FROM KNOWING A SECRET TO RECOGNIZING A PERSON: CONTINUOUS, BEHAVIOR-BASED IDENTITY ASSURANCE

AI for Identity and Access Management: Behavioral Authentication

Executive briefing — IAM and identity security leaders | ~20 min

The Password Is a Liability, Not a Control

- Passwords are shared, reused, and phished — they authenticate a string, not a person
- Credential theft remains a leading initial access vector across industry breach reporting (industry-reported range, not a specific unnamed study)
- Password resets and lockouts remain a persistent help-desk cost center
- Rotation and complexity policies push users toward predictable workarounds
- Static credentials offer a single point-in-time check, then trust the session indefinitely

Traditional MFA Raises the Bar, But Has Limits

- SMS OTP and push MFA are vulnerable to SIM swap, push fatigue, and adversary-in-the-middle proxies
- MFA is typically enforced only at login, leaving the rest of the session unverified
- Fatigue-based approval ('MFA bombing') has become a known bypass pattern
- Hardware tokens and FIDO2 improve phishing resistance but add cost, provisioning, and support overhead
- Even strong MFA cannot detect a session hijacked after successful authentication

Behavioral Biometrics: A Continuous Identity Signal

- Captures how a user interacts, not just what they know or possess
- Typing cadence: keystroke timing, rhythm, and pressure patterns unique to an individual
- Mouse and pointer movement: velocity, curvature, and click patterns
- Mobile gait and device-handling signals: how a phone is held, tilted, and carried
- Generates a running behavioral profile rather than a single login-moment check

Risk-Based and Adaptive Authentication

- Scores each access attempt in real time using device, location, network, and behavioral signals combined
- Low-risk sessions proceed with minimal friction; elevated risk triggers additional verification
- Moves authentication from a binary gate to a continuous, weighted decision
- Risk scores can factor in time-of-day, resource sensitivity, and historical access patterns
- Enables policy tuning per application or data classification rather than one-size-fits-all rules

Anomaly Detection Across Login and Session Behavior

- Machine learning baselines normal behavior per user and per peer group
- Flags deviations: unusual login times, impossible travel, atypical navigation paths within an application
- Extends detection beyond login into session-long monitoring for account takeover
- Can surface lateral movement or privilege misuse that credential checks alone would miss
- Requires sufficient historical data to establish a reliable baseline before alerts are trustworthy

AI-Driven Bot and Credential-Stuffing Defense

- Automated login attempts and credential-stuffing campaigns often exhibit non-human timing and interaction patterns
- Behavioral and device-fingerprint signals help distinguish scripted traffic from genuine users
- Complements rate limiting and CAPTCHA rather than replacing them
- Reduces reliance on CAPTCHA, which degrades user experience and is increasingly bypassed by automated solvers
- Should be paired with credential-breach monitoring to catch stuffing attempts using valid, stolen passwords

Balancing Friction and Security

- Step-up authentication applies stronger checks only when risk signals warrant it
- Goal is proportionate friction: invisible for routine access, robust for high-risk or high-value actions
- Over-triggering step-up prompts erodes user trust and increases support burden
- Under-triggering leaves genuine risk unaddressed — threshold tuning is an ongoing discipline, not a one-time setup
- Best practice: pilot risk thresholds on a limited population before organization-wide rollout

Privacy Considerations of Behavioral Data

- Behavioral telemetry (keystroke timing, mouse movement, device handling) is personal data and, in some interpretations, biometric-adjacent data
- Requires clear purpose limitation: data collected for authentication should not silently expand to other uses
- Data minimization and retention limits reduce exposure if telemetry stores are breached
- Transparency to employees and customers about what is collected and why supports trust and compliance
- On-device or edge processing of raw behavioral signals can reduce centralized data risk

Integrating with Existing IAM and SSO Infrastructure

- Behavioral and risk signals should feed existing IdP and SSO decision points, not create a parallel identity stack
- Standard integration points: risk-score inputs to conditional access policies, step-up triggers via existing MFA providers
- Vendor evaluation should weigh support for open standards (SAML, OIDC, SCIM) alongside proprietary behavioral APIs
- Phased integration reduces disruption to existing access policies and application dependencies
- Plan for fallback authentication paths if behavioral signal quality degrades or is unavailable

Managing False Positives and User Experience

- Behavioral models are probabilistic; some legitimate users will be flagged as anomalous
- Injuries, new devices, unfamiliar locations, or shared accounts can all trigger false step-up challenges
- Clear, low-friction fallback and appeal paths are essential to avoid locking out legitimate users
- False positive rates should be tracked as a first-class metric alongside detection rates, not an afterthought
- User communication about why a challenge occurred improves acceptance and reduces help-desk escalations

Regulatory Considerations for Biometric-Adjacent Data

- Behavioral biometrics may fall under biometric privacy statutes in some jurisdictions, depending on how data is derived and stored
- Requirements can include explicit notice, consent, and defined retention/destruction schedules
- Cross-border data transfer rules apply if behavioral telemetry is processed or stored outside the user's jurisdiction
- Legal and privacy counsel should review vendor data-handling terms before procurement, not after deployment
- Regulatory landscape is evolving; treat compliance review as a recurring checkpoint, not a one-time approval

Illustrative Scenario: Phased Adoption Pattern

- Illustrative scenario, not a verified case study: a mid-size financial services firm piloting behavioral authentication
- Phase 1: passive monitoring only, no enforcement, to build behavioral baselines and validate signal quality
- Phase 2: risk scoring feeds step-up authentication for high-privilege accounts and sensitive transactions
- Phase 3: broader rollout with tuned thresholds informed by Phase 1–2 false positive data
- Illustrates the general shape of low-risk adoption rather than a specific vendor outcome

Governance and Success Metrics

- Track authentication friction (challenge rate), detection efficacy, and false positive rate as a balanced scorecard
- Establish a cross-functional review board: IAM, security operations, privacy/legal, and user experience
- Define explicit rollback criteria if user impact or false positive rates exceed agreed thresholds
- Require vendor transparency on model logic and data handling as a procurement condition
- Revisit risk thresholds and baselines on a fixed cadence as user populations and threats evolve

Recommended Next Steps

- Commission a passive-monitoring pilot on a limited user population to validate signal quality before any enforcement
- Engage privacy and legal counsel early to scope data handling, consent, and retention requirements
- Shortlist vendors that integrate with existing SSO/IdP infrastructure via open standards
- Define success metrics (false positive rate, step-up rate, detection efficacy) before pilot launch
- Set a decision checkpoint at 90 days to evaluate expansion, threshold adjustment, or rollback