New Cyber Risk Forecast: AI Is Eroding BEC Controls and Driving Non-Linear Financial Loss

An executive brief on how AI is eroding business email compromise controls, accelerating the hybrid spray-to-target model, and driving non-linear financial loss.

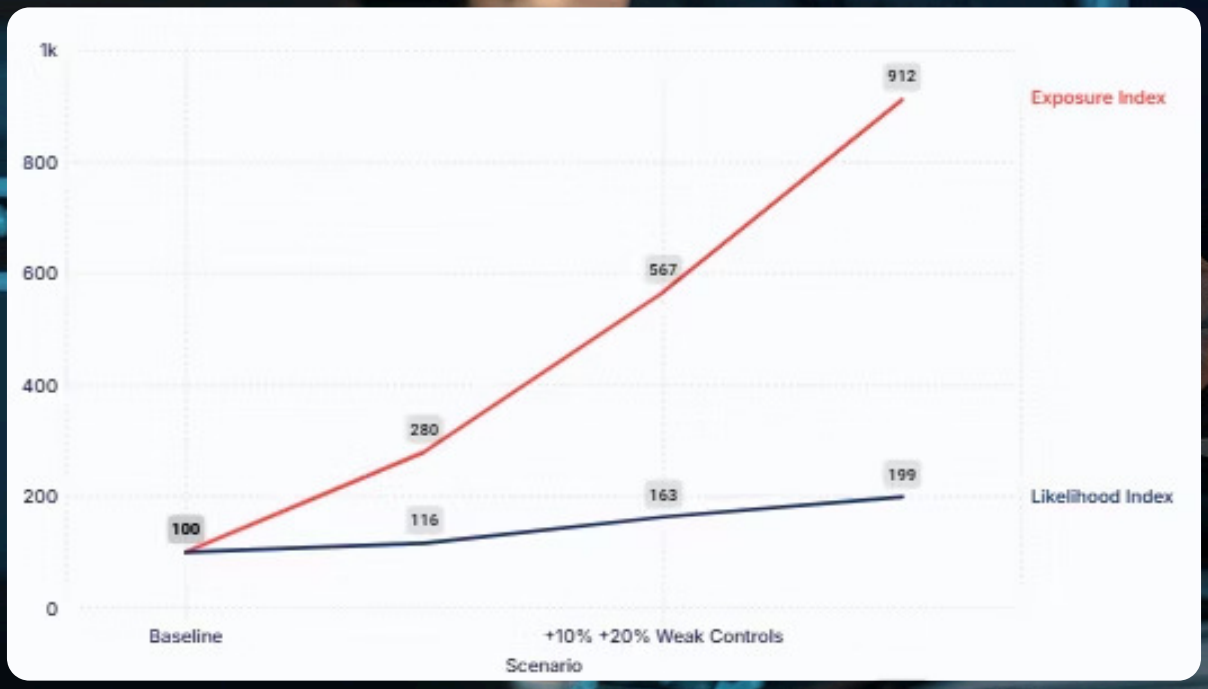

Executive Summary

We are in the middle of a transition of both the methodology and impact of business email compromise (BEC) attacks driven by AI. There is no timeline for this transition — it can happen at any moment in any business. That's because **a single click changes the attack**. This change in attack methodology **reduces the effectiveness of current protections** while simultaneously **increasing the likelihood of multiple compromise events**.

⚠ CISOs and CEOs need to understand: the primary mitigation is **human verification at the decision point** in their financial process — not technology.

Key Findings

AI Eroding Controls

AI is measurably reducing the real-world effectiveness of BEC controls across industries.

Credibility Shift

The main risk shift is more credible, payment-relevant interactions that bypass normal human judgment.

9× Financial Exposure

Model results suggest **financial exposure could increase by up to 9×, while attack likelihood may double.**

Execution Layer Gap

Verification and execution-layer controls now matter more than technological defenses.

What We Are Measuring

Our cyber risk forecast modeling uses public signals to identify the top attacks across industries and to measure baseline control effectiveness at the business level. We then develop scenarios with explicit likelihood and impact values to quantify expected financial exposure. **Our most recent analysis indicates** a **measurable reduction in BEC control effectiveness**, which we attribute to AI-enabled campaigns.

Defining Control Effectiveness

In this report, "control effectiveness" is not a maturity score — it is the **observed probability that controls prevent a BEC attempt from progressing to a financial outcome**.

What Baseline Reflects

- Historical performance against known attack methods
- Typical (not optimal) implementation within the industry
- Outcomes observed through reported incidents up to the baseline period

When a Control Is Degraded

A control is considered **ineffective or degraded** if current attack methods bypass or neutralize it. Based on recent business-level cyber risk forecasts, we are seeing reduced effectiveness of BEC controls.

ⓘ 86% of phishing is now AI-generated, according to current statistics — the primary driver of control degradation.

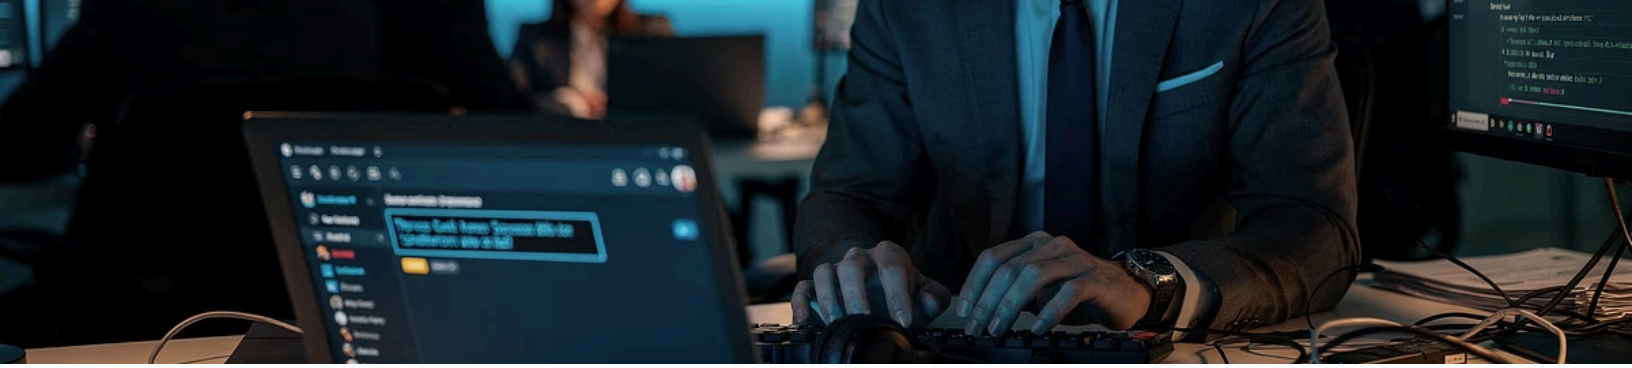

How BEC Is Accelerating

AI has been used to gather intelligence from open-source information for over a decade, beginning with spear phishing in the 2010s. However, the ability to automatically collect, correlate, and transform that data into deep organizational intelligence — and directly integrate it into phishing or BEC campaigns at scale — has only emerged in the past few years.

1 — 2010s

AI used for open-source intelligence gathering; early spear phishing campaigns emerge.

2 — Early 2020s

Automated correlation of organizational data begins enabling targeted BEC at limited scale.

3 — 2022–2024

Deep organizational mapping from open data becomes scalable and operational in cybercrime.

4 — 2025–2026

Hybrid spray-to-target model fully operational; **AI-driven BEC indistinguishable from legitimate communications.**

The Hybrid BEC Environment

We are currently in a hybrid BEC environment: attackers begin with broad outreach, but they do not invest targeted effort until they have proof that a transaction pathway is reachable. **The shift to targeted BEC is triggered by user interaction and exposure** — signals that reduce attacker uncertainty and provide the context needed to craft a payment-relevant pretext.

- The **trigger can be any interaction**: a reply, a forwarded invoice thread, a signature block revealing role and contact paths, an out-of-office delegate, or evidence of urgency and authority.

The Spray-to-Target Pivot Point

The trigger is the pivot point of the entire attack model. Before the trigger, defenses are largely probabilistic. After the trigger, the attack becomes targeted and process-centric.

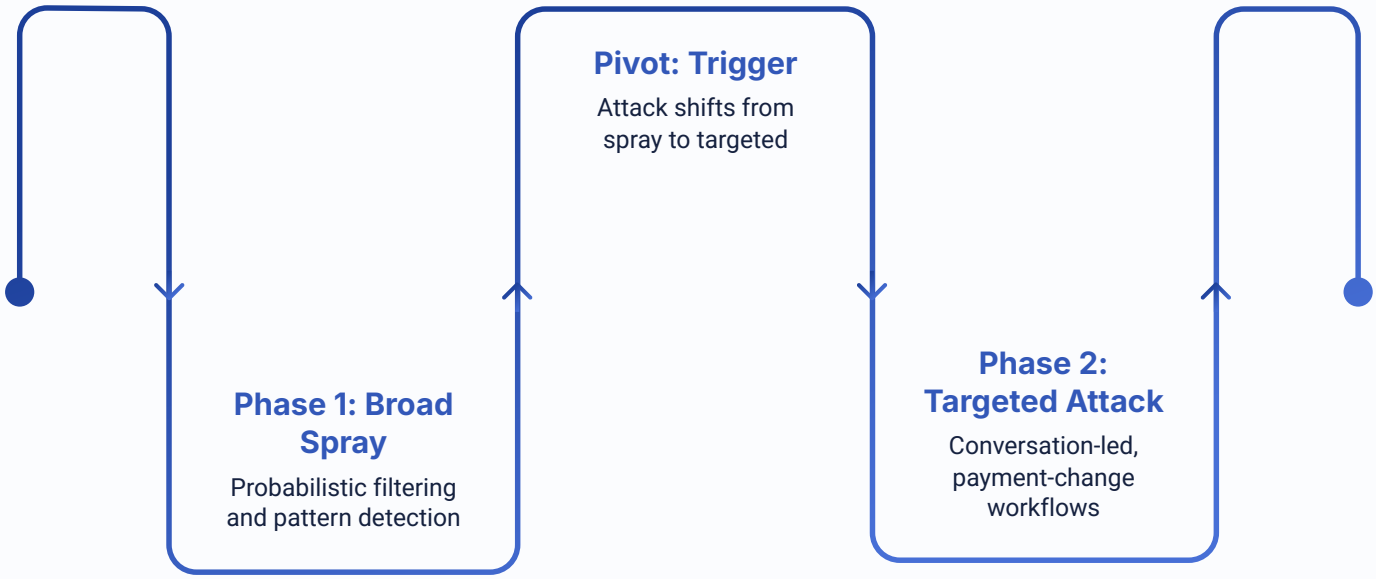

This pivot is where control effectiveness degrades most noticeably — because the attack transitions from pattern-based detection to conversation-based manipulation. The business's payment controls and verification discipline, not the inbox, determine the outcome after this point.

Before the Trigger

- Filtering and detection are active
- User caution provides a layer of defense
- Pattern-based detection catches broad attempts
- Probabilistic defenses are largely effective

After the Trigger

- Attack becomes targeted and process-centric
- Low detection footprint bypasses technology controls
- Informal human cues are neutralized
- Payment controls and verification discipline determine outcome

Why Traditional Controls Fail

In an AI-driven BEC environment, the primary point of failure is no longer the inbox — it is **human verification under pressure**. Email security and awareness reduce volume, but once a request is plausible and conversational, outcomes depend on whether people follow an independent verification path before money moves.

1

Prevention-Only Controls Are Insufficient

Filtering and detection will never be perfect in a high-variation, malware-free fraud problem. In the hybrid model, some attempts will reach users. The decisive question is what happens after engagement — especially for payment changes and exceptions.

2

The Click-Through Rate Problem

With active anti-phishing, click-through rates could be as low as **5–7%**. **AI-driven targeted messages have a 45% click-thru rate** which is the same as human-written messages. That's because AI removes obvious "tells," and the targeted messages are indistinguishable from valid communications.

3

Identity Assumption Exploitation

BEC exploits identity assumptions: "it looks like them" (sender name, thread, voice) is treated as proof. In an AI era, the control must be an **enforced, independent verification path** — known-good call-back, dual approval, enforced hold — that cannot be satisfied by replying to the same thread.

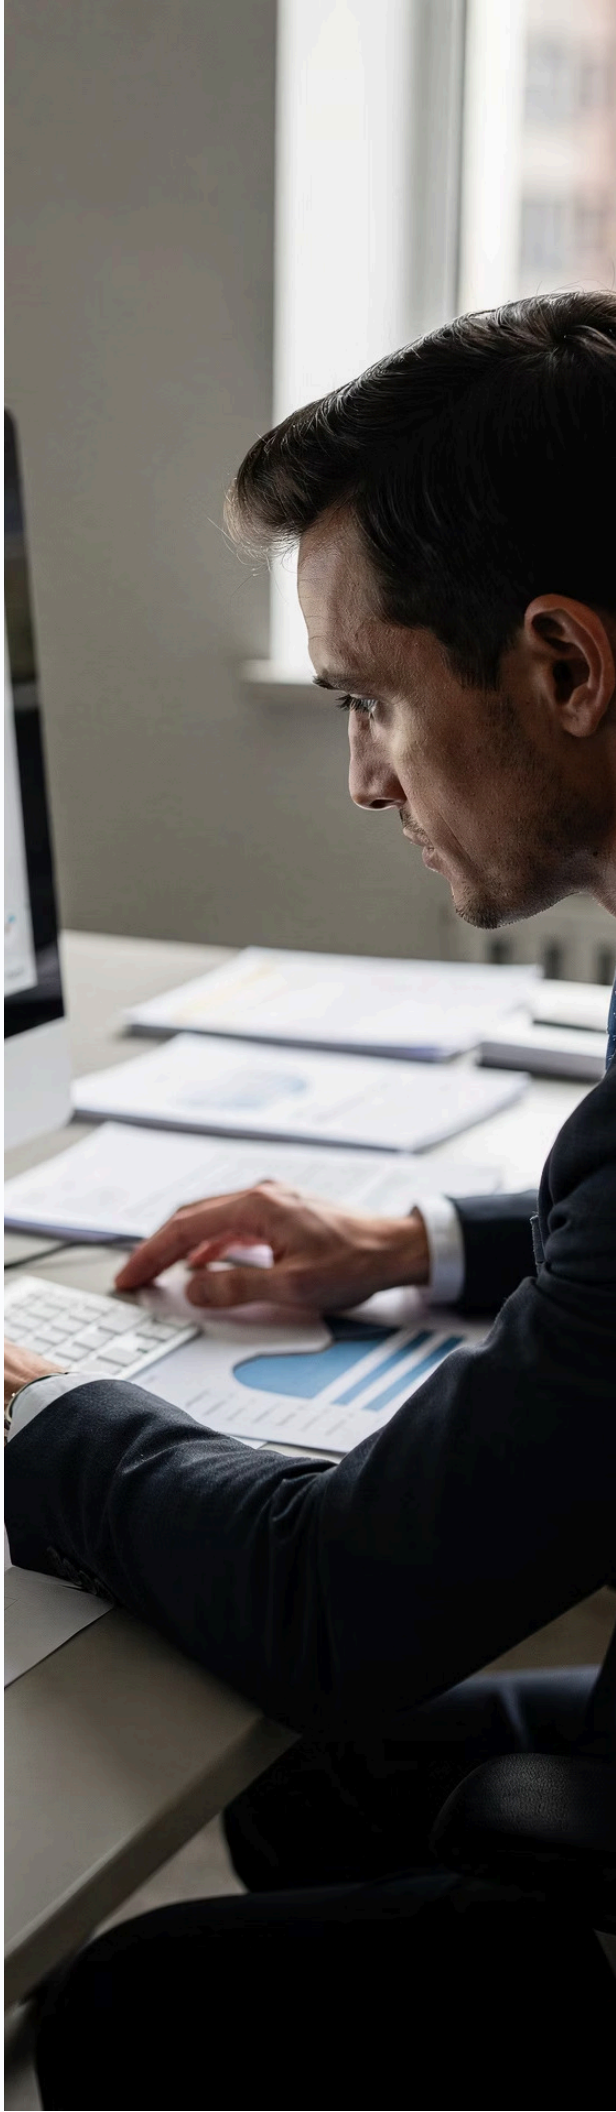

The Missing Control: The Execution Layer

BEC loss occurs when the organization executes an irreversible action — wire/ACH transfer, vendor master change, payroll redirect — based on a communication. This is **the execution layer, and it is the most critical and most overlooked control surface.**

Irreversible Actions Are the Target

Attackers specifically target first-time payments, beneficiary and bank changes, and exception approvals — places where urgency and ambiguity are common and reversibility is low.

Enforced Holds and Dual Control

The highest-leverage mitigation is execution-layer control: **enforced holds, dual control, and independent verification using known-good paths that do not rely on the initiating message or channel.**

Independent Verification Paths

Verification must use a known-good contact path — not contact details provided in the request. Calling back on a number from the original vendor record, not the email, is the standard.

The execution layer is the **missing control**. Enforced verification at the point of financial action is the primary mitigation in an AI-driven BEC environment.

Modeling Accelerated Financial Impact

We use a scenario-based normalization approach anchored to a baseline business profile. Rather than predicting exact future conditions, the model evaluates controlled variations from baseline so **we can isolate AI-driven shifts in likelihood, realized control effectiveness, and multi-event exposure.**

Below is an excerpt from the **Accounting and Audit business profile**. It was **selected because BEC is a top risk for this sector and recent analysis indicates reduction in control effectiveness for BEC**. The top row is the business-level baseline "typical" control implementation. The next three rows represent the scenarios we model for AI impact on financial outcome.

	Forecast Likelihood	Scenario Likelihood	Baseline Exposure	Multiplier x 2	Multiplier x 3	Multiplier x 4	Updated Exposure
Baseline	High (61–80%)	30.5% – 40%	$0 M – $1 M				$0 M – $1 M
Increased Likelihood (10%)	High-Very High (71–90%)	35.5%-45%	$0M - $1.4M	$0M - $2.8M			$0M - $2.8M
Increased Likelihood (10%) and weak controls (20%)	High-Very High (71–90%)	49.7% - 63%	$0M - $1.89M		$0M - $5.67M		$0M - $5.67M
Increased Likelihood (15%) and weak controls (30%)	Very High (76–95%)	60.8%-76%	$0M - $2.28M			$0M - $9.12M	$0M - $9.12M

Definitions

- **Forecast Likelihood** is the industry-level likelihood of the event in the next 90 days
- **Scenario Likelihood** is the business-level likelihood based on business profile
- **Business Profile** is a combination of 3 characteristics that determine how likely a business is to be a target; exposure, data attractiveness, and control maturity.
- **Baseline Exposure** is the 'typical' financial impact based on the last 90 days reported data
- **Event Multiplier** is a factor representing multiple independent financial outcomes

Rather than applying a single fixed multiplier, we evaluate a range based on structural conditions such as transaction volume, number of approval paths, workflow fragmentation, and the friction created by verification and holds.

- **2× → limited multi-event exposure**
- **3× → moderate operational scale**
- **4× → high-volume or fragmented workflows**

$$\text{Updated Exposure} = \text{Likelihood} \times \text{Event Multiplier} \times \text{Exposure}$$

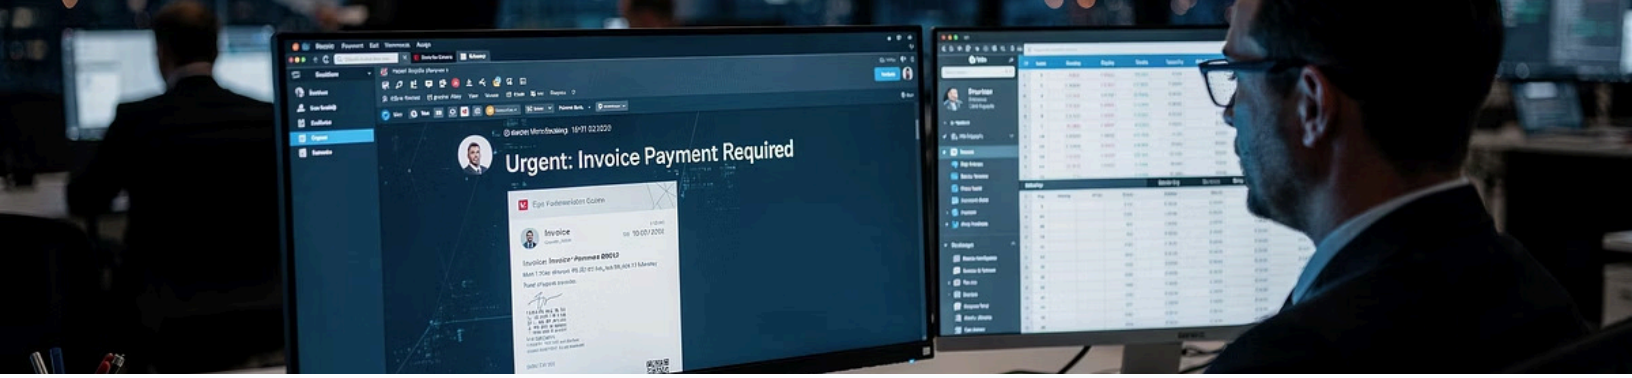

Normalizing And Indexing The Data

At a high-level normalization is putting everything on the same scale so comparisons are meaningful. The scenarios and multiplier combinations preserve interpretability by separating independent AI effects (likelihood uplift, control degradation, and multi-event exposure) so we can attribute changes in expected exposure to specific drivers and avoid double counting.

To put each scenario on the same scale, we will index the data. Indexing means converting values relative to a baseline (usually set to 100). The typical equation is **Index = (Current Value / Baseline Value) × 100.**

For the ease of modeling (graph) we will use single values; the low end of likelihood and high end of exposure. This encompasses the full range of each scenario.

Scenario	Description	Likelihood (%)	Updated Exposure ($M)
Baseline	Current-state conditions	30.5%	$1M
+10% Likelihood	Increased industry likelihood	35.5%	$2.8M
+10% +20% Weak Controls	Likelihood increase + control degradation	49.7%	$5.67M
+15% +30% Weak Controls	Combined maximum impact scenario	60.8%	$9.12M

Indexed Value = (Current Value ÷ Baseline Value) × 100

Below are the calculations as we index the data.

Likelihood Index Calculations

- Baseline: (30.5 ÷ 30.5) × 100 = **100**
- +10%: (35.5 ÷ 30.5) × 100 = **116**
- +10% +20%: (49.7 ÷ 30.5) × 100 = **163**
- +15% +30%: (60.8 ÷ 30.5) × 100 = **199**

Exposure Index Calculations

- Baseline: (1 ÷ 1) × 100 = **100**
- +10%: (2.8 ÷ 1) × 100 = **280**
- +10% +20%: (5.67 ÷ 1) × 100 = **567**
- +15% +30%: (9.12 ÷ 1) × 100 = **912**

Scenario	Likelihood Index	Exposure Index
Baseline	100	100
+10% Likelihood	~116	280
+10% +20% Weak Controls	~163	567
+15% +30% Weak Controls	~199	912

Indexed Growth Chart

By indexing both likelihood and exposure to a baseline of 100, the divergence between attack frequency and financial impact becomes starkly visible. The indexed growth chart below reveals the core finding of this analysis. **While we may expect the impact of AI driven BEC to increase likelihood by 2x the financial impact could increase by as much as 9x (as a conservative estimate)**.

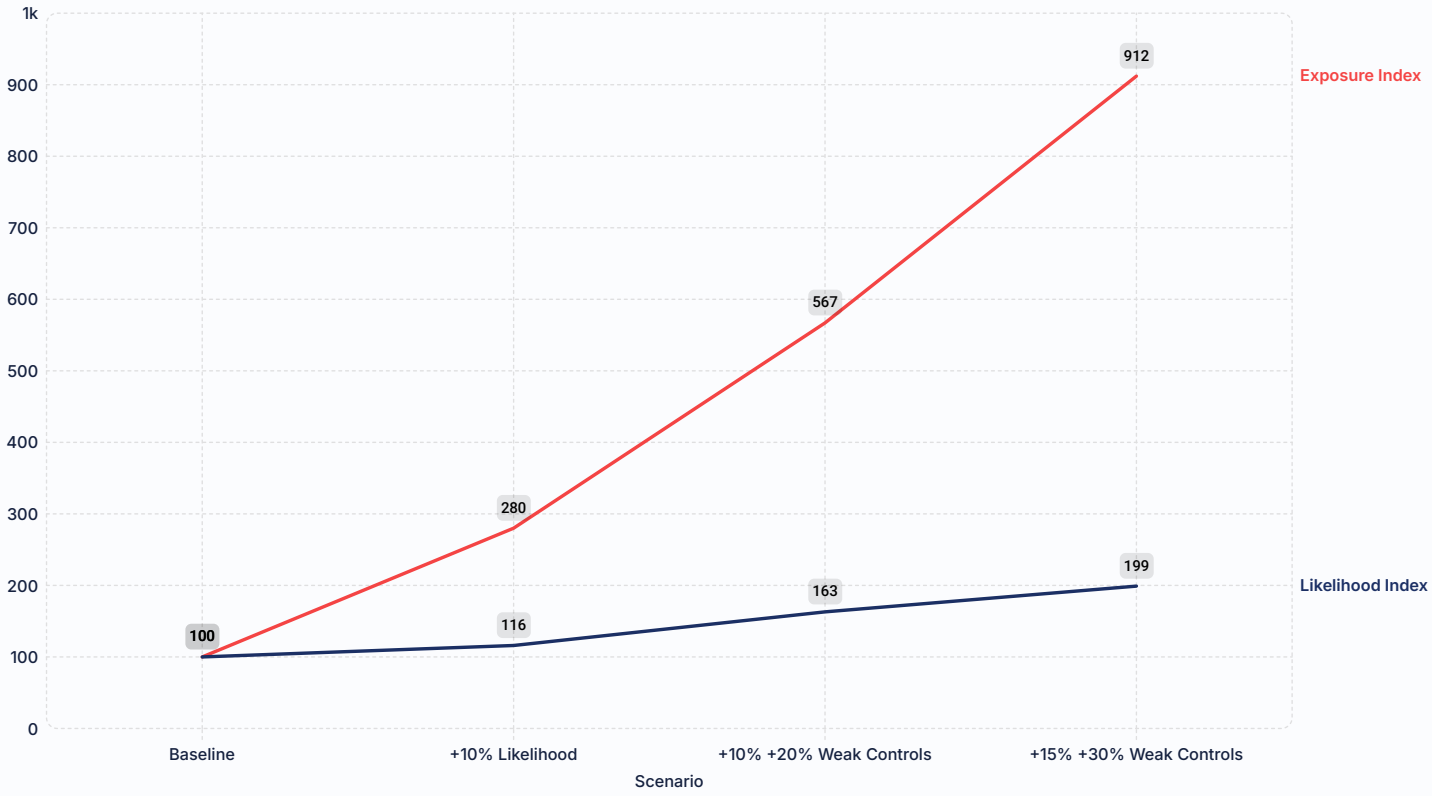

Likelihood Growth

Likelihood index grows from **100 → ~199** — roughly a **2× increase**. Attack frequency increases incrementally as AI expands targeting scale and reach, then stabilizes as adversaries shift toward precision over volume.

Exposure Growth

Exposure index grows from **100 → ~912** — roughly a **9× increase**. Even small increases in attack likelihood, combined with weakened controls, lead to exponentially higher financial risk due to multi-event outcomes within normal workflows.

⊗ The visual takeaway: while likelihood roughly doubles, financial exposure grows approximately **9 times**. This non-linear relationship is the defining characteristic of AI-driven BEC risk.

Findings & Takeaways For Executives

This model provides a realistic framework for forecasting accelerated financial impact from AI-driven BEC. The findings reframe how organizations should think about BEC risk and control investment.

1 AI Does Not Affect Risk Linearly

Likelihood increases early then stabilizes as adversaries shift toward precision targeting and optimize for high-confidence, high-value opportunities rather than maximum volume.

2 Control Degradation Becomes the Primary Driver

As AI improves realism and persistence, the dominant driver shifts from "more attempts" to "higher success per attempt" — especially once threads enter payment-change workflows and exception pathways.

3 Financial Impact Is Driven by Multiple Outcomes

The most significant shift is not simply more attacks or larger single losses — it is more financially valid actions occurring within normal workflows, **allowing losses to accumulate through repeated or parallel payment changes.**

4 Operational Environment Determines Risk Amplification

Risk increases most where payment processes are high volume, verification is manual or fragmented, multiple approval paths exist, and vendor relationships are active — conditions under which one compromise can enable multiple independent events.

> BEC should be treated as a process-driven risk embedded in normal operations. AI primarily amplifies success and follow-through rather than only increasing volume. **Verification and execution-layer controls are now the primary security control surfaces for BEC.**

CyberRiskModels.com provides business-level cyber risk forecasts, updated and delivered to your inbox monthly.

Charlene Deaver-Vazquez | Charlene@CyberRiskModels.com