WHITE PAPER

Persistent Identity Management: Infrastructure for an AI World

Connecting and continuously managing verified identities to protect organizations against today's fraud risks and tomorrow's agentic AI threats.

Prove

2026

Executive Summary

Identity Is No Longer a Checkpoint, It's a Continuous Signal
For decades, digital identity verification has operated as a gate: verify a user once at onboarding or login, then trust them throughout the session. This model was designed for a world of human-paced, human-initiated interactions. That world is disappearing.

The convergence of two forces is rendering moment-in-time identity models structurally obsolete. First, AI-driven fraud including deepfakes, synthetic identities, and automated credential attacks, has eroded the reliability of the signals on which legacy systems depend. Second, the emergence of autonomous AI agents acting on behalf of users introduces a category of interaction that legacy KYC and authentication architectures were never designed to handle.

This paper describes Prove's approach to both challenges:
A persistent, device-anchored identity layer that continuously verifies trust across every interaction, not just at entry points. Central to this architecture is the Prove ID, a durable, unique identifier that remains stable even as consumer phone numbers change. Prove ID serves as the backbone of Prove's real-time identity management capabilities.

Prove's Identity Platform operationalizes this model through a unified architecture that connects verified identities to a durable graph, authenticates them at device level, and monitors them in real time across the full customer lifecycle.

Key Thesis
Trust cannot be stamped once and assumed to hold. In an AI-driven environment, identity must become a live, continuously validated signal, not a credential stored at onboarding.

The Failure Modes of Legacy Identity Architecture

Legacy identity and fraud systems share a foundational assumption: that the entity initiating an interaction is a human, operating at human speed, verifiable through a document or knowledge-based credential. Each of these assumptions is now under systematic attack.

Velocity Overwhelms Verification

AI agents can initiate thousands of account creations, login attempts, and transactions simultaneously. Rule-based and manually-tuned systems were never designed to process this volume in real time. Systems relying on sequential checks, static thresholds, or batch processing introduce latency that fraud exploits.

The Signal-to-Noise Collapse

Traditional fraud signals such as device fingerprints, IP reputation, and behavioral heuristics depend on human behavior patterns as baselines. Generative AI can now synthesize those patterns with sufficient precision to defeat most behavioral detection systems.

Document-based verification platforms relying on computer vision are increasingly vulnerable to deepfake attacks that produce photorealistic but fabricated identity artifacts.

Point-in-Time Identity Breaks Down

Legacy systems verify identity at onboarding or login and then assume continuity. But identity is not static. Sessions can be hijacked mid-stream. AI agents can act on behalf of legitimate users in ways that shift risk profiles in real time. A user who was low-risk at login may be high-risk sixty seconds later. A system that does not continuously evaluate trust cannot detect the change until after damage is done.

Rules-Based Systems Cannot Adapt at Machine Speed

Fraud patterns driven by AI evolve in hours, not weeks. Static rule sets and traditional machine learning models require human intervention to update, and by the time a new attack pattern is detected, codified, and deployed, it has already been exploited at scale.

The Legacy KYC Gap and AI Agents

Legacy KYC was built to answer a single question: is this a real person?

It assumes a stable human identity, a one-time onboarding check, and human-scale interaction volumes. AI agents violate all three assumptions. They act across users, contexts, and sessions; they operate at machine speed and volume; and they can be granted, stolen, or spoofed access to verified human credentials.

KYC can tell you who enrolled but it cannot tell you whether the entity acting right now is authorized, low-risk, and genuinely human-backed.

The Core Problem
Legacy systems treat identity as a moment. AI turns identity into a continuous, high-speed stream. The gap between those two models is where fraud lives.

Phone Numbers as Persistent Consumer Identifiers

Phone numbers have become the preferred personal identifier for consumers. Over 90% of the 5 billion adults worldwide have a phone number, and 90% of those adults have maintained the same number for 5 to 20 years. Due to their widespread adoption and the unique bridge they create between the digital and physical realms, phone numbers are increasingly pivotal for critical services such as customer service, multi-factor authentication, identity verification, and fraud prevention.

SIM cards, inherently linked to phone numbers, bolster this security by offering robust cryptographic authentication and are typically associated with a single user. However, relying solely on phone numbers as long-term identifiers carries inherent risks. Without proper safeguards, phone numbers can change hands. These changes occur either through legitimate consumer actions (upgrades, carrier switches, moves) or malicious activities like SIM swapping. These change events jeopardize the integrity of the phone signal.

Persistent Identity

Prove's solution to the instability of raw phone numbers is the Prove ID, a unique identifier that remains constant even when phone numbers change. This allows Prove clients to maintain a reliable connection with their customers, regardless of phone number changes.

What Is a Prove ID?

At the core of the Prove platform is the Prove ID. Prove ID is a unique, persistent digital anchor for a human being. Unlike an account number or a simple username, a Prove ID is a tokenized, non-PII identifier that links users to their verified identity.

Ensuring Persistence: Managing Phone Number Changes

Approximately 120 million phone numbers are permanently disconnected in the US each year. Many disconnects occur when consumers ask their carrier to change their number, often following a data breach or to manage debt. Disconnected numbers become eligible for reassignment to new consumers. The following principles govern identity persistence:

- When a consumer changes their phone number but not their carrier account, Prove persists the Prove ID and updates it to point to the new number. Prove's Identity Established Date, or IED, persists.
- When the old phone number is reassigned to a new consumer, it receives a new Prove ID, and the IED is set to the reassignment date.
- If Prove has no evidence of a new phone number, the identifier is revoked. If the consumer opens a new carrier account with a new number, a new Prove ID and IED are assigned accordingly.

Phone Number Porting

Approximately 14 million phone numbers are disconnected from carriers each year due to porting when consumers switch carriers. In most cases, the phone number remains with the consumer and the Prove ID and IED remain intact. However, porting introduces important complexity. Porting is not a single event but a state, and original carriers may reclaim ported numbers in what are known as port deletions.

Illustrative Case Studies

The following scenarios illustrate how Prove manages identifier persistence through real-world porting situations:

Consumer A

Port and Return (Win Back): A consumer originally obtained their number from Carrier X ten years ago. Five years ago they ported to Carrier Y. Yesterday, they switched back to Carrier X for a promotional offer. While this appears to be another port, it is actually a port deletion and a "win back" — Carrier X owns the number and reclaimed it. The Prove ID and original IED remain fully intact.

Consumer B

Port and Account Closure: A second consumer also ported their number from Carrier X to Carrier Y five years ago. Yesterday, that consumer closed their Carrier Y account entirely. Once the account closes, the number "snaps back" to Carrier X, the original owner, and becomes eligible for reassignment to a new consumer. In this case, the Prove ID is revoked, and when the number is reassigned, a new Prove ID and IED will be set at the date of reassignment.

Managing False Positives

False Positives from Carrier Ports

Many disconnect notifications are misleading, implying that a disconnected phone number no longer belongs to the given individual. Since a disconnect notification is generated for each of the ~14 million account ports per year, Prove carefully differentiates permanent disconnects from false positives. False positives are suppressed and not acted upon. A port notification alone is not a sufficient signal, given the complexity of port deletions.

False Positives from Temporary Account Suspensions

Carriers also disconnect phone lines for non-payment. Because disconnect events are generated in these cases, vendors without proper suppression logic will treat them as permanent disconnects. These are effectively account suspensions, when the consumer pays their outstanding balance, the number is typically reactivated. If false positive disconnects are not suppressed, they create significant inaccuracies that erode the durability of phone numbers as consumer identifiers over time.

Permanent Disconnects and Fraud Mitigation

Prove manages these issues by suppressing false positives, generating a risk event only when a true permanent disconnect has occurred. This event signals that the identifier is considered inactive after the event timestamp.

Fraud mitigation is a key benefit of this approach. Fraudsters can claim phone numbers that have been permanently disconnected and use them to impersonate consumers at banks or other institutions.
If an organization does not know a consumer surrendered their phone number, it may reset passwords via SMS to that number, effectively handing access to the fraudster. Prove's real-time signaling prevents this scenario.

Coverage Across ~3,000 US Carriers

In addition to accuracy, Prove's carrier coverage is a key differentiator. Prove covers nearly all ~3,000 US carriers, whereas most vendors cover only the three major mobile network operators and their subsidiaries. This leaves significant blind spots, consumers can port numbers to any carrier, including non-fixed VoIP carriers that cannot receive SMS messages. By covering the full carrier ecosystem, Prove manages persistence across all ~500 million active US phone numbers, including critical events like permanent disconnects.

A New Architecture: The Persistent Identity Graph

Addressing the failure modes of legacy systems requires not just incremental improvement but an architectural shift in how identity is conceived, stored, and evaluated.

Prove's Identity Platform is built on this reconception: identity as a persistent, continuously updated graph rather than a credential verified at a point in time.

The Prove ID serves as the durable anchor within this graph, enabling the platform to maintain continuity of identity across all phone number lifecycle events.

Connect: Anchoring Identity to a Durable Graph

The first step in persistent identity management is creating a durable, tokenized representation of a verified individual. Prove connects a real person's national identifier with the Prove ID.

By anchoring identity to this graph at the point of verification rather than relying on credentials that can be duplicated or spoofed, Prove creates a foundation that is fundamentally cannot be faked: the token has no raw PII attached to it, but it carries a cryptographically verifiable link to a real, authenticated individual.

This tokenization architecture is specifically significant in AI environments. When AI agents or downstream systems need to act on identity, they reference a trusted token rather than handling raw PII. Sensitive data is not passed through every model, API, or agent in a pipeline. The token travels; the underlying identity stays protected.

Authenticate: Device-Bound Trust

Once an identity is connected to the graph, Prove silently generates an authentication key on the user's device during initial interaction, with consent. This token, or Prove Key, creates a hardware-bound link. This is not a statistical inference, not a behavioral heuristic, but a cryptographic binding between a verified identity and a specific physical device via carrier telemetry.

The significance of this approach is that it is deterministic rather than probabilistic. Competitors relying on behavioral models or document-based biometric verification are making statistical inferences. Prove establishes a physical, hardware-anchored fact, specifically resistant to deepfake attacks, since a synthesized document or manipulated video cannot spoof a carrier-verified SIM binding.

For returning users, this means silent, passive authentication. There are no OTPs, no re-uploads, no friction. This is done while maintaining cryptographic-grade assurance. The user experience improves as the security posture strengthens.

Verify: Layered Trust, Not Binary Pass/Fail

Prove's verification layer cross-references identity attributes against authoritative data sources to produce a graduated signal of how much trust can be placed in an identity at a given moment. Trust is established through the combination of verified data, authentication events, and real-time risk signals.

This multi-signal approach distinguishes persistent identity management from traditional KYC.

Rather than asking 'is this document valid?' Prove asks: 'given everything we know about this identity across its history in our graph, how confident are we that this interaction is authorized, low-risk, and genuinely human?'

Real-Time Phone Number Management with Identity Manager

The most significant architectural innovation in Prove's platform is not verification, it is what happens after verification.

Traditional systems onboard a user and then largely ignore that identity until the next interaction. Prove's Identity Manager is the operational heart of the platform's Manage process: a proactive, continuously active monitoring and management system that consolidates fragmented account data into a single trusted identity and keeps it accurate across its entire lifecycle. Prove IDs are the indexing mechanism that makes this possible.

Initial Upload and Ongoing Data Updates

Phone Number Management
Prove's Identity Manager provides data related to the ~500 million active phone numbers in the US across nearly all mobile/eSIM, landline, VoIP, and most non-fixed VoIP line types. Updates are delivered in real time as they occur, indexed by Prove.

When a consumer changes their phone number but not their phone account, the new phone number is provided in updates. The Prove ID and IED do not change. The new phone number inherits the tenure of the old one. Queries against the old identifier or old phone number will return associated history, including any permanent disconnect event.

Line Type Management
Line type is an essential attribute for a phone number. If a number ports from a mobile carrier to a landline, SMS OTP is no longer viable. Prove manages over 80 million intra-Service Provider Identification (SPID) change events across the ~3,000 US carriers.

For example, a number may appear to be an AT&T landline but, if AT&T leases it to a mobile MVNO, it is actually a mobile line type. Prove's real-time SPID change management ensures the most accurate line type for every number, well beyond the Number Portability Administration Center (NPAC) accuracy.

Identity Management
Identity Manager also provides data on ~280 million identities, representing the majority of the US adult population. These identities are provided for matching phone numbers to name and address data, with updates delivered monthly.

Role and Function of Identity Manager

Identity Manager sits within the fourth process of the platform's orchestrated flow:

Manage is where persistent trust is operationalized. Its primary functions are:

- Consolidating fragmented account data from multiple channels into a single, trusted identity view.
- Providing real-time webhook-based alerts tied to each customer's Prove ID when material lifecycle changes occur.
- Maintaining database integrity so downstream systems, fraud engines, authentication flows, customer experience layers, are always operating on accurate, current identity data.

Binding Customer IDs to Prove IDs

To ensure the ongoing relationship between Prove IDs and phone numbers, a unique customer ID should be bound with one or more identifiers.

Prove offers several binding processes:

- **Batch**: Prove can run the ~280 million identities through and append the appropriate Prove ID to each identity.
  - Verified identities are enrolled automatically via the Connect flow.
- **Real-time:** Each identity can be enrolled at the point of transaction.
  - Authenticated sessions through Unified Authentication enroll identities upon successful authentication.
  - Verification events via the v3/verify API trigger automatic enrollment.
  - Real-time standalone enrollment APIs are available for organizations enrolling existing user bases directly.

Enrollment is scalable to enterprise volumes, approximately 15 million records per week in batch mode. For organizations requiring real-time individual enrollment, the same v3 Verify API is available for direct integration.

What Identity Manager Monitors

Once an identity is enrolled, Identity Manager watches continuously for events that signal elevated risk or material change. Current live alerts include phone number changes and carrier port activity which are two of the most common precursors to account takeover. Additional lifecycle events, including name and address changes, monitored death indicator indexes, and more are on Prove's strategic roadmap.

Name and address change alerts extend this monitoring beyond fraud precursors into the compliance domain. Because Identity Manager notifies a subscriber the moment a bound identity attribute changes, an institution can move from calendar based KYC refresh to event driven ongoing due diligence, re-verifying a customer when their data actually changes rather than on a fixed review cycle. The notification is a provenance backed trigger, anchored to a verified Prove ID, that a subscriber then acts on by fetching the updated attribute and running it through its own re-verification and screening.

Every major AML regime, from the US BSA and Canada's PCMLTFA to the UK MLRs and the EU AMLR, requires regulated institutions to keep customer information current and monitor relationships on an ongoing, risk based basis. As these regimes shift toward risk-based and effectiveness-based supervision, event-driven attribute change alerts let institutions meet that duty by re-verifying when data actually changes, rather than on a fixed calendar.

The practical implication is that organizations are no longer limited to detecting fraud at the point of transaction. A customer whose phone number was ported to a new carrier in the hours before a high-value withdrawal attempt is flagged before the transaction is initiated, not discovered in a post-mortem.

Technical Architecture: Webhooks, Tokens, and Data Handling

Identity Manager is architected around real-time webhooks rather than batch polling. Webhook authentication uses JWT, no IP whitelisting required. Organizations implementing Identity Manager are expected to develop internal relational databases to store webhook notification history for auditing, compliance, and downstream analytics.

Identity proxying relies on tokenized, privacy-preserving identifiers throughout. Raw PII is not exposed through webhook payloads or monitoring APIs. The 2-hour token expiration applied during Authenticate and Verify processes ensures tightly scoped access windows in high-volume or automated workflows.

The Compounding Value of Managed Identities

Each enrolled identity generates a longitudinal record that compounds in value over time. An identity continuously monitored for twelve months with zero anomalous events carries a materially different risk profile than one enrolled yesterday. These signals inform Prove's Assurance Level capabilities which reflect this history: confidence and trust compound rather than resetting at each interaction.

This compounding dynamic also unlocks deep personalization, streamlined authentication for known low-risk users, proactive friction for anomalous sessions, in ways that are not possible with point-in-time verification systems.

Legacy Identity Model	Prove Identity Manager
Verify once at onboarding; ignore thereafter	Continuously monitor across full lifecycle
Fraud detected after the fact	Real-time pre-transaction alerting via webhook
Fragmented data across systems and channels	Consolidated into single trusted identity (Prove ID)
Credential-based, static risk assessment	Live, event-driven signals linked to device-bound identity
Friction scales with perceived risk	Passive authentication; friction decoupled from verification
Vulnerable to account takeover between sessions	Carrier port and phone change events trigger instant alerts
No lifecycle intelligence post-verification	Longitudinal identity record compounds in confidence over time
Phone number treated as static identifier	Identifier persists through number changes, ports, and reassignments

The Emerging Frontier: Agentic AI and Identity

The challenges described above – deepfake fraud, credential theft, account takeover – are present-day problems. A second, emerging challenge will require even more fundamental architectural adaptation: the rise of autonomous AI agents.

What AI Agents Introduce

Large language model-based agents are increasingly deployed to take actions on behalf of users such as executing transactions, initiating account changes, interacting with third-party services, and managing digital workflows. These agents operate autonomously, at machine speed, and across multiple sessions and contexts simultaneously.

From a human-to-agent identity architecture perspective, this creates a new category of legacy systems that are unprepared to answer important questions:

- How do we prove the identity of the originating human?
- How do we prove the agent is owned/controlled by that human?
- How do we prove the agent is acting within the explicit scope of its permissions?

The Tokenization Solution for Agent Interactions

Prove's tokenization architecture provides a foundational solution. Rather than exposing raw PII through every layer of an agentic pipeline, organizations can issue verified identity tokens that carry trust without exposing the underlying data. An AI agent operating on behalf of a verified user presents a Prove ID token that carries cryptographically-verifiable assurance of the underlying identity without requiring the agent to handle or transmit sensitive personal information.

Continuous Verification for Non-Human Actors

The deeper requirement for agentic AI is the ability to continuously validate whether an agent's actions are consistent with the scope and risk profile of the identity it represents. An agent granted access to a financial account should not be able to initiate high-value transfers after a credential compromise event that occurred three minutes ago.

Prove's real-time monitoring and webhook architecture provides the infrastructure for this continuous validation. When risk-relevant changes occur such as carrier port activity, phone number changes, or anomalous authentication patterns, those signals are propagated to downstream systems in real time, enabling organizations to dynamically adjust what AI agents are permitted to do on behalf of a given identity.

The Agentic Suite

Prove's Agentic Suite, currently enrolling design partners, is purpose-built to extend the platform's identity management capabilities to AI agent workflows. It reflects the recognition that the next wave of fraud and security challenges will not be human-initiated. The identity infrastructure required to address them must be built now, before those challenges fully materialize at scale.

Forward-Looking Observation
Organizations that build identity infrastructure capable of handling AI agents today will not only be better protected against emerging threats but they will also be positioned to deploy agentic AI capabilities with greater speed and confidence than competitors who must retrofit trust infrastructure after the fact.

Business Benefits: What Persistent Identity Management Delivers

The architectural case for persistent identity management is compelling on its own terms. The business case is equally strong.

Reduced Fraud Losses

Real-time monitoring and pre-transaction alerting enable organizations to intercept fraud before it occurs. Proactive signals, not reactive forensics.

Higher Conversion Rates

Passive, device-bound authentication eliminates friction for returning users. No OTPs. No re-uploads. Assurance improves as friction decreases.

Lower Authentication Costs

Eliminating repeated SMS OTP costs at scale represents significant savings for high-volume organizations. Prove Key authentication requires no carrier call-out once bound.

Faster Time to Value

A single platform implementation unlocks multiple solutions. Implement once; activate capabilities as needs evolve without rework or re-integration.

Regulatory Readiness

Continuous monitoring gives institutions an auditable, event-driven trail that supports the ongoing customer due diligence obligations common to every major AML regime, from the US BSA and Canada's PCMLTFA to the UK MLRs and the EU AMLR, each of which requires customer information be kept current and relationships monitored on a risk-based basis, without exposing raw PII.

Scalable Global Coverage

The Identity Graph spans 227 countries and territories. A single integration delivers consistent identity assurance across geographies, channels, and use cases.

Conclusion: From Checkpoint to Continuous Trust

The identity verification challenges organizations face today are not technical edge cases. They are structural consequences of building on an architecture designed for a world that no longer exists. Legacy systems treat identity as a moment. Fraud, and increasingly AI, treats it as a continuous stream.

At the center of Prove's architectural response is a durable thread that connects a consumer's identity across phone number changes, carrier ports, and lifecycle events. By indexing the full identity management lifecycle against tokenized identifiers rather than raw phone numbers, Prove eliminates the blind spots that fraudsters exploit and enables organizations to maintain accurate, persistent relationships with their customers regardless of underlying phone number changes.

Prove's Identity Platform represents foundational infrastructure for this shift. By connecting verified identities to a persistent graph anchored by the Prove ID, authenticating them at hardware level via carrier telemetry, and continuously monitoring them for risk-relevant changes, the platform transforms identity from a one-time gate into an always-on trust layer.

For organizations evaluating their identity infrastructure, the relevant question is not whether this architectural shift is coming. It is whether their current systems will be able to adapt when it arrives at scale. The organizations that build persistent identity management capabilities now will be better protected against today's fraud, better positioned to manage tomorrow's AI agents, and better equipped to deliver the seamless, trustworthy digital experiences their users expect.

For More Information
To learn more about Prove's Identity Platform, Identity Manager, and the Agentic Suite, visit prove.com or contact your Prove account representative.

About Prove

Prove makes identity work – verifying real people, businesses and agents in real time without friction or guesswork. Trusted by 19 of the top 20 U.S. banks and more than 2,000 leading brands worldwide, Prove helps the biggest names in banking, fintech, crypto, gaming, commerce, insurance, and healthcare grow with confidence.

Prove's identity verification and authentication solutions streamline onboarding, prevent fraud, and deliver seamless customer experiences across channels. With Prove, identity is no longer a question, it's proven.

Prove

www.prove.com