Abnormal

The Essential Guide to Human Risk Management

A Practical Framework for Measuring and Reducing Human Risk

Human Risk in the Age of AI-Driven Attacks

99%

Of organizations experienced a security incident linked to preventable user actions in the past year

2025 State of Security Awareness Training

83%

Of security leaders say security awareness training requires substantial effort to operate and maintain

2025 State of Security Awareness Training

53%

Say the effort required outweighs the impact it delivers

2025 State of Security Awareness Training

Human error continues to be one of the most consistent and costly drivers of security incidents. Despite widespread adoption of security awareness training (SAT), organizations struggle to meaningfully reduce risk posed by the human layer.

These findings point to a fundamental issue: **traditional awareness programs are not designed to keep pace with how modern attacks operate**. Today's AI-powered threats are highly personalized, behavior-driven, and continuously evolving to blend into legitimate business communications and workflows. Static training modules, generic phishing simulations, and periodic campaigns engineered to check compliance boxes cannot adequately prepare employees for real-world conditions.

For many organizations, human risk management remains limited in practice, often relying on static programming, engagement/completion-based metrics, and manual strategies that have constrained traditional SAT efforts. As a result, organizations continue to invest significant time and resources without gaining clear visibility into whether risk is actually decreasing.

The same technology that enables more powerful threat campaigns at scale is also reshaping how organizations prepare their workforce to respond. Advances in AI and behavioral analysis now make it possible to move beyond static training toward a continuous, adaptive model—one grounded in real threats and real user behavior. Real attacks can be safely repurposed into realistic simulations, training can be tailored to individual roles and behaviors, and feedback can be delivered the moment risky behavior occurs.

This guide examines:

- Why traditional security awareness programs struggle to deliver meaningful risk reduction
- What modern human risk management requires in today's threat landscape
- How organizations can move from awareness activities to measurable behavior change

The objective is to outline how organizations can more effectively manage human risk as a core component of their security strategy—shifting from merely checking compliance boxes to directly influencing user behavior and building a formidable security culture prepared for AI-enabled attacks.

Table of Contents

Why Security Awareness Training Falls Short

Security awareness training (SAT) is deeply embedded in modern security programs. Most organizations have established recurring training cycles, supported by phishing simulations, policy reminders, and role-based content. These programs are widely accepted as necessary and required for compliance.

Yet despite this widespread adoption, questions about their effectiveness persist.

Incidents Remain Closely Tied to Human Behavior

A significant portion of security incidents continues to involve user actions, whether clicking a malicious link, responding to a fraudulent request, or unintentionally compromising accounts by sharing credentials.

These incidents are not edge cases. Nearly half of organizations report that at least a quarter of their security incidents are linked to avoidable user actions.

This suggests that while traditional SAT programs are in place, they are not meaningfully translating into reduced exposure at the human layer.

Traditional Approaches Emphasize Coverage Over Effectiveness

Many SAT programs are structured around coverage: ensuring that all employees complete required training and are periodically tested through simulations. Common methods include phishing simulations (81%), topic-based modules (80%), and email-based reminders (71%)—all widely used across organizations.

These approaches are effective at delivering training content at scale, but they do not necessarily ensure that individual employees can recognize and respond to threats in real-world conditions: a phishing email disguised as an urgent request from the boss, or a sternly worded reminder about a nonexistent invoice coming from a vendor that appears legitimate. In fact, organizations that rely more heavily on these “less sophisticated” training methods often report higher rates of incidents tied to user actions.

The distinction is subtle but important: delivering training consistently is not the same as actually improving security decision-making under pressure.

Measurement Remains a Persistent Challenge

Another limitation lies in how training effectiveness is evaluated. 42% of organizations **report that reliably measuring the impact of their training programs is a challenge.**

Security teams tend to rely on metrics such as phishing simulation results, training completion rates, and employee feedback. While useful for tracking participation and short-term outcomes, these measures offer only a partial view of actual risk.

Phishing simulations, for example, provide a controlled snapshot of behavior in a low-stakes test environment, where employees know they're simply passing a quiz, but may not reflect how employees respond to real attacks. Similarly, completion rates indicate that training has been delivered, but not whether it has influenced behavior or delivered anything beyond information the employee discards in minutes.

This creates a measurement gap. Security teams are often tracking multiple indicators, yet still lack a reliable way to determine whether human risk is increasing, decreasing, or remaining unchanged.

Training Cannot Keep Pace With Evolving Threats

Organizations are actively working to improve their security awareness programs, but those improvements are typically introduced on fixed cycles—new modules, updated simulations, or refreshed content delivered periodically over time.

Attackers, however, do not operate on those timelines.

Modern social engineering tactics evolve continuously. New impersonation techniques, attack formats, and delivery channels emerge rapidly, often changing faster than training content can be updated or deployed. In fact, many organizations cite the inability to keep training aligned with current threats as a key limitation of their programs.

This creates a structural gap. Training reflects what threats looked like when content was created, while employees are confronted with attacks as they exist in the moment. Even well-designed programs can quickly become outdated, leaving users underprepared for the types of messages they actually receive.

As a result, incremental updates to training programs do not keep pace with the speed at which risk evolves.

From Training Delivery to Risk Management

Traditional Security Awareness Training	Modern Human Risk Management
Periodic training cycles (annual or quarterly)	Continuous, always-on adaptation to evolving threats
Generic, one-size-fits-all content	Training tailored to individual roles, behavior, & exposure
Simulated scenarios based on templates	Real attacks adapted into realistic, relevant simulations
Scheduled campaigns and modules	Just-in-time coaching delivered at the moment of risk
Manual program management and administration	Autonomous execution driven by behavioral signals
Measured by completion rates and simulation results	Measured by behavior change and reduction in human risk

The Human Attack Surface Has Changed

While many of the underlying cyberattack techniques—phishing, impersonation, and social engineering—are not new, the way they are executed has fundamentally changed.

Attackers are no longer relying on broad, easily identifiable tactics. Instead, they are targeting individuals with precision, using workplace context, timing, and behavioral cues to increase the likelihood of success. The human layer has become an active and dynamic attack surface.

Attacks Are Increasingly Personalized and Context-Aware

Modern social engineering attacks are designed to blend seamlessly into everyday business communication. Messages are crafted to reflect real relationships, ongoing workflows, and familiar patterns of interaction.

Credential phishing emails may replicate login portals with near-perfect accuracy. Business email compromise (BEC) attacks often impersonate executives or trusted partners, using urgency and authority to prompt action. Vendor email compromise (VEC) goes a step further, leveraging legitimate accounts and real transaction history to carry out fraud over extended periods of time. 44% of vendor email compromise (VEC) messages receive engagement, showing how effectively attackers replicate legitimate workplace communications and exploit pressure and expectations.

Signals of Malicious Intent Are Less Obvious

Historically, many awareness programs focused on identifying obvious warning signs such as poor grammar, suspicious links, or unfamiliar senders. Today, those indicators are far less reliable.

Advances in generative AI have made it easier for attackers to produce messages that are grammatically correct, contextually relevant, and stylistically consistent with legitimate communications. In many cases, malicious messages contain no clear red flags, making them difficult to distinguish from routine business activity.

At the same time, some attacks avoid traditional indicators altogether. Payloadless techniques, for example, rely on convincing users to take actions—such as sharing information or changing configurations—without the use of links or attachments.

As a result, identifying threats increasingly depends on judgment and context rather than pattern recognition alone.

The Attack Surface Extends Beyond Email

While email remains a primary entry point, attackers are expanding into additional channels that employees use every day.

Collaboration platforms such as Slack and Microsoft Teams are now used for social engineering, often mimicking internal conversations or urgent requests. QR code phishing, spoofed meeting invitations, and even deepfake voice messages and video are becoming more common as attackers experiment with new methods of engagement.

Attackers follow the user across the org, targeting all environments where communication and decision-making take place.

Behavior, Not Knowledge, Determines Outcomes

In this environment, the outcome of an attack often depends less on what an employee knows and more on how they respond in a specific moment.

Employees are frequently operating under time pressure, managing high volumes of communication, and making rapid decisions to keep business processes moving. Attackers are deliberately exploiting these conditions, designing messages that create urgency, familiarity, or a sense of obligation.

Even well-informed employees can make mistakes when a message appears legitimate and aligns with their expectations. This is particularly true when attacks are tailored to an individual's role, responsibilities, or current projects.

The implication is clear: awareness alone is not sufficient if it is not reinforced by context, timing, and relevance.

A Dynamic and Continuously Evolving Risk Surface

Taken together, these changes have transformed the human layer into a continuously evolving attack surface requiring an equally dynamic risk management culture.

This creates a moving target for security teams. The challenge is no longer limited to educating the workforce about known threats, but preparing employees for how risk can manifest based on individuals' roles and behavior.

As the human attack surface becomes more complex and adaptive, approaches that rely on static content and periodic reinforcement struggle to keep pace. The question is no longer simply how to increase awareness, but how to account for the variability and dynamics of human behavior in the face of evolving threats.

How Modern Human Risk Management Drives Measurable Change

As organizations mature their approach to human risk management, the focus shifts from program design to outcomes. The question is no longer how to deliver training at scale, but how to influence behavior in ways that can be observed, measured, and improved over time. Nearly all surveyed security leaders (98%) **agree that training tailored to individual behavior and delivered in real time would significantly improve security outcomes.**

In practice, this shift is not driven by a single capability, but by a set of coordinated changes in how training is generated, delivered, and evaluated. The ideal system is behavioral signal-driven, using real attack data and user interactions to continuously adjust how risk is addressed.

A Continuous Human Risk Management Loop

Modern approaches operate as a continuous system rather than a series of campaigns. At a high level, this system follows a repeatable loop:

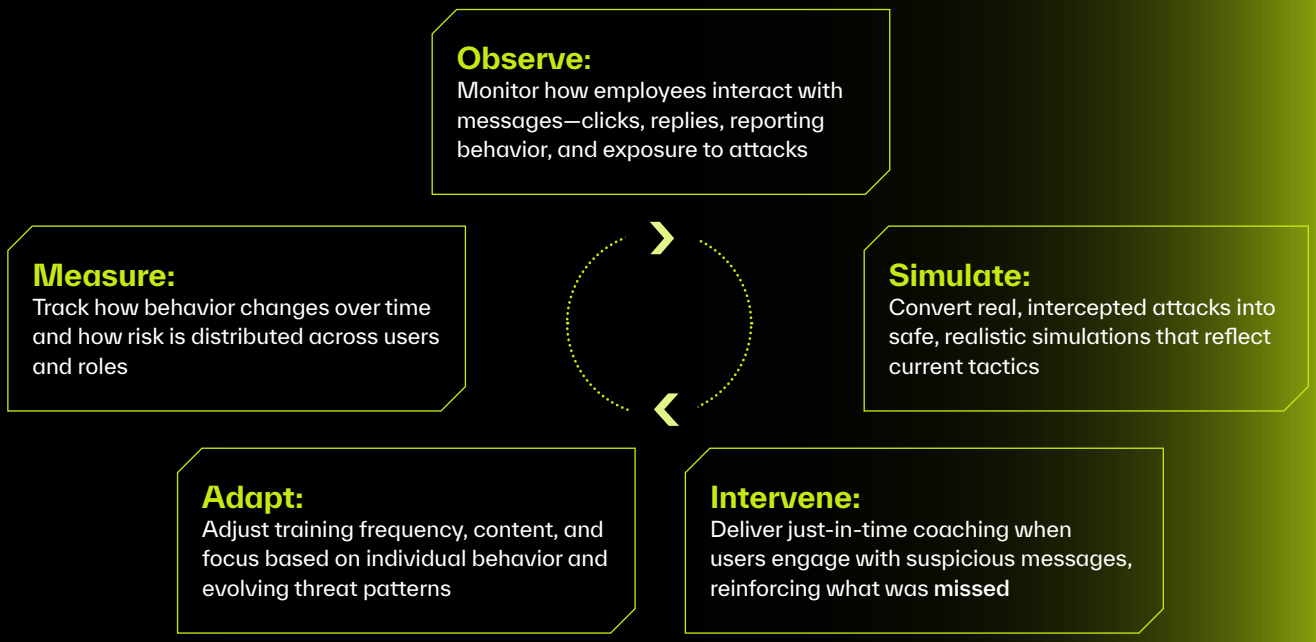

This loop allows organizations to move from periodic training to continuous influence over how decisions are made in real-world conditions.

From Simulated Scenarios to Real-World Conditions

One of the most significant changes is the shift toward realism.

Traditional simulations often rely on predefined templates or recognizable patterns. Over time, employees learn to identify these as tests rather than evaluate them as genuine threats, reducing their effectiveness.

Modern approaches use real attacks—safely neutralized and adapted for training—to mirror the types of messages employees actually receive at their organization. These simulations reflect current attacker tactics, realistic timing, and familiar business context.

For example, if finance teams are frequently targeted with invoice fraud, simulations can replicate those exact patterns—language, formatting, and urgency—rather than relying on generic phishing templates made to fit teams whose environments differ drastically.

This alignment ensures that employees are practicing decision-making in conditions that closely resemble their daily experience.

Continuous, Behavior-Driven Adaptation

Rather than delivering the same content to all users, modern programs adjust dynamically based on each employee's unique behavior and the types of threats targeting the organization.

Employees who frequently engage with risky messages or operate in higher-risk roles may receive more targeted simulations or more frequent reinforcement. Those who demonstrate stronger judgment may receive less frequent intervention.

At the same time, training evolves to reflect current attack tactics, ensuring simulations remain aligned with the threats employees are most likely to encounter.

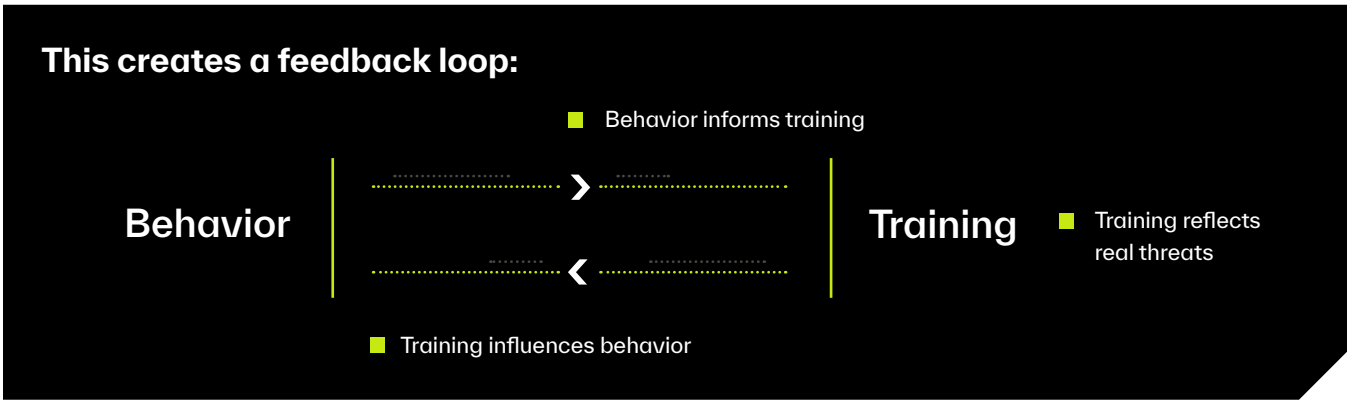

Over time, this allows organizations to prioritize high-risk users and adapt to changing exposure across the organization—rather than treating all employees the same.

Reinforcement at the Moment of Risk

Timing is critical.

When an employee clicks, replies, or reports a message, that interaction creates a high-value learning moment. Modern approaches use this moment to deliver in-the-moment, contextual coaching, explaining what signals were missed and

how similar threats can be identified in the future.

This just-in-time guidance is grounded in the specific message the user interacted with, not a generic example. It reinforces decision-making in context, rather than relying on recall from prior training sessions.

Over time, this approach builds pattern recognition and improves how employees respond under real conditions.

Autonomous Execution at Scale

Operationally, this model also reduces manual effort.

Traditional programs require teams to:

- Select simulations
- Schedule campaigns
- Manage follow-up training

Modern systems allow for the autonomous execution of these processes. Simulations are generated from real attack data, delivered continuously, and adjusted automatically based on user behavior.

This ensures consistency and allows programs to scale without increasing administrative overhead and adding manual work. Security teams can focus on understanding risk and outcomes, rather than managing training logistics.

Measuring What Actually Changes

The final shift is in how success is measured.

Instead of focusing solely on completion rates or isolated simulation results, organizations can track how behavior evolves over time. This includes:

- Changes in how users interact with suspicious messages
- Reduction in high-risk behaviors (e.g., clicking, replying, or sharing information)
- Differences in risk levels across roles or departments
- Improvements in response patterns to real and simulated threats

These metrics provide a more direct view of whether human risk is increasing or decreasing, and where intervention is needed.

A Direct Link Between Effort and Outcome

Taken together, these changes create a clearer connection between program activity and security outcomes.

The result is not simply more efficient training, but a system that actively shapes how employees respond to threats and makes that impact continuously visible.

Training reflects real threats.

Intervention happens in real time.

Programs **adapt** automatically.

And behavior is **measured** continuously.

Human Risk Management Evaluation Checklist

Not all approaches labeled “human risk management” operate at the same level. Many programs extend traditional awareness training with new terminology, while retaining the same underlying model.

The following criteria can be used to evaluate whether a vendor meaningfully reduces human risk, or rehashes the same old legacy security awareness gimmicks.

1. Are simulations based on real attacks?

- Are training scenarios generated from real attacks targeting the organization?
- Do simulations reflect current attacker tactics, context, and timing by continuously learning from new threats?
- Are simulations built from static templates that users learn to recognize and tune out?

Why it matters:

Training must reflect the same psychological ambiguity and workplace context as real attacks to build effective judgment.

2. Are simulations delivered in a way that mirrors real email?

- Do simulated messages behave like normal inbox traffic, without requiring safelisting or special configuration?
- Are delivery methods aligned with how real messages reach users (e.g., direct inbox delivery rather than routed campaigns)?
- Do simulations contain SAT artifacts that make them easy to identify as tests?

Why it matters:

If a simulation does not look and behave like a real message, it trains users to spot tests, not seriously assess risk.

3. Is training tailored to individual behavior and exposure?

- Is training adjusted based on how a specific user interacts with messages over time?
- Are high-risk users identified based on role, behavior, and exposure?
- Is the same generic content delivered uniformly across the organization?

Why it matters:

Human risk is uneven. Effective programs focus attention where it is most needed, and pay attention to individual user behaviors.

4. Is reinforcement delivered at the moment of risk?

- Do users receive immediate feedback when they engage with a suspicious message?
- Is that feedback tied to the specific signals they missed?
- Is training delivered separately from real-world interactions?

Why it matters:

Behavior changes in context. Delayed, vague feedback can't prepare the workforce for the next novel attack.

---

5. Does the program adapt continuously to threats and behavior?

- Are new attack patterns incorporated into training as they emerge?
- Is training adjusted based on user behavior over time?
- Or is content updated only on fixed schedules?

Why it matters:

Static programs fall out of alignment with evolving threats. Bad actors equipped with AI improve their methods by the hour; human risk management must keep up with that pace.

---

6. Is execution automated and operationally scalable?

- Are simulations generated and delivered without manual campaign setup?
- Can training operate continuously without heavy administrative overhead?
- Or does the program rely on manual scheduling and content management?

Why it matters:

Manual processes limit consistency and scale. Autonomous human risk management ensures that security culture is enforced without increasing operational burden for the SOC.

---

7. Are outcomes measured using reliable behavioral signals?

- Are metrics based on real user interactions rather than inferred activity?
- Does the system avoid distorted signals (e.g., automated "ghost clicks")?
- Or are outcomes primarily measured through completion rates and simulation results?

Why it matters:

Inaccurate data leads to incorrect conclusions about risk. Accurate reporting ensures behaviors are truly changing and lowers the chance of the workforce merely learning how to pass the test.

Conclusion

Security awareness training remains an important component of modern security programs. But in the age of AI-powered threats, it has become clear that awareness alone is not sufficient to reduce risk at the human layer.

Today's attacks are designed to exploit behavior, context, and trust. They are personalized, adaptive, and often indistinguishable from legitimate business interactions. In this environment, the effectiveness of security programs depends not only on what employees know, but on how they act in specific moments under real-world conditions.

This has led to a broader shift in how organizations approach human risk.

Rather than focusing solely on delivering training, security teams are placing greater emphasis on understanding how risk manifests across individuals, roles, and interactions. This includes gaining visibility into behavior, aligning training with real threats, reinforcing decisions at the moment they occur, and accurately measuring how outcomes change over time.

Human risk management, when operationalized in this way, becomes more than an extension of awareness programs. **It becomes an ongoing discipline reflecting the dynamic nature of both human behavior and the threat landscape.**

Organizations that adopt this model move beyond training as a compliance activity and begin managing human risk as a system defined by real attacks, in-the-moment coaching, and continuous, behavioral signal-driven adaptation. While this continues to fulfill (and improve) compliance expectations, it also brings the crucial benefit of measurably changing how employees respond to novel threats.

About Abnormal AI

Abnormal AI is the leading AI-native human behavior security platform, leveraging machine learning to stop sophisticated attacks and detect compromised accounts across email and connected applications. Our anomaly detection engine leverages identity and context to analyze normal behavior and assess the risk of every cloud email event—detecting and stopping sophisticated, socially-engineered attacks that target your organization's most valuable cybersecurity asset: your people.

You can deploy Abnormal in minutes with an API integration for Microsoft 365 or Google Workspace and experience the value of the platform instantly. Additional protection from Abnormal is available for Slack, Workday, ServiceNow, Zoom, and multiple other cloud applications. Abnormal is currently trusted by thousands of organizations, including more than 25% of the Fortune 500, as it continues to redefine how cybersecurity works in the age of AI.

Reduce human risk with adaptive coaching grounded in real attacks and user behavior

Schedule a Demo ›

See how one organization replaced fragmented SAT workflows with unified, AI-native human risk management

Read More ›