D I S C U S S I O N P A P E R

Cyber Insurance and the Compliance-Reality Gap

Examining Why Regulated Organizations Experience Claim Denials and Methods for Assessing Actual Exposure

Including the Cyber Insurance Survivability Index (CISI)
Claim Payability Scoring Framework

Cybantage • Healthcare & Regulated Industries Practice • 2025

Executive Summary

A dangerous assumption has taken root across regulated industries — healthcare providers, digital health platforms, financial services firms, and SaaS companies processing sensitive data. That assumption is simple, seductive, and increasingly expensive: that a bound cyber insurance policy is a functioning safety net.

It is not. Cyber insurance is better understood as a contractual audit conducted at the worst possible moment — under forensic conditions, after systems are down, data is exposed, and leadership is searching for financial relief. And that audit has a striking failure rate.

What makes this more than a headline statistic is what the data reveals when examined by company size and claim amount: the denial problem is not distributed uniformly. Smaller companies face dramatically worse odds of payment, and the financial exposure from a denied claim scales in ways that can threaten organizational survival. Understanding those dimensions is now a CFO-level obligation.

40%+	**3:1**	**20:1**
Claims Denied or Partially Paid	**Claims Closed Without Payment vs. Paid**	**Unpaid vs. Paid in Excess Layer Policies**
Slingshot/Fitch 2025; multiple sources [21]	*NAIC 2024: 28,555 unpaid vs. 9,941 paid* [2]	*NAIC 2025 Cybersecurity Insurance Report* [3]

These are not edge cases. They represent the systematic outcome of a market that has corrected — carriers who absorbed record ransomware losses now verify, under forensic conditions, whether your security program matches what you told them it was. In most cases, it does not.

This paper examines the structural anatomy of that failure: why claims collapse, what the market now requires, how claim size and company size fundamentally alter exposure, and how regulated organizations can measure their actual claim payability before a breach forces the question.

It introduces the Cyber Insurance Survivability Index (CISI) — a scoring framework developed by Cybantage to translate that measurement into a structured, actionable assessment that answers the one question that matters: Would a carrier actually stand behind us when it matters?

Defining Question for Every Executive Team

If a breach occurred today and every control we claim to have was examined in forensic detail — would our insurance claim be paid? If the answer is not an unqualified yes, supported by evidence, the gap between that answer and yes represents an unquantified financial liability that exists today, regardless of whether a breach has yet occurred.

1. The Market Reality: What the Data Actually Shows

The cyber insurance denial problem resists clean quantification — and that is not accidental. The industry is structurally opaque. Carriers do not publish claim outcome statistics, and the NAIC reporting framework captures aggregate market data without separating denial rates from other closure categories. But when credible sources are triangulated, a consistent and alarming picture emerges.

1.1 The Denial Rate Range

The most defensible range, supported across multiple 2024–2025 industry analyses, is that 40 to 44% of cyber insurance claims are denied or only partially paid.[1] This figure represents claims that go through the full submission and investigation process and are then rejected or reduced based on coverage conditions.

More damning is the NAIC's own 2024 data: nearly three times as many claims were closed without payment (28,555) as those that resulted in a payout (9,941).[2] In the excess policy segment — the layer most relevant to large regulated organizations — unpaid claims outnumber paid claims by more than 20 to 1.[3]

Headline Statistic

Roughly one in three to nearly one in two cyber insurance claims do not result in full payment. In excess layer policies, unpaid claims outnumber paid claims by more than 20 to 1.

1.2 What the Data Looks Like by Claim Size

The aggregate 40% denial figure, while alarming, obscures a more nuanced — and more actionable — reality. When claim outcomes are examined by size, a clear stratification emerges that fundamentally changes the financial calculus for executive leadership.

Claim Type / Size	Average Loss Amount	Payment Likelihood	Key Driver
BEC – Standalone	**$35,000**	Moderate	**Policy exclusions for employee error; notification gaps**
BEC + Funds Transfer Fraud	**$106,000**	Moderate–High	**Swift action enables clawback; control evidence critical**

[1]Slingshot Information Systems / Fitch Ratings, December 2025. Cyber insurance claim rejection rate over 40% across 2024 and 2025 industry data. https://www.slingshotis.com/blog/cyberinsurance-claim-denied/

[2]NAIC 2025 Cybersecurity Insurance Market Report. Of 38,496 total claims closed in 2024, 28,555 were closed without payment vs. 9,941 paid.

[3]NAIC 2025 Cybersecurity Insurance Market Report. Excess policy segment data.

Claim Type / Size	Average Loss Amount	Payment Likelihood	Key Driver
Ransomware (avg. paid)	**$100,000 (Coalition avg.)**	Higher — if controls proven	**MFA enforcement, backup integrity determinative**
Ransomware (median ransom)	**$2,000,000**	Variable — policy limits often insufficient	**Coverage cap frequently exceeded; partial payment common**
Full recovery w/o ransom	**$2.7M–$3.0M**	At Risk — limits rarely adequate	**Business interruption, forensics, legal, notification compound**
Data breach (IBM avg.)	**$4.88M**	Partial at best	**Multi-year exposure; excess layer denial common**
Enterprise wiper/destructive	**$3B+ (Change Healthcare)**	Contested — nation-state exclusions	**Policy language, nation-state exclusions, scale of loss**

Sources: Coalition 2024–2025 Cyber Claims Report; IBM Cost of a Data Breach 2024; Sophos 2024 State of Ransomware; NAIC 2025 Cybersecurity Insurance Report

The financial pattern is clear: smaller, simpler claims face denial for administrative and control-attestation reasons, while larger claims face the dual problem of control failure and coverage limits that were set before recovery costs reached current levels. An organization with a $2M ransomware event may discover that its $1M policy limit is already inadequate — and then lose that partial coverage when forensic investigation finds MFA gaps.

1.3 What the Data Looks Like by Company Size

Company size is the most significant predictor of claim payment probability that is not widely communicated in the market. The NAIC's 2024 data reveals a structural disparity that CFOs in small and mid-size organizations need to confront directly.[4]

Company Size	Claim Payment Rate	Typical Policy Limit	Avg. Breach Cost	Coverage Gap Risk
Enterprise (>$1B revenue)	Higher — forensic readiness generally stronger	**$10M–$100M+**	**$4.88M avg.**	Low–Moderate
Mid-Market ($100M–$1B)	Moderate — controls variable, limits often adequate	**$2M–$10M**	**$1.5M–$3M avg.**	Moderate

[4]National Association of Insurance Commissioners (NAIC), 2025 Cybersecurity Insurance Market Report, based on 2024 annual statement filings. https://content.naic.org/sites/default/files/inline-files/2025_Cybersecurity_Insurance%20Report.pdf

Company Size	Claim Payment Rate	Typical Policy Limit	Avg. Breach Cost	Coverage Gap Risk
SMB ($10M–$100M)	Lower — only ~1 in 3 claims paid (NAIC 2024)	$500K–$2M	$500K–$1.5M avg.	High
Micro (<$10M revenue)	Lowest — excess policies 20:1 unpaid	<$500K	$150K–$500K avg.	Critical

Sources: NAIC 2024 Cybersecurity Insurance Report; Coalition 2024 Cyber Claims Report; NetDiligence 2025 Cyber Claims Study; citybiz NAIC analysis Feb 2026

The NAIC's 2024 data makes the SMB reality explicit: only approximately one in three cyber insurance claims filed by small to mid-size businesses resulted in payment.[5] Insurers have tightened policy language and raised security standards in response to years of losses, and smaller organizations — which generally have less mature control environments, less documentation discipline, and less legal sophistication during the claims process — bear the consequence disproportionately.

1.4 The Financial Impact of a Denied Claim by Company Size

The financial consequence of a denied claim is not simply the loss of the claim amount. It is the compounding of every breach-related cost category onto an organization that sized its financial resilience against an assumed coverage backstop. When that backstop fails, the full cost of the event lands on the balance sheet.

Cost Category	SMB Impact (Denied Claim)	Enterprise Impact (Denied Claim)
Forensic investigation	$50K–$500K (uninsured)	$500K–$2M+ (uninsured)
Legal counsel + notification	$100K–$300K	$500K–$5M
Regulatory fines (HIPAA, etc.)	$100K–$1.9M	$1M–$16M+
Business interruption losses	Often existential for SMB	Significant but survivable
Reputational/customer loss	Can trigger client contract defaults	Manageable with IR
Class action defense	Typically not self-fundable	$2M–$50M+ over years
Total uninsured exposure	**$350K–$3M+ — often > annual revenue**	**$5M–$75M+ — manageable with reserves**

The SMB Survivability Threshold

[5]KY3/citybiz analysis of NAIC 2024 data, published February 2026. 'Only one-third of cyber insurance claims filed by small- to mid-size businesses in 2024 resulted in payment.' https://www.citybiz.co/article/805423/

For a small business with $5M in annual revenue, a $500K uninsured breach cost represents 10% of annual revenue absorbed in a single event — before accounting for lost business, customer notification costs, and the operational disruption of a weeks-long recovery. This is why roughly 60% of small businesses that experience a significant cyber incident do not survive the following 12 months.

1.5 Why This Is Happening

Claims are not denied because attacks were not real. They are denied because the organization failed a test it did not know it was taking — a forensic verification of whether its security controls matched its attestations. The specific failure categories, drawn from consistent reporting across carriers and incident response data:

- MFA was not universally enforced across all covered access points
- Privileged access controls were inconsistently governed or could not be evidenced
- Documentation did not prove what was claimed — attestation and operational reality diverged
- Policy language exclusions around 'failure to maintain controls' gave carriers legal grounds to deny
- Notification requirements were not met within required timeframes
- Application information was inaccurate — even unintentionally — and the carrier argued the policy was void

Coalition's 2024 Cyber Threat Index found that 82% of claims involved organizations without universal MFA.[6] Marsh McLennan's 2024 data found 41% of applications are denied on first submission, with missing MFA and inadequate endpoint protection as the top two reasons.[7]

1.6 The Market Has Corrected — Permanently

Early cyber insurance operated on a questionnaire model: self-attested controls, minimal validation, broad coverage. Carriers absorbed record losses through 2020–2022 and responded with structural changes that are not temporary market conditions — they represent the permanent architecture of how cyber risk is underwritten.

Then (Pre-2022)	Now (2024–2026)
Self-attested questionnaires taken at face value	**Technical validation — screenshots, log exports, policy reports**

[6]Coalition, 2024 Cyber Threat Index / Claims Report. 82% of claims involved organizations without universal MFA enforcement.

[7]Marsh McLennan, 2024 Cyber Insurance Market Analysis. 41% of applications denied on first submission; missing MFA and inadequate endpoint protection cited as top two reasons.

Then (Pre-2022)	Now (2024–2026)
MFA as best practice recommendation	MFA as binding coverage condition — phishing-resistant preferred
EDR optional for most policies	EDR on all endpoints mandatory for most carriers
Backups assumed adequate	Immutable, isolated backups with tested recovery required
IR plan accepted as documentation	IR plan must be tested — tabletop evidence required
Third-party risk acknowledged	Vendor security mapped and monitored
Premium driven by revenue and industry	Premium driven by demonstrated control maturity
AI-driven underwriting minimal	Automated scanning of public-facing assets pre-binding

2. The Quiet Signal: What Coverage Denial Before a Breach Means

There is a phenomenon that occurs earlier than a breach — quieter and more revealing — that most organizations misread or ignore entirely. It is the moment a broker cannot place coverage.

When a broker says 'the market is tight right now' or 'we're having trouble finding a carrier' or 'we need to revisit some controls before placement,' most organizations hear a procurement problem. They should hear a survivability warning.

2.1 What Is Actually Being Said

Brokers do not get paid to place bad risk. Underwriters do not get paid to accept it. When coverage cannot be placed, it is frequently because the carrier has already run the mental simulation of how a claim would go — and concluded it would not hold.

- The gap between stated controls and actual controls is too wide for the carrier to underwrite
- The organization cannot produce defensible evidence of control operation
- The risk profile suggests a high probability of post-claim dispute
- Missing baseline requirements — MFA, EDR, isolated backups, tested IR — make the application non-bindable

Survivability Warning

If you cannot get coverage, it is not just an insurance problem. It is a survivability problem. The same gaps preventing coverage are the ones that lead to breaches, break response efforts, collapse recovery, and leave you funding the entire event yourself.

2.2 Three States of Insurability

Every regulated organization exists in one of three states — and most believe they are in a better state than they actually occupy:

State	Status	What It Means
STATE 1	Insurable and Defensible	Coverage binds. Controls verified. Claims likely to pay. The organization can produce evidence that matches its attestations under forensic scrutiny.
STATE 2	Insurable but Fragile	Coverage binds. Claims questionable. Controls documented but not consistently operated. A forensic investigation will find gaps — the outcome depends on how significant they are.
STATE 3	Uninsurable	Coverage does not bind — because claims would not pay. The broker's 'market difficulty' is a pre-claim denial. The same gaps

State	Status	What It Means
		that prevent coverage will drive breach severity and collapse recovery.

Most organizations believe they are in State 1. Current data suggests a significant portion are in State 2, and a material number of those who receive quiet coverage declines are in State 3. The CISI framework in Section 5 provides a structured mechanism for determining which state your organization actually occupies.

3. The Forensic Anatomy of Claim Failure

Denied claims rarely fail on exotic grounds. They fail on fundamentals — control categories that appear repeatedly in applications, frameworks, and regulatory requirements. The following represents a composite of post-breach forensic investigation findings and carrier denial rationale, consistent across the 2022–2025 period.

3.1 The Six Primary Failure Modes

Control Domain	Common Attestation	Forensic Reality
Multi-Factor Authentication	**MFA enforced across all remote access**	**Legacy systems, VPNs, partner portals, SaaS apps excluded. SMS-based MFA accepted where phishing-resistant required.**
Privileged Access	**Privileged access managed and monitored**	**Stale accounts, shared credentials, no access review cadence. Service accounts unmanaged.**
Endpoint Protection	**EDR deployed enterprise-wide**	**Agents outdated, excluded directories, BYOD/contractor devices unmanaged. No active monitoring.**
Logging & Detection	**Comprehensive logging and SIEM in place**	**Cloud workloads, SaaS platforms, OT systems not ingested. Alerts unreviewed. Log retention inadequate (<90 days).**
Incident Response	**IR plan documented and maintained**	**Plan untested. Key personnel unaware of roles. No external counsel pre-engaged. Communication protocols missing.**
Control Evidence	**Controls operating as documented**	**Cannot produce evidence of continuous operation. Audit reports from prior period submitted — insufficient for claim period.**

3.2 The Evidentiary Problem

Underlying all of the above is a systemic evidentiary gap: the inability to produce documentation proving controls were functioning as stated at the time of the incident. This can affect organizations whose controls were genuinely adequate. If you cannot prove a control was operating, the insurer's forensic team will proceed as though it was not.

Specific evidentiary requirements that carriers now routinely apply during claim investigation:

- **Admin screenshots showing active policies for email, VPN, and privileged groups — not just that the tool was purchased.** MFA enforcement
- **Dashboards listing protected endpoints by OS and policy, with evidence of active monitoring and response.** EDR coverage reports
- **Documented successful restore tests, immutability settings, storage location diagrams, and recovery time measurements.** Backup integrity

- **Change tickets and vulnerability scan summaries showing the age of critical findings at time of incident.** Patch cadence
- **Dated playbook, call tree, and tabletop summary with remediations — not just a document with a date.** Incident response
- **Evidence that logs were retained for the full required period (typically 90 days minimum) and protected from tampering.** Log retention

The Compliance Theater Problem

SOC 2, HITRUST, and ISO 27001 reports are snapshots. They confirm controls existed at a point in time. A claim investigation operates in continuous time: it determines whether controls were consistently enforced and effective from policy binding through the date of incident. An organization can hold a current SOC 2 Type II report, a bound cyber policy, and still receive a denied claim — because the audit and the investigation are asking fundamentally different questions.

3.3 The Structural Drivers of Misalignment

The compliance-reality gap is not random. In regulated industries, four structural forces predictably widen it:

Growth-Stage Scaling

Rapid growth introduces access sprawl — a proliferation of identities, credentials, and permissions that exceeds the organization's capacity to govern. Security controls designed for the organization at 50 people are not automatically adequate at 500.

Engineering Velocity

Technology organizations prize speed. Security review creates friction. The result is a production environment that continuously diverges from the documented security program.

Third-Party Dependency Proliferation

In 2024, 35.5% of all data breaches originated from third-party compromises.[8] Modern regulated platforms operate within ecosystems of dozens to hundreds of vendor integrations.

Security Team Under-Resourcing

Security teams in growth-stage regulated organizations are structurally under-resourced relative to the environments they govern, forcing prioritization toward systems evaluated in audits rather than those most likely to be probed in an actual attack.

[8]Verizon, 2025 Data Breach Investigations Report (DBIR). Third-party involvement in breaches doubled year-over-year to approximately 30%.

3.4 The Full Taxonomy of Claim Denial: Claimant-Side vs. Insurer-Side

The industry narrative around cyber insurance denials heavily emphasizes policyholder failures — and the data supports that claimant-side gaps are the largest single denial category. But the picture is materially incomplete without acknowledging that 20–30% of denials originate from insurer-side policy exclusions that no security control can address. Understanding both dimensions is essential to accurate claim payability assessment.

Category 1: Security Control Gaps (claimant-side, ~35–44% of denials). MFA not universally enforced, EDR coverage gaps, backup integrity failures, insufficient log retention. These are the failures addressed by the CISI's first nine domains. A 2023 industry report found 44% of claims denied because organizations did not meet security requirements attested to at binding. The Travelers v. ICS case (2022) established that a single server missing MFA — one exception in an otherwise compliant environment — was sufficient grounds for full claim denial.

Category 2: Misrepresentation at Application (claimant-side, ~10–15% of denials). Inaccurate or incomplete application responses — even unintentional ones — can void coverage retroactively. Pre-existing vulnerabilities not disclosed at underwriting fall into the same category. Marsh McLennan's 2024 data found 41% of initial applications denied on first submission, with missing MFA and inadequate endpoint protection as the top two reasons.

Category 3: Procedural Failures Post-Incident (claimant-side, ~10–12% of denials). Late notification to the carrier, failure to preserve forensic evidence, unauthorized use of non-approved vendors, or failure to follow the policy's incident response protocols. An organization can have excellent security and still lose a claim by notifying its insurer three weeks after discovery rather than within the required timeframe.

Category 4: Insurer-Side Policy Exclusions (~20–30% of denials). This is the category most often absent from market assessments — and the one most relevant to regulated industries facing nation-state threat actors. Three sub-categories define the primary exposure:

- **Nation-State / War Exclusions.** Lloyd's of London mandated that all standalone cyber policies include state-backed attack exclusion clauses effective March 2023. Healthcare, financial services, and critical infrastructure are primary targets of Iranian, North Korean, and Russian state-affiliated threat actors. The Merck v. ACE Insurance litigation over NotPetya — ultimately settled in 2024 after New Jersey courts rejected the war exclusion — forced the market to codify explicit exclusion language that now appears in most policies renewed after 2022. An organization in a high-risk sector that has not reviewed its policy for this language may have coverage void for the most likely class of attack it faces.
- **Third-Party / Supply Chain Exclusions.** Standard cyber policies frequently exclude downstream losses when a vendor or cloud provider is breached — even when the policyholder had no control over the vendor's vulnerability. The Change Healthcare ransomware attack (February 2024) demonstrated this at scale: tens of thousands of healthcare organizations — hospitals, physician practices, pharmacies — suffered operational and financial losses from a breach at a clearinghouse they depended on but did not control. Total disclosed costs exceeded $3.09 billion. Third-party breach coverage must be explicitly confirmed in writing; it cannot be assumed from general policy language.

- **Systemic / Correlated Event Exclusions.** The July 2024 CrowdStrike incident — a faulty software update that crashed 8.5 million systems globally — generated between $400 million and $1.5 billion in insured losses and accelerated insurer movement toward excluding non-malicious correlated events. As organizations integrate more third-party technology infrastructure, business interruption coverage for software update failures, cloud provider outages, and widespread correlated events must be explicitly reviewed. Lloyd's itself acknowledged that cyber losses "have the potential to greatly exceed what the insurance market is able to absorb" — systematic exclusion of these scenarios is a market capacity decision, not a policyholder failure.

The coverage gap implication

A CISI score of 186/215 from an organization that has not reviewed its policy exclusion exposure reflects high claimant-side preparedness but unknown insurer-side risk. For a healthcare organization heavily dependent on a billing clearinghouse operating under a policy with unreviewed state-backed attack and third-party exclusion language, the effective claim payability may be far lower than the security control score suggests. Domain 10 of the CISI framework — and the three independent flags it generates (D10-NS, D10-TP, D10-SY) — exists specifically to surface this gap regardless of total score.

4. What Carriers Now Require: The 2025–2026 Baseline

The underwriting requirements that determine insurability — and more importantly, claim payability — have been codified across the market. What were once 'best practices' are now binding conditions.

4.1 The Non-Negotiable Controls

Multiple carrier applications, including Beazley, Coalition, Chubb, and others, now explicitly require evidence of the following. Absence of any of these is sufficient grounds for application denial or claim denial:

- **Required on email, remote access (VPN, RDP, cloud access), all administrative and privileged accounts, and cloud platform access. Phishing-resistant MFA (FIDO2, hardware keys) increasingly preferred.** Universal MFA
- **EDR on all endpoints:** 65% of carriers now require EDR.[9] Traditional antivirus does not qualify. Must be deployed with active monitoring and host isolation capability.
- **Immutable, isolated backups:** A third of carriers now require offline or air-gapped backups. The backup must include tested recovery — 94% of ransomware incidents involved attackers targeting backup systems.[10]
- **Not just a document — a plan with evidence of testing. Carrier applications specifically ask whether the IR plan has been exercised.** Documented and tested IR plan.
- **PAM required for business-critical systems at larger organizations, increasingly across the board.** Privileged access management
- **Continuous scanning and remediation with documented SLAs, not ad hoc responses.** Vulnerability management

4.2 Advanced Requirements for Regulated Industries

- SIEM and 24/7 SOC monitoring for active threat detection with defined response SLAs
- Third-party risk program with periodic re-assessment beyond onboarding
- Data encryption at rest and in transit with documented key management
- Network segmentation isolating patient data or sensitive regulated data from general operational environments
- Regulatory compliance evidence (HIPAA Security Rule, state privacy laws) demonstrating ongoing compliance
- Evidence-ready documentation packages: MFA enforcement reports, EDR coverage dashboards, backup restore test results, and IR artifacts

4.3 The AI Underwriting Shift

Some carriers now deploy automated tools that scan public-facing infrastructure before binding coverage. Discrepancies between claimed controls and observable infrastructure can trigger application

[9]Zero Networks / industry survey, 2024. 65% of carriers now require EDR as a binding coverage condition.
[10]Coalition, 2024 Cyber Claims Report. 94% of ransomware incidents involved attackers targeting backup systems.

denial or warranty exclusion. Carriers using AI-driven underwriting are moving toward continuous monitoring post-binding, with policy conditions tied to real-time security posture.

The Implication

The gap between your attestation and your actual environment is increasingly measurable from outside your organization. The question is whether you measure it first.

5. Cyber Insurance Survivability Index (CISI)

The Cyber Insurance Survivability Index (CISI) is a claim payability scoring framework developed by Cybantage to help regulated organizations measure the actual alignment between their security operations and the conditions their cyber insurance coverage requires.

It is not a compliance checklist. It is not an audit instrument. It is a forensic simulation: a structured assessment of whether your organization could withstand the level of scrutiny a claim investigation would trigger, applied before a breach forces the question.

What CISI Measures

CISI answers a single question: If a breach occurred today and every control we claim to have was examined in forensic detail — would our insurance claim be paid? It scores across eight domains weighted by their relative frequency as drivers of claim denial. A high CISI score does not mean you will not be breached. It means that if you are, coverage will respond.

5.1 Framework Structure (Updated March 2026 — Revised Domain Weightings)

The CISI comprises 10 domains, 34 assessment questions, and a total score of 215 points. Domain point values were refined in March 2026 to reflect updated forensic weighting based on real-world claim investigation patterns. Domains are weighted by their forensic significance — the frequency with which each appears as a driver of claim denial or partial payment in post-breach investigations and carrier reporting. The framework explicitly scores both claimant-side security control gaps (which drive approximately 60–70% of denials) and insurer-side policy exclusion risks (which drive 20–30% of denials) — a dimension absent from most market assessment tools.

#	Domain	Points	Primary Denial Risk
1	Identity & Access Control	30	Highest-frequency breach vector and denial trigger
2	Endpoint Protection & Detection	25	Coverage gaps expose control misrepresentation
3	Backup & Recovery Integrity	25	Where survivability and payouts break under ransomware
4	Logging, Monitoring & Evidence	20	Evidentiary failure makes all other controls unprovable
5	Control Reality vs. Documentation	30	Compliance theater — where audit and forensics diverge
6	Incident Response Capability	25	Operationalization gap: documented vs. functional
7	Third-Party & Supply Chain Risk	15	35% of 2024 breaches originated from third parties
8	Policy Alignment & Attestation Accuracy	20	Silent denial trigger — application inaccuracy voids coverage

#	Domain	Points	Primary Denial Risk
9	Data Classification & Exposure Scope	10	Carriers control beach scope narrative without this
10	Policy Exclusion & Coverage Gap Risk	15	Insurer-side denials that no security control can prevent — 20–30% of all claim denials
TOTAL		215	

5.2 Domain Assessment

Each domain is assessed through a structured set of questions with fixed-value scoring: 0 (not in place), partial credit, or full credit. Scoring is based on demonstrated operational practice, not documentation alone. Evidence requirements are implicit in each criterion.

DOMAIN 1: IDENTITY & ACCESS CONTROL · 25 Points

Primary breach vector and denial trigger. Coalition data: 82% of claims involved organizations without universal MFA.

#	Assessment Question	Score	Scoring Criteria
1.1	Is MFA enforced universally across all access points: email, VPN, remote access, all admin/privileged accounts, SaaS platforms, and cloud consoles? Is phishing-resistant MFA (FIDO2, hardware key) used for privileged access?	0/5/10	**0** — Partial or inconsistent enforcement; gaps in coverage **5** — Broadly enforced with documented exceptions; SMS-only on some privileged accounts **10** — Universal enforcement, no exceptions; phishing-resistant for privileged access
1.2	Are privileged accounts fully governed? Are shared credentials eliminated? Is a formal access review conducted at defined intervals (quarterly minimum)? Are all stale/orphaned privileged accounts removed within defined windows?	0/5/10	**0** — Unmanaged or shared; no formal review process **5** — Partially controlled; reviews occur but not consistently **10** — Fully governed; no shared credentials; quarterly review documented
1.3	Does an automated joiner/mover/leaver (JML) process govern identity lifecycle? Are departing employees' accounts deprovisioned within 24 hours? Are access rights reviewed and right-sized on role change?	0/5	**0** — Orphaned accounts present; manual or delayed deprovisioning **5** — Automated JML process with documented SLAs; no orphaned accounts in evidence

1.4	Are session token lifetimes restricted and subject to revocation? Is conditional access policy enforced based on device compliance, location, and risk signals? Is there evidence that compromised tokens can be invalidated in real time?	**0/5**	**0 —** No session controls; tokens long-lived or non-revocable; no conditional access **5 —** Token lifetimes defined and enforced; conditional access operational; revocation demonstrable on demand

DOMAIN 2: ENDPOINT PROTECTION & DETECTION · 20 Points

65% of carriers now require EDR. Deployment without active management is a coverage gap.

#	Assessment Question	Score	Scoring Criteria
2.1	Is EDR/XDR deployed on all endpoints including servers, workstations, laptops, and contractor devices? Are agents actively managed, updated, and monitored? Can you produce a coverage report showing all in-scope devices?	**0/5/10**	**0 —** Partial deployment; legacy AV only on some devices **5 —** Majority of endpoints covered; some gaps in BYOD or contractor devices **10 —** Full coverage; active monitoring with documented response SLAs; current coverage report available
2.2	Are EDR alerts actively reviewed and responded to? Is there a defined escalation path for detections? Is there evidence of actual alert triage and response actions — not just tool deployment?	**0/5**	**0 —** Alerts generated but not managed; no defined response process **5 —** Active monitoring with documented response times; alert-to-action evidence available
2.3	Is vulnerability scanning conducted on a continuous or at minimum quarterly basis? Are critical/high vulnerabilities remediated within defined SLAs (e.g., critical within 30 days)? Is patch cadence documented and evidenced?	**0/5**	**0 —** Ad hoc or reactive patching; no scanning cadence **5 —** Continuous or quarterly scanning; defined and evidenced SLAs for remediation
2.4	Is the network segmented to limit lateral movement between critical systems, user endpoints, and administrative infrastructure? Are east-west traffic controls in place? Can segmentation boundaries be demonstrated to a forensic investigator?	**0/5**	**0 —** Flat network architecture; no meaningful segmentation; lateral movement unrestricted **5 —** Documented segmentation with enforced controls; east-west traffic monitored; demonstrable to forensic examiner

DOMAIN 3: BACKUP & RECOVERY INTEGRITY · 20 Points

94% of ransomware events involve attackers targeting backups. Backup integrity is where survivability and payouts break.

#	Assessment Question	Score	Scoring Criteria
3.1	Are backups fully isolated from the production environment? Are backups immutable (cannot be modified or deleted by ransomware)? Are they stored offline or in air-gapped infrastructure? Is the 3-2-1 rule implemented?	0/5/10	**0** — Backups accessible from production network; no immutability **5** — Partially segmented; some isolation but not fully air-gapped **10** — Fully isolated, immutable backups; offline/air-gapped copy confirmed
3.2	Are backup restores tested on a defined schedule? Is there documented evidence of successful full-system restore tests (not just incremental)? When was the last restore test conducted and documented?	0/5	**0** — Backups untested; no restore test documentation **5** — Regular (minimum quarterly) tested restores with documented evidence and sign-off
3.3	Are RTO and RPO targets formally defined, documented, and validated through testing? Are they aligned with coverage requirements? Is the organization confident it can meet stated recovery timelines under actual attack conditions?	0/5	**0** — RTO/RPO assumed or undocumented; never tested **5** — Formally defined, tested, and validated; documented evidence of meeting stated targets
3.4	Are encryption keys for backup data stored independently of the production environment and protected with hardware security modules or equivalent controls? Is key recovery feasible without access to the compromised primary environment?	0/5	**0** — Keys stored in production environment; would be compromised alongside data in a ransomware event **5** — Keys independently stored with HSM or equivalent; key recovery process documented and tested

DOMAIN 4: LOGGING, MONITORING & EVIDENCE · 15 Points

Evidentiary failure makes all other controls unprovable. Log retention and forensic readiness are investigated in every denied claim.

#	Assessment Question	Score	Scoring Criteria
4.1	Is logging centralized across all critical systems including cloud environments, identity providers, SaaS platforms,	0/5	**0** — Fragmented logging; critical systems not ingested; cloud/SaaS excluded

	network devices, and endpoints? Are all systems ingested into a SIEM? Are gaps documented with compensating controls?		**5** — Centralized SIEM with comprehensive coverage; documented scope and gap remediation
4.2	Are logs retained for a minimum of 90 days (12 months preferred for regulated environments)? Are they protected from tampering or deletion? Is retention aligned with carrier requirements and regulatory obligations?	**0/5**	**0** — Incomplete or unreliable retention; logs below 90-day minimum or unprotected **5** — Full retention (90+ days) across all in-scope systems; integrity protected
4.3	Can the organization reconstruct a complete event timeline from log data in the event of an incident? Has forensic readiness been tested? Is there a defined log review and alert response process with evidence of operation?	**0/5**	**0** — Cannot reconstruct events; no forensic readiness **5** — Full event traceability demonstrated; forensic exercise conducted; documented chain of custody
4.4	Are identity plane events specifically captured from all identity providers, including cloud directory services, SSO platforms, and privileged access management systems? Is retention for these logs aligned with the policy period, not just the 90-day minimum?	**0/5**	**0** — Identity plane events absent from SIEM; cloud directory and SSO logs not ingested **5** — All identity provider logs ingested; retention meets or exceeds policy period; forensic investigators can trace identity events to source

DOMAIN 5: CONTROL REALITY vs. DOCUMENTATION · 30 Points

The compliance theater domain. Where SOC 2 passes and claims investigations diverge. Highest weight because it underlies every other domain.

#	Assessment Question	Score	Scoring Criteria
5.1	Are security controls enforced technically, not just documented in policy? Is there evidence that controls are operating — not just that they are written down? Can the gap between documented policy and operational reality be demonstrated to be minimal?	**0/5/10**	**0** — Controls primarily on paper; significant gap between policy and operation **5** — Controls partially enforced; some operational evidence available **10** — Fully enforced; continuous validation; operational evidence readily available for all primary controls

5.2	Is there a continuous monitoring process that validates control effectiveness in real time — not just at audit cycles? Are deviations from policy detected and remediated between audits? Is control drift measurable?	0/5	**0 —** Audit-only validation; no visibility between assessment periods **5 —** Continuous monitoring with defined alerting on control degradation; documented evidence
5.3	Is ownership assigned for every critical control? Is accountability clear at the individual level? Are owners informed of their responsibilities and able to evidence performance against them?	0/5	**0 —** Unclear or unassigned ownership; no individual accountability **5 —** Defined ownership for all critical controls; accountability documented and evidenced
5.4	Can the organization produce, on short notice, a defensible evidence package that proves controls were operating at the time of an incident? Does this evidence exist independently of audit artifacts?	0/5	**0 —** Evidence is manual, inconsistent, or primarily audit-cycle artifacts **5 —** Continuous evidence capture; defensible package assembler; available on demand
5.5	Has penetration testing or red team assessment been conducted within the last 12 months against production systems? Are findings tracked to remediation? Does the organization use external testing results as input to control attestation, not just internal validation?	0/5	**0 —** No external adversarial testing in the past 12 months; controls validated by self-assessment only **5 —** Annual penetration test or red team exercise against production; findings remediated with tracked closure; results feed attestation process

DOMAIN 6: INCIDENT RESPONSE CAPABILITY · 25 Points

The plan is not the capability. Carriers distinguish organizations with documented IR from those with tested, operationalized IR.

#	Assessment Question	Score	Scoring Criteria
6.1	Has the incident response plan been formally tested through a tabletop exercise within the last 12 months? Are the results documented with identified gaps remediated? Is the test scenario representative of likely attack vectors (ransomware, BEC, third-party compromise)?	0/5	**0 —** Plan never or not recently tested; no tabletop documentation **5 —** Tested within 12 months; documented findings; remediation tracked and closed
6.2	Are all key personnel aware of their roles and able to execute them without	0/5	**0 —** Roles unclear; personnel unaware; plan not operationalized

	reference to the document? Are runbooks operationalized for the most likely attack scenarios? Is the plan reviewed and updated on a defined schedule?		**5 —** Operational plan; roles confirmed with personnel; scenario-specific runbooks maintained
6.3	Are external response partners pre-engaged and documented: legal counsel, forensic IR firm, PR/communications, regulatory notification counsel? Are retainers in place? Are insurer-preferred vendors identified?	**0/5**	**0 —** No pre-engaged external partners; relationships ad hoc **5 —** All key partners pre-engaged with retainers; contacts documented and tested
6.4	Is there a defined decision-making and escalation framework for incident response? Are thresholds for notifying carriers, regulators, and leadership clearly defined? Are notification timelines aligned with policy and regulatory requirements?	**0/5**	**0 —** Ad hoc decision-making; notification thresholds undefined **5 —** Structured escalation; notification timelines documented and aligned with policy/regulatory requirements
6.5	Is there a documented decision authority matrix for high-severity incidents, including who can authorize containment, ransom payment consideration, and public disclosure? Is the carrier notification requirement embedded in the decision workflow with defined time thresholds?	**0/5**	**0 —** No documented authority matrix; ransom and disclosure decisions made ad hoc; carrier notification requirements unknown or untested **5 —** Authority matrix documented and exercised; carrier notification SLAs embedded in workflow; decision chain confirmed with named individuals

DOMAIN 7: THIRD-PARTY & SUPPLY CHAIN RISK · 10 Points

35.5% of 2024 breaches originated from third-party compromises. Supply-chain breaches now account for nearly 30% of all incidents.

#	Assessment Question	Score	Scoring Criteria
7.1	Is there a formal third-party risk management program with defined assessment criteria, periodic re-assessment (not just onboarding), and escalation processes for high-risk vendors? Is vendor risk ownership assigned?	**0/5**	**0 —** No formal program; onboarding-only assessment; no ongoing monitoring **5 —** Formal program with periodic reassessment; risk tiering; documented ownership
7.2	Are all critical vendors and their access pathways mapped and monitored? Is third-party access governed with least	**0/5**	**0 —** Critical vendor dependencies unknown or unmapped; access ungoverned

	privilege and MFA? Are vendor access rights reviewed and deprovisioned when no longer required?		**5 —** Critical vendors mapped with access governance; monitored; access reviews documented
7.3	Does the organization assess the security posture of critical vendors' own subprocessors and downstream dependencies for the highest-risk vendor relationships? Are contractual security requirements flowed down to subprocessors handling the organization's data?	**0/5**	**0 —** No visibility into vendor subprocessors; downstream dependencies unknown; security requirements end at the direct vendor relationship **5 —** Subprocessor assessment required contractually for critical vendors; downstream dependencies mapped for highest-risk relationships; flow-down provisions enforced

DOMAIN 8: POLICY ALIGNMENT & ATTESTATION ACCURACY · 15 Points

The silent denial trigger. Application inaccuracy — even unintentional — allows carriers to argue the policy was issued on false premises.

#	Assessment Question	Score	Scoring Criteria
8.1	Does the cyber insurance application accurately reflect the current operating environment? Has the application been reviewed against actual control state — not documentation — before submission? Are all attestations verifiable with evidence?	**0/5**	**0 —** Application reflects assumed or hoped-for controls; gap with reality likely **5 —** Application verified against actual environment; all attestations evidenced before submission
8.2	Have controls attested to at policy binding been continuously maintained through the policy period? Is there a process for identifying control drift and updating carrier notification where material changes occur?	**0/5**	**0 —** No monitoring of control state against policy attestation; drift likely **5 —** Continuous alignment monitoring; material changes identified and carrier notification process defined
8.3	Is there a governed change management process that evaluates security control impact before technology or organizational changes are implemented? Are significant changes that affect covered controls reviewed with the broker before execution?	**0/5**	**0 —** Changes implemented without security impact assessment; controls broken by change **5 —** Change management includes security control review; broker consulted on material changes

8.4	Has the organization experienced any material changes since the policy was bound — including significant technology changes, workforce reductions affecting security functions, or new systems in scope — and were these evaluated for carrier notification obligations? Is there a recurring process for this evaluation?	0/5	**0 —** No process to evaluate mid-term material changes against notification obligations; carrier unaware of significant environment changes **5 —** Defined process for evaluating material changes; recurring review against policy terms; carrier/broker notified where obligations triggered

DOMAIN 9: DATA CLASSIFICATION & EXPOSURE SCOPE · 10 Points

Controls claim payability. Without a data classification program, carriers control the scope narrative — always to the insured's disadvantage.

#	Assessment Question	Score	Scoring Criteria
9.1	Does the organization maintain a current data classification inventory that maps where PHI, PII, and other regulated data is stored, processed, and transmitted? Is this inventory updated on a defined schedule and validated through discovery tools — not manual documentation alone?	0/5	**0 —** No data classification inventory; regulated data location unknown or assumed; scope of any breach must be reconstructed forensically at significant cost **5 —** Current data classification inventory maintained with automated discovery validation; regulated data mapped across cloud, SaaS, and on-premises environments; updated on defined schedule
9.2	In the event of a breach, can the organization rapidly determine what data was in scope — including cloud storage, SaaS platforms, and third-party environments — without relying on attacker-disclosed information or forensic reconstruction alone? Is there a defined data breach scope assessment process?	0/5	**0 —** No rapid scope determination capability; breach scope dependent on attacker disclosure or extended forensic investigation; carrier controls the narrative **5 —** Defined breach scope assessment process; data inventory enables rapid determination; third-party and cloud data locations documented and verifiable within 24 hours of incident declaration

DOMAIN 10: POLICY EXCLUSION & COVERAGE GAP RISK · 15 Points

Insurer-side denials that no security control can prevent. 20–30% of all claim denials originate here. Lloyd's March 2023 mandate requires state-backed attack exclusions in all standalone cyber policies. Change Healthcare ($3.09B) established the existential scale of third-party coverage gaps.

#	Assessment Question	Score	Scoring Criteria
10.1	Does your organization operate in a sector designated as a primary target of state-sponsored threat actors (healthcare, financial services, critical infrastructure, defense)? Has your current policy been reviewed by qualified counsel for state-backed cyberattack exclusion language — including Lloyd's-mandated exclusions effective March 2023?	0/3/5	**0** — High-risk sector; policy contains state-backed attack exclusion; scope not reviewed with counsel **3** — High-risk sector; exclusion present but scope and implications not confirmed with broker or counsel **5** — Exclusion language reviewed and confirmed with counsel; coverage scope understood; sector risk factored into limit adequacy
10.2	Does your policy explicitly cover losses from third-party vendor breaches — including clearinghouses, cloud EHR platforms, billing processors, and SaaS infrastructure? Has coverage scope been confirmed in writing with your broker, and have critical vendor dependencies been mapped against your policy's third-party provisions?	0/3/5	**0** — Standard policy; third-party coverage not confirmed; significant vendor dependencies exist **3** — Partial or sub-limited third-party coverage; vendor dependencies not fully mapped against policy scope **5** — Explicit third-party coverage confirmed in writing; critical vendor dependencies mapped against policy; gaps identified and addressed
10.3	Does your policy address non-malicious systemic outage events — cloud provider failures, software update failures such as the July 2024 CrowdStrike incident, or widespread correlated outages? Has business interruption coverage been confirmed for scenarios where the triggering event originates outside your organization?	0/5	**0** — No data classification inventory; regulated data location unknown or assumed; scope of any breach must be reconstructed forensically at significant cost **5** — Current data classification inventory maintained with automated discovery validation; regulated data mapped across cloud, SaaS, and on-premises environments; updated on defined schedule

5.3 Scoring Interpretation

Score Range	Claim Outcome	What It Means
198–215	EXEMPLARY — CLAIM DEFENSIBILITY OPTIMIZED	Security controls validated, evidence defensible, and policy exclusion exposure reviewed. Both claimant-side and insurer-side denial risks have been assessed. Carrier forensic investigation is highly likely to confirm compliance. This tier represents a negotiating advantage at renewal.
169–197	CLAIM LIKELY PAID	Strong operational security with validated data governance. Controls verified, evidence defensible. Insurer's forensic investigation likely to confirm compliance.
134–168	AT RISK — PARTIAL PAYMENT LIKELY	Coverage probable but not certain. One or more high-impact gaps. Claim may be reduced or subject to negotiation. Remediate critical domains.
105–133	HIGH DISPUTE / DENIAL RISK	Multiple control gaps. Forensic investigation will expose material misalignment. Expect challenge on coverage. Significant remediation required.
<105	CLAIM LIKELY DENIED	Critical control failures across multiple domains. Policy attestation likely inaccurate. Coverage will be contested. Organization absorbs breach cost.

CRITICAL THRESHOLD WARNING A score below 105 should be treated as an organizational emergency — not because of the insurance implications alone, but because the control failures that produce a score below 105 are the same failures that make breaches more likely, more severe, and more operationally devastating. The CISI score below 105 is a survivability signal, not just a coverage signal.

5.4 CISI Score Mapped to Company Size and Financial Exposure

The following matrix connects CISI scoring outcomes to company size and estimated uninsured financial exposure in a denied claim scenario. This is the table that belongs in a CFO briefing.

CISI Score	Company Size	Payment Probability	Typical Claim Amount	Uninsured Exposure if Denied
198–215	Any	High	Any	Minimal — coverage responds
134–163	Enterprise	Moderate–High	$500K–$5M	$0–$500K partial gap
134–163	SMB	Moderate	$100K–$500K	$50K–$300K partial denial

CISI Score	Company Size	Payment Probability	Typical Claim Amount	Uninsured Exposure if Denied
105–133	Enterprise	At Risk	$1M–$10M+	$500K–$3M+ exposed
105–133	SMB	Low	$100K–$1M	$100K–$1M fully uninsured
<105	Any	Denied	Any	Full breach cost on balance sheet
<105	Micro/SMB	Denied	$150K–$500K	Existential — potential insolvency

Note: Uninsured exposure estimates incorporate forensic investigation, legal, notification, regulatory, and business interruption costs as documented in NetDiligence 2025, IBM 2024, and Coalition 2025 reporting.

5.5 Domain Priority Matrix: Where to Focus First

Not all gaps are equally consequential for claim payability. The following matrix reflects the relative impact of domain improvement on expected claim outcome:

Domain	If Score is Low…	Priority	Immediate Action
1	**Identity & Access**	CRITICAL	Eliminate shared credentials. Enforce MFA universally. Audit privileged accounts this week. Restrict session tokens and enforce conditional access.
5	**Control Reality**	CRITICAL	Implement continuous control monitoring. Assign control owners. Build evidence capture. Schedule penetration test within 90 days.
10	**Policy Exclusion & Coverage Gap Risk**	CRITICAL	Engage insurance counsel immediately to review: (1) state-backed attack exclusion language; (2) third-party/supply chain coverage scope; (3) systemic event BI coverage. These gaps cannot be resolved through security investment — only through policy review, endorsements, or supplemental coverage.
3	**Backup & Recovery**	HIGH	Isolate backups immediately. Conduct restore test. Document RTO/RPO results. Verify encryption key independence.
8	**Policy Alignment**	HIGH	Re-review application against actual environment. Evaluate mid-term material changes. Fix discrepancies before next incident.

4	**Logging & Evidence**	HIGH	Expand SIEM ingestion to cloud/SaaS. Capture identity plane events. Verify log retention meets 90-day minimum.
9	**Data Classification**	HIGH	Initiate data discovery. Map PHI/PII across cloud and SaaS. Build rapid breach-scope determination capability.
6	**Incident Response**	MEDIUM	Schedule tabletop within 60 days. Pre-engage legal/IR counsel. Document decision authority matrix. Define ransom-payment authorization chain.
2	**Endpoint Protection**	MEDIUM	Audit EDR coverage. Close BYOD/contractor gaps. Implement network segmentation. Enable active monitoring.
7	**Third-Party Risk**	MEDIUM	Map critical vendor access and subprocessors. Implement periodic reassessment. Flow security requirements to downstream vendors.

6. The Path to Defensible Coverage: Operational Disciplines

Organizations that consistently survive breaches and recover their insurance claims share a common characteristic: they treat security as an operational discipline rather than a compliance exercise.

6.1 Continuous Control Validation (Not Audit-Cycle Compliance)

The single most consequential organizational shift is moving from assumed compliance to demonstrated compliance. This means implementing mechanisms that verify, on a continuous basis, whether controls are functioning as designed. Not whether they are documented. Not whether they passed an audit eighteen months ago. Whether they are operating in the current environment, today.

6.2 Identity Governance as a Continuous Practice

The forensic record of breach investigations points to identity as the primary attack pathway. Stolen credentials, over-privileged accounts, and ungoverned service accounts appear in the vast majority of significant breach investigations. As AI agents proliferate in regulated environments, this discipline extends beyond human identities — non-human identities face increasing scrutiny from underwriters.

6.3 Shared Accountability Across Engineering and Security

Security programs that operate as isolated functions are structurally limited. Integrating security requirements into engineering processes, making security visibility available to operational teams, and establishing metrics that reflect the actual security health of the environment — not just the security team's activity — is what closes the compliance-reality gap at its source.

6.4 Pre-Underwriting Readiness as an Ongoing Practice

The CISI is most powerful when it functions not as a one-time assessment but as an ongoing practice. Organizations that conduct CISI assessments as part of their annual renewal preparation position themselves to negotiate from a position of demonstrable strength, reduce premium exposure through proven control maturity, and resolve discrepancies between their attestation and their environment before a carrier does.

7. Breach Case Studies: The Compliance-Reality Gap in Practice

Abstract frameworks for understanding claim risk become viscerally clear when examined through the lens of real incidents. The following case studies illustrate precisely where the gap between documented controls and operational reality becomes the difference between organizational survival and organizational crisis.

7.1 Change Healthcare (2024): When Infrastructure Becomes the Attack Surface

The Change Healthcare breach of February 2024 remains the most consequential healthcare infrastructure event in U.S. history. A ransomware attack by the ALPHV/BlackCat group exploited a single ungoverned access point: a Citrix remote access portal that did not require multi-factor authentication. Attackers moved laterally through the network over nine days before deploying ransomware, ultimately forcing the shutdown of payment processing systems handling roughly one-half of all U.S. healthcare claims.

UnitedHealth Group ultimately disclosed breach-related costs exceeding $3.09 billion.[11] The downstream cascade extended far beyond the immediate breach victim — hospitals, physician practices, pharmacies, and health systems lost access to claims processing, eligibility verification, and prior authorization systems simultaneously.

Three insurance-critical observations:

- The breach entry point — an MFA-exempt Citrix portal — is precisely the control gap that cyber insurance applications require attestation on and that claim investigations examine first.
- The third-party cascade exposed how organizations had acknowledged third-party cyber risk without operationalizing governance of it.
- The scale of financial exposure illustrates what happens when insurance coverage is uncertain at exactly the moment it is needed most: the financial backstop fails simultaneously with the operational failure it was designed to address.

7.2 Stryker Corporation (March 2026): The MedTech Sector's Defining Incident

On March 11–12, 2026, Stryker Corporation confirmed it was responding to a global network disruption resulting from a cyberattack.[12] The breach, claimed by the Iran-linked threat actor Handala, involved attackers gaining Intune administrator or global administrator privileges within Stryker's Microsoft environment and weaponizing Microsoft Intune to push remote wipe commands across enrolled systems. Handala claimed to have wiped approximately 200,000 corporate systems and exfiltrated 50 terabytes of data.

Products operating on architecturally separate systems — including Mako surgical robots, Vocera communications platforms, and LIFEPAK monitors — were confirmed unaffected because Stryker's product security architecture had maintained separation from corporate IT systems.

[11] UnitedHealth Group, SEC filings and earnings disclosures, 2024–2025. Breach-related costs from the Change Healthcare incident disclosed as exceeding $3.09 billion.

[12] Check Point Research and public incident disclosures, March 2026. Stryker Corporation 8-K filing with the U.S. Securities and Exchange Commission, March 12, 2026.

The CISI failure map for this event:

- **The attack required Intune administrator or global administrator credentials. Whether those credentials were properly governed — least privilege, MFA-protected, regularly reviewed — will be the central forensic question.** Domain 1 — Identity & Access Control:
- **Domain 3 — Backup & Recovery Integrity:** When 200,000 systems are simultaneously wiped, the distinction between recoverable and catastrophic depends entirely on whether immutable, isolated backups existed independent of the compromised MDM environment.[13]
- **Mass wiping of enrolled devices eliminated forensic evidence on those endpoints. Centralized, independently protected log repositories become determinative.** Domain 4 — Logging & Evidence:
- **Stryker's behavior — SEC filing within 24 hours, activation of pre-existing response plans, CEO communication, CISA collaboration — represents the operationalized IR that carriers require evidence of.** Domain 6 — Incident Response Capability:

The Wiper Attack and Policy Coverage Question

Many cyber insurance policies are not clearly written to address destructive wiper attacks conducted by geopolitically motivated state-affiliated actors. Organizations should verify whether their policies cover: data destruction events (not just encryption); system wipe scenarios across managed device fleets; and activity claimed by nation-state-affiliated groups. Whether Handala constitutes a 'state-sponsored actor' under policy exclusions is a question that will be litigated.

7.3 Cross-Industry Implications

Industry Vertical	Primary Attack Surface	Regulatory Cascade on Breach	Critical CISI Domains
Healthcare Providers	**Identity, third-party infrastructure (clearinghouses, EHR, billing)**	**HIPAA, OCR, state AGs, class action**	**1 (Identity), 5 (Control Reality), 7 (Third Party)**
MedTech Manufacturers	**Privileged access, MDM/UEM platforms, supply chain ordering systems**	**FDA cybersecurity, HIPAA (if PHI), SEC 8-K (if public)**	**1 (Identity), 3 (Backup), 6 (IR), 7 (Third Party)**
FinTech / Payments	**Third-party integrations, API keys, open banking access chains**	**DORA, SEC, state regulators, PCI-DSS**	**1 (Identity), 5 (Control Reality), 7 (Third Party), 8 (Attestation)**

[13]Forrester Research analysis of Stryker/Handala incident, March 2026. Cited in multiple cybersecurity trade publications including SC Media and CybersecurityDive.

Industry Vertical	Primary Attack Surface	Regulatory Cascade on Breach	Critical CISI Domains
Digital Health SaaS	Cloud environments, PHI data stores, identity providers	HIPAA, FTC Health Breach Rule, state privacy law	1 (Identity), 4 (Logging), 5 (Control Reality)
Regulated Data Processors	Cloud integrations, vendor access chains, SaaS platform dependencies	GDPR, CCPA/state privacy, SEC (if public)	5 (Control Reality), 7 (Third Party), 8 (Attestation)

8. The Regulatory Compounding Effect

For healthcare organizations, digital health platforms, and regulated SaaS companies, the insurance denial problem does not exist in isolation. The same gap between documented compliance and operational reality that drives claim denial also drives regulatory exposure — and both are triggered simultaneously by a breach.

The organization that suffers a breach with a compliance-theater security program faces a convergence of consequences: the denied insurance claim, the regulatory investigation, the class action litigation, and the operational disruption all arrive simultaneously. Each amplifies the others. The denied claim removes the financial backstop precisely when it is needed most. The regulatory investigation proceeds without the benefit of counsel fees being covered.

For Healthcare Organizations: Patient data exposure carries regulatory consequences under HIPAA, OCR enforcement, and state privacy law simultaneously. Operational disruption can affect care delivery in ways that create additional regulatory and legal exposure.[14] The cost of a denied insurance claim in healthcare is not just financial — it is operational, regulatory, legal, and reputational in a compounding cascade that can threaten organizational viability.

[14]EY / KLAS Research, MedTech Cybersecurity Survey, late 2025. More than 70% of MedTech organizations reported moderate to severe financial effects from a cyber incident in the prior two years; nearly 60% reported clinical impacts.

9. Recommendations for Executive Leadership

The following five recommendations are governance imperatives, not technical prescriptions. They represent the actions that fall squarely within executive accountability and that most directly address the compliance-reality gap at the organizational level.

#	Recommendation	Action
1	**Conduct a CISI Assessment Before the Next Renewal**	**Commission a structured CISI assessment against the current operating environment — not the documented security program, but the actual one.**
2	**Require Continuous Control Visibility, Not Audit Reporting**	**Require evidence of continuous control monitoring: MFA enforcement rates, privileged access review cadence, EDR coverage percentages, backup restore test results.**
3	**Bridge Engineering and Security Accountability**	**Integrate security requirements into engineering processes. Establish shared metrics. Hold both functions jointly accountable for the security posture of the production environment.**
4	**Pre-Engage Your Crisis Response Team**	**Legal counsel, forensic IR firm, PR/communications support, and regulatory notification counsel should be engaged before a breach — not selected during one.**
5	**Ask the Defining Question at Every Board Meeting**	**If a breach occurred today and every control we claim to have was examined in forensic detail — would our insurance claim be paid?**

Conclusion

The cyber insurance market has undergone a permanent structural correction. The era of insurance as a reliable backstop for organizations with nominal security programs has ended. What has replaced it is a system in which coverage is genuinely conditional — conditional on controls being implemented, maintained, and operated in the manner attested to.

The NAIC's own 2024 data makes the outcome visible: nearly three times as many claims were closed without payment as with payment.[15] That ratio does not represent insurer misconduct. It represents the systematic divergence between what organizations believe they have built and what forensic investigation reveals.

When that data is examined through the lens of company size and claim amount, the picture sharpens considerably. Smaller organizations face the harshest odds of payment — roughly one in three claims paid for SMBs[16] — while simultaneously facing claim amounts that represent existential financial exposure when denied. Larger organizations face a different problem: claim amounts that routinely exceed policy limits that were set before recovery costs reached current levels.[17]

The path forward does not require wholesale reinvention of the security program. It requires a shift in orientation — from compliance as destination to security as operational discipline — and the implementation of continuous verification mechanisms that close the gap between documented policy and operational reality.

The Cyber Insurance Survivability Index exists to make that gap measurable before a breach makes it expensive. Organizations that conduct this assessment and address what it reveals will find that their security posture strengthens, their breach probability decreases, their regulatory position improves, and — when they need it most — their insurance actually pays.

The Cost of Getting This Wrong

The cost of getting this wrong is no longer theoretical. It is being realized, quietly, consistently, and across every regulated industry. The question is not whether your organization has cyber insurance. The question is whether it has defensible coverage — and whether that coverage will scale to the actual financial exposure your organization faces if a breach occurs today.

About Cybantage

Cybantage is a cybersecurity firm specializing in the healthcare and regulated industries sector. Our practice focuses on the intersection of security operations, compliance frameworks, and risk management — helping organizations build security programs that perform under real-world

[17] SC Media / Sophos, 2023–2024 survey data. 63% of respondents denied full compensation cited coverage limits as primary reason. https://www.scworld.com/resource/why-your-cyber-insurance-may-not-cover-everything

conditions, not just in structured evaluations. The CISI framework is a Cybantage proprietary methodology.

This paper is intended for informational and strategic planning purposes. It does not constitute legal, insurance, actuarial, or compliance advice. Organizations should consult qualified legal and insurance counsel when evaluating their specific coverage environments, policy terms, and regulatory obligations. CISI scores represent an internal assessment framework and do not constitute a guarantee of coverage or claim outcome.