low-risk user activities—which may constitute 95% or more of all traffic—are allowed to proceed seamlessly. The "friction" of a step-up challenge, like an MFA prompt, is reserved only for the small percentage of actions that are genuinely anomalous or suspicious. This targeted application of security is the essence of a modern, user-centric approach.[52]

5.3 Achieving the Win-Win: A User-Centric Security Model

The ideal CIAM program is one that achieves a "win-win" state: it delights customers while satisfying the rigorous demands of security and risk teams. It enhances the user experience by simplifying the most common user journeys, such as registration and login, through features like social sign-on or passwordless authentication, thereby reducing friction and abandonment.[67]

Simultaneously, this modern approach provides security teams with a far more robust and adaptive defense than traditional, static models. By leveraging multiple layers of intelligence—device reputation, behavioral analysis, real-time risk scoring—it creates a security posture that is resilient against the evolving threat landscape.[74] This successful alignment of what were once seen as competing goals—customer satisfaction and enterprise security—is the definitive hallmark of a mature, AI-driven CIAM strategy.

This paradigm shift reframes the role of security within the business. Traditionally, security has often been perceived as the "Department of No," a necessary function that introduces cumbersome controls that can impede business agility and degrade the customer experience. The deeper truth of the modern CIAM era is that AI-driven, user-centric security inverts this relationship. A secure, yet frictionless, identity experience is no longer a security tax; it has become a powerful competitive differentiator that actively builds customer trust, increases conversion rates, and drives top-line business growth. Security becomes a feature that attracts and retains customers, rather than an obstacle they must overcome.[69]

This also leads to a more nuanced and intelligent understanding of the concept of "friction." The old paradigm viewed all friction as inherently negative. The new, AI-enabled model distinguishes between "bad friction" (unnecessary, universal obstacles applied to all users, like mandatory MFA on every login) and "good friction" (intelligent, targeted challenges applied only in response to a specific, contextual risk). The goal of a modern CIAM program is not to achieve zero friction, but to achieve zero unintelligent friction. The AI-powered risk engine provides the critical intelligence to distinguish between the two, applying security controls only when and where they are truly needed, thus preserving a seamless experience for the trusted majority.[68]

Section 6: The North Star: Future Trends and Strategic Vision

As organizations navigate the complexities of the current identity landscape, it is imperative to look toward the future and align strategies with the emerging trends and visionary concepts that will define the next era of CIAM and fraud prevention. This final section synthesizes the report's findings through the lens of industry thought leadership, explores the dual-use nature of Generative AI, and outlines the strategic path forward.

6.1 Lessons from Identiverse: Andre Durand's Vision for the Future of Identity

The annual Identiverse conference serves as a barometer for the identity industry, and the keynote addresses from leaders like Andre Durand provide a "North Star" for future strategy. His recent vision highlights several critical themes that will shape the years to come.

- **From Assumed Trust to Verified Trust:** The central tenet of Durand's message is a direct response to the AI-driven threat landscape. In a world where deepfakes can create perfect audio-visual impersonations, the foundational human assumption that "seeing is believing" is broken. This necessitates a paradigm shift in our security models away from "trust but verify" and toward a more rigorous default state of "verify, then trust".[5] This concept reinforces the absolute necessity for the continuous, contextual, and often invisible verification methods discussed throughout this report. Trust can no longer be a static state granted at login; it must be a dynamic attribute that is continuously earned and re-verified at every step of the digital journey.
- **The Age of Agentic Automation:** A critical and imminent challenge is the proliferation of non-human identities. In the near future, users will increasingly delegate tasks to personal AI agents, assistants, and bots that will interact with digital services on their behalf, primarily through APIs rather than graphical user interfaces.[5] This trend fundamentally disintermediates the direct relationship between a human customer and a brand, reducing engagement to a series of programmatic calls.

It also creates a profound new security challenge: how does an organization authenticate, authorize, and securely manage entire fleets of bots that are acting with the authority of a human user? The authority of a human user? This requires a re-architecting of CIAM to evolve into what might be better termed "Entity and Agent Management," capable of governing the complex one-to-many relationships between a human principal and their digital deputi es.[77]

- **The Call to Action:** Becoming Guardians of Authenticity: Durand's vision culminates in a powerful call to action for the identity industry. The role of the identity professional is evolving from that of a technical administrator or gatekeeper into a far more profound societal function: a "Guardian of Authenticity." In an information ecosystem polluted by AI-generated misinformation and impersonation, the identity and security community is tasked with building and maintaining the bedrock of verified trust upon which the entire digital economy and civil discourse depend.[5]

6.2 The Duality of Generative AI: Fraud Copilot vs. Defender's Assistant

Generative AI stands as the ultimate dual-use technology in the identity space, offering powerful new capabilities to both attackers and defenders.

- **GenAI as a "Fraud Copilot":** As established in Section 1, malicious actors are already leveraging LLMs as a powerful assistant to augment their capabilities. They use these tools to generate flawless phishing text, write polymorphic malware code that evades signature-based detection, and create convincing scripts for complex social engineering scams. GenAI acts as a force multiplier, enhancing the efficiency and effectiveness of their fraudulent operations.[13]

- **GenAI as a "Defender's Copilot":** Conversely, security teams are beginning to harness the same technology for defensive purposes, transforming how fraud is investigated and prevented:

- AI-Powered Investigation: For a human fraud analyst, GenAI can serve as a powerful investigative assistant. It can summarize complex security alerts in natural language, query vast and disparate log sources through simple conversational prompts, identify correlations between events that a human might miss, and even suggest appropriate remediation steps based on organizational playbooks. This dramatically accelerates the investigation and response lifecycle.[37]
- Synthetic Data Generation: One of the biggest challenges in training fraud detection models is the scarcity of high-quality fraud data. GenAI can solve this problem by creating vast quantities of high-fidelity synthetic data that realistically mimics novel and emerging attack patterns. This allows organizations to train and validate their ML models on a much richer and more diverse dataset without using or compromising real customer data, leading to more robust and predictive detection capabilities.[30]
- Proactive Threat Modeling: Defensive AI can be used to simulate attacker behavior and model potential attack paths through an organization's systems. By thinking like an attacker, these AI models can proactively identify and highlight potential vulnerabilities and security gaps before they can be exploited by real adversaries.[13]

This evolution suggests a future where the primary role of the human security analyst shifts from being a manual, reactive investigator to becoming a proactive "AI trainer." Their deep domain expertise will be used not to solve one case at a time, but to curate high-quality training data, design realistic synthetic fraud scenarios, and provide critical feedback to fine-tune the defensive AI models. This elevates the analyst's role, allowing them to scale their individual expertise across the entire organization through the AI system they manage and improve.data privacy regulations like the GDPR and CCPA will require robust governance over how AI models are trained and deployed, ensuring that customer data is protected at all stages[26]

6.3 The Road Ahead: Decentralization, Privacy, and the Future of Verification

Looking further ahead, several key trends will continue to shape the identity landscape:

- **Decentralized Identity:** There is a growing movement towards decentralized identity models and verifiable credentials. In this paradigm, users take control of their own identity data, storing it in a personal digital wallet. They can then present cryptographically verifiable "claims" (e.g., "I am over 21") to a relying party without having to share the underlying sensitive data (e.g., their full date of birth). AI will play a crucial role in verifying the authenticity of these claims while preserving user privacy.[5]

- **Ethical AI and Privacy:** As AI becomes more deeply embedded in security and identity systems, the focus on ethical considerations will intensify. Addressing challenges such as algorithmic bias in facial recognition or risk scoring models will be paramount to ensure that AI-powered systems are fair, transparent, and equitable. Furthermore, adherence to

Conclusion: Future-Proofing Your Identity Strategy

The digital identity landscape is undergoing a period of rapid and profound transformation, driven by the dual forces of AI-powered attacks and AI-powered defense. The key takeaways from this analysis are clear: the threat is real, accelerating, and fundamentally challenges our traditional notions of trust. A passive or reactive security posture is no longer sufficient. The path forward requires a strategic embrace of a new paradigm—Verified Trust—built upon a foundation of intelligent, adaptive, and interconnected technologies.

The analysis has demonstrated that the only viable response to AI-driven threats is a defense that leverages the same principles of AI and cloud scale. Practical, powerful tools such as anomaly detection, behavioral biometrics, and adaptive authentication are available today, enabling organizations to build a defense that is both robust and user-centric. Furthermore, the emergence of open standards like the Shared Signals Framework provides the blueprint for a future where security tools no longer operate in silos but as a cohesive, intelligent ecosystem capable of responding to threats in milliseconds.

Crucially, this technological evolution resolves the long-standing conflict between security and user experience. By making security invisible and applying friction intelligently, organizations can create digital experiences that are both seamless for customers and formidable for adversaries. This transforms CIAM from a necessary security cost center into a strategic business enabler that builds trust, enhances loyalty, and drives growth.

Actionable Recommendations

To translate these insights into action, professionals across different functions should consider the following strategic priorities:

- **For Security Leaders (CISOs and Security Architects):**
  - Champion a Unified Platform Strategy: Advocate for investment in a modern, unified identity platform that natively integrates AI-powered fraud detection and adaptive authentication capabilities. Move away from a collection of point solutions toward a cohesive platform that provides a single, contextual view of identity risk.
  - Drive Adoption of Open Standards: Begin planning for and demanding support for the Shared Signals Framework (SSF) from your security vendors. Fostering an interconnected ecosystem is the most effective way to maximize the return on investment (ROI) of your entire security stack and enable true, real-time response.
  - Rethink Identity Governance: Expand your governance models to include non-human identities. Develop strategies now for authenticating, authorizing, and managing the lifecycle of the AI agents and bots that will soon represent your customers.

- **For IT and Digital Transformation Leaders:**
  - Prioritize a Frictionless User Experience: In all CIAM procurement and development decisions, prioritize solutions that enable a seamless, passwordless user journey. View identity not as a security hurdle for customers to overcome, but as the front door to your digital brand and a core component of the overall customer experience.
  - Leverage Identity as a Growth Engine: Recognize that a secure and user-friendly CIAM platform is a competitive differentiator. Use the data and insights gathered from the identity platform (with appropriate consent) to personalize customer journeys, build trust, and increase engagement and retention.
- **For Fraud Prevention and Risk Management Teams:**
  - Evolve Beyond Static Rules: Transition fraud detection strategies away from brittle, static rule-based systems and toward dynamic, AI/ML-based models that can adapt to evolving threats.

- Embrace AI as an Augmentation Tool: Begin exploring how Generative AI can be used as an investigative "copilot" to augment your team's capabilities. Pilot tools that can automate alert analysis, data correlation, and reporting to free up your analysts for more strategic threat hunting and model improvement.

- Focus on Proactive Defense: Shift from a purely reactive posture of investigating past fraud to a proactive one. Leverage AI to model potential threats and use synthetic data to train your systems against the attacks of tomorrow, not just the attacks of yesterday.

Is your Identity Strategy Future-Proof?

The shift from "Assumed Trust" to **"Verified Trust"** requires a clear understanding of your current gaps. Indigo Consulting offers a **Free IAM Assessment t**o help you:

- Identify vulnerabilities in your current CIAM workflow.
- Evaluate your readiness for AI-driven threats.
- Create a roadmap for a proactive, interconnected ecosystem.

Get Your Free Assessment

By embracing these strategies, organizations can move beyond a state of constant reaction and begin to build a future-proof identity program that is resilient, intelligent, and prepared for the challenges and opportunities of the AI era.

info@indigoconsulting.ca

Works cited

1. The Impact of AI on Major Cyberattacks in 2024 | Appgate, accessed August 21, 2025, https://www.appgate.com/blog/the-impact-of-ai-on-major-cyberattacks-in-2024
2. Staying ahead of threat actors in the age of AI | Microsoft Security Blog, accessed August 21, 2025, https://www.microsoft.com/en-us/security/blog/2024/02/14/staying-ahead-of-threat-actors-in-the-age-of-ai/
3. Microsoft Digital Defense Report 2024, accessed August 21, 2025, https://www.microsoft.com/en-us/security/security-insider/threat-landscape/microsoft-digital-defense-report-2024
4. 8 Types of Identity-Based Attacks | CrowdStrike, accessed August 21, 2025, https://www.crowdstrike.com/en-us/cybersecurity-101/cyberattacks/identity-attack/
5. Identiverse 2025 Highlights | Ping Identity, accessed August 21, 2025, https://www.pingidentity.com/en/resources/blog/post/identiverse-2025-highlights.html
6. AI is helping hackers automate and customize cyberattacks - Cybersecurity Dive, accessed August 21, 2025, https://www.cybersecuritydive.com/news/ai-automate-cyber-threats-crowdstrike/756694/
7. 2025 Identity Fraud Report - Entrust, accessed August 21, 2025, https://www.entrust.com/sites/default/files/documentation/reports/2025-identity-fraud-report.pdf
8. 5 Ways Cybercriminals Are Using AI in Cybercrime in 2024 - BlinkOps, accessed August 21, 2025, https://www.blinkops.com/blog/using-ai-in-cybercrime
9. Generative AI fraud friend or foe? - RSM UK, accessed August 21, 2025, https://www.rsmuk.com/insights/advisory/generative-ai-fraud-friend-or-foe
10. 10 Ways That Cybercriminals Are Weaponizing AI - ID Agent, accessed August 21, 2025, https://www.idagent.com/blog/10-ways-that-cybercriminals-are-weaponizing-ai/
11. How Cybercriminals Are Leveraging AI to Build Better Attacks - NexusTek, accessed August 21, 2025, https://www.nexustek.com/blog/how-cybercriminals-are-leveraging-ai-to-build-better-attacks
12. FBI Warns of Increasing Threat of Cyber Criminals Utilizing Artificial Intelligence, accessed August 21, 2025, https://www.fbi.gov/contact-us/field-offices/sanfrancisco/news/fbi-warns-of-increasing-threat-of-cyber-criminals-utilizing-artificial-intelligence
13. Self-defending systems are no longer a vision, but a reality in the making, accessed August 21, 2025, https://www.dqindia.com/interview/self-defending-systems-are-no-longer-a-vision-but-a-reality-in-the-making-9678647

14. How cybercriminals are using gen AI to scale their scams - Okta, accessed August 21, 2025, https://www.okta.com/newsroom/articles/how-cybercriminals-are-using-gen-ai-to-scale-their-scams/
15. LLMs + fraud: How criminals use large language models to commit fraud - Persona, accessed August 21, 2025, https://withpersona.com/blog/llm-fraud
16. AI-Enhanced Social Engineering Will Reshape the Cyber Threat Landscape | Lawfare, accessed August 21, 2025, https://www.lawfaremedia.org/article/ai-enhanced-social-engineering-will-reshape-the-cyber-threat-landscape
17. Fraud in financial services: Leaning on generative AI to protect a rapidly expanding attack surface | Elastic Blog, accessed August 21, 2025, https://www.elastic.co/blog/financial-services-fraud-generative-ai-attack-surface
18. What is Generative AI Fraud? - Ironscales, accessed August 21, 2025, https://ironscales.com/glossary/generative-ai-fraud
19. Leveraging AI LLMs to Counter Social Engineering: A Psychological Hack-Back Strategy, accessed August 21, 2025, https://www.tripwire.com/state-of-security/leveraging-ai-llms-counter-social-engineering-psychological-hack-back-strategy
20. Identity Fraud: 3 Rising Trends To Watch in 2025 - Incode, accessed August 21, 2025, https://incode.com/blog/identity-fraud-3-rising-trends-to-watch-in-2025/
21. Deepfake Deception in Digital Identity - Identity Management Institute®, accessed August 21, 2025, https://identitymanagementinstitute.org/deepfake-deception-in-digital-identity/
22. How to Stop Deep Fake Impersonation with Advanced Identity Verification - Authenticate, accessed August 21, 2025, https://authenticate.com/resources/blog/stop-deep-fake-impersonation-with-advanced-identity-verification
23. How Deepfakes Are Undermining Biometric Identity Checks in 2025 - Mea: Digital Integrity, accessed August 21, 2025, https://www.mea-integrity.com/how-deepfakes-are-undermining-biometric-identity-checks-in-2025/
24. I Deepfaked My Boss: What It Reveals About Identity Verification in the Age of AI, accessed August 21, 2025, https://www.intellicheck.com/resource-library/i-deepfaked-my-boss-what-it-reveals-about-identity-verification-in-the-age-of-ai
25. 2025 Fraud Trends: Protecting Against Emerging Threats | FinTalk - Jack Henry, accessed August 21, 2025, https://www.jackhenry.com/fintalk/2025-fraud-trends-protecting-against-emerging-threats
26. AI-Powered Fraud Detection in CIAM - Deepak Gupta, accessed August 21, 2025, https://guptadeepak.com/customer-identity-hub/ai-powered-fraud-detection-in-ciam