Secure Every Moment of Truth

How Identity-Centric Security Must Operate in the Age of AI

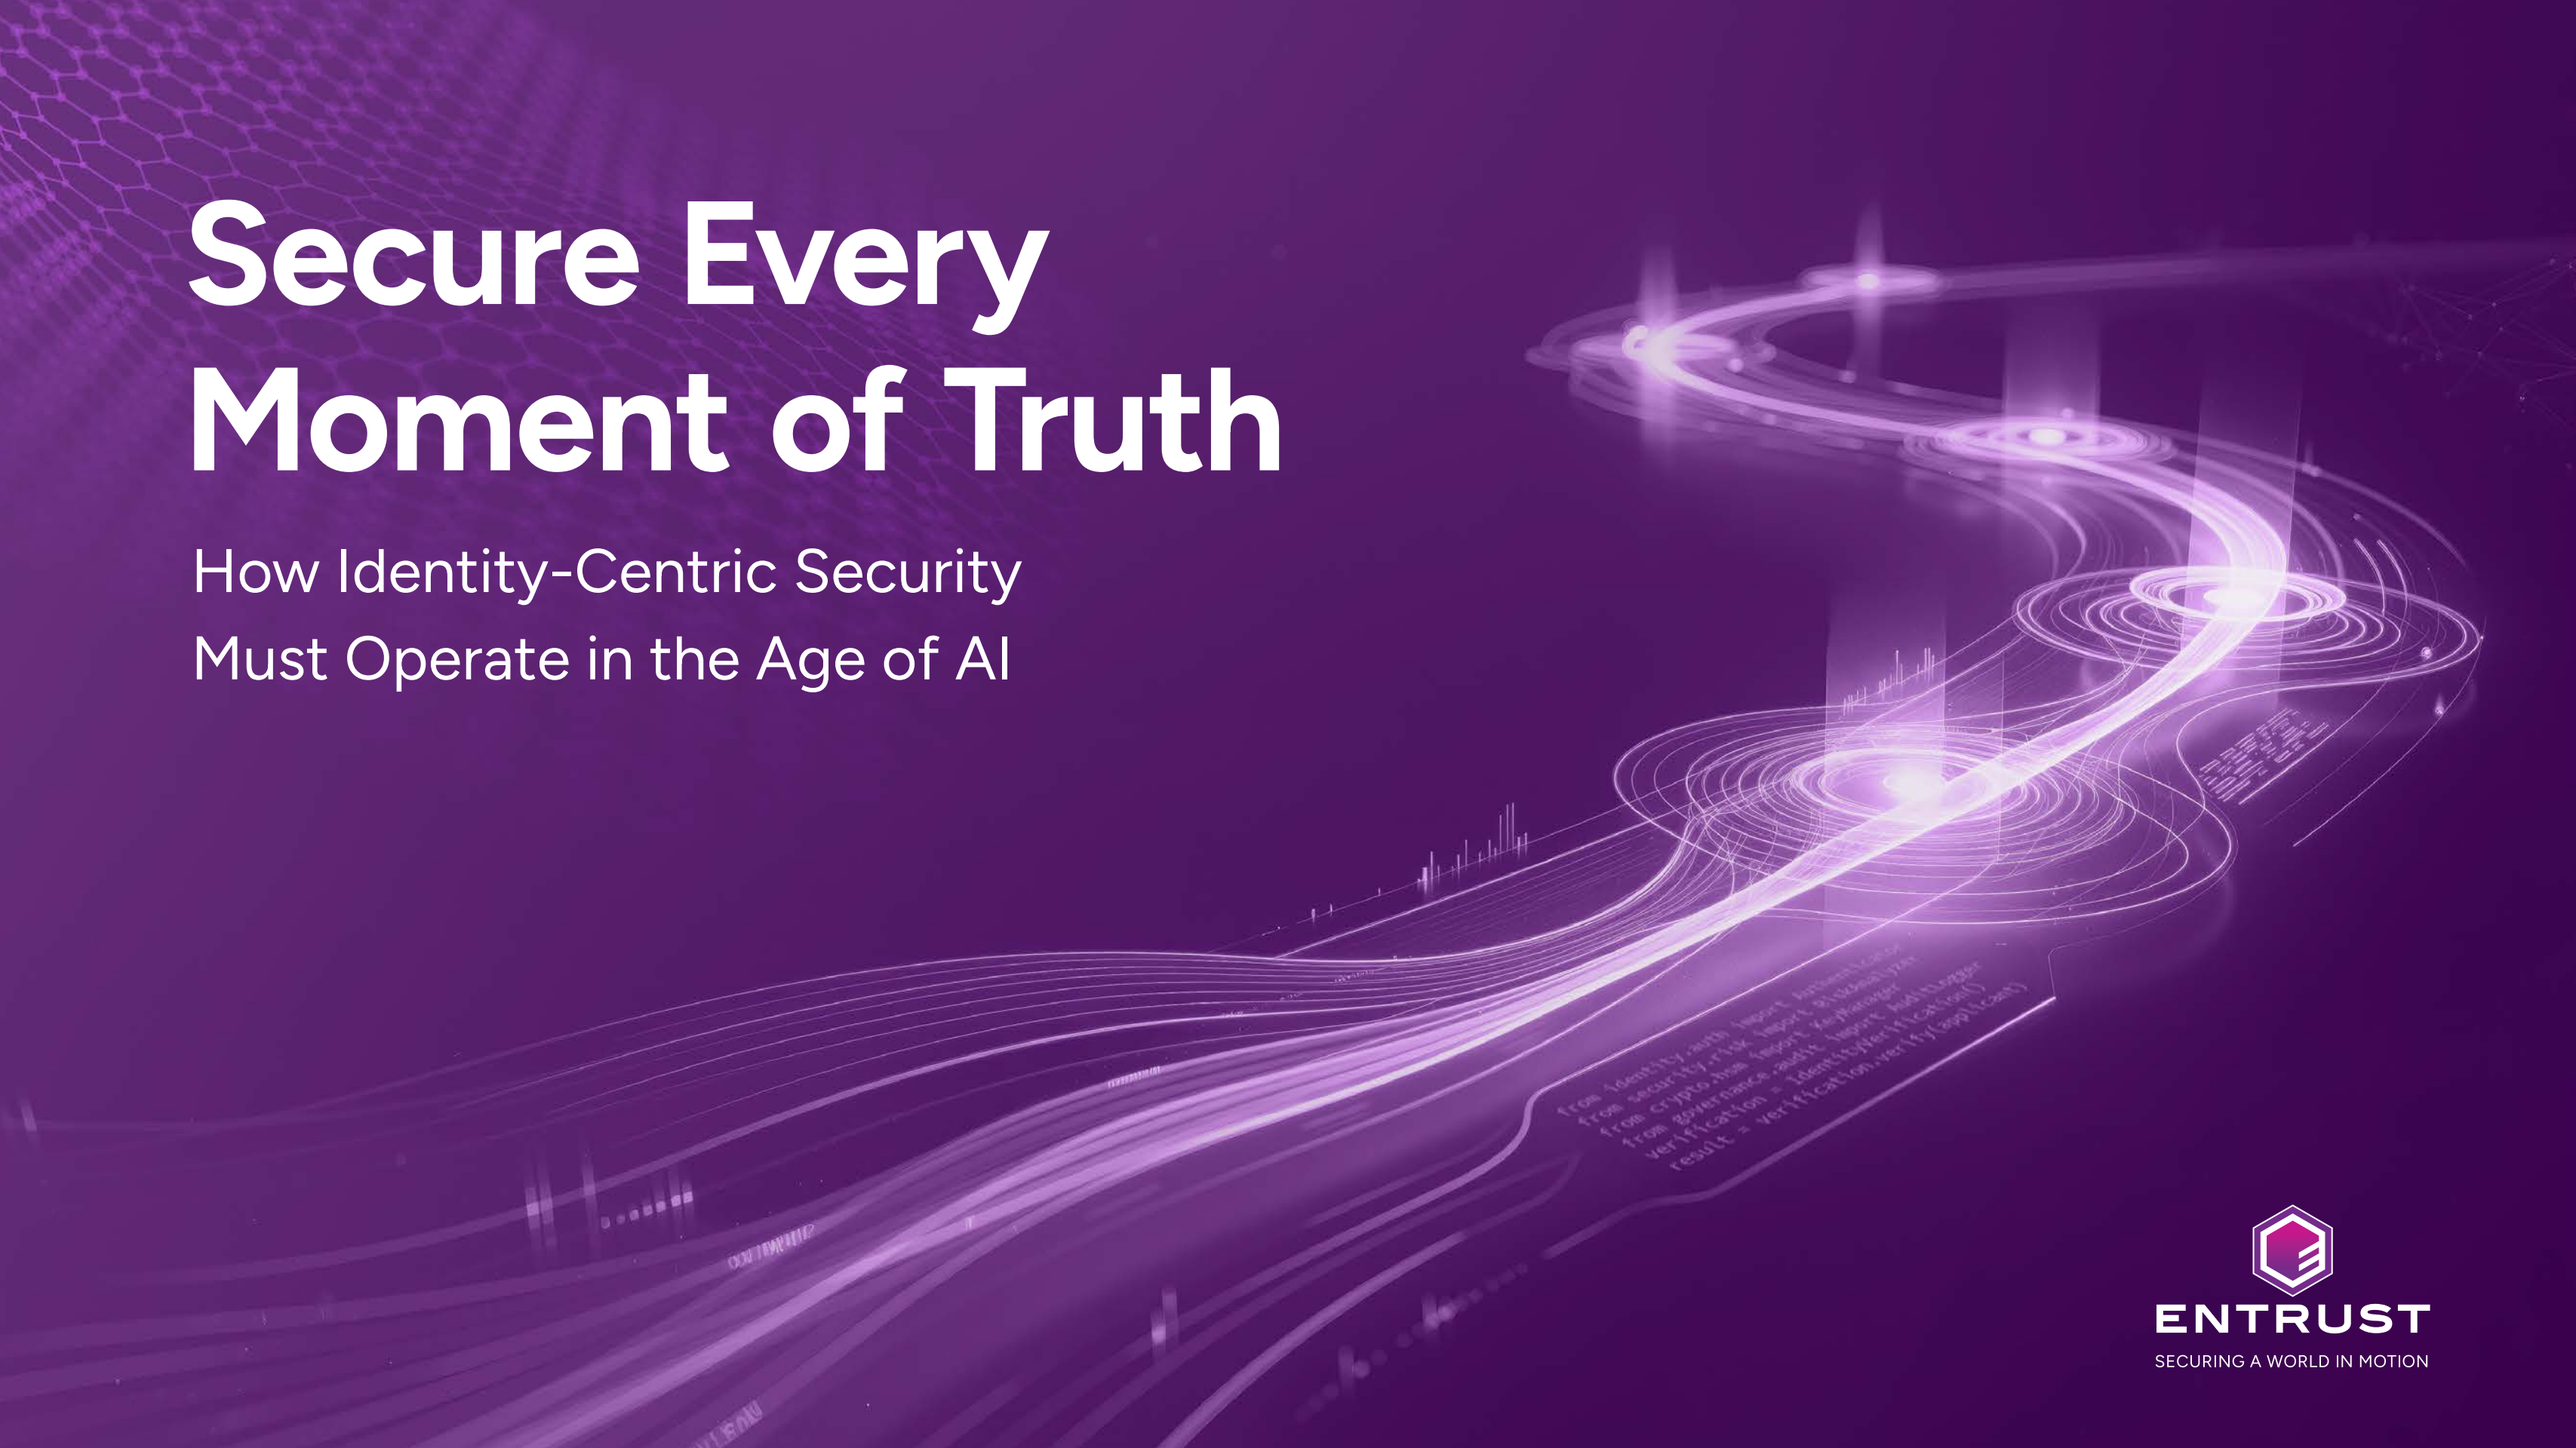

ENTRUST

SECURING A WORLD IN MOTION

EXECUTIVE SUMMARY

Trust Is the Missing Layer in AI

Today, every digital action brings both opportunity and risk. To improve security, identity must be continuously validated. Every interaction across the identity lifecycle represents a Moment of Truth – a point where trust is granted, authorized, monitored, and upheld. At each of these moments, security works to prevent disruption, reduce fraud, and protect people, devices, data, and agents – all without slowing the business and the people it serves.

Fraud is a persistent threat in today's world, with bad actors constantly developing new methods and attack paths. Entry points focus on high opportunity areas: New user onboarding, high-value transactions, and gaps in visibility, such as with missed or expired certificates. Additionally, AI-assisted scammers try relentlessly to manipulate trusted users into granting unauthorized access, risking account takeover. All of this together creates blind spots, vulnerabilities, and expands security risk.

The key is to start with verified identity – of people, devices, machines, and agents – so that every action can be traced back to a known, accountable identity, and every agent identity is bound to a human principle.

Table of Contents

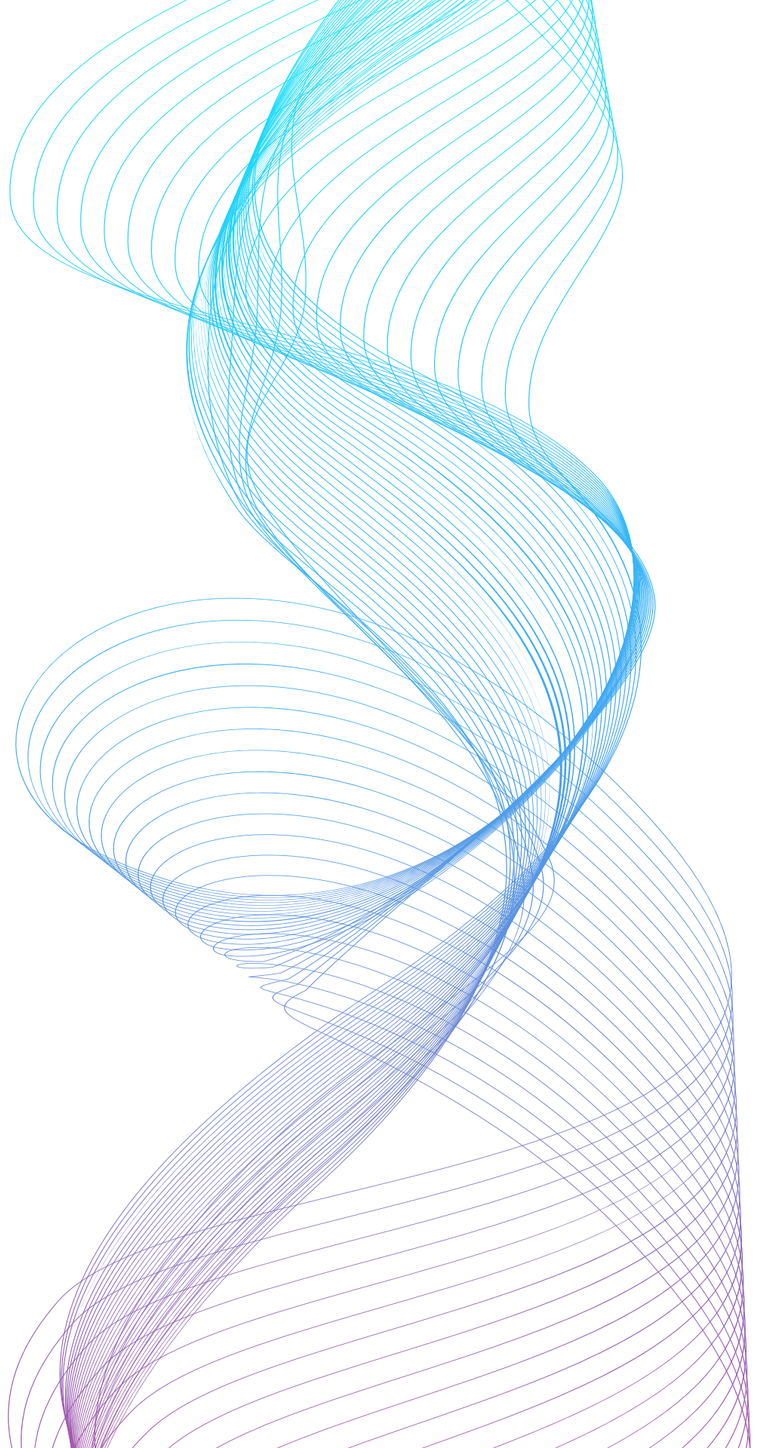

What is a Moment of Truth?

Identity moves in stages. Risk moves in moments.

Trust is a decision made, or missed, at every critical interaction across the identity lifecycle. These critical interactions are Moments of Truth: Points where identity is tested and security is either confirmed or compromised.

The identity lifecycle spans three core stages: Onboarding, transactions and access approvals, and ongoing monitoring for security and compliance. Within each stage are inflection points where trust is most exposed and the consequences of failure are immediate. To your customers, these moments are not seen – they're part of everyday life. New customers are expecting fast, secure access at the first interaction and every interaction thereafter, making failure not an option.

However, if a financial account is opened under a stolen identity, it can enable a significant amount of fraud before it's detected. A compromised user account can grant a threat actor persistent access to systems and data. Even operational gaps, such as an expired digital certificate, can cause significant damage.

As the identity surface expands, so does the definition of trust. Every identity – human and non-human – must be continuously verifiable, governed, and auditable across every moment that matters.

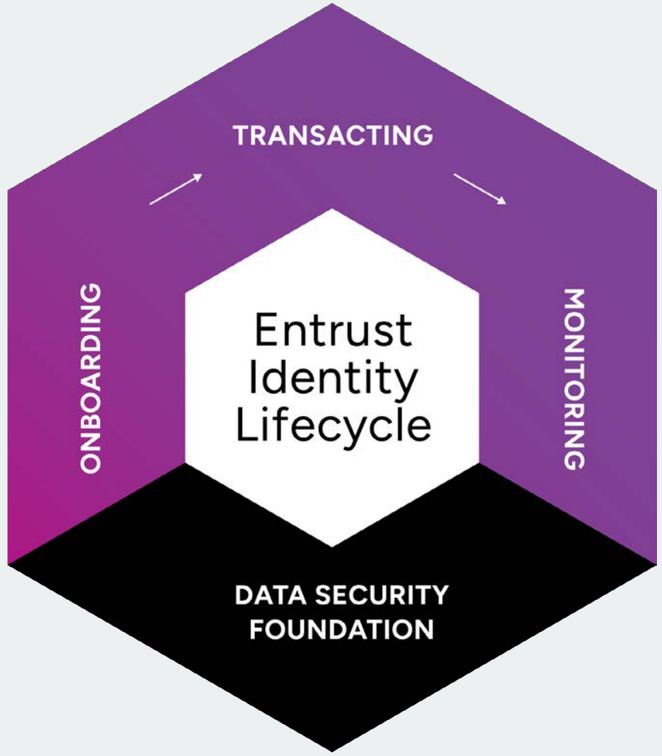

Every interaction across the identity lifecycle represents a Moment of Truth.

Why Digital Trust Matters in the Age of AI

Digital engagements are the norm. Threats are evolving, and so is the attack surface.

Three forces are converging to make Moments of Truth more consequential than ever.

First, every customer interaction is a test of trust. Digital engagement must be fast, frictionless, and secure. When interactions are slow or uncertain, users abandon. Organizations that don't validate trust invisibly and instantly weaken security, and risk their customers and reputations.

Second, bad actors are using AI to enhance and accelerate fraud and cyberattacks. Deepfakes, synthetic identity fraud, injection attacks, and AI-powered phishing now move faster than security controls designed for a pre-AI world. Common entry points, such as account opening, high-value transactions, and credential issuance, are targeted with increased sophistication. The window to detect and stop a fraudulent interaction is shrinking.

Third, enterprises themselves are embedding autonomous AI into core workflows, whether executing payments, managing procurement, initiating financial transactions, or committing organizations to external agreements. As these systems begin to make decisions at machine speed, the identity surface expands beyond people to include machines, services, and autonomous agents. An April 2026 **Deloitte** survey found that 80% of IT and security leaders lack mature governance capabilities for agentic AI, even as 74% expect their organizations to be using it at least moderately.

The Agent Autonomy Spectrum

Not all agents carry equal levels of risk. Required controls must scale with the level of autonomy an agent is granted. As agents move from read-only to fully autonomous operation, the governance model must move with them.

Autonomy Level	Key Risk	Required Controls
Read only	Data exposure	Access boundaries, visibility
Recommendation only	Unsafe guidance	Human accountability
Human-approved execution	Approval bypass	SoD, policy enforcement
Time-delayed execution	Drift, missed review	Queued approval, cancel windows
Fully autonomous	Expanded blast radius	Least privilege, monitoring, kill switch

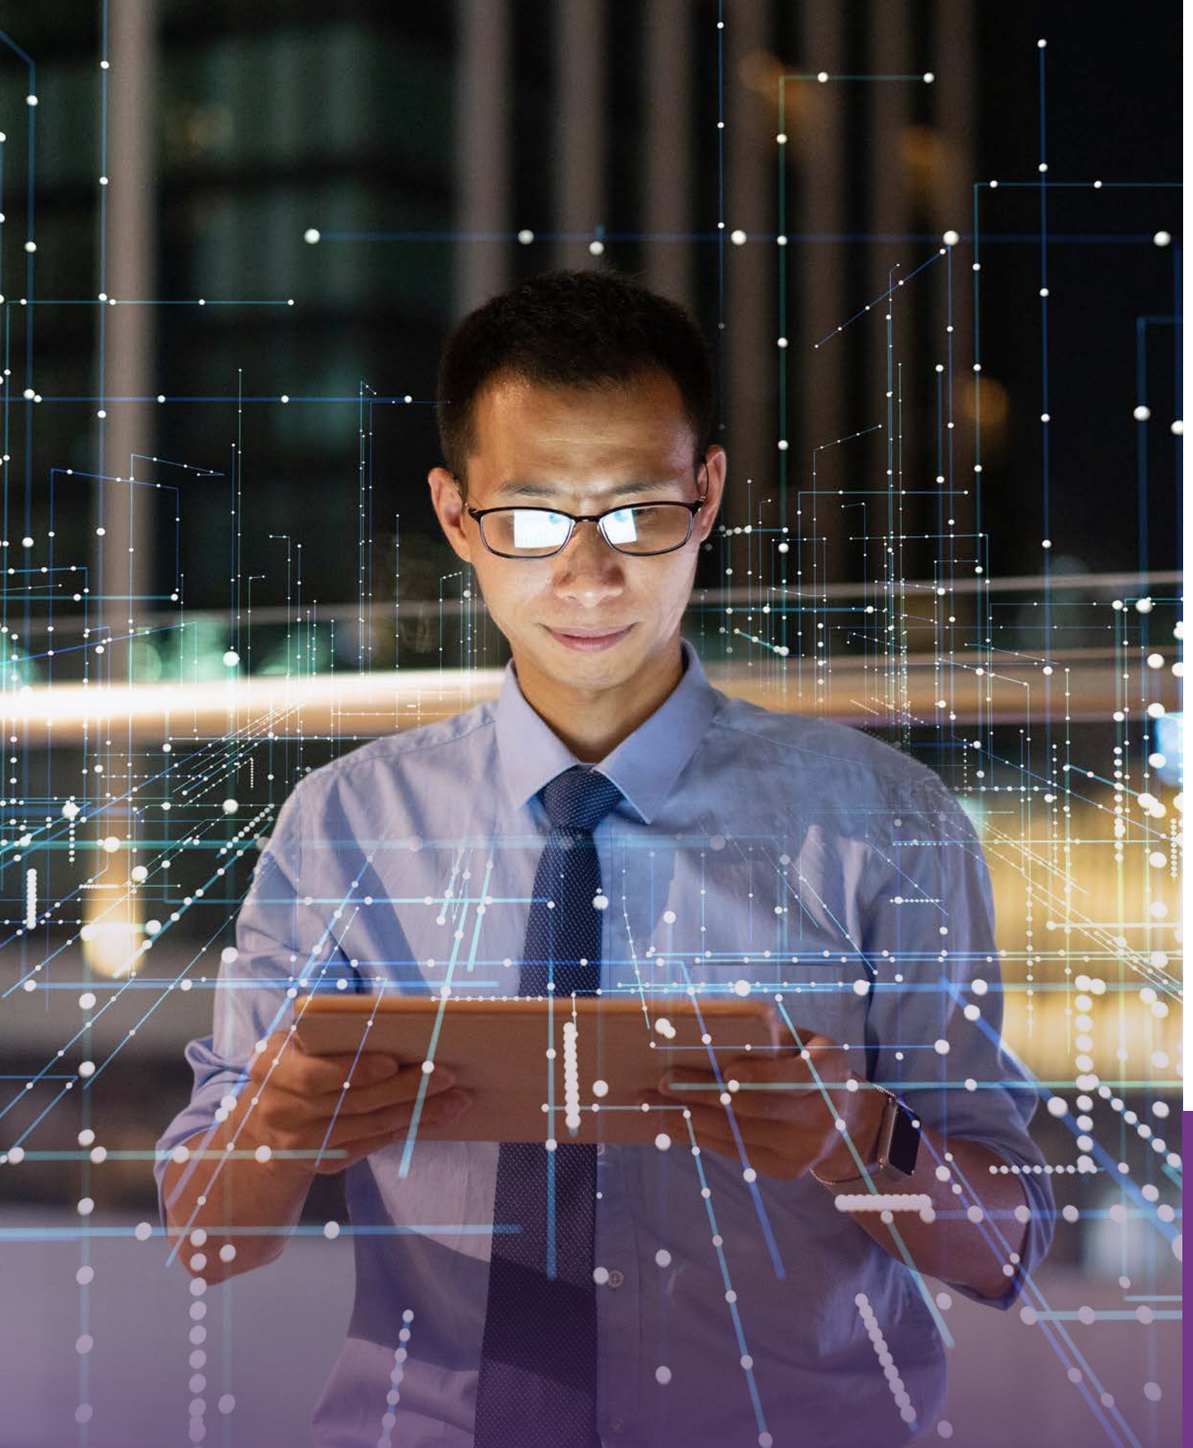

This shift changes the risk environment. As **Gartner** notes, compliance is moving from a point-in-time exercise to a continuous requirement as both AI systems and the regulations governing them evolve in real time. Move too quickly without a consistent way to verify identity, enforce policy, or prove what happened, and an organization can introduce massive operational, financial, and regulatory exposure.

These forces give rise to a singular need for identity-centric security to verify who's requesting it, whether it's authorized, and what the action is.

Autonomous AI is accelerating the movement toward what is essentially the full realization of Zero Trust – never trust, always verify. But as regulated organizations look to deploy autonomous agents across internal and external platforms, they find that existing control models aren't designed for this environment. Platform-native controls secure activity within a single ecosystem but create fragmentation and lock-in when agents operate across multiple environments. Traditional IAM-based approaches govern access, but not autonomous action at scale. Meanwhile, emerging AI-native vendors bring fresh innovation but often lack the operational depth and track record needed to demonstrate auditable identity, authorizations, actions, and outcomes under regulatory scrutiny.

What's missing is an independent trust layer anchored in verifiable cryptographic identity that can hold up under audit, regulation, and real-world complexity. As AI platforms evolve, this architectural layer sits beneath applications and models, enabling enterprises to consistently verify every identity, enforce policy in real time, and create auditability across every interaction.

IT and security leaders need a blueprint for trust in every Moment of Truth.

Identity Verification: Establishing Trust at Onboarding

Onboarding sets the identity that every decision depends on.

The onboarding stage is the foundation for all future trust decisions. It's where organizations determine whether a new identity is real, verified, and safe to operate within their environment.

Organizations must get identity verification right at the start, otherwise the risks are real and immediate. The identity established at onboarding becomes the baseline used throughout every transaction, access request, or high-risk moment thereafter.

- **The first interaction sets the tone for trust.** The security posture established at onboarding defines every interaction that follows.
- **High assurance verification must happen instantly, or the business pays later.** From synthetic identity fraud to deepfakes and impersonation attacks, the cost shows up as fraud losses, chargebacks, failed audits, or breaches traced back to an identity that should never have been verified in the first place.
- **It's about onboarding people first, then connect devices and agents to that identity.** Human identity verification is the anchor, and then devices and agents connect through cryptography to that human identity. As enterprises expand into AI-driven workflows, every device and autonomous agent must be cryptographically bound to a verified human identity from day one.
- **The link across humans, devices, agents, and organizations is what makes the entire chain of trust traceable, auditable, and accountable.**

AI AT PLAY
Deepfake and synthetic identity attacks are increasingly used to defeat onboarding checks. Verification must be resilient enough to detect AI-generated documents, images, and biometric spoofing, in addition to validating that a document was submitted.

If a wrong identity is created at onboarding, that fraudulent identity persists across systems, increasing risk and expanding compliance exposure throughout the entire lifecycle.

Onboarding Through Activation: Issuing Trust at Scale

Trust becomes usable the moment a credential enters the world.

Activation is where trust becomes usable – the moment a verified identity is issued as a credential. Typical activation scenarios include: A financial institution issuing a payment card, a government entity issuing a national ID or driver's license, a business granting badge access to an employee, or an educational or healthcare institution granting a badge to a student or patient. Each one of these scenarios requires activation as the point where that identity becomes operational.

Physical and digital credentials now coexist, and attackers look for opportunities to exploit the gaps between them.

> Trust embedded at issuance is the architecture that makes everything else governable.

Transaction Security: Verifying Trust in Real Time

Every transaction is a trust decision made in real time.

After a new identity has been successfully onboarded and activated, that identity carries through as trusted anchor for all downstream activity. The ongoing challenge for organizations is ensuring that trust is continuously validated across every interaction and not assumed after initial verification.

During transactions, organizations must dynamically evaluate whether access or action should be granted based on context, risk, and verified identity. In healthcare, for instance, doctors and patients need access to specific records, such as test results, treatment histories, and financial data, based on their specific and verified roles. Providing a “blanket login” risks giving access to information that they shouldn’t have under applicable regulations.

In financial services, transactions often require step-up authentication at defined, often higher-risk or time-based thresholds. These thresholds trigger biometric verification or risk-based authentication, especially at times of unusual behavior, high-value transfers, or anomalous access patterns. Government services similarly rely on verified identity checks at sensitive access points to ensure proper authorization.

At the transaction stage, trust is continuously evaluated using the identity established at onboarding and reinforced through adaptive authentication methods such as biometrics and risk-based verification.

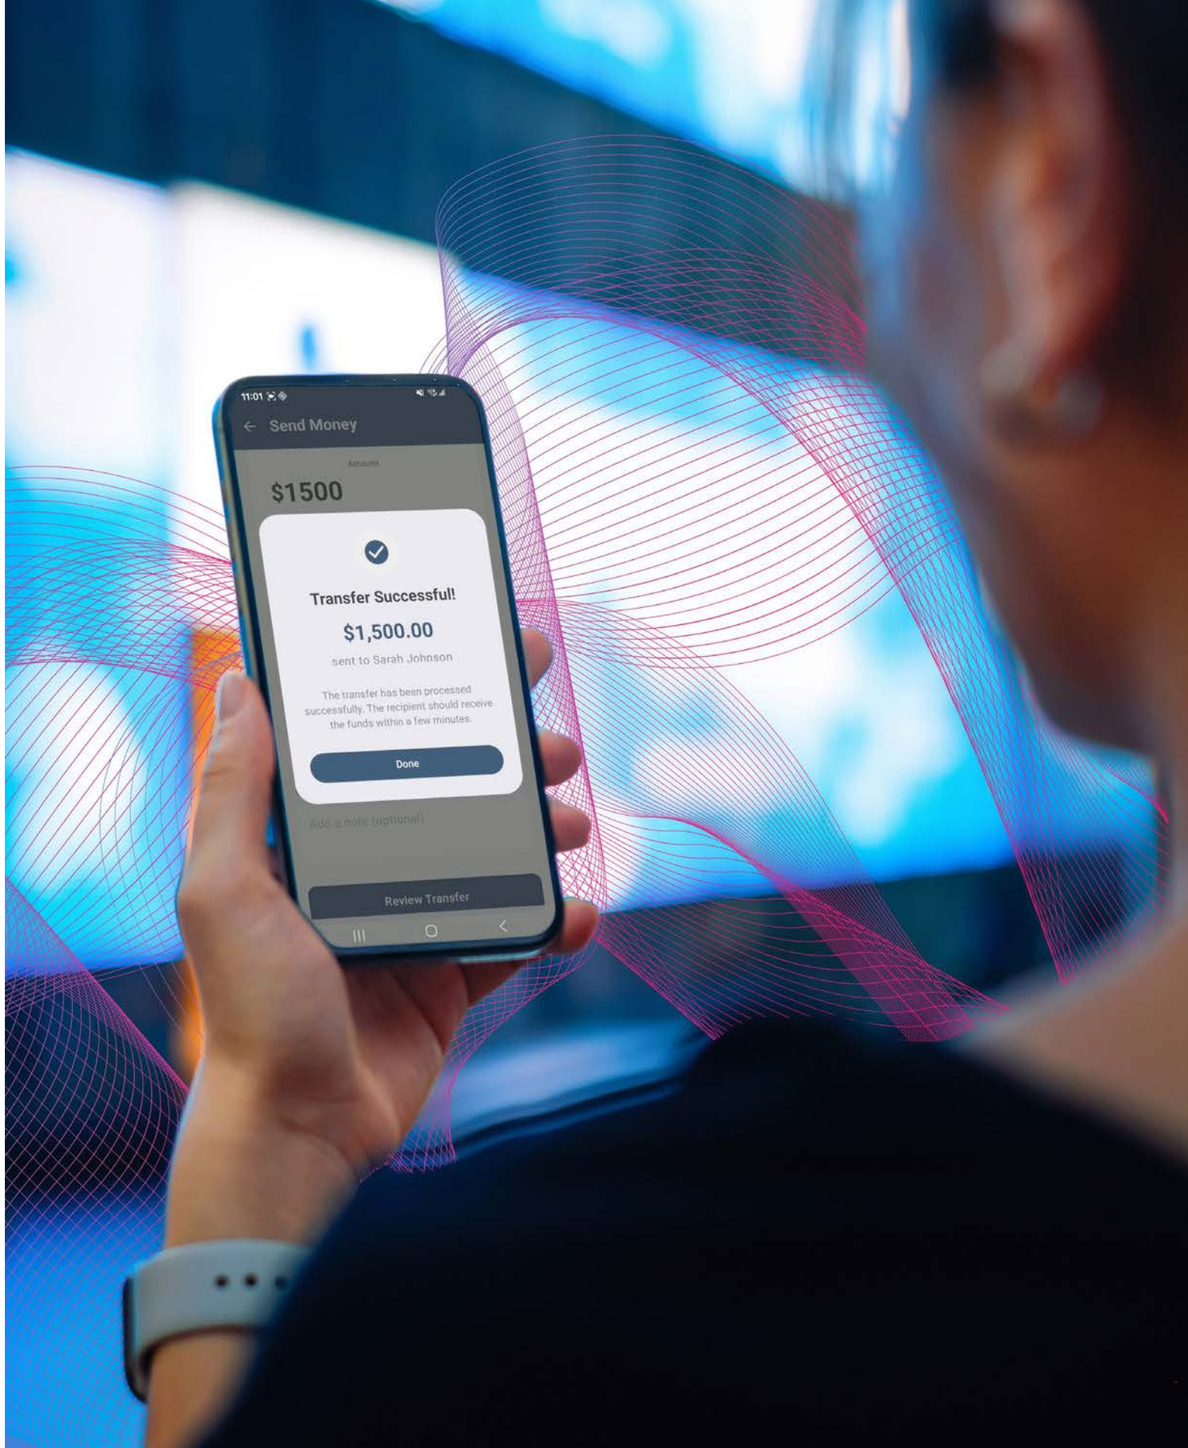

AI AT PLAY

AI-driven attacks, including account takeovers, automated bot attacks, and synthetic identity exploitation, are becoming more common. This makes real-time validation essential across every interaction.

Therefore, trust decisions must be:

Continuous, not one-time, extending verification throughout the interaction lifecycle in line with Zero Trust principles: Never trust, always verify

Contextual, evaluating both the identity and what initiated the action in real time

Verifiable, confirming every action or decision aligns with policy before access or execution is granted

AI agents introduce a key extension to the trust model. Guardrails need to define what an agent is allowed to do, under which policies, and on whose behalf. Every autonomous action must be documented and traceable through a chain of trust that links a verified human identity to a cryptographically bound agent identity. In other words, every autonomous transaction must maintain a clear, auditable lineage: A verified human initiator, a governed agent, and a policy-enforced action. Organizations need to answer these three questions in real time:

- Who, or what, initiated this action?
- What is it authorized to do?
- And on whose behalf is it operating?

Without this accountability layer, agent autonomy shifts from an efficiency gain into a systemic business liability. This is why human-in-the-loop remains critical during high-risk authentication. No autonomous system should execute irreversible or high-consequence actions without a verified human checkpoint.

A missing chain of trust leaves autonomous agents operating outside accountability.

Continuous Trust: Governing Human, Machine, and AI Identities

Continuous assurance requires visibility, automation, accountability, and resilience.

Security teams are responsible for continuously monitoring the health and integrity of keys, certificates, secrets, devices, workloads, and identities across their environments. Expired certificates, improperly authenticated devices, and unmanaged credentials can disrupt operations, expose sensitive systems, and create compliance gaps.

Organizations need automated capabilities that can continuously discover, track, renew, and report across their entire certificate and identity estate – not just at audit time, but in real time.

Monitoring is where governance becomes visible. The identities being observed – human, machine, device, workload, and agent – have already been established, credentialed, and governed earlier in the lifecycle. Proper monitoring offers continuous assurance that those identities remain trusted, authorized, and compliant, and that every non-human identity has both a beginning and an intent. Understanding who or what created it, what it's intended to do, and how it behaves over time is important to maintaining trust.

As organizations scale AI agents and connected systems, automation is needed to manage certificates, rotate secrets, enforce policy, and maintain audit readiness. Compliance requires continuous evidence, accountability, and control across the identity lifecycle.

Cryptographic infrastructure has an expiry date, and most organizations are closer to that date than they realize. Harvest now, decrypt later attacks mean bad actors are collecting encrypted data today with plans to decrypt it once quantum computing power reaches maturity.

Organizations need to build crypto-agility now. This requires protecting customers against long-term data exposure, managing trust across an ever-growing identity estate, and visibility and control to support algorithmic transactions. All of this is built on a cryptographic root of trust – identities and data protected by keys, certificates, and hardware security modules (HSMs).

AI AT PLAY

As AI agents multiply across environments, monitoring must cover autonomous identities as thoroughly as human ones, tracking who created an agent and how its behavior evolves over time.

You can't protect what you can't see, and you can't move at machine speed without automation. The organizations that define trust infrastructure now will shape the next decade of digital business.

Source: Why trust is your greatest asset in a digital-first world

Trust at Stake: Industry Perspectives

Moments of Truth look different depending on the industry. For highly regulated industries, the ability to enforce policy and drive compliance, and prove it in an audit, is just as important as stopping fraud in the moment.

Industry	Moments of Truth	The Consequences	What Winning Organizations Get Right
Financial Services	Account opening, payment card issuance, high-value transactions, account changes, access recovery	Account takeover (ATO), financial loss, regulatory penalties, customer churn, reputational damage. **90% of organizations say failure to increase use of AI fraud prevention will increase their losses.** Source: Mastercard Report	Continuous, contextual authentication that stops ATO fraud without adding friction to legitimate users and transactions
Government and Public Sector	Issuance of national IDs, driver's licenses, citizen access to services, border control, eID/eWallets	Fraud at scale, national security exposure, non-compliance, public distrust. **The U.S. Government Accountability Office found that 23 out of 24 major U.S. federal agencies failed to fully address previously identified cybersecurity weaknesses in 2024.** Source: U.S. GAO Cybersecurity Overview	Identity infrastructure that scales credential issuance across physical and digital channels without gaps
Healthcare	Patient record access, clinical device authentication, provider credentialing, provider badge access to restricted areas and patient records	Patient safety, HIPAA violations, clinical disruption, ransomware vulnerability. **Healthcare is the most expensive industry for data breaches for the 14th consecutive year.** Source: IBM Cost of Data Breach 2025 Report	Role-based access controls tied to verified provider identity, updated in real time as roles and contexts changes
Education	Student and faculty onboarding, badge access, digital credentials, campus ID issuance	Unauthorized access, credential fraud, compliance gaps, campus safety risks. **80% of school districts reported phishing attacks in 2025.** Source: MS-ISAC's 2025 K-12 Cybersecurity Assessment	Unified identity governance across physical badge access and digital systems from day one of onboarding
Telecom	SIM swaps, subscriber onboarding, network device authentication, certificate management at scale	Account fraud, network outages, subscriber trust erosion, regulatory exposure. **In 2025, losses tied to telecommunications fraud cost $41.82 billion.** Source: Communications Fraud Control Association (CFCA) Fraud Loss Survey Report	High-assurance onboarding, fraud prevention, cryptographic device trust, automated CLM that prevents outages before they happen

Secure Every Moment of Truth

The lifecycle is yours to protect. Entrust helps you own it.

Identity is the trust layer. Entrust helps make it continuous, governable, and provable.

Trust is the gating factor to realizing the full value of the identity lifecycle as the identity surface expands. Enterprises need a consistent foundation for establishing trust, enforcing policy, and proving accountability across every interaction.

At the center of this is human-in-the-loop step-up authentication. No matter how autonomous a system becomes, high-risk actions require a verified human checkpoint before execution. Trust that scales with automation, without losing the human accountability regulators and customers expect, is a key differentiator.

For CIOs and CISOs, the stakes are concrete: Fraud that bypasses identity controls, identity attacks that compromise systems from the inside, and network outages triggered by expired certificates or unverified devices. Enterprises that govern who can act, what they can do, and prove outcomes continuously and in real time will become tomorrow's leaders.

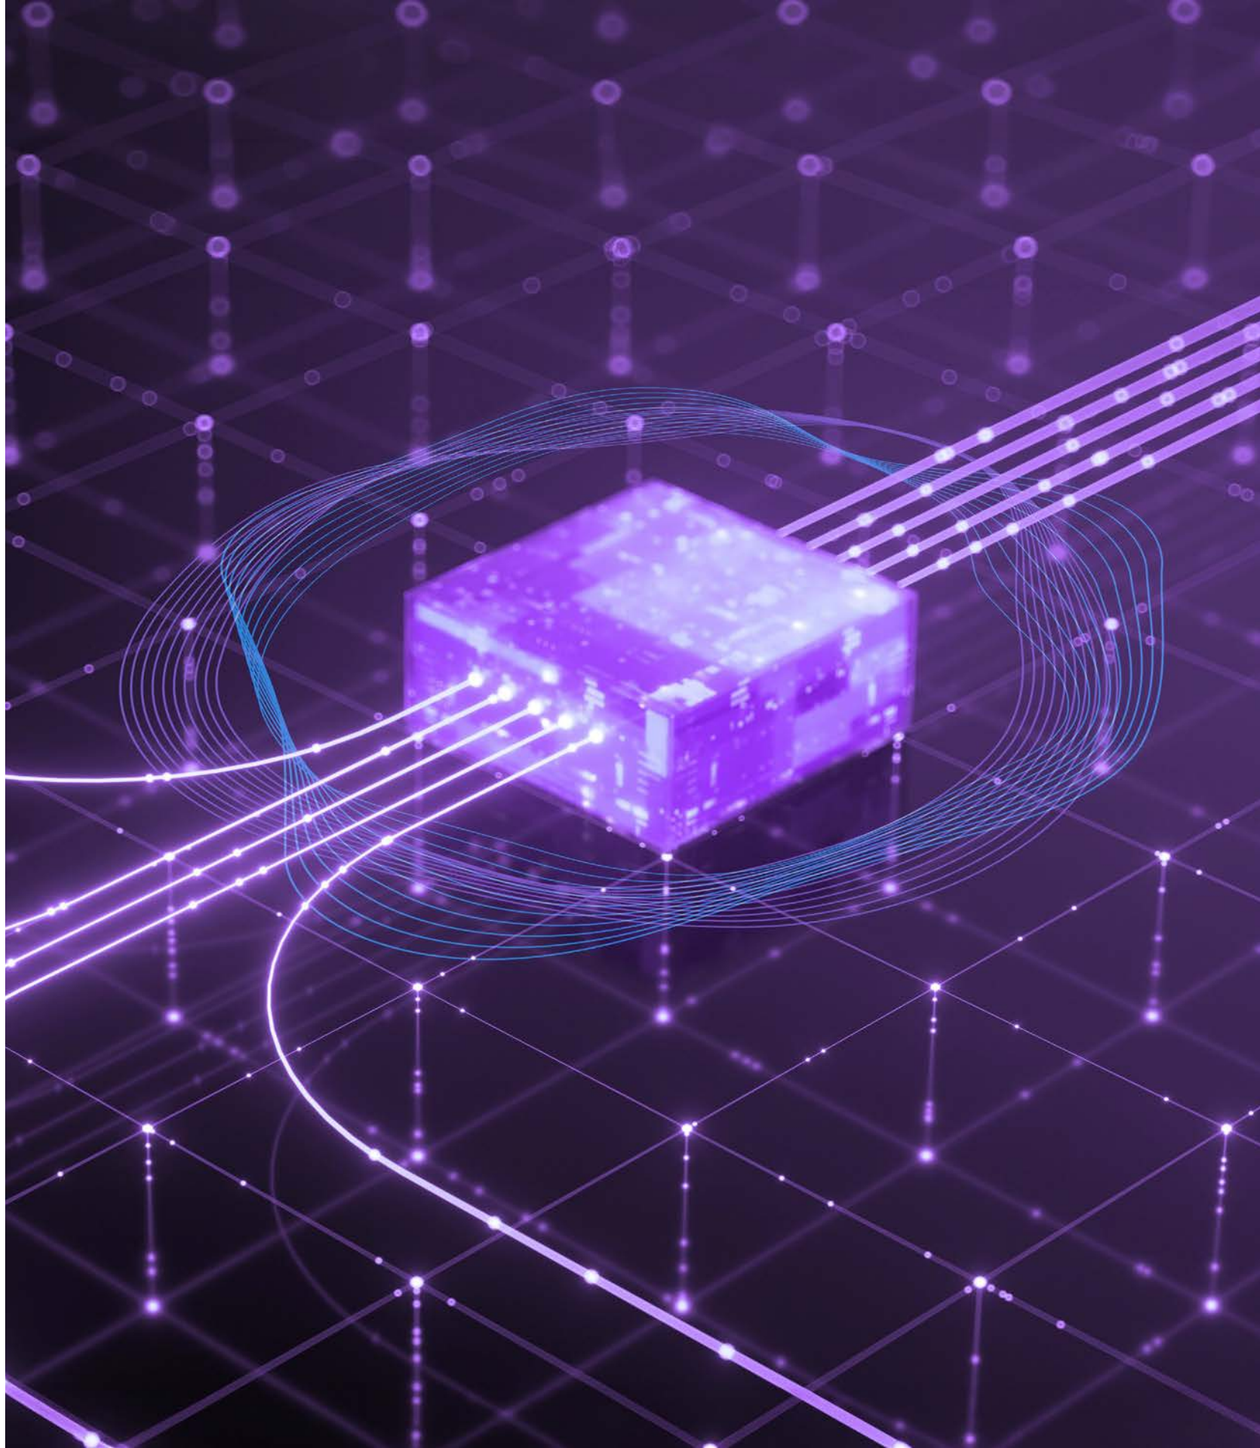

Every Moment of Truth is an opportunity to build trust or lose it.

For organizations managing an expanding identity surface, the trust layer is a critical component of enterprise governance and compliance:

Identity:
A verifiable connection between actions and accountable entities.

Authorization:
Delegated authority that remains governed, restricted, and aligned to policy.

Accountability:
The evidence needed to understand what happened, why it happened, and on whose behalf.

Entrust brings decades of experience helping organizations establish trust across identities, credentials, certificates, and cryptographic infrastructure, including biometric identity verification, adaptive authentication, certificate lifecycle management, and hardware-rooted cryptographic trust. Entrust provides the building blocks organizations rely on to verify users, secure devices, protect machine identities, and maintain accountability across ecosystems.

Building this trust infrastructure early gives organizations the footing they need to adopt new technologies with confidence. They can move faster without sacrificing control, demonstrate continuous compliance rather than reconstructing it after the fact, and maintain clear accountability across people, devices, agents, and systems.

As autonomous systems take on more responsibility, the same foundations that secure human identity must extend to devices, workloads, and AI agents, keeping identity, authorization, and accountability resilient across every Moment of Truth.

Whether you're managing certificate lifecycles, scaling digital identity programs, deploying AI agents, or preparing for post-quantum cryptography, Entrust is the trust layer organizations need to operate securely at every moment that matters.

Learn more at entrust.com/solutions/ai-security

ABOUT ENTRUST

Entrust fights fraud and cyber threats with identity-centric security that protects people, devices, and data. Our comprehensive solutions help organizations secure every step of the identity lifecycle, from verifying identity at onboarding to securing connections and fighting fraud in everyday transactions. Ongoing monitoring supports compliance and safeguards keys, secrets, and certificates. With a foundation of identity-centric security, our customers can transact and grow with confidence. Entrust has a global partner network and supports customers in over 150 countries.

For more information, visit entrust.com.

entrust.com Toll-Free: 888.690.2424 | International: +1.952.933.1223 | sales@entrust.com

ENTRUST

SECURING A WORLD IN MOTION