AETOS ONE

MANAGED SECURITY OPERATIONS

Signal Over Noise

What Twenty-Five Leading Security Reports Agree On, and Which Eight Areas Deserve Your Budget

A cross-report analysis of threat intelligence, claims, board-governance, and survey data for executives who fund cybersecurity, mapped to the Common Sense Security Framework v2.0

AETOS ONE | WHITE PAPER | SEPTEMBER 2026

CONTENTS

Table of Contents

Executive Summary

Eight areas of control reduce the length and cost of a cyber disruption more than anything else you might fund this year. This finding comes from twenty-five security reports, all but one published in 2025 and 2026. The set spans incident response caseloads, insurance claims, product telemetry, board-governance research, and practitioner surveys. This paper analyzes them together and reads them with skepticism. It names the eight areas, shows the evidence behind each, and maps them to the Common Sense Security Framework (CSSF) v2.0, a framework built from the same evidence base.

Attackers enter through two doors. Stolen identities open one, and exposed internet-facing systems open the other. Sophos attributes **roughly two-thirds of 2025 incident root causes to identity compromise**, and Red Canary measured identity threats climbing to 53 percent of its 2025 detection volume from 20 percent the year before. Verizon, Mandiant, and IBM each rank exploitation of edge and public-facing systems first among initial access vectors in the same period. ReliaQuest's 2025 report, covering 2024, put external remote services first among successful intrusion techniques at 45 percent.

Losses concentrate in mundane events, not headline ransomware. Business email compromise and funds transfer fraud produced 58 percent of Coalition insurance claims in 2025, and AI voice cloning now drives a growing share of that fraud. For CEOs and boards, this fraud exposure carries a second cost beyond the direct loss. It creates regulatory exposure, reputational damage, and diligence findings that surface during a sale or a financing round. Speed has compressed decision windows below human reaction time, and defenders still take **a median of 14 days to notice an intrusion already underway**. Confidence is not evidence. Most leaders express confidence in their recovery plan, although few have tested it. NACD's board research finds a parallel gap in the boardroom itself, where fewer than half of directors feel their board's cyber expertise or oversight structure is where it needs to be.

No single control substitutes for another. A company with strong identity controls and no recoverable backups still faces a shutdown. CSSF v2.0 scores this directly. A failed load-bearing control **caps reported readiness at 60 percent**, no matter what else scores well. This paper applies the same discipline to the evidence below.

How This Analysis Weighs the Evidence

Not every report in this set deserves equal weight. This paper applies three tiers of skepticism throughout.

Tier One: Incident and Claims Data

The Verizon Data Breach Investigations Report, Mandiant M-Trends, IBM X-Force, the Sophos Active Adversary Report, the Coalition Cyber Claims Report, the NetDiligence Cyber Claims Study, the FBI's 2025 Internet Crime Report, and ReliaQuest's Annual Cyber-Threat Report. These sources measure what happened to real victims, through forensic investigation, paid loss, or logged incident response engagements. They carry the most weight in this paper.

Tier Two: Product Telemetry

CrowdStrike, Red Canary, Microsoft, Varonis, VulnCheck, GreyNoise, KnowBe4. Each measures large activity volume through the lens of its own product or portfolio. An identity vendor finds identity threats surging, while an email vendor finds email threats surging. Although the framing sells the product, the measurements are real. This paper uses the numbers and sets the framing aside.

Tier Three: Survey and Governance Data

Sophos State of Ransomware, Veeam, the World Economic Forum's cybersecurity outlook and Global Risks Report, SecurityScorecard, ProcessUnity, the NACD Director's Handbook on Cyber-Risk Oversight, and parts of the Cisco AI Security report. These sources measure belief, board practice, and self-reported experience. They earn their place where the gap between belief and outcome is itself the finding. NACD's board research sits here for the same reason. It reports what directors say about their own oversight, not an independently audited outcome.

One source falls outside the 2025 and 2026 window. The Marsh McLennan Cyber Risk Analytics Center study of April 2023 appears in the Devices section, flagged, because the same research center published an updated and industry-adjusted analysis in 2025 that supersedes its headline effect size. This paper reports the 2025 figures and cites the 2023 study only to show how far the estimate moved.

The Denominator Problem

Ransomware appears in 48 percent of breaches Verizon analyzed in 2026. Separately, ransomware claims frequency among Coalition policyholders was 0.32 percent in 2025, against 1.54 percent for cyber claims of every type. Both numbers are correct. The first describes what a breach looks like once it happens. The second describes your odds in a given year. Confusing the two leads executives to either panic or under-invest.

Correlation Discipline

Control-effectiveness figures throughout this paper are univariate correlations drawn from claims and survey data. Company size, industry, and overall security maturity travel together with any single control. Read every weight in this paper as an evidence-based priority, not a guarantee. CSSF v2.0 narrows its own claim on the same reasoning. It targets the duration and cost of a disruption rather than its likelihood, because the measured likelihood effects available today are too small for a mid-market company to detect against a 1.54 percent annual base rate.

Perception Is Not Incidence

The World Economic Forum's Global Risks Report ranks cyber insecurity sixth among global risks over a two-year horizon and eighth over ten years, based on executive perception surveys, not incident counts. This paper treats that ranking as a governance signal. It shows that cyber risk holds a durable place on the executive agenda. It does not measure how often or how badly incidents occur. Claims and forensic data set the operating priorities in this paper. Perception data explains why boards still under-invest despite them.

Eight Areas at a Glance

The Common Sense Security Framework v2.0 organizes 38 controls into these eight areas. The list below maps the strongest evidence in this paper to each one, and names the Aetos One service tier built to close the gap.

Area	Top-weight control	First question to ask
Govern Your Risk	GV3, one authoritative asset inventory	Who owns cyber risk, and when did they last report to the board?
Protect Your Identities	ID1, phishing-resistant authentication	What share of admin and remote accounts use phishing-resistant sign-in?
Protect Your Devices	DV1/DV2, EDR coverage and patch SLA	What percent of known-exploited vulnerabilities did we fully remediate last quarter?
Protect Your Networks	NW1, validated internet-facing footprint	When did we last scan our external footprint rather than rely on memory?
Protect Your Data	DA1, offline or immutable tested backups	When did we last restore from an offline copy, and does our payment callback use an independently sourced number?
Protect Your People	PE1/PE2, training and exercised response	When did executives last exercise the incident response plan together?
Protect Your Partnerships	PA1, verified critical-vendor posture	Which vendor is a single point of failure with no tested fallback?
Protect Your Uptime	UP1/UP2, off-hours response and proven recovery	When did we last prove our recovery time objective by full-scope test?

Eight Fundamentals, Mapped to the Framework

The sections below follow the eight areas of CSSF v2.0 in order. Each closes with the question to put in front of your security leader, your IT provider, or your managed security partner. A no, or an honest I do not know, is the signal to act.

1. Govern Your Risk

Every control in this paper scores against a baseline, and the baseline comes from an asset inventory most companies do not maintain. CSSF v2.0 places this control, GV3, in Tier 1 at weight 9 for exactly this reason. A coverage percentage against an unknown asset count is a guess dressed as a metric.

The Asset Inventory Gap

This gap is distinct from the denominator problem described in the introduction. That problem concerns how to read two correct but differently scaled statistics. This gap concerns whether the company can count its own assets at all. Without that count, every coverage figure in this paper, endpoint protection, patch remediation, vendor oversight, rests on a number nobody can verify.

Governance gaps show up before any technical failure does. The World Economic Forum's 2026 outlook finds CEOs and CISOs rank different threats first. CEOs name fraud, CISOs name ransomware and supply chain. For a CEO or a board, fraud carries a cost beyond the direct loss. It triggers regulatory inquiry, breach-notification liability, reputational damage with customers and partners, and a red flag in due diligence during a sale or a financing round. Divergence at the top means no one owns the full risk in terms every executive in the room understands. Those outcomes map to the risk categories your own filings already use: operational and strategic disruption, legal and regulatory compliance exposure, financial and economic loss, and market and competitive damage. The WEF's Global Risks Report reinforces the same point from a different angle. Cyber insecurity holds a top-ten spot on both the two-year and ten-year executive risk outlook, a durable presence that boards keep underfunding relative to its rank.

NACD's Director's Handbook puts a number on that underfunding. Only 34 percent of public company directors call it very or extremely important to improve their own

cybersecurity expertise, and 40 percent say the same about clarifying which board committee owns cyber oversight. Private company directors report a wider gap, at 45 percent and 49 percent respectively, on samples of 158 public and 85 private directors. The largest gap directors name is not their own expertise but the quality of what management reports to them, called very or extremely important by 43 percent of public and 57 percent of private directors. Boards are not ignoring cyber risk. They are, by their own admission, not yet equipped to govern it. NACD also finds 86 percent of Fortune 100 boards now disclose cybersecurity expertise as a sought or existing director qualification, evidence that the market has priced in the gap faster than most boards have closed it.

NetDiligence data shows why tiering matters. Four percent of small and mid-size claims exceeded 1 million dollars, and large companies, meaning those above 2 billion dollars in revenue, at 2 percent of claims drove 51 percent of total incident cost. A flat budget protects the wrong processes when loss concentrates this unevenly.

Insurance is a governance decision, not an IT purchase. NetDiligence's global figure shows insurer payouts covering 32 percent of total incident cost across all organization sizes. That blended average splits sharply once you segment by size. Small and mid-size companies recovered 69 percent of total incident cost through insurance, while large companies recovered only 27 percent. The larger a company gets, the more of its own incident cost it carries directly, since large-loss events increasingly exceed policy limits and sub-limits. Coalition puts average loss at 116,000 dollars globally and 262,000 dollars for ransomware specifically. A policy with a waiting period longer than your tested recovery time pays out on paper and nothing in practice.

Blind spot: the WEF figures measure opinion and executive perception, not incident outcomes. NACD's board figures measure director self-assessment, not an audited oversight standard. Treat both as a warning about governance readiness, not a ranking of true risk.

Ask this: *Who owns cyber risk in this company, with authority to halt operations during an incident, and when did that person last report to the board or ownership group? Does our board have a documented process for assessing its own cyber-oversight expertise?*

2. Protect Your Identities

Attackers log in more often than they break in. Sophos attributes roughly two-thirds of 2025 incident root causes, 67 percent across 661 investigations, to identity-related tactics: compromised credentials, brute force, phishing, and token theft. Red Canary's separate telemetry puts identity threats at 53 percent of its 2025 detection volume, up from 20 percent. CrowdStrike reports 82 percent of its 2025 detections were malware free, up from 51 percent in 2020. That figure describes how attackers get in, not how the incident ends. The same attacker who moves on stolen credentials still deploys a ransomware payload after reaching the systems worth encrypting. Malware-free intrusion and ransomware event are two stages of one attack.

Standard multifactor authentication no longer closes this door. Session hijacking through adversary-in-the-middle phishing kits **bypassed MFA in every successful business email compromise incident its analysts investigated in 2024. Not most. All.** KnowBe4 documents the same kits capturing session cookies after a victim approves a legitimate MFA prompt, and reports a 31.4 percent rise in phishing that evades secure email gateways, though it published that figure without a stated base or window. Mandiant, cited in the People section, shows plain email phishing losing ground as an infection vector. Together the two describe one shift: simple phishing volume is dropping, and what remains is built to beat the defenses most companies already have.

Phishing-resistant authentication answers this directly. Microsoft reports it blocks over 99 percent of unauthorized access attempts. Marsh McLennan measures a higher signal strength for phishing-resistant MFA than for standard MFA, but publishes no percentage gap between the two, so read the direction as the finding and not the size. Marsh's underlying survey puts MFA adoption at 90 to 100 percent across almost all organizations, without breaking out which systems carry it. ReliaQuest fills that gap from the attacker's side, naming insecure VPN configuration, meaning VPNs lacking MFA or device certificates, as one of five recurring control failures behind the breaches it investigated. Primary user and email accounts have largely closed this gap. VPNs and service accounts remain the softer target, so the remaining risk is method and scope, not headline adoption.

Cutting standing privilege is the second load-bearing control here, also weighted at 9. Verizon finds 65 percent of the techniques it examined name restricting permissions as a primary mitigation. Its attack-graph data is sharper: in 16 percent of organizations, an attacker starting from a low-privilege foothold had an 80 percent or better chance of

reaching a key administrative account. Sophos measures the same exposure in time: attackers reach Active Directory a median of 3.4 hours after initial access, 70 percent faster than last year. Standing admin rights turn one stolen password into the whole environment. The FBI names reducing them a Winter SHIELD defense.

The newest identity to govern is not human. Varonis finds 99 percent of organizations have sensitive data exposed to AI tools and 98 percent running unsanctioned applications, including shadow AI. Cisco reports 83 percent planning agentic AI deployments against 29 percent who feel ready to secure them.

Blind spot: CrowdStrike's malware-free figure describes its own endpoint agent, built to stop malware. Independent sources corroborate the identity-compromise trend. None corroborates any single vendor's number.

Ask this: *What share of admin and remote-access accounts use phishing-resistant authentication today, how many standing administrator accounts do we still carry, and do we inventory and govern the AI tools with access to our data?*

3. Protect Your Devices

Coverage, not deployment, is the number that matters. A tool installed on most of the fleet still leaves the uncovered fraction open, and attackers land there. CSSF v2.0 scores endpoint detection and response on measured coverage rather than a yes or no, placing it in Tier 1 at weight 9. The top weight of 10 goes to backups and off-hours response.

Endpoint configuration carries its own measurable gap, distinct from the network-edge exploitation in the next section. Verizon's assessment data on workstation configuration checks found 97 percent of assessed devices, spanning desktops and servers, failed the check for limiting failed login attempts before lockout, and 90 percent failed the check requiring passwords of 15 characters or more. Desktops carry wider exposure than servers to credential-theft techniques targeting system memory, reflecting how often administrators log into end-user machines with elevated credentials.

The strongest control finding in the evidence set concerns configuration, not detection. Marsh McLennan's 2025 analysis ranks network hardening first among twelve tracked control categories, ahead of EDR, logging and monitoring, awareness training, and incident response planning. Signal strength is the rate of breach claims among organizations answering no to a control, divided by the rate among those answering yes. Hardening scores 1.179 and EDR 1.174, so each correlates with roughly 15 percent fewer claims. An earlier Marsh study from April 2023 put hardening near six times, and that figure still circulates. Marsh has since adjusted its method for industry threat variation and says the two editions should not be compared one to one, so this paper treats the 2023 figure as superseded. CSSF v2.0 still holds hardening at Tier 2, because a claims correlation carries unseparated confounders. Budget, staffing, and sector maturity travel with hardening adoption.

1.179

Marsh McLennan's 2025 signal strength for network hardening, highest of twelve tracked controls. Organizations without it file breach claims at roughly 1.18 times the rate of those with it.

1.174

Marsh McLennan's 2025 signal strength for endpoint detection and response, second of twelve. Signal strength is a claims correlation, not a measured causal effect.

Blind spot: incident response datasets skew toward severe intrusions at organizations able to hire responders, inflating the apparent share of exploit-driven attacks at large enterprises.

Ask this: *What percent of our fleet reports healthy endpoint detection coverage against our full asset inventory, and what percent of workstations pass baseline hardening checks for lockout and password length?*

4. Protect Your Networks

You defend the perimeter you know about, and four independent incident datasets now agree the perimeter is where most intrusions start.

Verizon places exploitation of vulnerabilities at 31 percent of initial access, up 55 percent year over year, overtaking credential abuse for the first time. Mandiant ranks exploits first for the sixth consecutive year at 32 percent of investigations. IBM recorded a 44 percent rise in exploitation of public-facing applications, reaching 40 percent of its caseload, and found 56 percent of tracked vulnerabilities required no authentication to exploit. ReliaQuest's incident data narrows the picture further. Among techniques that succeeded in 2024, abuse of external remote services such as VPN portals, RDP, and virtual desktop infrastructure led at 45 percent, with public-facing application exploitation second at 23 percent. Its 2026 update moves drive-by compromise ahead of phishing as the leading initial access technique, so read the ranking as moving rather than fixed. Listings offering stolen VPN access on criminal marketplaces surged 250 percent between 2023 and 2024, a supply-side signal that the demand for this door is rising. The mid-market carries this exposure disproportionately. Verizon reports that roughly 96 percent of the ransomware victims it could size were small and mid-size businesses, entered most often through compromised credentials at 38 percent or unpatched edge-device vulnerabilities at 29 percent. Sophos finds 40 percent of attacks in the 501 to 1,000 employee segment started with an exploited vulnerability.

Patching the edge is not achievable on every device, and the data proves the point rather than excusing it. Verizon found organizations fully remediated only 26 percent of known-exploited vulnerabilities in 2025, down from 38 percent the year before, with a median remediation time of 43 days. Sophos found a 322-day median gap between patch release and confirmed exploitation across all confirmed exploited vulnerabilities in its 2025 casework, not edge devices alone. It identified the specific CVE in 52 of those cases, and one SonicWall firewall flaw, CVE-2024-40766, accounted for 67 percent of that subset. Treat it as a narrow but pointed base. VulnCheck attributes 42.5 percent of 2025 exploited edge vulnerabilities, drawn from 181 tracked edge CVEs, to end-of-life or likely end-of-life products, rising to 65 percent among vulnerabilities targeted by botnets.

322 days

Sophos's median gap between patch release and confirmed exploitation, measured across all confirmed exploited vulnerabilities in its 2025 casework.

Hours

GreyNoise's reported window between vulnerability disclosure and active scanning, stated by its research team without a published figure or method. Reconnaissance is fast. Compromise waits for a forgotten device.

These two speeds are not in conflict. GreyNoise puts the window between a vulnerability disclosure and active scanning at hours rather than weeks, though it states this as a research-team observation rather than a published measurement with a method behind it. Sophos separately measures a 322-day median gap between patch release and confirmed exploitation. Both are true because they measure different stages of the same problem. Automated scanning finds every exposed device within hours of a new disclosure. The severe, forensics-worthy breaches come later, from the older devices nobody patched, the SonicWall appliance still running last year's firmware, the VPN concentrator nobody remembered to decommission. Reconnaissance is instant. Compromise waits for a forgotten device. CISA responded to the same pattern with Binding Operational Directive 26-02 on February 5, 2026, requiring federal agencies to inventory, decommission, and replace end-of-support edge devices on a two-year schedule that opened with an inventory deadline of May 5, 2026. CISA, the FBI, and the UK's National Cyber Security Centre have each urged non-federal organizations to follow the same sequence, which makes the directive a usable benchmark outside government.

Once inside, a flat network turns one foothold into the whole environment, and log gaps compound the problem. Sophos names missing logs a leading impediment to investigation, citing firewall retention defaults as short as seven days, and in some cases 24 hours. ReliaQuest reaches the same conclusion from its own incident review. Inadequate logging was a contributing factor in most of the customer breaches it analyzed, alongside insecure VPN configuration, unmanaged devices, weak help-desk verification, and external exposure.

Email authentication is the cheapest control in this paper and the one most often left half finished. SPF, DKIM, and DMARC set to enforcement, rather than parked at monitoring, is the test. CSSF v2.0 puts NW2 at the entry stage of its business email compromise loss scenario, which scores at its weakest link, so a weak answer here caps that scenario no matter how disciplined your payment approvals are. Marsh ranks email filtering and web security among its twelve tracked control categories at a signal strength of 1.111, with enforcing sender policy framework named as one of the specific signals. The FBI lists

email authentication and malicious content protection as one of the ten Winter SHIELD defenses.

Blind spot: GreyNoise measures scanning volume, not successful compromise. ReliaQuest's figures come from its own incident-response caseload, which skews toward customers already engaged for a live incident. The urgency of patching an exposed asset does not change either way.

Ask this: *When did we last validate our external footprint by scan rather than by memory, and how long did we keep an exposed remote-access port open before closing it?*

5. Protect Your Data

Two very different controls live in this area. DA1, offline or immutable backups covering data, systems, identity services, and virtualization management, carries weight 10. That is the highest weight in CSSF v2.0, shared only with off-hours response, and it sits in Tier 1. Mandiant documents ransomware crews deliberately destroying backup infrastructure, identity services, and virtualization management before making any demand, which is what converts a recoverable outage into an existential one. The FBI names offline immutable backups with tested restoration as a Winter SHIELD defense for the same reason. The proof of that control is a restoration test, so its evidence sits in the Uptime section below. What follows here is the loss you are far more likely to file a claim on.

The most probable loss in your business is financial, not technical. Coalition's 2025 claims data, drawn from more than 100,000 policyholders, puts business email compromise at 31 percent of claims and funds transfer fraud at 27 percent, together 58 percent of claims, averaging 27,000 and 141,000 dollars respectively. NetDiligence confirms the pattern across 10,402 claims from 2020 through 2024. Ransomware and business email compromise led losses for small and mid-size companies at close to 55 percent of claims in 2024.

58%

Share of Coalition's 2025 insurance claims from business email compromise and funds transfer fraud combined.

$141,000

Average severity of a funds transfer fraud claim, the costlier of the two categories.

The control here is procedural, and a simple callback no longer holds by itself. The FBI's 2025 Internet Crime Report documents over 30 million dollars in business email compromise losses involving AI, mostly voice cloning used to request a wire payment. The same report logs over 632 million dollars in investment-fraud losses with a reported AI nexus in 2025, against more than 8 billion dollars in total investment-fraud losses that year. Read those shares carefully, because IC3 confirms nothing. It records what victims report, and it draws the opposite inference this paper might tempt you toward: the FBI states that the gap between the two investment figures shows many victims never learn AI was involved. The reported share is a floor, not a ceiling. The callback number must come from a source independent of the request, such as a maintained vendor directory,

because a callback fails against a cloned voice if the attacker supplied the number being called.

Beyond payment fraud, data exposure runs broader than most companies assume. Varonis reports 90 percent of organizations have sensitive cloud data exposed, and 100 percent hold at least one account able to export all data.

Blind spot: claims data only measures the insured population, which passes underwriting, and Coalition values complex claims at six-month maturity, which undercounts slow-developing litigation cost. Treat claims severity as a floor, not a ceiling.

Ask this: *Does our payment-verification callback use a number pulled from an independent, maintained source, when did we last test the procedure against a simulated fraud attempt, and when did we last restore from an offline copy no attacker could reach?*

6. Protect Your People

Attackers moved to the phone, the help desk, and now the video interview. Mandiant recorded email phishing falling from 22 percent of initial infection vectors in 2022 to 6 percent in 2025, while voice phishing rose to second place at 11 percent. That drop describes plain phishing volume, not phishing risk overall. KnowBe4 separately reports a 31.4 percent rise in phishing built to evade secure email gateways, using adversary-in-the-middle kits that bypass MFA outright, a trend detailed in the Identities section above. Read together, simple phishing is losing ground as an entry method while sophisticated, evasion-built phishing is gaining ground. ReliaQuest attributes 14 percent of 2024 breaches to social engineering against help desks and users, naming weak identity verification at the service desk as the enabling gap, the same gap the ransomware group Black Basta exploited by impersonating IT support inside Microsoft Teams to reach employees directly.

AI has moved impersonation from theoretical to reported. The FBI's 2025 report logs **over 22,000 complaints reporting AI-related information, with adjusted losses exceeding 893 million dollars**, much of it voice cloning aimed at finance and executive-adjacent staff. A newer thread targets hiring directly. The FBI documents voice and video spoofing during remote job interviews, aimed at network access rather than a paycheck, with audio and lip-movement misalignment as an observable warning sign. The same report finds over 19 million dollars in confidence and romance-scam losses with a likely AI nexus, including voice-cloned distress calls impersonating a family member, a tactic that translates directly to impersonating an executive on a support call.

Rehearsal is the control that speed demands. The Marsh McLennan Cyber Risk Intelligence Center ranks incident response planning fifth among twelve tracked control categories by signal strength, behind network hardening, endpoint detection and response, logging and monitoring, and awareness and phishing testing. Marsh publishes no percentage reduction for this control. Its own reading of the result is the useful part. The effect likely runs through secondary benefits, because organizations that rehearse tend to fix what the rehearsal exposes. NACD's board guidance places this same discipline in the boardroom. Leading boards participate directly in tabletop and crisis simulations rather than reviewing results after the fact, on the reasoning that decision-making under stress is a skill directors need as much as the security team does.

Blind spot: KnowBe4 classifies 85.8 percent of recent phishing as AI-driven, and no vendor, KnowBe4 included, has published a validated method for that label. Treat AI-attribution percentages as directional until a source discloses its criteria.

Ask this: *When did we last test our callback-verification procedure and run a synthetic-media awareness module for finance and executive-adjacent staff, and when did executives last exercise the incident response plan together?*

7. Protect Your Partnerships

Third parties multiply exposure faster than most vendor programs can track it. Verizon found third-party involvement in 48 percent of confirmed breaches, up 60 percent year over year. This is the harder number in the evidence set, drawn from independent forensic investigation rather than self-report. ProcessUnity's survey separately reports 90 percent of organizations experienced at least one third-party breach or security incident in the past twelve months, averaging twelve each. That figure comes from a vendor-sponsored survey of self-reported incidents, a broader category than a confirmed breach, and it should be read as a measure of third-party exposure rather than a count of confirmed compromises. Verizon's 48 percent stands as the number this paper relies on for breach involvement.

The assessment machinery meant to manage this risk is failing on its own terms. ProcessUnity finds 63 percent of assessments consume more than 40 hours of staff effort, 27 percent of vendors never respond, and 66 percent of organizations onboard a vendor before remediation completes. SecurityScorecard reports 78 percent of organizations cover less than half their vendor ecosystem with internal oversight, while 90 percent of leaders remain confident operations would continue seamlessly through a critical vendor incident. That gap between confidence and coverage is the finding, not a reassurance.

22%

Organizations covering half or more of their vendor ecosystem with internal oversight, per SecurityScorecard. The other 78 percent do not.

90%

Leaders confident operations would continue seamlessly through a critical vendor incident. Confidence is not coverage.

Concentration, not headcount, is the real exposure. A single vendor supporting a critical process with no tested alternative converts that vendor's bad day into your outage, and no security control you own will shorten it. Contract terms matter as much as monitoring. SecurityScorecard reports 55 percent of organizations still coordinate breach response with vendors by phone and email, and 60 percent take eight days or more to remediate a high-severity vendor issue. ReliaQuest's incident data adds a related access path. Abuse of trusted third-party relationships and abuse of cloud accounts each enabled 9 percent of the successful intrusions its analysts investigated in 2024, 18 percent between them.

That is smaller than the edge and identity doors, level with voice phishing, and a route attackers use rather than a theoretical concern.

Blind spot: the third-party survey data comes from vendors selling third-party risk platforms. Keep the finding a skeptic accepts, that the breach-involvement rate is Verizon's independent incident data, and discount the framing around the remedy each survey recommends.

Ask this: *Which vendor is a single point of failure for a critical process, and does that vendor's contract carry a breach-notification deadline measured in hours?*

8. Protect Your Uptime

Speed and detection move in opposite directions, and together they compound into the same failure. CrowdStrike reports an average eCrime breakout of 29 minutes in 2025, with the fastest observed at 27 seconds. ReliaQuest's February 2026 incident telemetry puts average breakout at 34 minutes, down from 48 minutes in 2024, with its fastest observed case at 4 minutes. In its 2024 data, Remote Desktop Protocol led lateral-movement techniques and legitimate administrative tools appeared in 60 percent of hands-on-keyboard incidents. Mandiant documents access hand-offs between criminal groups completing in under 30 seconds. Set against that speed, Mandiant's global median dwell time before detection sits at 14 days. The two figures are not in conflict. They describe a single compounding failure. Attackers finish lateral movement in well under an hour, and the median defender takes two weeks to notice. The gap between those two numbers is the entire window an incident response plan must work with. Mandiant's split inside that median matters more than the headline: 9 days when the organization finds the intrusion itself, 25 days when someone outside makes the call. Sophos finds 88 percent of ransomware deployment and 79 percent of data exfiltration now happen outside business hours, when most companies run their thinnest coverage.

Ransomware operators have changed their target. Mandiant documents a deliberate shift toward destroying backup infrastructure, identity services, and virtualization management before demanding payment, converting a recoverable incident into an existential one. The outcome data confirms the danger. Veeam finds 90 percent of security leaders confident they will recover within their objectives, yet only 28 percent of organizations hit by ransomware fully recovered their data, and 38 percent suffered extended downtime of critical systems.

90%

Security leaders confident they will recover within their objectives, per Veeam.

28%

Organizations hit by ransomware that actually recovered their data in full. Confidence and outcome point in opposite directions.

Recovery statistics need a careful read, because not every recovery is the same recovery. Sophos reports 97 percent of encryption victims recovered data by some means, but that figure includes the 49 percent of victims who paid the ransom to get it,

alongside the 54 percent who restored from backup. Veeam's 28 percent full-recovery figure measures something narrower and more useful: the share of ransomware victims who recovered their data completely, without that outcome depending on paying an attacker. That is the number that should set your recovery target, not the 97 percent. Sophos reports median ransom payments falling 50 percent to 1 million dollars in the same period, and Verizon reports 69 percent of ransomware victims paid nothing at all. Two figures set the financial shape of the event, and they measure different things: the **median ransom paid was $1M**, while the **mean recovery cost excluding ransom was $1.53M**. One is a median and the other a mean, so do not net them against each other. Refusal is a capability built on a tested restore, not a policy decided in the moment. NACD's board guidance frames this as a resilience metric directors should track directly. Time to detect, contain, and recover belongs on the same board dashboard as financial and legal risk, not buried in a technical appendix.

Blind spot: Sophos and Veeam both rely on self-reported outcomes from organizations that survived to answer a survey, so recovery numbers likely run better than the true population. ReliaQuest's breakout-time figures come from its own incident-response caseload rather than a random sample of all intrusions. The direction still holds, because it matches Mandiant's independent forensic observations from the attacker side.

Ask this: *When did we last prove our recovery time objective with a full-scope restoration test that did not involve paying a ransom, and do we have a named responder with pre-delegated authority to isolate systems outside business hours?*

Where the Numbers Disagree

Which Door Comes First?

Sophos ranks identity compromise as the root cause of two-thirds of incidents. Verizon, Mandiant, and IBM all rank exploitation of internet-facing systems first, at 31 to 40 percent, with credentials counted lower. ReliaQuest's successful-intrusion data splits the difference again. External remote services and public-facing applications together account for 68 percent of successful access, with voice phishing against identity-verification gaps at 14 percent. The conflict is largely definitional. Sophos measures root cause, while Verizon, Mandiant, and ReliaQuest measure the first observed access step. Verizon's own analysis reconciles the two, finding credential abuse appears somewhere in 39 percent of attack chains regardless of the opening move. One conflict survives the definitions. Sophos reports vulnerability exploitation as a root cause at its lowest level since 2021, while Verizon reports a 55 percent rise in the same vector. Both cannot be describing the same population. **Fund identity and the edge together and stop arbitrating the ranking**. No dataset puts a third door close.

Are Speed and Detection in Conflict?

Mandiant's global median dwell time rose from 11 to 14 days in 2025, while every speed metric elsewhere fell: a 29-minute average breakout in CrowdStrike's telemetry, a 34-minute lateral-movement window in ReliaQuest's 2026 incident data, access hand-offs completing in under 30 seconds. These numbers describe one compounding problem rather than two contradictory ones. Attackers move within an hour, while defenders notice within two weeks. The median dwell-time figure also rose on long-dwell espionage cases even as ransomware crews got faster and are found sooner within their own population, which is why the global median and the individual speed metrics can move in different directions at once. **Ask your security leader for detection time on the intrusion types most likely to hit your business, not the global median**.

How Long Before You Notice?

This is the sharpest disagreement in the evidence set, and it sits between two Tier 1 sources. Mandiant reports a global median dwell time of 14 days, up from 11. Sophos reports a median of three days, and falling. Nearly a five-fold gap on an identical metric.

The populations explain it. Mandiant's caseload carries espionage and North Korean IT worker intrusions with median dwell times of 122 days, which drag the global median upward, and 41.5 percent of its cases are still found inside a week. Sophos leans on managed detection engagements, where the median runs two days against five for its all-cause incident response work. Neither number describes your business. The useful question is not which report is right. **It is what your own detection time looks like on the intrusion types you actually face, and how often the first notice arrives from outside**.

How Big Is the AI Threat?

Survey respondents rank AI the top driver of change, and KnowBe4 puts AI-driven phishing at 85.8 percent of attacks without disclosing how it makes that call. The World Economic Forum's Global Risks Report finds adverse AI outcomes the fastest-rising risk on its ten-year horizon, moving from thirtieth to fifth place. Incident data tells a narrower story in the meantime. The FBI's own 2025 figures show reported AI-nexus losses are real and growing, and they remain a minority of total fraud losses in every category the FBI tracks, though the FBI treats that share as an undercount rather than a ceiling. Paid losses still ride on stolen credentials, unpatched edge devices, and fraudulent payment instructions, with AI lowering attacker cost and raising lure quality rather than opening an entirely new door. **Let reported losses, not projected fear, set your spending order today, while tracking the trend line for tomorrow**.

Does the Board Agree with the Security Team?

NACD finds a meaningful share of directors rate their own cyber expertise and committee oversight structure as needing improvement, at rates roughly ten points higher at private companies than public ones. Meanwhile 86 percent of Fortune 100 boards already disclose cybersecurity as a sought director qualification. The gap between disclosure and self-assessed readiness suggests boards are adding the credential faster than they are building the muscle. **Ask whether your board's cyber literacy comes from a named director's background or from a standing practice of education and testing across the full board**.

Right-Sizing the Program

Two numbers from claims data start the sizing exercise: frequency and severity. Coalition's 2025 figures put all-cause claims frequency at 1.54 percent of policyholders, with ransomware alone at 0.32 percent, and average severity at 116,000 dollars, down 19 percent year over year. Adjust for your revenue, industry, and data sensitivity, but the shape of the estimate matters more than its precision. Expected annual loss for most mid-market companies lands in the five-to-six figure range. The tail is a different number. NetDiligence records eight claims above 100 million dollars across its five-year dataset, two of them at organizations under 700 million dollars in revenue. Rare, and a solvency question rather than an operating expense.

1.54%

Coalition's 2025 all-cause claims frequency among policyholders. Ransomware alone ran at 0.32 percent.

$116,000

Coalition's 2025 average claims severity across all policyholders.

Median ransom demands tell a different part of the story than average claims severity does. Sophos reports a median ransom payment of 1 million dollars in 2025, well above Coalition's average claims severity of 262,000 dollars for ransomware. The two figures describe different populations. Coalition's number averages every insured ransomware claim, including the smaller and mid-size events that make up most claim volume. Sophos surveys organizations that were hit by ransomware and reports what they paid, a population that skews toward larger, less-insured, or self-negotiated cases. Both numbers are correct. Neither should be read as the other's contradiction. A third figure makes the point harder. Verizon measures the same metric, median ransom paid, on its own breach dataset and reports 139,875 dollars. The seven-fold spread between Verizon and Sophos on one definition is the clearest evidence in this paper that no single ransom number should anchor a budget.

Spending more does not equal reducing more risk, but the fix is not simply spending less. Coalition's 2026 report opens with what it calls the cyber protection paradox: security spending projections rose 24 percent in two years while incident likelihood roughly doubled over five, and tool sprawl itself now generates alert fatigue and visibility

gaps. Sophos points to the same root cause from the operator side. Victims most often cite lack of expertise, unknown gaps, and lack of capacity, each near 40 percent, ahead of any missing product. The corrective move is not to buy fewer tools and stop there. It is to redirect the budget those redundant tools consume toward the visibility infrastructure that makes existing tools worth having: longer log retention, expanded storage, and the analyst hours to review what the logs show. A tool nobody reviews produces the same blind spot as no tool at all.

Independent validation for this list exists from three directions with no commercial stake in each other. Insurance carriers weight the same controls in underwriting because their claims experience prices them. Law enforcement named a similar list in January 2026, when the FBI Cyber Division launched Operation Winter SHIELD with ten voluntary defenses: risk-based vulnerability management, protected and retained security logs, phish-resistant authentication, third-party risk management, end-of-life technology retirement on a defined schedule, offline immutable backups with tested restoration, email authentication and malicious content protection, internet-facing asset inventory, reduced administrator privileges, and an exercised incident response plan. Every one of the ten maps to a CSSF v2.0 control. Winter SHIELD carries no regulatory force and published no effect data of its own, so CSSF v2.0 treats it as corroboration rather than independent proof. Board-governance research adds a third source. NACD's handbook tells directors to track the same resilience metrics, detection time, containment time, recovery time, and vendor coverage, that this paper derives from claims and incident data. The insurance and claims convergence carries the weight here. Law enforcement and board-governance guidance corroborate the direction.

Insurance is the transfer leg of resilience, not a substitute for the eight areas above. Carrier-required controls and these eight areas overlap almost entirely, so meeting them shortens a disruption and improves coverage terms at the same time. CSSF v2.0 stops short of claiming they cut breach frequency. Stack the published control effects against a 1.54 percent annual base rate and the absolute reduction lands under half a percentage point, which no mid-market company will detect and no CFO should fund a program on. Duration and cost are the defensible claim. Every statistic in this paper carries a date, and the one figure older than 24 months is flagged where it appears and in Appendix A.

From Findings to a Working Session

Twenty-five reports, three source tiers, and one conclusion: **eight areas of control, funded and verified, reduce breach duration and cost more than any expansion of the tool stack**. What changed this year is the strength of the convergence, the cost of ignoring it, and clear evidence that the gap now runs through the boardroom as well as the security operations center.

Reading this paper changes nothing by itself. Review the eight questions above with your security leader, your IT provider, or your managed security partner, and score the answers honestly. A no, or an honest I do not know, is not a failure. It is the prioritized to-do list this analysis exists to produce.

The Common Sense Security Framework v2.0 turns this exercise into a scored assessment built on the same evidence base analyzed in this paper. Eight areas, 38 controls, each a governance question paired with the evidence to demand. Scoring multiplies coverage by weight by verification tier, so an unmeasured control scores zero and a control resting on an executive's word alone scores a fraction of what full testing would earn. Ten of the 38 controls are load-bearing, carrying weight 9 or 10 in Tier 1. These ten sit inside the eight areas, distributed by evidence rather than spread one per area, so some areas carry two load-bearing controls and others carry none. Any one of them caps reported readiness at 60 percent on three triggers: answered no, running below 50 percent coverage, or resting on an executive's assertion with no evidence behind it. Controls do not substitute for each other. Five loss scenarios score at their weakest link, so a strong composite cannot hide one open path.

Run the assessment with your leadership team in a single afternoon. The gaps it surfaces should set next quarter's agenda, and your board's, not this paper.

A NOTE ON WHAT COMES NEXT

Getting the score is the milestone, but expert operational coverage is the mandate.

A scored gap does not close itself, and most mid-market teams do not have the headcount to close it alone. Ask us how Guardian, Bastion, and Citadel provide named security leadership and expert operational coverage matched to the areas this assessment flags, so next quarter's review is a progress update, not a fire drill.

cssf@aetos.one | aetosone.com

Appendix A: Source Quality Reference

Tier 1 carries the most weight (incident and claims data), Tier 2 is telemetry, Tier 3 is survey and governance research. Dates refer to publication. Statistics older than 24 months from the date of this paper are flagged in this table and at the point of use in the text.

Source (year)	Tier	Primary blind spot
Verizon DBIR (2026)	1	Contributor mix shifts year to year. Vector definitions changed.
Mandiant M-Trends (2026)	1	Caseload skews toward severe intrusions at firms that hire responders, and toward long-dwell espionage, which lifts its median dwell time well above other incident datasets.
IBM X-Force (2026)	1	Same severe-caseload skew. Regional mix follows IBM's client base.
Sophos Active Adversary (2026)	1	Missing victim telemetry forces a broad credential-compromise catchall.
Coalition Cyber Claims (2026)	1	Insured population passes underwriting. Six-month claim maturity undercounts slow losses.
NetDiligence Claims Study (2025)	1	Global payout figure of 32 percent blends small and large companies. Segment by size before use.
FBI IC3 2025 Annual Report	1	Complaint-based and reliant on victim self-report. IC3 records reported or likely AI involvement, never confirmed involvement, and treats its own AI figures as an undercount.
ReliaQuest Annual Cyber-Threat Report (2025, 2026)	1	Incident-response caseload from ReliaQuest's own customer base, not a random sample. Its 2024 technique shares are superseded in part by the February 2026 edition, which moves drive-by compromise ahead of phishing.
CrowdStrike GTR (2026)	2	Malware-free share reflects what its own endpoint agent blocks best.
Red Canary TDR (2026)	2	Identity-threat growth partly reflects expanded detection coverage and broader identity product adoption, so read the direction rather than the multiple.
Microsoft DDR (2025)	2	Sees the Microsoft ecosystem. MFA block rate measured on its own stack.
Varonis Data Security (2025)	2	Customer base already buys data-security tooling.
VulnCheck Edge Report (2026)	2	Covers network edge devices only, not endpoints. End-of-life labels partly inferred from update silence.
GreyNoise Edge Report (2026)	2	Measures reconnaissance volume, not successful compromise. States its disclosure-to-scanning window as a research-team observation, with no published figure or method.

Source (year)	Tier	Primary blind spot
KnowBe4 Phishing Trends (2026)	2	85.8 percent AI-driven classification rests on an undisclosed method. Gateway-evasion figure published without a stated base or time window.
Marsh McLennan CRIC (2025)	2	Signal strength is a claims correlation with unseparated confounders, not a causal effect. Self-assessment responses paired against claims, with no control group.
Sophos State of Ransomware (2025)	3	The 97 percent recovery figure includes ransoms paid. Self-reported outcomes from a survey of hit organizations.
Veeam Resilience Report (2026)	3	Backup-vendor framing. Confidence-gap finding still holds against outcomes.
SecurityScorecard Supply Chain (2026)	3	Vendor sells the remedy it recommends.
ProcessUnity TPRM Report (2026)	3	Self-reported incidents, a broader category than confirmed breach. Same category bias as other vendor surveys.
WEF Cybersecurity Outlook (2026)	3	Measures sentiment, not incident ground truth.
WEF Global Risks Report (2026)	3	Executive perception survey. Ranks risk salience, not incident frequency or severity.
NACD Director's Handbook (2026)	3	Director self-assessment of board practice. Not an independently audited oversight standard.
Cisco State of AI Security (2026)	2/3	Early-stage field. Incident base rates for AI attacks remain thin.
Marsh Cyber Risk Analytics Center (2023)	2	Superseded by the 2025 CRIC analysis, which adjusts for industry threat variation. Past the 24-month refresh rule. Cited only to show how far the estimate moved.

About the Authors

This analysis was prepared by Aetos One, a managed security services provider (MSSP) with a mission to simplify cybersecurity.

Jerod Brennen is CEO and Co-Founder of Aetos One, with more than 25 years in cybersecurity across banking, healthcare, retail, manufacturing, and private-equity-backed companies. He built the first PCI compliance program at Abercrombie & Fitch, protecting 4 billion dollars in annual transactions, and has served as fractional CISO across businesses managing nearly 5 billion dollars in revenue.

Robert Vaile is CTO and Co-Founder of Aetos One, with more than 30 years in enterprise-scale information security, risk, privacy, and governance leadership, and a law-trained background in technology and compliance. As a Fractional CISO, he advises organizations on security strategy, cyber risk management, incident response, and regulatory compliance. Robert spent most of his career as a CISO in enterprise organizations where he built information security or risk functions from the ground up, including at Grange Insurance, Kansas State University, Consumers Energy and Idaho Power. His career also includes leadership roles at Gartner, GuidePoint Security / Deepwatch, Deloitte, and UBS/Warburg Dillon Read.

Aetos One delivers fractional CISO leadership through **Guardian**, AI-driven security operations through **Bastion**, and automated compliance through **Citadel**, enabling mid-market organizations to achieve enterprise-grade security without enterprise-grade headcount. The firm's work rests on one principle: security is a business function, and its job is enabling the business to deliver its mission without disruption.

Learn more at aetosone.com.