● CYBANTAGE | HEALTHCARE & REGULATED INDUSTRIES 2026

EXECUTIVE RISK BRIEFING

The Assumption *Stack*

Why Your Cyber Insurance Safety New Has a 40% Rejection Rate – and What That Means to Your Balance Sheet

40%
CLAIMS DENIED OR PARTIALLY PAID

31.3%
BREACHED ORGS DID NOT SURVIVE

~20%
ANNUAL GROWTH IN CYBER LOSSES

EXECUTIVE RISK BRIEFING

The Assumption Stack

Why Your Cyber Insurance Safety Net Has a 40% Failure Rate — and What That Means for Your Balance Sheet

The Financial Risk No One Is Measuring

Most organizations have built their cyber risk management strategy on a three-part assumption: that compliance certifications validate their security posture, that their security team is managing the right risks, and that cyber insurance will pay when a breach occurs.

All three assumptions are structurally wrong. This paper quantifies the exposure.

Published

March 2026

Authors

Rod Andes, CISSP, CCISO, CGEIT

Cybantage

Healthcare & Regulated Industries Practice

Research basis: 1,478 breach events (Jan 2023 – Feb 2026) | SOC 2 and HITRUST framework analysis | Cyber insurance market data (NAIC, Aon, Marsh, Coalition, Beazley, Chubb) | Multi-carrier underwriting requirement compilation

Executive Summary

Every year, finance and risk officers at SMB and mid-tier organizations sign off on cyber budgets, insurance premiums, and compliance program investments on the basis of a shared assumption: that the combination of security certifications, a capable IT security team, and a bound cyber insurance policy constitutes a functioning risk management strategy.

It does not. What it constitutes is an assumption stack — a sequence of individually plausible beliefs that compound into a single catastrophic miscalculation when tested by an actual breach event. This paper examines each layer of that stack, documents its structural failure mode, and quantifies the financial exposure that results when organizations confuse the appearance of cyber risk management with its substance.

The conclusion is direct: the cyber safety net that most SMB and mid-tier organizations believe they possess does not exist in the form they imagine. Approximately 40% of cyber insurance claims are denied or only partially paid. Compliance certifications evaluate governance documentation, not adversarial durability — a company can pass a compliance audit and fail the forensic audit that follows a breach. Security teams governed by IT timelines optimize for uptime and compliance, not risk reduction, and the gap between those objectives is exactly where modern attackers operate.

Meanwhile, the industry continues spending heavily on perimeter defenses against a threat that has largely moved elsewhere. Attackers are no longer breaking in. They are logging in — using stolen, phished, or inadequately protected credentials to operate inside organizational systems as legitimate users. The controls designed to prevent this are frequently rolled out incompletely, undermined by user friction, or not enforced at the policy level. And every gap in that enforcement becomes, under forensic review at claim time, the evidence a carrier needs to deny payment.

For the CFO and risk officer, this is not a technology problem. It is a financial exposure problem that has been mislabeled as a technology problem — and that mislabeling is costing organizations their existence.

~$12B+	**40%**	**31.3%**	**98%**
projected global cyber losses in 2026, growing ~20% annually	of cyber insurance claims denied or only partially paid	of breached organizations did not survive as independent entities	of claims originate from companies under $2B revenue

The central argument of this paper

The compliance audit and the forensic audit are not the same test. Your organization may pass the SOC 2 or HITRUST audits and fail the forensic. The difference is that you discover the latter result after systems are down, data is exposed, and the carrier's forensic team is on-site — at the moment when financial relief is most urgent and least likely to arrive.

Section 1: The Assumption Stack — How Risk Accumulates in Layers

Risk management failures in cyber security rarely result from a single bad decision. They result from a chain of individually defensible assumptions that are never examined together. Each layer of the stack seems reasonable on its own. Combined, they constitute an unquantified and largely unacknowledged liability sitting on the balance sheet of every organization that has built its cyber risk strategy on compliance and insurance alone.

1.1 Layer One — The Compliance Assumption

The first assumption is that achieving and maintaining compliance certifications — SOC 2, HITRUST, ISO 27001, or equivalent — validates the organization's security posture. This assumption is understandable. Compliance programs require real effort, real investment, and real organizational change. They are externally validated by credentialed assessors. They are recognized by customers, regulators, and business partners as evidence of security maturity.

But compliance frameworks evaluate a specific and bounded thing: whether controls are documented, implemented, and evidenced according to a defined set of criteria during a defined review period. They do not evaluate whether those controls hold under adversarial pressure, they do not test whether they can withstand a modern-day attack. These assessments are intended to evaluate whether you have defensive architectures in place, not will they defend. They do not ask whether your identity architecture can withstand credential abuse, token replay, or privilege escalation.

SOC 2 attests that controls were "suitably designed and operating effectively relative to the Trust Services Criteria during the review period." HITRUST certifies that control maturity requirements are documented, implemented, and scored above threshold. Neither framework makes any claim about what happens when a sophisticated attacker attempts to traverse the organization's actual environment.

Compliance audit vs. forensic audit: different tests, different verdicts

A compliance audit asks: Do documented controls exist and is there evidence they were applied? A forensic audit asks: Were controls actually enforced at the moment of attack — and can you prove it? The first question is answered in a conference room. The second is answered after the breach, by the carrier's forensic team, with your organization's survival on the line.

The Change Healthcare breach — $3.09 billion in damages occurring while the organization held active HITRUST r2 certification — is not an anomaly. It is the proof of concept. Certification had increased organizational confidence. That confidence, never subjected to adversarial validation, had closed the feedback loop that would otherwise have surfaced the gaps. The attackers found those gaps. The forensic team documented them afterward.

1.2 Layer Two — The Security Team Assumption

The second assumption is that the organization's security function is managing the right risks. In most SMB and mid-tier organizations, this assumption conceals a structural problem: security has been delegated to IT.

This is not a personnel quality issue. The problem is structural. IT operations and security operations require fundamentally different orientations. IT governance is organized around availability, continuity, and efficiency — keeping systems up, keeping users productive, keeping projects on schedule. Security governance requires adversarial thinking — modeling how an attacker would traverse the environment, abuse identity infrastructure, and operate inside systems as a legitimate user while evading detection.

When the same executive — typically a CIO or CTO — governs both domains, the security function inherits IT's incentive structure. Security investments are evaluated through an operational lens. Controls that disrupt workflow face organizational resistance that security-specific leadership would be positioned to manage; IT leadership has neither the authority nor the incentive to absorb that friction. The result is a characteristic spending pattern documented consistently in breach data:

- Heavy investment in perimeter and endpoint controls — visible, auditable, and operationally familiar to IT
- Compliance-driven spending that produces documentation and certifications aligned with audit timelines
- Systematic underinvestment in identity governance, detection engineering, and incident response capability — the domains that address the actual vectors responsible for the vast majority of breach events
- Minimal spend on adversarial validation: penetration testing scoped beyond compliance requirements, red team exercises, or assumed-breach simulations

The consequence, visible in organizational breach data, is that security budgets are concentrated on a narrow slice of the actual attack surface. Network servers and email systems — both structured around identity and authentication — account for 88% of breach entry points in the analyzed population. Physical security perimeter and traditional endpoint controls account for less than 3%. Organizations directing the majority of their security investment at the 3% are not making a random error. They are making the error their governance structure almost guarantees.

1.3 Layer Three — The Insurance Assumption

The third and most financially consequential assumption is that the cyber insurance policy constitutes a functioning safety net. For the CFO, this assumption represents a specific balance sheet calculation: the organization has transferred the financial risk of a cyber event to its carrier, and in the event of a breach, the policy will pay.

This calculation is wrong at a rate of approximately 40%. Across multiple credible analyses of 2024–2025 claim outcomes, the most defensible range for claims denied or only partially paid is 40 to 44%. This is not a marginal edge-case rate. It represents the systematic outcome of a market that has corrected.

Carriers absorbed record ransomware losses in 2020 and 2021. They responded by raising premiums, tightening coverage conditions, and investing heavily in forensic investigation capability. They now verify, under forensic conditions after a breach, whether the organization's security program matches what was represented on the insurance application. In most cases involving a denied claim, if it does not.

The insurance policy is not a security instrument. It is a financial instrument with specific performance conditions.

Those conditions are evaluated forensically, after the breach, at the worst possible moment. The question the carrier is answering is not whether you had good intentions or reasonable efforts. The question is whether the documented controls on your application were actually enforced at the time of the event. If they were not — and in most SMB and mid-tier environments, critical controls have enforcement gaps — the claim faces denial.

Section 2: How the Stack Fails — The Mechanics of Claim Denial

Understanding how the assumption stack fails requires understanding the mechanism by which cyber insurance claims collapse. It is not random. It follows a predictable pattern rooted in the gap between what organizations believe their security posture is and what forensic investigation reveals it to be.

2.1 The Application Misalignment Problem

Every cyber insurance application asks detailed questions about security controls: Is MFA enforced across all remote access and email systems? Is EDR deployed on all endpoints? Are backups immutable and tested? Does the organization maintain a documented and exercised incident response plan?

Organizations answer these questions based on what they believe to be true about their environment — which is typically based on IT reporting filtered through a governance lens rather than adversarial testing. The gap between believed state and actual state of these controls is, in most cases, substantial.

Multi-factor authentication is the clearest example. An organization may have MFA enabled across its primary systems and report full enforcement with complete accuracy — because IT has measured enrollment rates and policy settings. What IT has not measured, and what the forensic investigation will surface, is the legacy VPN endpoint that was exempted during rollout due to operational friction, the contractor access portal that was never brought into scope, or the administrative account that was flagged as too complex to migrate. Any of those gaps is sufficient to deny coverage for a breach that entered through that vector.

2.2 The Enforcement Gap

The enforcement gap between policy and practice is the primary mechanism of claim denial, and it originates almost entirely from the same source: user friction and operational pushback on security controls.

Security controls that disrupt workflow generate complaints that travel upward through operational channels to IT leadership, where they are resolved through exemptions, delayed rollouts, and policy modifications that reduce friction. In an IT-governed security program, these resolutions are operationally rational. The security team reports strong enrollment rates; the compliance audit confirms policy documentation; the insurance application reflects the policy as written.

The forensic investigation does not examine the policy. It examines what was enforced at the moment of the event. The exemption that IT approved eighteen months ago, the phased rollout that never completed, the policy that exists in documentation but not in technical enforcement — these are the artifacts that populate the forensic report and form the basis of claim denial.

Denial Ground	Forensic Finding Pattern
MFA not universally enforced	82% of denied claims involved organizations without enforced MFA on all critical access vectors. Having MFA on most systems is not sufficient — a single unenforced access point is the entry vector.
EDR incomplete or misconfigured	65% of carriers require true EDR with active monitoring. Traditional antivirus does not qualify. Partial deployment — servers covered, contractor endpoints not — fails forensic review.
Backup integrity unverified	94% of ransomware incidents involved attackers targeting backup systems. Carriers require tested, immutable, isolated backup with documented restore validation. A backup that exists but has never been tested is not a backup for claim purposes.
IR plan never exercised	Documentation of an incident response plan is not the same as evidence it was exercised. Carriers now require tabletop exercise records or equivalent evidence of plan testing.
Representations not current	AI-driven underwriting tools scan public-facing infrastructure at application and renewal. Discrepancies between claimed controls and observable posture can void warranty provisions in the policy.

2.3 The Billion-Dollar Math Problem

For the CFO, the denial rate is not the only number that matters. The trajectory of the market is equally significant. Global cyber losses have grown at approximately 20% annually over recent years, with projections that place the market well above $10 billion in annual losses within the current planning horizon. The cyber insurance market has responded with premium increases, coverage restrictions, and exclusions — but the loss growth rate has consistently outpaced the market's capacity to absorb it.

The practical implication for SMB and mid-tier organizations is a safety net that is simultaneously shrinking in coverage and increasing in cost, while the underlying loss exposure continues to grow.

The organizations in this segment — which generate 98% of cyber insurance claims — are paying increasing premiums for decreasing certainty of payment.

~20%	**+40%**	**40–44%**	**3–5x**
annual growth in global cyber losses	increase in reported claims in 2024 alone	claims denied or partially paid — systematic, not exceptional	uninsured loss typically exceeds the insured payout even on paid claims

The final column in that table deserves emphasis. Even when claims are paid, the total economic cost of a cyber event typically runs three to five times the insured payout. Uninsured costs — business interruption not covered by the policy, reputational damage, customer attrition, regulatory penalties, leadership time — are not captured in claim statistics and not recoverable from the carrier. The cyber insurance policy, even when it pays, covers a fraction of the actual loss.

Section 3: The Spend Paradox — Billions Aimed at the Wrong Target

One of the most consequential failures in organizational cyber risk management is not a failure of effort or investment. It is a failure of targeting. The industry — including vendors, consultants, compliance frameworks, and regulatory guidance — spent the better part of two decades building a comprehensive defensive architecture around a single threat model: the attacker who is trying to break in.

That attacker still exists. But the dominant threat has moved. Modern attackers are not breaking in. They are logging in.

3.1 The Credential Economy

Credential compromise has become the primary initial access vector for the vast majority of significant cyber events. Phishing campaigns harvest credentials at scale and at low cost. Password reuse across personal and professional accounts means that a breach of a low-value third-party platform can yield valid credentials for high-value organizational systems. Credential stuffing tools automate the testing of harvested credentials across thousands of targets simultaneously. MFA bypass techniques — SIM swapping, adversary-in-the-middle phishing, push notification fatigue attacks — have matured to the point where organizations that believe MFA is protecting them may be wrong.

The forensic data from the 1,478 organizational breach events analyzed for this paper is direct on the point. Email systems and network servers — both environments organized around identity and authentication — account for 88% of breach entry points. Attackers are not defeating perimeter controls. They are authenticating past them.

> **Modern attacks do not look like intrusions. They look like logins.**
>
> An attacker operating inside a network with a valid set of stolen credentials generates activity that looks, to monitoring tools calibrated for intrusion detection, identical to legitimate user behavior. The attacker is not breaking rules the system enforces. They are following rules — as a ghost resident of the environment, moving laterally, escalating privilege, and positioning for impact with no technical anomaly to detect.

3.2 Where Security Budgets Are Going Instead

Against this threat landscape, organizational security budgets remain heavily weighted toward the perimeter and endpoint layer — the layer that defends against an attacker trying to break in. Firewall management, intrusion detection, perimeter monitoring, endpoint protection, and network security tooling represent the majority of discretionary security spend in most SMB and mid-tier organizations.

These are not wasted investments. The perimeter must be maintained. But their concentration reflects a risk model calibrated to a past threat, not the present one. The identity and access management layer — the environment where 88% of breach events originate — remains systematically underfunded. Detection engineering that would identify credential abuse patterns within normal-looking authentication activity receives minimal investment. Incident response capability built for the assumed-breach scenario, where the attacker is already inside, is rare.

The result is a security program that produces excellent compliance documentation — auditable, certifiable, reportable — while leaving the actual attack surface largely undefended. The spending is visible. The results are not.

3.3 The Control Enforcement Deficit

The spend paradox is compounded by a control enforcement deficit that operates at every level of the security program. Controls are purchased. Controls are deployed. Controls are documented. Controls are not enforced uniformly — and the gap between documented enforcement and actual enforcement is where claims go to die.

The enforcement deficit originates in a structural tension between security and operations that is never resolved in organizations where IT governs security. Every security control that touches user behavior generates resistance. MFA adds authentication steps. Password requirements create complexity. Access controls restrict workflow flexibility. EDR monitoring slows endpoints. Each of these frictions generates complaints that travel up the operational chain to IT leadership, where the path of least resistance is an exemption, a phased rollout delay, or a policy modification.

Security-specific leadership — a CISO or equivalent with the authority and mandate to absorb that friction — is the mechanism by which enforcement is maintained against operational pressure. In organizations where that role does not exist independently of IT, enforcement consistently loses to operations over time. The policy says one thing. The technical enforcement state says another. The forensic investigation reads the technical enforcement state.

Section 4: The Forensic Audit — Where the Safety Net Gets Tested

The moment a cyber insurance claim is filed, the organization faces a different kind of audit than any it has previously experienced. The compliance audit is conducted in advance, with preparation time, by assessors whose mandate is to evaluate whether controls meet a defined standard. The forensic audit is conducted under crisis conditions, with no preparation time, by investigators whose mandate is to determine exactly what happened and whether the organization's security program matched the representations made in the insurance application.

These are fundamentally different examinations. Understanding the forensic audit — what it looks for, how it differs from compliance assessment, and what it finds in most SMB and mid-tier environments — is essential context for any financial officer making cyber risk decisions.

4.1 What the Forensic Team Is Looking For

Carrier forensic teams are not looking for evidence of effort or good faith. They are building a factual record of what controls were and were not in place at the time of the event, and comparing that record against the warranty representations embedded in the policy. The investigation is adversarial in the precise legal sense: it is building a factual record that will be used to support a coverage position.

The investigation systematically examines the same control domains that the insurance application addressed: authentication enforcement, endpoint protection state, backup integrity and isolation, incident response documentation and exercise history, and vulnerability management practice. For each domain, the investigation seeks the same artifact: contemporaneous technical evidence of the state of enforcement at the time of the breach.

Policy documentation is not contemporaneous technical evidence. A written MFA policy does not establish that MFA was enforced on the specific access vector the attacker used. A backup policy does not establish that backups were tested and immutable. An incident response plan does not establish that the plan was exercised and that the organization knew how to execute it under pressure. The forensic investigation requires technical artifacts — configuration exports, enforcement logs, testing records, monitoring data — not policy documents.

4.2 What the Forensic Team Typically Finds

In organizations where claims are subsequently denied or reduced, forensic investigations consistently surface the same categories of finding:

Finding Category	Forensic Observation	Denial Basis
Authentication gaps	MFA policies exist but enforcement is incomplete across all access vectors; legacy systems or contractor portals are excluded from scope	Denial basis: warranty misrepresentation on MFA enforcement scope
Endpoint coverage gaps	EDR is deployed on primary endpoints but contractors, remote workers, or server classes are outside coverage; traditional AV remains in place on some systems	Denial basis: EDR representation materially overstated
Backup integrity failures	Backups exist but have not been tested for restoration; backup systems were accessible to the attacker and encrypted or deleted in the attack	Denial basis: backup representation not substantiated by evidence
Unexercised IR plans	Incident response plans are documented but no tabletop or live exercise record exists; response during the event was disorganized and delayed containment	Denial basis: IR plan representation not supported by exercise evidence
Vulnerability management gaps	Critical vulnerabilities in internet-facing systems were known but unpatched beyond the SLA documented in policy; the breach entered through a known unpatched vulnerability	Denial basis: vulnerability management representation contradicted by forensic evidence

The pattern visible in this table is not coincidental. The forensic investigation is structured to evaluate exactly the controls the application required. The denial basis is derived directly from the gap between application representation and forensic finding. In most cases, the organization did not intentionally misrepresent its posture. It accurately reported what it believed to be true — and what IT had reported to it. The enforcement gap was real; it simply had not been measured by anyone who would have reported it adversarial.

4.3 The AI Underwriting Escalation

Forensic review now begins before the claim is filed. Several major carriers have deployed automated tools that scan public-facing organizational infrastructure at the time of application and at policy renewal. These tools identify observable security posture discrepancies — open ports on sensitive services, expired certificates, misconfigured authentication endpoints, publicly visible system information — and flag them against application representations.

Discrepancies identified at this stage can result in application denial, warranty exclusion for specific control claims, or premium surcharges. More significantly, AI-driven continuous monitoring post-binding — tied to real-time security posture assessment — is the market's stated direction of travel. Organizations that maintain their current approach of declaring controls at application and not revisiting enforcement state until the next compliance audit cycle will find themselves increasingly unable to obtain coverage, or holding policies riddled with exclusions that preclude payment for the most likely loss events.

Section 5: The Survival Dimension — What Happens When the Stack Fails

Abstract risk quantification has limited impact on executive decision-making. Specific outcomes are more clarifying. The organizational survival analysis conducted as part of the research underlying this paper provides the most direct answer to the question finance and risk officers should be asking: when the assumption stack fails, what happens to companies like ours?

5.1 The Survival Finding

Of the 1,478 healthcare provider and business associate organizations that reported major breaches between January 2023 and early 2026, 31.3% did not survive as independent entities. This figure encompasses permanent organizational closure, bankruptcy filings, forced merger or acquisition under financial distress, and material operational incapacitation that effectively ended the organization's capacity to function.

Nearly one in three. Not as a projection or a modeled estimate. As a documented outcome from a population of real organizations that experienced a major breach event during the study period.

31.3% is not a risk probability. It is a historical frequency.

For a CFO building a risk model, the distinction matters. This is not the probability that a breach will occur, which is a separate and arguably more tractable question. This is the probability that an organization which has already experienced a major breach will not survive it as an independent entity. One in three of your peers who faced this scenario are gone.

The survival data intersects with the insurance denial data in a way that is financially important. An organization that experiences a major breach and receives full insurance payment is in a materially different financial position than one that experiences an equivalent breach and faces claim denial. The denied claim does not simply reduce recovery — it removes the primary financial mechanism through which the organization was expecting to fund its response, recovery, and continuity operations.

For an SMB or mid-tier organization, a denied claim combined with the full economic cost of a major breach — direct costs, business interruption, regulatory penalties, legal fees, reputational damage — is frequently not a setback from which the organization recovers. It is an extinction event.

5.2 The Scale Asymmetry

The survival risk is not evenly distributed. Large enterprises have resources, legal capacity, and balance sheet depth that allow them to absorb even a significant denied claim and continue operating. SMB and mid-tier organizations do not. Yet this is precisely the population from which 98% of cyber insurance claims originate.

The market structure of cyber risk is therefore deeply asymmetric. The organizations most exposed to organizational non-survival from a combination of breach and claim denial are the organizations with the least capacity to self-fund recovery and the least leverage in disputes with carriers. The safety net, such as it is, is most likely to fail precisely where the fall is most dangerous.

This asymmetry has a specific financial implication: the value at risk on a CFO's balance sheet is not the probability of a breach multiplied by the claim amount. It is the probability of a breach multiplied by the probability that the claim is paid, multiplied by the partial recovery that a paid claim represents, offset against the full economic cost of the event. When that calculation is run with accurate inputs — including a 40–44% denial rate and a 3–5x uninsured loss multiplier — the residual financial exposure in most SMB and mid-tier organizations is substantially larger than the number sitting in the risk register.

5.3 The Business Associate Dimension

For organizations that function as vendors, technology partners, or service providers to larger enterprises, the survival risk carries an additional dimension. A breach event affecting your organization can trigger notification obligations, contractual penalties, and client attrition that operate independently of whether your insurance claim is paid. Business associates in the study population faced disproportionately severe outcomes — not because the breach itself was worse, but because the downstream contractual and reputational consequences compounded the direct financial impact.

An organization that carries HITRUST certification or SOC 2 attestation and experiences a significant breach faces a particular form of reputational exposure: the certification that was supposed to provide assurance to clients becomes, in the breach narrative, evidence that certifications do not predict breach prevention. Managing that narrative while simultaneously managing incident response, regulatory notification, carrier interaction, and operational recovery requires resources and leadership bandwidth that most mid-tier organizations cannot sustain simultaneously.

Section 6: The Financial Officer's Risk Reframe

The argument of this paper has a direct implication for how CFOs and risk officers should model and manage cyber risk. The current dominant model — compliance program investment plus cyber insurance premium as risk transfer — is not functioning as modeled. The actual financial exposure is different in kind and magnitude from what the model represents.

6.1 Reframing the Risk Register

A cyber risk register entry that reads "Risk: major breach event. Mitigation: cyber insurance policy. Residual risk: policy deductible" is not an accurate representation of the organization's exposure. An accurate entry would read:

> **Risk: major breach event.**
>
> Mitigation: cyber insurance policy (estimated payment probability: 56–60%, based on 40–44% industry denial rate).
>
> If claim is paid: recovery of approximately 20–33% of total economic loss (based on 3–5x uninsured cost multiplier).
>
> **If claim is denied: full economic loss with no financial backstop.**
>
> Probability of organizational non-survival: approximately 31% based on documented outcomes in comparable breach events.
>
> **Current residual risk on balance sheet: material and unquantified.**

The gap between the first version and the second version of that risk register entry is the unquantified liability that this paper is designed to surface. Finance officers who manage risk for a living should find the first version of that entry professionally uncomfortable.

6.2 The Questions the Board Should Be Asking

Based on the structural failure modes documented in this paper, the following questions represent a minimum due diligence threshold for boards and risk committees evaluating their organization's cyber risk position:

- What percentage of our critical access vectors have MFA technically enforced — not enrolled, not policy-compliant, but enforced such that there is no path to those systems without completing a second factor?

- When was the last time our incident response plan was exercised under simulated breach conditions, with documented results? Does the carrier's underwriting team have evidence of that exercise?
- Are our backup systems isolated from our production environment in a manner that would survive a ransomware event designed to destroy backup capability? When were those backups last successfully restored in a tested recovery exercise?
- If our security controls were examined forensically today — by an investigator whose objective was to identify gaps in documented representations — what would they find?
- What is the specific basis on which we believe our cyber insurance claim would be paid? Has that belief been tested against the carrier's current underwriting standards, not the standards that applied when the policy was bound?
- Does our security leadership have adversarial thinking capacity — the ability to model how an attacker would traverse our environment — or is security governed by people whose primary expertise is operational IT?

These are not rhetorical questions. They have specific, verifiable answers. An organization that cannot answer them with specificity and evidence has not managed its cyber risk. It has documented its cyber risk management activity — which is a different thing, and not the thing that will matter at claim time.

6.3 The Direction the Market Is Moving

The market is not moving in a direction that favors organizations that continue on the current trajectory. Several structural shifts are already underway that will tighten the gap between compliance-based security postures and insurable security postures:

- Carrier underwriting standards that were aspirational best-practice guidance two years ago are now binding policy conditions. The pace of that migration is accelerating.
- AI-driven pre-binding infrastructure scanning is already deployed by multiple major carriers. Post-binding continuous monitoring tied to real-time posture assessment is the stated direction of development across the market.
- Regulators in healthcare and financial services are strengthening enforcement frameworks in ways that create personal liability for board members and executives who cannot demonstrate adequate cyber risk oversight — independent of insurance status.
- The frequency of large cyber events is increasing. The probability that any given SMB or mid-tier organization experiences a significant breach event in a five-year window is higher today than at any prior point in the market's history.

Organizations that wait for the market to force the issue will face it after the breach — in the forensic investigation, in the claim denial, or in the survival statistics. Organizations that confront it before the breach have the opportunity to close the gap between their assumed posture and their actual posture before the test that matters.

Section 7: What Actually Needs to Change

The purpose of this paper is not to argue that compliance certifications are worthless, that cyber insurance should be abandoned, or that IT security teams are incompetent. None of those conclusions follows from the evidence. Compliance programs build governance discipline. Cyber insurance provides meaningful financial protection when coverage conditions are met. IT security teams do important work.

The argument is more specific: these elements have been assembled into a risk management architecture that does not function as its designers believe, at a price point that has become financially material, against a threat that has evolved faster than the architecture has adapted. Three structural changes address the core failure modes.

7.1 Distinguish Between What Compliance Measures and What Security Requires

Compliance frameworks measure governance maturity. They are designed to do that, they do it well, and organizations should continue pursuing and maintaining appropriate certifications. But governance maturity and adversarial resilience are different properties, and conflating them is the origin of the assumption stack.

The practical change is the addition of adversarial validation as a distinct and independent layer of security assurance. This means penetration testing scoped to actual threat scenarios — not the compliance-scoped annual exercise that satisfies HITRUST requirements — and assumed-breach simulations that answer the question: if an attacker were already inside our environment with valid credentials, how long before we detect them, and what would they be able to reach before containment?

These exercises are not expensive relative to the financial exposure they address. They are, however, uncomfortable — because they produce findings that compliance audits do not. The organization that has never subjected its identity architecture to adversarial testing is the organization that will discover its gaps in the forensic investigation.

7.2 Align Security Leadership With Risk Rather Than Operations

The structural conflict between IT timelines and security risk management cannot be resolved by asking IT leadership to manage it differently. It requires governance separation — a security leadership function with independent authority and reporting lines that do not flow through the operational IT structure.

This does not require a full internal CISO function in every SMB organization. It requires that whoever governs the security program has adversarial thinking capacity, is evaluated against risk outcomes rather than compliance outputs, and has the organizational authority to sustain control enforcement against operational friction without being overruled by the operational team whose friction is being generated.

The consequence of not making this change is not abstract. The forensic record will reflect the enforcement decisions that IT made under operational pressure. Those decisions, not the policy documents that preceded them, will be the basis of the carrier's coverage position.

7.3 Test the Safety Net Before the Breach Does

Every organization with a cyber insurance policy should conduct a pre-breach forensic simulation: an examination of the organization's actual control enforcement state, conducted from the perspective of a forensic investigator building a record for a carrier's coverage determination.

This examination asks different questions than a compliance audit. It asks not whether the MFA policy exists, but whether MFA is technically enforced on every access vector that would be material to a breach event. It asks not whether backups are documented, but whether those backups could be restored today and whether they are isolated in a manner that would survive a ransomware event. It asks not whether the incident response plan is current, but whether the organization can demonstrate that it has been exercised.

The output of that examination is not a compliance artifact. It is a pre-breach gap analysis that answers the only question that will matter if the organization experiences a significant breach: would the carrier pay?

An organization that can answer that question with evidence and confidence is in a materially different risk position than one that cannot. The insurance policy in both cases may read identically. The probability of payment does not.

The test you do not conduct before the breach, the forensic team will conduct after it.

The compliance audit tells you whether your documentation passes a governance standard. The forensic audit tells you whether your controls held against an actual attack. For the CFO, only one of those audits has a material effect on the balance sheet.

Conclusion: The Safety Net Cannot Be Counted On

The central financial reality this paper documents is this: the cyber safety net that most SMB and mid-tier organizations believe they possess does not perform as modeled. The compliance certifications that underpin confidence in the security program do not evaluate adversarial resilience. The security team that manages the program is governed by incentives that optimize for compliance output rather than risk reduction. And the insurance policy that is supposed to catch the organization when the other elements fail has a 40–44% probability of not paying, at a moment when the organization's financial continuity depends on it.

These are not predictions about the future. They are documented outcomes from the recent past. One in three organizations in the studied population that experienced a major breach did not survive it. The ones that fared worst were the ones whose assumption stacks were tallest — the most certain, before the breach, that their compliance posture, their security program, and their insurance policy had transferred the financial risk to someone else.

For the CFO and risk officer, the appropriate response to this paper is not a technology investment decision. It is a risk framework question: does the organization's current cyber risk management posture accurately represent the actual financial exposure on the balance sheet? If the answer is not clearly and evidentially yes, the assumption stack deserves examination before the forensic team arrives to examine it instead.

The compliance audit and the forensic audit are not the same test.

The assumption that they are is the most expensive assumption in cyber risk management.

Every organization will eventually take one of these two tests. The choice is which one to prepare for.

Appendix: Data Sources and Research Basis

Organizational Survival Analysis

Primary breach dataset: 1,478 healthcare providers and business associates reporting major breaches affecting 500 or more individuals between January 1, 2023 and early 2026. Source: U.S. Department of Health and Human Services Office for Civil Rights public breach notification repository. Excluded categories: health plans and payors, government-operated entities, educational institutions with affiliated healthcare operations, and cases where required organizational data was not accessible for analysis.

Survival determination: organizational status assessed as of March 2026 through public records, regulatory filings, acquisition announcements, and closure notices. Non-survival categories include confirmed closure, bankruptcy filing, material acquisition under financial distress, and operational incapacitation sufficient to end independent organizational function.

Cyber Insurance Market Data

Denial rate range (40–44%): derived from triangulation across Network Assured industry reporting, Marsh McLennan cyber claim analysis (2024–2025), and Coalition Cyber Threat Index findings. The industry does not publish standardized claim outcome statistics; the range represents the most defensible synthesis of available non-carrier-published analyses.

Premium and loss data: NAIC Cyber Supplement annual reporting, Aon U.S. Cyber Market Updates, and Marsh McLennan annual cyber market analyses. Claims volume and frequency data from NAIC 2024 preliminary reporting and Marsh mid-year 2025 assessment.

Carrier underwriting requirements: compiled from publicly available application materials and underwriting guidelines from Coalition, Beazley, Chubb, AIG, Munich Re, Zurich, AXA, CNA, Arch, AXIS, QBE, and Liberty Mutual, cross-referenced against Marsh McLennan's 2025 application analysis and Coalition's 2024 Cyber Threat Index.

Compliance Framework Analysis

SOC 2 analysis: based on AICPA attestation standards (AT-C 205) and Trust Services Criteria documentation. HITRUST analysis: based on HITRUST CSF v11 and Assessment Handbook, PRISMA maturity methodology documentation, and HITRUST Trust Report analysis. Change Healthcare breach financial figures: derived from UnitedHealth Group regulatory filings and public statements, March 2024 through March 2026.

Attack Vector and Breach Entry Data

Entry point analysis (88% email and network servers): derived from primary analysis of the 1,478 breach event study population, cross-referenced against CISA advisory data, Verizon Data Breach Investigations Report 2024, and CrowdStrike Global Threat Report findings for the relevant period.

About Cybantage

Cybantage is a healthcare and regulated industries cybersecurity advisory practice specializing in the gap between governance compliance and operational resilience. Our work spans security program governance, adversarial validation, HIPAA and HITRUST compliance architecture, cyber insurance posture assessment, and the CISI (Cyber Insurance Survivability Index) — a 215-point, 10-domain claim payability scoring framework that evaluates both claimant-side control failures and insurer-side policy exclusion risk independently. The CISI surfaces three independent Domain 10 flags (D10-NS, D10-TP, D10-SY) that identify nation-state exclusion, third-party coverage, and systemic event coverage gaps regardless of an organization's security posture score. We work with CFOs, boards, and risk committees who need an adversarial perspective on their organization's actual security posture — not a compliance attestation.

cybantage.com

© 2026 Cybantage. All rights reserved. This paper is intended for executive and board-level audiences. Reproduction or distribution requires written permission from Cybantage. Statistical findings are based on the cited research corpus and should be read in conjunction with the methodological notes in the Appendix.