IDENTITYWEEK AMERICA

GLOBAL • TRUSTED • VISIONARY

Beyond Authentication

Continuous Identity Assurance in the Age of AI

September 2, 2026

Dr. Sarbari Gupta, CISSP, CISA

Founder & CEO, Electrosoft

Electrosoft

ABOUT THE SPEAKER

Two Decades in the Trenches of Federal Identity

FIPS 201 (2005)

Supported NIST in developing FIPS 201, the standard issued in response to HSPD-12 — still the basis for the PIV credential used across federal agencies, and still the gold standard.

SP 800-63 Co-Author

Continues to serve as a subject matter expert supporting NIST's Digital Identity Guidelines, most recently as a named co-author of SP 800-63-4.

Electrosoft

Founder and CEO of Electrosoft, delivering cybersecurity and identity management solutions to federal agencies for over two decades.

Electrosoft

TWO CALLS, TWO YEARS APART

2024

Arup Engineering Firm

$25.6 Million

A finance employee joined a video conference in which the CFO and other apparent colleagues were represented using deepfake video and audio. The employee subsequently authorized 15 transfers totaling approximately $25.6 million

2026

Singapore Government Impersonation Scam

At least SGD $4.9 Million

A business professional joined a fabricated Zoom meeting that appeared to include Singapore's prime minister and other senior officials. Deepfake video helped establish trust, after which the victim transferred funds through a series of transactions.

Familiar faces and a plausible setting did not authenticate who was actually participating.

THE MODEL WE INHERITED

Verify Once. Trust for the Session.

Digital identity didn't invent this model — it borrowed it from the physical world.

The Physical World

Checked once at the gate of a facility. Once inside, free to roam — unless an inner door demands proof again.

The Digital World

Authenticated once at login — PIV/CAC, MFA, passkey. Once inside, the session remains valid until timeout, logout, reauthentication, or a detected risk event.

THE MODEL WE INHERITED

Establish. Authenticate. Grant Access.

For decades, digital identity has relied on this straightforward model.

Establish Identity → **Authenticate the User** → **Grant Access**

MFA, biometrics, passkeys have made authentication far stronger.

But it is primarily a single decision, made at a single point in time.

THE FEDERAL REALITY

Strong Authenticators Answer One Question

"Who are you right now?"

Derived PIV credentials and high-assurance authenticators answer this with real rigor.

"Are you still you, five minutes from now?"

No authenticator, however strong, answers this on its own. That's a gap — not a flaw.

THE NEW PRESSURE

AI Isn't Attacking the Lock. It's Attacking the Assumption.

Synthetic Identities

Convincing, AI-generated identities built to pass proofing checks.

Voice Cloning

Realistic clones from just seconds of audio.

Deepfakes

Video and audio forgeries used in remote proofing, impersonation and social-engineering fraud.

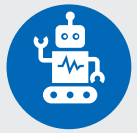

Automated Social Engineering

AI-scaled, highly targeted attacks at unprecedented pace.

THE FRAMEWORK

IAL, AAL, FAL — In Plain Language

IAL — **Identity Assurance Level**

The strength of the identity-proofing process used when a person sets up an authentication credential.

AAL — **Authentication Assurance Level**

Robustness of the authentication process and the binding between authenticators and the subscriber.

FAL — **Federation Assurance Level**

The rigor of the federation process that ensures the person who authenticated is the one who obtains access to a federated resource.

IT LOOKS LIKE THIS

When a Bot Wears Your Badge

1

User authenticates

PIV card, MFA — properly verified, AAL met.

2

A bot or AI agent takes over

Operating under that already-authenticated identity.

3

Full access, potentially unnoticed

Bot inherits privileges available through that authenticated session

4

Continuous signals flag it

Unusual activity for that user, in that context, raises the alarm.

Electrosoft

THE SHIFT

From Authentication to Continuous Assurance

NIST SP 800-63B-4, published in July 2025, explicitly recognizes session monitoring as a risk-management mechanism for detecting possible fraud during authenticated sessions.

> *An important boundary: session signals are risk-reduction indicators — not authentication factors. They do not change the AAL and do not substitute for periodic reauthentication.*

The question shifts from “was this identity authenticated?” to “can this session still be trusted?”

THE FRAMEWORK

A Practical Four Pillar model for Continuous Assurance

Behavioral Analytics

Typing, navigation, and usage patterns consistent with the user.

Device Posture

Endpoint health, patch level, and management status.

Contextual Risk Signals

Location, network, and privilege-change anomalies.

AI-Powered Threat Detection

Correlating signals across users, devices, and entitlements.

Together, these let an organization continuously ask whether trust in an active session remains justified.

Privacy Guidance: Collect only signals needed for defined security purposes; assess privacy impacts, limit retention and access, provide appropriate transparency, and establish redress for consequential decisions.

PILLAR 1 OF 4

Behavioral Analytics

Signals worth watching

- Typing cadence, navigation patterns, normal working hours
- Sudden bulk downloads, privilege escalation, unusual command sequences

Adaptive response ladder

Establish a baseline at the user, peer-group, and role level. Evaluate sequences of activity, not isolated events.

Increase monitoring → restrict sensitive functions → step-up authentication → suspend or terminate

PILLAR 2 OF 4

Device Posture

Signals worth watching

- Enrollment and management status, patch level, EDR status
- Secure boot, disk encryption, jailbreak/rooting status

How to use it

Bind the session as closely as practical to the device used at authentication. Reevaluate posture continuously — not only at login.

Revoke or restrict sessions when EDR, MDM, or vulnerability tools report a critical change.

PILLAR 3 OF 4

Contextual Risk Signals

Signals worth watching

- Geolocation and travel feasibility, IP reputation
- New device or browser, recent password reset or privilege change

A useful way to frame it:

Session Risk = f(Identity, Authenticator, Session, Device, Network, Behavior, Resource, Threat)

Weigh both likelihood of compromise and consequence of the action — an anomalous login to a public portal isn't the same as one preceding a privileged export.

PILLAR 4 OF 4

AI-Powered Identity Threat Detection

What AI adds

Correlates behavioral, device, contextual, and threat signals that rule engines evaluate separately — UEBA, session-hijack detection, graph analysis across entitlements.

The guardrail that matters

AI should not be the sole, opaque authority for high-impact decisions.

Use deterministic controls for known-bad conditions; require human review for ambiguous cases; retain evidence and model version for auditability.

An Example Graduated Response Model

LOW	Expected behavior, compliant device, normal location	**Continue the session**
GUARDED	New network or moderate behavioral deviation	**Increase monitoring; restrict sensitive ops**
ELEVATED	Multiple correlated anomalies or posture degradation	**Require step-up authentication**
HIGH	Token theft indicators, malware, impossible concurrent activity	**Suspend transaction; revoke session**
CRITICAL	Confirmed compromise or active identity attack	**Revoke sessions/tokens; contain account; initiate IR**

Electrosoft

As AI transforms both cyber offense and defense, we must move beyond asking whether an identity has been authenticated — and begin asking whether it can be continuously trusted.

Sarbari Gupta | Founder & CEO, Electrosoft
https://www.linkedin.com/in/sarbari-gupta/

Thank you.