



30-MINUTE SESSION • CYBERSECURITY & COMPLIANCE

When **AI Attacks**, the Rules Change

Why “having the tools in place” no longer satisfies your cyber insurance underwriter — or the attackers.

Derek Boczenowski – CISM, CISA, CDPSE, PCI QSA, CCA

Chief Architect

Compass IT Compliance

30-MINUTE SESSION

Q&A INCLUDED

Two Cases, One Hard Lesson



JANUARY 2024 — ARUP (HONG KONG)

\$25.6M

wired in 15 transactions after a finance employee joined a video call where every other participant — including someone who looked and sounded exactly like the CFO — was an AI-generated deepfake.

The employee had doubts — and was talked out of them by “colleagues” on the call who weren’t real.



MARCH 2026 — FEDERAL COURT RULING (TEXAS)

\$874,863 lost → \$250,000 paid

Perry & Perry Builders wired two payments to a fraudster impersonating its steel vendor. A federal court held the loss capped at the policy’s single \$250,000 sublimit — regardless of how many transfers went out. Case No. 6:25-cv-00106 (W.D. Tex.).

The policy existed. It just didn’t pay what most people in this room would assume it would.

AI Removed the Old Warning Signs



< 1 hour

BREAKOUT TIME

The gap between initial breach and an attacker moving laterally through the network has collapsed to under an hour in many cases — leaving almost no window for a manual response.



RaaS

DEMOCRATIZED ATTACKS

AI-powered ransomware-as-a-service platforms let novice criminals rent enterprise-grade extortion tools — no longer requires a skilled crew.



Hyper-personalized

NO MORE RED FLAGS

Generic, typo-ridden phishing is gone. AI crafts polished, individually tailored messages, fake websites, and cloned voices that bypass the cues employees were trained to spot.

The Scale Is Not Small Anymore

700%+

Year-over-year increase in deepfake fraud

U.S. SECRET SERVICE DATA

93%

Of ransomware incidents began with a phishing interaction

IBM COST OF A DATA BREACH REPORT

146%

Quarter-over-quarter jump in QR-code (“quishing”) phishing, Q1 2026

MICROSOFT THREAT INTELLIGENCE

Translation for the room: the controls and training that worked five years ago are being outrun — which is exactly why insurers rewrote the rulebook.

From Checkbox to Technical Audit

THEN

- Self-attested yes/no questionnaire
- A signature and a premium quote
- Rarely verified unless a claim was filed
- “Do you have a firewall?”-level questions

NOW

- Evidence-based verification — screenshots, logs, reports
- External attack-surface scans before they’ll quote you
- Proof controls were active on the day of the loss
- Claims disputed — or policies voided — when attestation doesn’t match reality

A Sublimit Isn't a Typo — It's the Real Limit



Social engineering and funds-transfer-fraud coverage is almost always sublimited well below the overall policy limit — commonly around \$250,000 within a \$1M (or larger) policy, no matter how the loss occurred or how many transfers were involved.

Perry & Perry Builders lost **\$874,863**. The court held the policy's **\$250,000** cyber-crime limit applied once — the two transfers, sent one minute apart, were a single loss under one limit, not two.

A second, related trap:

Many carriers treat a deception-based transfer as a “voluntary parting” of funds — no system was technically breached, so some policies exclude it from cyber coverage and push it toward commercial crime coverage, which most organizations carry inadequately, or not at all.

Ask your broker: what is my actual sublimit for social engineering / funds transfer fraud — and does it come anywhere near a realistic loss for my organization?

Denial and Rescission Are Not the Same Thing

CLAIM DENIAL

This specific claim doesn't get paid. The rest of the policy — and any other coverage in the same term — stands.

RESCISSION



The carrier voids the entire policy retroactively, back to day one. The current claim is denied — and payouts already made in the same policy term can potentially be clawed back.

Travelers v. International Control Services (2024): a carrier sought — and won — rescission after the insured represented MFA was deployed across all systems when it was not. Courts have held the carrier doesn't need to prove the misrepresentation caused the loss; the misrepresentation itself is enough.

The Three Hard Gates



MFA

Multi-factor authentication enforced on all remote access and email accounts — not just offered, enforced.



EDR

Endpoint Detection & Response on every endpoint and server. Traditional antivirus alone no longer qualifies with most carriers.



Immutable Backups

Backups that are isolated, immutable, and restore-tested on a documented schedule — with proof of the last successful restore.

“FORMAL” HAS A SPECIFIC MEANING

What Insurers Now Ask For



Written plan, named owner

A documented incident response plan with clear roles — who decides, who communicates, who calls outside help — not tribal knowledge.



Tested, not just written

A tabletop exercise within the last 12 months, documented. Some carriers now require this as a condition of renewal.



Continuous monitoring

Evidence of active monitoring and log retention — carriers want proof controls were running before the incident, not switched on after.



Pre-approved response team

Your breach counsel and forensics vendor confirmed as approved on the policy before an incident — using an unapproved vendor mid-crisis can create its own coverage dispute.

AI Governance Is Now Its Own Underwriting Category

42%

of companies now carry an AI-related exclusion somewhere in their cyber policy — language that can apply not only when the insured deploys AI, but when the attacker used AI to execute the loss.

DELINEA SURVEY, 2026

86%

report premium discounts or credits for deploying AI-based security tools

INDUSTRY SURVEY, 2026

23+DC

states have adopted the NAIC AI Model Bulletin shaping how insurers evaluate AI risk

AS OF APRIL 2026

The dividing line at renewal isn't "do you use AI" — it's whether you can document how, and whether your policy still covers you when the attacker is the one using it.

What You Should Be Able to Produce



RESPONSE & RESILIENCE

- Written incident response plan with named ownership
- Documented tabletop exercise within the last 12 months
- Immutable, restore-tested backups — with dated logs



ACCESS & GOVERNANCE


- MFA enforced everywhere — not just available
- An AI usage inventory: every tool, every model, what data it can reach
- Visibility into critical vendors' security posture



COVERAGE LITERACY

- Your actual funds-transfer-fraud sublimit vs. realistic exposure
- Your notification deadline — often 24–72 hours to preserve coverage
- Confirmation your breach counsel is pre-approved on the policy

The 3-2-1-1-0 Rule



Zero unverified restores is the one insurers ask about most — and the one most organizations can't actually answer.

Four Things to Do This Month

1

Pull your last renewal questionnaire

Go line by line. For each item, ask: could I produce evidence today, or only a verbal “yes”?

2

Ask your broker two blunt questions

Is there an AI exclusion in my current policy — and what is my actual funds-transfer-fraud sublimit versus my realistic exposure?

3

Confirm your breach counsel is on the panel

A five-minute email this week. Don’t find out who’s approved in the middle of an incident.

4

Schedule a tabletop exercise

If it’s been more than 12 months — or never — that’s the single highest-value hour you can put on the calendar this quarter.

THE REFRAME

This was never really about the insurance form.

The same documentation that satisfies an underwriter — a tested plan, verified backups, real visibility, an honest sublimit conversation — is exactly what shortens a real incident and protects the people counting on you.

Questions & Discussion