THEMATIC RESEARCH · CYBERSECURITY

The Referee Has a Book on the Game

Cybersecurity has no scoreboard except the insurers who pay when the products fail. Read carefully, their filings kill the spend-versus-breaches argument both sides cite, show the products working where deployed and failing as the front door, and show the scorekeepers selling, crediting and investing in the equipment they score — conflicts of interest the industry has yet to sort out. Whether the arms race or the ratchet is the better account is the reader's call, better informed on the way out than in.

Data as of the National Association of Insurance Commissioners' (NAIC) 2025 report (data year 2024), Aon's 2024 US Cyber Market Update, AM Best's 26 June 2026 report on the 2025 filings, and carrier and broker publications through April 2026. A companion TON618 note, The Survivor Premium, sorted the sector's stocks; this note is the sector's outcomes, and who measures them.

DOWNLOAD PDF

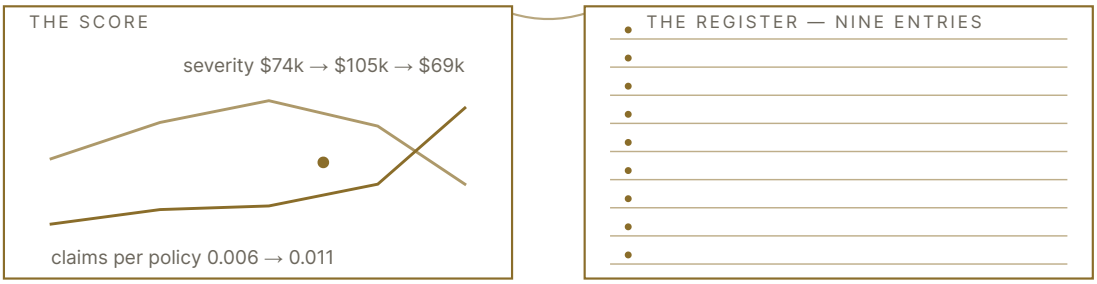

total — left blank on both questions

THE NUMBERS UP FRONT

0.006 → 0.011	insured cyber claims per policy, 2020 to 2024. Nearly doubled.
$105k → $69k	average insured severity, 2022 peak to 2024. Down a third.
73%	share of At-Bay's 2025 ransomware claims that entered through a VPN appliance, from 38% two years earlier.
5.6x	the largest measured effect of any single control on event likelihood (configuration hardening, Marsh McLennan). Not a product.
12.5%	the premium credit Coalition pays for CrowdStrike, SentinelOne, or its own managed service.
+0.09	rank correlation of carrier loss ratios from one year to the next. Last year's "best underwriter" predicts nothing about this year's.
$0	fees documented in the register flowing from any vendor to any carrier or broker for a credit, shortlist or designation. The book is an interest, not a payment.

Three parties get their strongest case before any evidence is weighed: the two accounts of the industry, and the insurers whose filings supply the evidence. **The cases, in their own terms — before the scorecard is read.** *Arms race:* attackers choose the

ground; entry migrating email → remote desktop (RDP) → VPN appliances (the network gateways companies use for remote access) → endpoints (individual computers and servers) without monitoring is what product success looks like from the claims side; controls measured inside one attacker population cut events 1.4–5.6x; severity falling while frequency rises is defense-in-depth working. *Ratchet:* the leading ransomware entry vector is a security product; the mandated endpoint layer was present in 60% of the year's worst victims; the strongest control is free; each failure becomes the next mandate, sold by the party issuing it. *The carriers as referees:* they pay the losses, so a credit for a control that does not work costs them money; every relationship in §6 is taken from the carrier's or vendor's own public release — the question is silence in the studies, not secrecy; no row shows a fee flowing from a vendor to a carrier for a credit, shortlist or seal (Marsh: "No charges, fees or expenses are payable"); the referee has published Google above its own investor three years running and named Cisco, Fortinet and SonicWall as the risk. **Tested here:** unit-level control effects (§5); the appliance share (§5); whether the counter-examples hold (§6); whether the studies disclose (§6). **Not testable in this frame:** the aggregate counterfactual behind either metaphor; whether any credit or resale was actuarially justified — the calendar-year loss ratio is too noisy at carrier level (§4) to vindicate or convict a carrier's crediting, so the carriers' strongest defence cannot be scored here; the ratchet's revenue claim (companion note); Marsh's 2025 successor study, not retrieved; Coalition's "73% fewer claims," which the frame cannot separate from selection.

FO. 1

The verdict

The scorecard is real, partial, and interested. Read it accordingly.

Two accounts of the cybersecurity industry compete. The **failure ratchet** says the industry is paid for failing: spend rises, breaches rise, and every breach is unpaid marketing for breach prevention. The **arms race** says attackers improve, defenders must improve in kind, and a rising breach count measures the attacker's progress, not the defender's failure. Both cite the same chart — spend and breaches, both rising — and both are consistent with it. The difference is a counterfactual no one has. On aggregate data the two accounts are observationally equivalent, and neither is falsifiable. That is the first thing this note establishes, because it is the thing the industry's own argument never concedes.

One party does not argue in anecdotes. Cyber insurers pay when a product fails and report the result to regulators every year. Their claims data is the only outcome data in this industry with a P&L behind it. Read carefully, it says seven things:

1. **Intrusions doubled; damage fell by a third.** Insured claims per policy rose from 0.006 to 0.011 between 2020 and 2024; average severity peaked at $105k in 2022 and fell to $69k. Prevention is losing and containment is winning. The front line moved from the perimeter to the blast radius.
2. **The products work where they are deployed.** Within one population facing the same attackers, Marsh McLennan measured configuration hardening at 5.6x fewer events, privilege management 2.9x, endpoint detection 2.2x, multi-factor authentication 1.4x when broad — and no effect when confined to administrator accounts. Vendor choice moves email and VPN claim rates two to seven times. The arms race holds at the unit level.
3. **The security products are the attack surface.** VPN and firewall appliances — Cisco, Fortinet, SonicWall, Citrix, Palo Alto — were the entry point for 73% of At-Bay's ransomware claims in 2025, up from 38% in 2023; Coalition's book says six in ten. The mandated endpoint product was present in 60% of the year's dominant ransomware strain's victims, who were encrypted anyway. The ratchet's sharpest evidence, in a form it never predicted.
4. **After each layer fails, the insurers require the next one, credit it, and increasingly sell it.** First email filtering; then multi-factor authentication; then endpoint detection software (which watches each computer and server for intrusions); then endpoint detection with a 24/7 monitoring service attached — each step visible in the carriers' own application forms, premium-credit schedules and product launches. The arms-race reading calls that underwriters mandating, against a moving attacker, the control that currently holds; the ratchet reading calls it each failure converted into the next product. The note prints both.
5. **The referee has a book.** Nine kinds of commercial relationship between the carriers that keep score and the vendors being scored, each documented from a party's own release, form or filing — none disclosed in the claims studies that rank vendors. Two counter-examples show the referee can still call it straight, and the note prints them too.
6. **The 2024 loss-ratio deterioration was price, not risk — and 2025 rose again.** The industry loss ratio — claims plus defense costs as a share of premium — rose from 42% to 49% in 2024; premium earned per policy fell 13% while frequency (+33%) and severity (−26%) nearly cancelled. US rates have fallen nine consecutive quarters and are accelerating, on a price level still roughly twice where the hard market — the 2020–22 run of steep rate rises — began. AM Best's first read of 2025, on the same statutory filings: 53.0, up 4.3 points (implying 48.7 for 2024), on premium up about 6% — a second straight rise into a price cut that was still accelerating. Whether price or risk carried it awaits the decomposition Aon publishes, which was not out at this writing.

7. **Concentration is credited, not surcharged.** Across the five endpoint shortlists in the register CrowdStrike appears on four, SentinelOne and Microsoft on two each. The carriers booked CrowdStrike's July 2024 outage as a catastrophe event in their 2024 results, then attached premium credits to more of the same. The companion note's tripwire on this — the monoculture cost gets priced — has not fired, and this note explains why: the party positioned to price it is building it.

The house's verdict is on none of the metaphors. It is on the scorecard: what it shows, what it cannot, and whose book it sits in.

FO. 2

The players — who keeps score, and who is scored

The note's two questions turn on relationships between parties most readers will not know by name. Players are listed because they appear in the register or the tables, not by market weight: of the twenty largest US cyber insurance groups, nine have a row in §6 (Chubb, Travelers, AIG, AXIS, Beazley, At-Bay, Sompo, Zurich, AXA); the other eleven, and roughly 200 smaller reporting insurers, have none.

The scorekeepers. The **NAIC** — the National Association of Insurance Commissioners, the standard-setting body of the state insurance regulators — designs the annual statement supplement every US insurer files on its cyber business and publishes a short report on the totals; it publishes no industry loss-ratio series. **Aon**, one of the three global brokers, computes the industry figures from those filings each year; the "49%, up seven points" number is Aon's. Aon also sells the security assessment carriers use in underwriting, and is paid by commission on the policies it places. **Marsh**, the largest broker, runs Cyber Catalyst (§6), published the control-effectiveness study in §5 from its own claims, and was "the first 'Cyber Insurance Partner' with AWS." **Howden**, a London broker, publishes the only chained cyber pricing index. The **CIAB** is the US brokers' trade association; its quarterly survey is the US rate series in §4, and it sells nothing to insureds.

The carriers — traditional, ranked by 2024 US-admitted cyber premium. **Chubb** (#1, $561M) joined Google's insurance program in 2025. **Travelers** (#2, $535M) acquired Corvus, whose application form names six endpoint detection and response (EDR) vendors. **AIG** (#8) insures CrowdStrike's $1M breach warranty. **AXIS** (#12) is a Cyber Catalyst insurer whose form asks for the EDR vendor by name. **Beazley** (#16 in the table of US-admitted insurers — those licensed state by state; a Lloyd's of London specialist) owns Beazley Security, sits in Cyber Catalyst, Google's program and Microsoft's incident-response partnership — the most connected traditional carrier in the register. **Hartford** (#15) runs an endorsement-heavy small-business book at an 11% loss ratio both years. **Starr** (#10) is a first-time reporter whose 96% loss ratio is a new book catching up on reserves (§4). **Allianz** co-designed Google's policy and, through Allianz X, co-led Coalition's $250M round and provides its US capacity. **Munich Re**, the world's largest reinsurer, is the node of the register: a Cyber Catalyst insurer, co-designer and underwriter of Google's policy, the "up to 20% off" Google endorsement through its HSB subsidiary, and an investor in At-Bay. Reinsurers — the insurers' insurers — sit above every carrier in the table; the top five hold 62% of cyber reinsurance.

The carriers — "InsurSec." A newer class that bundles insurance with its own security platform, and the source of most of §5's carrier data. **Coalition** — the largest, valued at $5B in 2022, Allianz-backed — sells a risk platform and its own managed detection, pays up to 12.5% premium credit for CrowdStrike, SentinelOne or its own service, and publishes an annual claims report. **At-Bay** — #7 by US premium in 2024, up from #32 the year before, the fastest-rising book in the table — counts Microsoft's venture fund and Munich Re Ventures among its investors, sells a managed service built on CrowdStrike's software and credits policyholders for it, publishes the reports that supply the appliance findings, and coined "InsurSec."

The vendors. *Endpoint* — the products every carrier's form names: **CrowdStrike** (on Corvus's shortlist, in Coalition's credit schedule, inside At-Bay's service, a Cyber Catalyst designee, warrantied by AIG; its July 2024 update disabled roughly 8.5 million Windows machines), **SentinelOne** (Coalition's partner; its own $1M warranty, underwriter unnamed), **Microsoft** (on two forms; its venture fund in At-Bay; its incident-response unit Beazley's partner; its Microsoft 365 ranked below Google Workspace in At-Bay's claims data three years running). *Perimeter* — the products the carriers' data names as the entry vector: **Cisco**, **Fortinet**, **Palo Alto Networks**, **Citrix**, **SonicWall**. None is resold or credited by any carrier in the register. *Hyperscalers* — the platforms that now co-design policies: **Google Cloud**, **AWS**, **Microsoft**.

The antagonist. Akira — a ransomware operation, not a company: more than 40% of At-Bay's ransomware claims in 2025, 86% of them through SonicWall appliances, 60% of its victims with a leading EDR deployed, demands averaging $1.2M. It is the reason the 2025–26 data reads the way it does.

In one paragraph: regulators collect the data; a broker computes it; brokers designate vendors and partner with clouds; traditional carriers write the premium, own security firms, back vendor warranties, invest in insurtechs and co-design policies with the hyperscalers; the insurtech carriers sell the endpoint vendors' products under their own brand

and publish the studies that recommend them; the endpoint vendors are shortlisted, credited, resold, designated and warrantied; the perimeter vendors are named as the risk and sold by no one in the chain; and one attacker drives a third of the year's losses through one perimeter vendor's device.

FO. 3

What the industry cannot settle — and the one test that can

Every threat report the industry publishes compounds the impasse rather than resolving it. Verizon's "22,000 incidents, the most ever reviewed" is a count of what its contributors submitted, not a rate. A vendor's "800% rise in stolen credentials" is that vendor's telemetry. Set beside rising spend, each count reads as "spend doesn't reduce failures" to the ratchet and "spend is why failures aren't far worse" to the arms race — the same observable, and no time series contains the counterfactual that would separate them. The aggregate argument is unfalsifiable, and the note declines to run it.

What insurer data can do that aggregate data cannot is *hold the attacker constant.* Every insured in a carrier's book faces the same year's attackers. If insureds with a control have fewer or smaller claims than those without, the control works, whatever the attacker did that year. That cross-sectional test is the one thing in this industry that can adjudicate efficacy, and the carriers run it because they price on it. It is also published by the carriers — which is why §6 exists.

FO. 4

The scorecard

What the data is

Since 2016 every US-domiciled insurer has filed a Cybersecurity Insurance Coverage Supplement with its annual statement: premiums written and earned, losses paid and incurred, defense costs, claim counts, policies in force. Three terms carry the rest of the note. A **loss ratio** is losses divided by premium — here direct losses plus defense and cost-containment expense (DCC), over earned premium, on a calendar-year basis, so it includes reserve changes on prior years and is restated as claims develop. **Alien surplus lines** are non-US insurers (mostly Lloyd's syndicates) writing US risks outside the admitted market; they file separately, and the NAIC's top-20 table excludes them, which matters for every Lloyd's-heavy carrier in it. **Endorsement, primary and excess** cover are, respectively, cyber bolted onto another policy with a small sublimit, a standalone policy, and a layer above one; endorsement is 55% of policies and a tenth of the loss ratio of primary. Cyber is not a separate line on the annual statement, so expenses are not reported and no true combined ratio can be computed.

Data year 2024

US direct written premium, excluding alien surplus lines, fell 2.3% to $7.08B — the first decline on record, after a run of +21.7%, +75.3%, +50.5% and −0.2%. Including alien surplus lines, total US premium was about $9.14B, also a first reduction. Policies in force were flat (−0.03%). Written premium per policy fell 1.8% — after falling 10.2% in 2023; the price cut was 2023's, and 2024 was the year softening eased. Earned premium fell 10%, because 2023's cuts earned through. Claims closed without payment ran 2.87 to one against claims closed with payment — 28,555 to 9,941.

The industry loss ratio, as Aon computes it from the filings: 67% in 2020 (the spike), 67% in 2021, 45% in 2022, 42% in 2023, **49% in 2024** — primary 49%, excess 54%, endorsement 10%. That leaves, in Aon's phrase, "51 points of margin for other expenses."

The dispersion that isn't a finding

The NAIC's top-20 table for 2024 runs from Beazley at 9.24% to Starr at 96.26%: an 87-point, 10.4-times spread that reads, on first sight, as underwriters disagreeing by a factor of ten about what a breach costs. It is not that. Both endpoints are artefacts — Starr a first-time reporter that booked *zero* loss on $128M of earned premium in 2023 and was included in 2024, so its 96% is a new book catching up on reserves; Beazley's 9.24% the US-admitted slice of a group whose whole cyber segment ran a 48.5% claims ratio in the first half of 2025. The spread does not persist: among the eighteen groups in both years' tables the median one-year move was 16.4 points, seven of eighteen moved more than twenty, and **the rank correlation of carrier loss ratios between the two years is +0.09** — a single calendar-year loss ratio on a $150–500M book is a noisy measurement, not an opinion. And much of the level spread is product

mix: endorsement books run near 10%, excess near 54%, and Hartford at ~11% both years is simply an endorsement book.

What survives is Aon's own finding: among writers over $50M, the whole distribution shifted up about ten points in 2024 — 5th to 95th percentile from 7–73% to 12–88% — with "more extreme outliers at the high end."

How much of 2024 was price — and the first read on 2025

Aon decomposes the seven-point rise from 42% to 49% into three legs: earned premium per policy −13%, frequency +33%, severity −26%. Loss cost per policy therefore moved (1.33 × 0.74) − 1 = −1.6%; the loss ratio implied is 42% × 0.984 / 0.87 = 47.5%, against 49% reported — rounding and mix. **The 2024 deterioration was essentially all price.** Per insured policy, the books did not get riskier; they got cheaper.

Three independent measures say what price has done since. Howden's global pricing index — the only chained series — sits "22% below its mid-2022 peak but remains 103% higher relative to pre-hard market levels": with the pre-hard-market level at 100, the index is at 203 and peaked at an implied 260 (both figures derived from Howden's stated anchors). Marsh's global renewal index has fallen twelve consecutive quarters, dating the global turn to 3Q23. The CIAB's survey of US brokers, read from each quarter's by-line table, has US cyber premiums at renewal falling nine consecutive quarters, dating the US turn to 2Q24 — and accelerating, not flattening:

	Q1	Q2	Q3	Q4
2021				**+34.3%** (all-time high)
2024		−1.7%	−1.5%	−1.8%
2025	−2.1%	−1.5%	−2.6%	−3.3%
2026	−3.5%	−3.2%		

Why, in the brokers' own words: 43% of respondents reported more cyber underwriting capacity in 3Q25 and 45% more demand; the capacity came from a favourable reinsurance market and the cyber catastrophe-bond market, which placed more than $750M of privately placed (144A) cyber catastrophe bonds in 2024, including Beazley's $210M PoleStar 2024-3. The same survey offered a second reason a year earlier: "Cyber resiliency may also have contributed, as industry cyber loss ratios have continued to fall from their peak in 2020 and 2021, despite a notable rise in ransomware attacks."

The 2024 loss ratio was earned on 2023's price. The US decline then accelerated through 2025 and 2026 while frequency was rising 20–30% a year. Beazley's segment, the only large-carrier mid-year read, showed the direction first: rate change −6.8%, claims ratio up three points to 48.5%, combined ratio up ten points to 78.8%, in the first half of 2025.

AM Best's report on the 2025 filings (26 June 2026) then gave the full-year read: direct written premium $7.5B against $7.1B; the industry loss ratio 53.0, up 4.3 points — implying 48.7 for 2024, against Aon's 49% on the same filings; surplus-lines writers, now about two-thirds of premium, at 55.9 against admitted carriers' 50.2; third-party claims rising. In AM Best's words, "this time the loss ratio increase is occurring as pricing is still declining and even accelerating the decline."

Two readings of that number are open. The carriers': 53 leaves 47 points before expenses, and AM Best kept its outlook on the global segment stable three weeks later (15 July 2026), citing "solid demand for coverage, even as the market pricing softens" and "favorable profitability over the intermediate terms." What the level cannot yet say: whether the rise was price or risk; that needs the earned-premium, frequency and severity split Aon publishes, which was not out at this writing. §10 arms a tripwire on it.

Frequency and severity: the front line moved

	2020	2021	2022	2023	2024
Claims per policy	0.006	0.007	0.007	0.008	**0.011**
Average severity	$74k	$94k	$105k	$90k	**$69k**
Closed without payment	70%	61%	62%	65%	**74%**

Frequency nearly doubled in four years. Severity peaked in 2022 and fell a third. More intrusions get through; each one costs less. Three confounders ride with that reading, and each cuts at it. The 2024 frequency point is inflated by notices, not losses: Aon reports "heightened precautionary loss notifications in 2024 due to the CrowdStrike and Change Healthcare events, many of which ultimately had negligible incurred loss" — claim count rose 40% and closed-without-payment jumped to 74% the same year; strip 2024 and frequency still rose a third over 2020–23. Severity's fall is not only products: the median ransom paid fell and 64% of organizations refused to pay; insurers mandated multi-factor authentication (MFA), endpoint detection and tested backups from about 2021, and Aon's own control data shows a 9% improvement in critical controls in 2024 — some of the decline is behaviour, and some is insurers forcing the products to be installed, which is evidence products work when mandated and not evidence the market buys them on merit. And the pool is selected and self-cleaning: the series measures outcomes among companies good enough to get insured. Coalition states its policyholders "experience 73% fewer claims than the industry average" — a figure the frame cannot separate from selection.

FO. 5

Question one — paid for failing, or keeping pace?

The scorecard says what happened. Whether it means the industry is paid for failing or keeping pace turns on the one test §3 named — hold the attacker constant and compare insureds with a control to those without. The carriers run it, and publish it.

The controls work — Marsh McLennan's signal strengths

Marsh McLennan joined the answers on its Cybersecurity Self-Assessment to a year of its US claims and notices (November 2020 to November 2021) and computed, for each control alone, the ratio of event probability without it to event probability with it. Each control is measured in isolation; the ratios are "not additive." Frequency only, not severity.

CONTROL	SIGNAL STRENGTH
Hardening — configuration tools enforce and redeploy settings (e.g. Active Directory group policy, Windows' central settings tool)	**5.58**
Privileged access — local administrator rights managed via endpoint privilege management	2.92
Endpoint detection and response	2.23
Own security operations centre (SOC) or managed provider, with alert thresholds and event monitoring (SIEM)	2.19
Patch high-severity vulnerabilities within 7 days	2.19
Internal phishing campaigns at least annually	1.76
Network intrusion detection/prevention	1.67
Patch critical vulnerabilities within 7 days	1.57
Email attachments sandboxed before delivery	1.56
24×7 SOC with containment and threat intelligence	1.56
Multi-factor authentication, all three questions together	**1.44**
MFA for administrator access, alone	**0.85 — "no discernible effect"**

Two things to read off it. The largest measured effect belongs to configuration discipline, which is not a product. And the control every carrier mandates first, MFA, is a 1.4x effect when broad and nothing when partial; only 24% of the organizations studied patched high-severity vulnerabilities within seven days.

The controls work — the carriers' own books

At-Bay (more than 100,000 policy years): Google Workspace users' email-claim frequency ran 41% below average and Microsoft 365 users' 18% above in the 2018–22 data, and on-premises Exchange more than 2.5 times Google; by the 2021–25 data Google was 29% below and Microsoft 13% above, with claim rates rising for both. Users of Cisco and Citrix remote-access VPN appliances were "nearly 7X more likely" to suffer ransomware than businesses with no VPN detected; any on-premises VPN nearly 4x versus cloud or none. **Coalition** (2023 data): internet-exposed Cisco ASA firewalls "nearly 5x more likely to experience a claim"; Fortinet boundary devices 2x; exposed remote desktop without a boundary device 2.5x.

Caveats on all of it: correlational; controls self-reported at application; carrier samples are small-business-weighted; a vendor's users differ from another vendor's users in ways the data cannot hold constant; Marsh's study is one claims year.

What this settles: within a single attacker population, controls cut claim frequency by 1.4 to 5.6 times and product choice moves claim rates by two to seven times. The efficacy half of the arms-race account holds at the unit level.

The security products are the attack surface

The same carriers' data then adds something neither account predicted. At-Bay's 2026 report, on calendar 2025: "nearly 3 in 4 ransomware attacks (73%) started with a VPN," up from 38% in 2023 and 66% in 2024. SonicWall appliances alone were the entry for 27% of ransomware claims and were "present in 86% of Akira's attacks"; Akira, one operation, accounted for more than 40% of all ransomware claims. Coalition's 2025 threat index: most ransomware "start[s] when attackers exploit virtual private networks (VPNs), remote desktop tools, and firewalls," and "Fortinet, Cisco, SonicWall, and Palo Alto Networks build the most frequently compromised products."

And the mandated endpoint layer: "More than half (60%) of Akira's victims had a leading Endpoint Detection & Response (EDR) solution in place and were still compromised." In the third quarter of 2025, 91% of Akira attacks ended in full encryption, and "every single company that avoided encryption had a professionally managed" EDR — the product with 24/7 human monitoring behind it, sold as managed detection and response (MDR).

Read together: the perimeter appliance — the firewall, the VPN concentrator — is the leading way in, and rising. The endpoint product carriers require on 95% of machines was present in most of the year's worst victims. What held was a service wrapped around the product. And the control with the strongest measured effect in the literature is sold by no one. The products sold to stop intrusions are where intrusions start. The ratchet would call that "failure as marketing" in its literal form; the arms race calls it one product category's patch discipline (§8).

73%

The share of At-Bay's 2025 ransomware claims that entered through a VPN appliance — up from 38% two years earlier; Coalition's book says six in ten. The products sold to stop intrusions are where intrusions start.

The mandate ladder

Put the findings in sequence and a mechanism appears, documented at every step. Email was the first vector; carriers began asking about filtering and training. By 2021–22 every ransomware application asks whether MFA secures all remote access — and Marsh's data says the partial version does nothing. The 2022 Corvus form lists six named EDR products to tick; AXIS asks the vendor's name and the share of endpoints covered. In 2023–24 the carriers move from requiring EDR to crediting and selling MDR. The 2025–26 reports name specific VPN vendors as the risk; the next form will ask which firewall.

Email → RDP → VPN appliance → EDR → managed EDR. Each layer's failure creates the next requirement; the requirement shapes the insured pool; the pool generates the next claims study; the study tests — and so far has validated — the requirement. That sequence is the observable. What it means is where the accounts part: for the arms race, competent underwriters responding to a moving attacker by mandating the control that currently holds; for the ratchet, each failure of a security product converted, by the party that should price the failure, into a mandate for the next product — which the same party increasingly sells.

Scoring it, both columns

FINDING	READS FOR THE ARMS RACE	READS FOR THE RATCHET
INTRUSIONS DOUBLED, SEVERITY −33%	The front line moved; containment is a defensive win	Prevention failed and was paid anyway
CONTROLS CUT CLAIMS 1.4–5.6X WHERE DEPLOYED	Products work; the attacker moved	Conceded — but the top control is free
73% OF RANSOMWARE ENTERS VIA SECURITY APPLIANCES	The attacker targets the strongest layer	The industry's own products are the vector
EDR PRESENT IN 60% OF AKIRA VICTIMS; ONLY MANAGED EDR HELD	Defense-in-depth requires the next layer	The product without the service is a mandate, not a control
THE MANDATE LADDER	Rational underwriting	Failure converted to the next product
TOTAL	*— left blank —*	

Verdict on question one: withheld. The unit-level studies favour the arms race; the perimeter numbers favour the ratchet; the aggregate proves nothing; the ladder reads both ways. The reader gets the table and the tests (§10), not a side.

FO. 6

Question two — referee, or backing the players?

Everything in §5 was published by carriers and brokers. Whether that evidence can bear the weight put on it depends on the answer here. The house pulled the relationships between the scorekeepers and the scored, by type, each from one of the two parties' own documents. No intent is asserted anywhere in this section.

OWNERSHIP

Lodestone, "a wholly owned subsidiary of leading specialist insurer Beazley," was merged with Beazley's cyber services team in 2024 to form **Beazley Security**, which sells managed detection "priced based on the number of endpoints" to "Beazley policyholders, while also driving organic growth outside of the Beazley client community." A Lloyd's of London specialist that is one of the twenty largest US cyber writers owns a managed-security firm that sells detection to its insureds and to the open market.

RESALE

At-Bay sells Stance MDR — "best-in-class EDR from CrowdStrike," run by At-Bay's own security team — and credits policyholders for it; its 2026 report states, "Not a single At-Bay MDR customer filed an Akira claim in 2025." **Coalition** sells its own managed detection, which earns the credit below.

PRICE SUPPORT

Coalition pays "up to 12.5% premium credit" for CrowdStrike Falcon Complete, SentinelOne Vigilance, or Coalition MDR — products "hand-selected by experts on Coalition's security and actuarial teams." **Munich Re's HSB** gives Google Cloud customers who run Google's diagnostic "up to 20% off their cyber insurance premium." **AWS's** insurance partners "can now reward customers that present a security posture that follows AWS best practices similar to 'safe-driver' discounts." A premium credit is a price cut on the vendor's product, paid by the carrier.

SHORTLIST

Corvus's 2022 ransomware application asks which endpoint technology is in place and offers, as EDR: CrowdStrike Falcon Insight, SentinelOne Singularity Complete, Sophos Intercept X with XDR, FireEye, Microsoft Defender for Endpoint, VMware Carbon Black, or none. **AXIS's** asks the applicant to "identify EDR solution(s) in place, including company names & product name," and whether Microsoft 365 Defender is enabled. A named-vendor form is a shortlist; the pool is shaped to it before any claims data exists.

DESIGNATION

Eight insurers — Allianz, AXA XL, AXIS, Beazley, CFC, Munich Re, Sompo, Zurich — run **Marsh's Cyber Catalyst**, designating products whose adopters "may be considered for enhanced terms and conditions." Designees include CrowdStrike, FireEye, Forescout, KnowBe4, Carbon Black, CyberArk, Varonis, Trend Micro. The vendor agreement cuts both ways: "No charges, fees or expenses are payable by Marsh to Vendor or by Vendor to Marsh" — and "Marsh may … reference those Approved Security Products and associated Vendors in marketing material relating to the Program." CrowdStrike's own blog told customers the designation could earn them "enhanced terms." Forrester's analysts called the program "a promotional campaign for cyberinsurers and security vendors to get their products into the doors of more businesses."

PLATFORM CO-BRANDED INSURANCE

In 2021 **Google Cloud, Munich Re and Allianz** launched Cloud Protection +, a policy Munich Re "designed … exclusively for Google Cloud customers"; the customer runs Google's tool and sends the report to the carriers, who "use the data from Google Cloud's Risk Manager reports to simplify the insurance application and risk assessment." **Chubb and Beazley** joined in 2025; Chubb's flyer promises "specialized terms tailored for Google Cloud customers" and "customized pricing … based on use and risk posture," with the policy issued "inclusive of broader coverage via Google Endorsement." **AWS** runs the same structure with At-Bay, Marsh, Cowbell, Resilience and Measured; **Microsoft** became Beazley's approved incident-response partner in December 2025. The three hyperscalers now each have an insurance channel in which the cloud vendor's own diagnostic generates the underwriting submission.

WARRANTY BACKING

CrowdStrike's $1M breach-prevention warranty: "CrowdStrike has purchased an insurance policy from an AIG insurance company to help protect their financial risk." SentinelOne's: "Our third-party insurance partner underwrites our warranty." The vendor's guarantee is a carrier's liability.

INVESTMENT

Microsoft's venture fund M12 joined At-Bay's cap table in December 2020, alongside Munich Re Ventures, and again in 2021. Allianz X co-led Coalition's $250M round at a $5B valuation while Allianz provides Coalition's US capacity.

CO-PUBLISHED RESEARCH

At-Bay's ranking of email-security vendors is served from Google's own servers; the release led with Google Workspace's result.

THE NODE

Munich Re appears four times — Cyber Catalyst insurer, co-designer of Google's policy, the 20%-off endorsement, At-Bay investor. Beazley four times. CrowdStrike five: Corvus's shortlist, Coalition's credit, At-Bay's resold service, a Cyber Catalyst designee, an AIG-backed warranty. A reinsurer of Munich Re's size would appear in any list of cyber programs; the count measures scale before it measures anything else.

TWO COUNTER-EXAMPLES, PRINTED ON PURPOSE

The register would be an insinuation without them. Coalition names Cisco ASA at 5x, Fortinet at 2x, and Fortinet, Cisco, SonicWall and Palo Alto as "the most frequently compromised products" — and sells endpoint services, not perimeter appliances, with no resale relationship with any of them. And Microsoft's venture fund has sat on At-Bay's cap table since 2020 while At-Bay has published, three years running, that Google Workspace users file fewer email claims than Microsoft 365 users. Whatever the relationships do, they did not stop that.

WHAT THE REGISTER SHOWS

The parties that keep cybersecurity's only money-backed score have commercial relationships with the vendors being scored, of at least nine kinds, none disclosed in the studies that rank vendors or controls. Those are conflicts of interest in the ordinary sense — a position, not a proven act — and the industry has yet to sort them out. **What it does not show:** that any study's numbers were altered, shaded or selected; the counter-examples are evidence they were not, in at least those cases.

THE CARRIERS' CASE, TYPE BY TYPE

Ownership and resale are vertical integration into loss control — the carrier absorbing the mitigation it would otherwise pay for on the claims side. Price support and shortlists are what a competent underwriter does with a credible control. The hyperscaler programs replace questionnaires with telemetry, which is better underwriting data. Warranty backing is a vendor buying insurance like any other company. Venture stakes are small and passive. Each is defensible on its own. The note's claim is narrower than any of them: **a reader of the referee's scorecard is entitled to know the referee's book.**

THE CONCENTRATION READ

The companion note's eighth tripwire fires on its insurer leg when "a major cyber insurer publicly attaches an aggregation surcharge or exclusion to single-vendor security estates." No row above shows a surcharge. Several show a credit paid to the concentrated vendors, a service built on one of them, and shortlists on which CrowdStrike appears four times in five and SentinelOne and Microsoft twice each. The carriers booked CrowdStrike's July 2024 outage as a catastrophe event — Aon: "Widespread events such as CrowdStrike, Change Healthcare and CDK struck the market, highlighting the potential materiality of catastrophe events" (Change Healthcare and CDK being the year's other widespread events, in Aon's term) — then continued to concentrate the insured estate on the same handful of platforms. The tripwire has not fired, and the register explains why: the party positioned to price the concentration is building it.

Verdict on question two: withheld as to motive. The structure is a fact; the note prints it and the tests (§10).

$0 Fees documented in the register flowing from any vendor to any carrier or broker for a credit, shortlist or designation. The book is an interest, not a payment.

FO. 7

Where the two questions meet

Question two governs the weight question one's evidence can bear. The same At-Bay report supplies both halves of the year's story — the appliance as the door and managed detection as what held — so the appliance number is disinterested only as to the appliance vendors; it is the first half of a finding whose second half sells the author's service. The weighting instruction applies finding by finding, not author by author: what survives best is the row that names a vendor the author does not sell. That is not a verdict; it is a reading instruction, and it is why the two questions belong in one note.

The narrower thing both questions support, without a metaphor: the insurance market has become the security industry's distribution channel as well as its scorekeeper. At-Bay's word for the model is "InsurSec." Every layer the attackers get through becomes the next thing the carriers require, credit and sell — and the claims data on the layer after that will be written by the sellers.

FO. 8

The case against this note — both directions

(Four sides are steelmanned here. Question one's verdict is withheld, so both accounts get their strongest statement. Question two's structure cuts against the carriers, so the carriers get theirs — and so does the skeptic the note declines to join.)

For the arms race, at full strength. Security is adversarial and the attacker chooses the ground. The claims data shows exactly what an adaptive attacker produces: entry migrating from email to remote desktop to VPN appliances to unmonitored endpoints, each shift *after* the previous layer was hardened. That is product success forcing the attacker to move. The unit-level effects are large and measured within a single year's attacker population, which is the only fair test. Severity falling a third while frequency doubles is the signature of defense-in-depth working. The perimeter finding is about one product category with a known vulnerability profile, sold by network-equipment companies, not about "the security industry"; and the response — segment, patch within seven days, monitor — is what the industry has been saying. The mandate ladder is the market learning.

For the ratchet, at full strength. The industry's revenue rises with its failures, and the claims data now supplies the mechanism. The leading entry vector is a security product. The endpoint product carriers require was present in most of the year's worst victims. The largest measured effect in the literature is configuration discipline that costs nothing, and the control every carrier mandates first is a 1.4x effect when done fully and nothing when done the way most buyers do it. The response to each failure is not a refund; it is a mandate for the next tier, sold by the party issuing it. Severity fell because victims stopped paying and insurers forced controls — not because the market chose products that work.

For the carriers as referees. A carrier that mandates, credits and even sells detection is doing what a loss-control function should: it has the claims data, it knows which controls hold, and a credit for the control that cuts its losses is price aligned with risk. Vertical integration is the carrier paying for mitigation up front instead of losses afterward. The hyperscaler programs replace self-reported questionnaires with machine-read telemetry. Every relationship in §6 comes from the parties' own public releases, and none involves a fee from a vendor for a credit, shortlist or seal. The counter-examples prove the referee calls it straight when the data points at a vendor it has no reason to protect — and, in At-Bay's case, at its own investor. Disclosure would be welcome; corruption is not shown.

For the skeptic of the carriers. Nine kinds of relationship, none disclosed in the studies that rank vendors. The carrier that publishes "only managed detection stopped the year's dominant ransomware" sells managed detection. The carrier that publishes a vendor ranking hosts it on the winning vendor's servers. The referee's shortlists recur on three names, its credit schedule pays toward them, and it booked the last failure of one of them as a catastrophe. No surcharge followed. Whether or not any number was shaded, the score and the sale are made by the same hand, and the reader is not told.

FO. 9

Testing it

On the arms race: holds where a controlled comparison exists — conceded in full. Where it overreaches: "one product category" is the category the industry has sold as the perimeter for twenty-five years, from four of its largest vendors, and its share of entry rose from 38% to 73% in two years while those vendors' revenues did not fall. "The attacker moves after the layer hardens" requires the layers to have hardened; 60% of Akira victims had the mandated layer.

On the ratchet: its strongest evidence — the appliance share — is real, sourced, and disinterested as to the appliance vendors (§7), and it survives. Where it overreaches: "revenue rises with failures" is a claim the claims data cannot test; the demand half was tested in the companion note (CrowdStrike's recurring revenue up 51% after its own outage) and held there, not here. "Severity fell but not because of products" is half true: Aon's control data and Marsh's signal strengths both say mandated controls reduce loss, which is a product working under compulsion — a fact the ratchet must own rather than wave off. And the mandate ladder is consistent with competent underwriting; the ratchet reads it as capture without evidence of capture.

On the carriers as referees: the vertical-integration defence is coherent and the counter-examples are strong. Where it fails: none of it is disclosed. A study that recommends a service should say the author sells it; a ranking hosted by the winner should say so. The defence is a defence of the practice, not of the silence, and the silence is what the note documents.

On the skeptic: the structure is proven; the effect on the numbers is not, and the two counter-examples are affirmative evidence against a strong form of capture. Correct as a disclosure claim, unproven as a corruption claim; the note holds it to the first.

Net. Neither metaphor is adjudicated and the note does not pretend otherwise. What is adjudicated: the products reduce claims in the insured pool; the products are also the leading vector into it; the scorekeepers have a book on the game and have not said so. Those three are the findings. The metaphors are the reader's.

FO. 10

What would change our view

The note rests on five readings: that unit-level control effects are real and not artefacts of selection; that the appliance share is a trend and not one carrier's book; that the frequency/severity split is a change in the world and not in notification behaviour; that the register is undisclosed rather than merely unremarked; and that concentration is being credited rather than priced. Each carries an observable that would flip it. Silence scores as silence throughout.

- **CI-1 — the controls stop working out of sample.** A carrier or broker publishes a control study with a comparison group *not* subject to the same mandate, or Marsh's next update shows the top controls' effects collapsing toward 1.0 → the unit-level case for the arms race weakens and the selection objection is confirmed. Mirror: a non-mandated comparison confirms effects of the same order → the efficacy claim is established outright.
- **CI-2 — the perimeter stops being the door.** The 2027 At-Bay and Coalition reports put the appliance share of ransomware entry below 50% → the finding was a two-year exploit cycle, not a structural fact. Mirror: a third carrier or the NAIC publishes entry vectors and confirms a majority → the ratchet's strongest evidence is no longer one book's.
- **CI-3 — frequency turns.** Claims per policy fall for two consecutive data years with closed-without-payment also falling → prevention is winning again and "the front line moved" retires. Mirror: frequency rises again in data year 2025 after the notification effect washes out → the 2024 step was real.
- **CI-4 — the price leg catches up with the loss leg.** Data year 2025's industry loss ratio rises to or through 55% with earned premium per policy again the dominant leg → the softening has consumed the margin faster than the carriers' commentary allows, and the pricing question overtakes both metaphors. Mirror: it holds at or below 49% → the carriers priced the decline correctly and the "referee" reading gains a point. First observation: AM Best's 53.0 on the 2025 filings (26 June 2026) sits between the two thresholds — not fired, not mirrored; the attribution leg waits on Aon's decomposition.
- **CI-5 — the referee discloses, or doesn't.** A carrier's claims study that ranks vendors carries a disclosure of the author's commercial relationship to them, or a carrier publishes a finding against a product it sells → the disclosure gap closes and the note's second claim is met. Mirror: a carrier is shown to have withheld or altered a finding about a vendor it resells → the skeptic's strong form is established and the note re-spines on it.
- **CI-6 — concentration gets priced** (inherits the companion note's CS-8, insurer leg). A major cyber insurer attaches an aggregation surcharge, sublimit or exclusion to single-vendor estates → the monoculture cost becomes a line item; scored in the companion note's ledger, never re-scored here. Mirror: by year-end 2028 no carrier has, and no second cross-customer single-vendor outage has occurred → the companion note retires the term and this note's "credited, not surcharged" becomes the description of a risk that did not materialize.

Two-directional by construction: CI-1 and CI-2 fire against the account that currently has the better evidence on each finding; CI-3 and CI-4 fire against the note's own reading of the scorecard; CI-5 and CI-6 fire for the carriers if met and against them if not.

FO. 11

What is actionable — honestly

Nothing in this note is a position, and the house takes none from it. The findings bear on how to read every other cybersecurity document, not on any security.

For a reader of vendor or carrier research: every carrier claims study is now evidence with a known interest; read it with §6 open. The finding that survives interested authorship best is the one that names a vendor the author does not sell.

For the coverage universe: the appliance finding lands on the network-security incumbents — Fortinet, Cisco, Palo Alto (Citrix and SonicWall are private) — as the vendors whose products are the documented entry vector. The endpoint finding lands on CrowdStrike, SentinelOne and Microsoft as the vendors whose products are mandated, credited and, in one case, resold by the scorekeepers, and present in most of the year's worst victims. Neither is a valuation claim; the companion note carries those. What this note adds to that note's sort is that the consolidation it priced now runs through the insurance channel, and the channel's scorekeepers have not priced the concentration they are building.

For the insurers: the 2024 loss ratio was earned on 2023's price; 2025's, on AM Best's first read, came in at 53.0 — up again, on cuts twice as deep; which leg carried it is not yet known, and CI-4 has not fired. Beazley's first half is the only large-carrier read and it shows the combined ratio up ten points. The reinsurance renewal for 1 January 2027 is being quoted now.

What the house will not do with this data: rank endpoint vendors against each other on carrier claims (the mandate, the resale relationships and the service wrapper all sit between product and outcome); state a security-industry key performance indicator from the aggregate loss ratio (its denominator is price); or read the mandate ladder as either metaphor's proof.

APPENDIX

Sources & Method (appendix)

Statutory data. NAIC, *Report on the Cybersecurity Insurance Market*, 2025 edition (data year 2024) and 2024 edition (data year 2023): Table 2 / Exhibit 1 top-20 groups, footnotes excluding alien surplus lines. Aon, *U.S. Cyber Market Update: 2024 U.S. Cyber Insurance Profits and Performance* (10th ed.): industry loss ratio and primary/excess/endorsement split; Exhibits 3, 15–17, 19–22; footnote 5 (the first-time reporter); actuarial addendum. The historical loss-ratio series was read from Exhibit 15 as rendered, not from text extraction. AM Best, *Best's Market Segment Report: U.S. Cyber Market's Premium Grows Slightly; Surplus Lines Writers Continue to Increase Role*, press release of 26 June 2026 (data year 2025, NAIC Cyber Supplement), and *AM Best Maintains Stable Outlook on Global Cyber Insurance Segment*, 15 July 2026; Aon's data-year-2025 update had not been published as of 11 September 2026. Beazley plc Interim Report 2025, Cyber Risks segment.

Pricing. Howden, *Rebooting growth: Howden's 2025 cyber insurance report* (September 2025), Figures 6, 7, 8 and 10; the 203/260 index levels are derived from Howden's two stated anchors. Marsh Global Insurance Market Index, quarterly, 4Q23–2Q26, and Marsh's US cyber market update (May 2025). CIAB Commercial Property and Casualty (P&C) Market Index, quarterly by-line tables 3Q24–1Q26 and press releases 4Q21, 2Q24, 4Q24, 2Q26; the 4Q25 figure is from the PDF, as the web article at that address carries the prior year's text. Marsh and CIAB measure rate change at renewal, which is already year-over-year; no chained index was built from them.

Control effectiveness. Marsh McLennan, *Using data to prioritize cybersecurity investments* (April 2023), Figures 1–2. At-Bay, *Ranking Email Security Solutions* (2023), *2025 InsurSec Rankings Report* (October 2025), *2026 InsurSec Report* (April 2026) — the latter two from At-Bay's own articles and releases, the full PDFs being gated. Coalition, *2024* and *2025 Cyber Claims Reports* and *Cyber Threat Index 2025*; the "six in ten" perimeter-device figure from Coalition's own announcement of the Threat Index. Aon Global 2025 Cyber Risk Report. Verizon and Flashpoint figures as cited by the NAIC, not independently retrieved.

Relationships. Every row from one of the two parties' own documents: Beazley releases (February and June 2024); At-Bay releases and articles (December 2020, October 2023, January 2024); Coalition (February 2024); SentinelOne (June 2021, July 2016, May 2024); Allianz (March 2021, July 2022); Google Cloud (March 2021); Munich Re HSB product page; Chubb flyer (March 2025); AWS (November 2023; 2024 partner deck); Microsoft Security blog (December 2025); Corvus Ransomware Supplemental Application (July 2022); AXIS Cyber Ransomware Supplemental (January 2022); Marsh Cyber Catalyst page and 2020 vendor agreement; CrowdStrike (June 2018; October 2019); Forrester (April 2019). Not used: trade-press multipliers attributed to Coalition's 2025 work; the platform vendor beneath Beazley Security's service; SentinelOne's underwriter, which it does not name.

Companion-note figures. CrowdStrike recurring revenue and the July 2024 outage: *The Survivor Premium*. CS-8 text: that note, §11.

Method. The 10x loss-ratio dispersion was tested (year-over-year swings, rank correlation, endpoint audit, mix) and rejected; the test and every computation are archived with the note's research files. The Aon decomposition reconciliation is arithmetic on Aon's stated means. The register asserts no intent; the counter-examples are printed in the same section as the rows.