

The Security Paradox

Why More Cybersecurity Does Not Necessarily Create More Certainty

A research-led examination of security, dependency, evidence and resilience

Author: Chris Els

Founder & CEO, P3 Audit | Resolut

Research paper | September 2026

Central proposition

The enterprise security problem may no longer be a shortage of controls. It may be a shortage of continuously verifiable relationships between controls, suppliers, software, obligations, evidence and liability.

Executive Summary

Across the research responses, the strongest common finding is not that cybersecurity is failing. It is that the way enterprises measure and govern cybersecurity is increasingly mismatched to the system they are trying to protect.

Medium-to-large organisations spend heavily on security. IANS/Artico's 2024 CISO benchmark reported an average security budget equal to 0.69% of revenue and 13.2% of IT spend; its 2025 study reported security spending at 10.9% of IT spend. These ratios vary materially by organisation, sector and methodology. Yet there is still no credible industry-wide "protection success rate" that converts expenditure into a measurable probability of preventing an attack. Security effectiveness is instead assessed through risk reduction, coverage, resilience, detection, response and recovery.^[1-2]

At the same time, enterprise defence has become a network: dozens of tools, numerous suppliers, cloud and SaaS dependencies, outsourced security services, third-party processors and transitive software components. After an incident, technical cause, contractual responsibility, governance responsibility and legal liability can point to different places. Complete root-cause reconstruction is not always possible.

The research therefore points toward a different control problem: not simply whether each component is "secure" or "compliant", but whether an enterprise can continuously establish what it depends on, what is known about each dependency, where that evidence came from, how current it is, what obligation it satisfies, and who becomes responsible when the status changes.

We can estimate the spend. We cannot calculate the protection rate.

Cybersecurity has become one of the few corporate functions in which increasing expenditure can coexist with increasing uncertainty. Boards approve larger budgets, CISOs deploy more sophisticated controls, enterprises buy managed detection, threat intelligence, identity platforms, cloud security and insurance—and still nobody can credibly state the percentage of attacks the organisation will prevent.

That is not an indictment of cybersecurity. It is a clue that we may be asking the wrong question.

Public benchmarking confirms substantial enterprise cybersecurity expenditure, while also showing why precise cross-industry averages should be treated cautiously. IANS/Artico reported average security spending of 13.2% of IT spend and 0.69% of revenue in its 2024 CISO benchmark; in 2025, security spending represented 10.9% of IT spend. Definitions, accounting boundaries, sector exposure and company scale can materially change these ratios.^[1–2]

The more important finding is what the evidence cannot tell us. There is no accepted denominator for “attacks attempted”, no stable threat population against which enterprises can be compared, and no simple binary outcome called security. A control can fail while the incident is contained. An attacker can enter a system without causing material harm. A breach can result from several individually survivable weaknesses interacting.^[3]

The sensible metrics are therefore different: exposure reduction, control coverage, dwell time, detection and response performance, blast radius, recovery, loss avoidance and resilience. Cybersecurity is less like a wall with a measurable failure rate and more like a continuously changing system of defences operating against a continuously changing adversary.

The enterprise is not a stack. It is a dependency network.

The traditional diagram of enterprise security is usually drawn as layers: identity, endpoint, network, cloud, application, data. Operational reality is less tidy. A business service may depend on an internal team, an outsourced SOC, a SaaS provider, a cloud platform, a fourth-party processor, an open-source component and a specialist incident-response firm. Each relationship carries different evidence, contracts, responsibilities and failure modes.

Published evidence confirms substantial security-tool and supplier complexity, although estimates vary widely by survey and definition. Recent studies report anything from more than ten tools in many security operations centres to averages above fifty technologies in some enterprise samples. The precise number is less important than the operational consequence: modern security depends on multiple products, providers and data sources that must work together.^[4–6]

After a breach, “which product failed?” may be the wrong forensic question.

It is tempting to assume that digital forensics can always identify a definitive root cause. The evidence is more nuanced. Investigations can establish strong evidence about control weaknesses and contributing factors without reconstructing every step of the causal chain. In complex incidents, evidence may be missing, retention may be incomplete, telemetry may never have existed, and several failures may have combined. ^[7]

That produces an important governance distinction. Technical cause is not automatically contractual responsibility. Contractual responsibility is not automatically regulatory accountability. Regulatory accountability is not automatically legal liability. A mature investigation must preserve those distinctions rather than collapsing them into a convenient search for a single failed vendor.

This suggests a more useful objective: construct an evidence-backed causal graph showing which safeguards were expected, what status they had, what changed, which dependencies were involved, who had responsibility, and how confidence in each fact was established.

Cyber insurance exposes the same architectural problem.

Cyber insurance increasingly examines concrete security controls and evidence, but the research also warns against treating underwriting practice as uniform. MFA, endpoint controls, backups, vulnerability management and incident response are common underwriting subjects; exact requirements, exclusions, supplier definitions and dependent-business-interruption terms vary materially by carrier and policy. ^[8–10]

Insurance therefore creates a revealing test. The enterprise may maintain one view of its critical suppliers, procurement another, security a third, regulators a fourth and the insurance policy a fifth. If those maps do not describe the same dependency architecture, the organisation may discover the mismatch only when something fails.

The question is no longer simply: “Do we have cyber insurance?” It is: “Does our evidence, our contractual chain, our technical dependency map and our insurance wording describe the same risk?”

Nth-party risk is real—but infinite visibility is not the answer.

The research supports two positions that are often presented as opposites. Third-party, software-supply-chain and downstream dependency incidents are material. At the same time, organisations already possess meaningful mechanisms for direct-supplier due diligence, external monitoring, SBOMs, contractual flow-downs, regulatory registers and supplier mapping. ^[11–12]

The unresolved problem is not total blindness. It is fragmentation and depth. Complete recursive visibility across every supplier, sub-processor and transitive software component is neither consistently achievable nor necessarily useful. The more credible objective is

selective, criticality-driven visibility: follow the dependency chain as far as the business consequence justifies, and attach provenance and confidence to what is known. ^[12–15]

That is a more demanding idea than “know your suppliers”. It means knowing which dependencies matter to a particular service, obligation or loss scenario—and being explicit about where verified fact ends and inference begins.

The market has many answers. It does not yet have one shared truth.

The landscape research found mature or fast-maturing solutions for individual layers: third-party workflow and GRC; external security ratings; multi-tier supply-chain intelligence; SBOM creation and analysis; software provenance; supplier insurance-certificate tracking; and regulatory dependency registers. ^[13–16]

It did not find convincing evidence of a single established platform maintaining the full causal chain from a deep-tier corporate dependency, through software components and vulnerabilities, into contractual responsibility, regulatory controls, evidence state and applicable insurance.

That gap should not automatically be interpreted as an opportunity for one enormous replacement platform. The stronger architectural conclusion from the research is almost the opposite: the future may be federated.

From systems of record to systems of verifiable relationships

A federated model would allow evidence to remain with authoritative sources while exposing a continuously current status, provenance, timestamp, confidence level and relationship to the obligation being tested. A fact could be identified as observed, declared, attested, audited, inferred or cryptographically verified rather than being flattened into a generic green tick. ^[15–16]

This reframes compliance. Traditional compliance systems ask whether evidence exists. A relationship-based infrastructure asks whether the relevant claim remains true now, who can prove it, who depends on it, and what should happen when it changes.

It also reframes cybersecurity. The security industry has spent decades improving individual controls. The next problem may be orchestrating certainty across the relationships between those controls.

A possible new category: Continuous Trust Infrastructure

The literature and market landscape use several adjacent labels - dependency intelligence, compliance relationship infrastructure, digital trust infrastructure, federated assurance networks and compliance status infrastructure. None is yet an established market category.

That ambiguity may be significant rather than inconvenient. Established categories describe the component being purchased: TPRM, GRC, security ratings, SBOM, supply-chain risk, cyber insurance. The emerging requirement described by this research is horizontal: a way to

connect evidence, status, dependencies, obligations and responsibility across those categories.

One candidate description is Continuous Trust Infrastructure: infrastructure that does not promise that every participant is safe, but continuously establishes what can be trusted, why, by whom, for which purpose and for how long.

The controversial proposition is therefore this: perhaps enterprises do not primarily need another security product. They need an infrastructure for knowing whether the complex network of products, suppliers, software, controls and obligations they already depend upon is still behaving within the conditions on which trust was granted.

The questions analysts and boards should now ask

- Can we identify the dependencies that are genuinely critical to each important business service - not merely our direct vendors?
- Can we distinguish verified facts from declarations, ratings, inferences and assumptions?
- Can we show when a compliance or security status changed, and who upstream became exposed by that change?
- Do our technical dependency maps, supplier contracts, regulatory registers and insurance policies describe the same risk architecture?
- Can evidence be reused and referenced without every participant repeatedly collecting and storing another copy?
- When an incident occurs, can we reconstruct the relevant relationship and responsibility chain without pretending that technical cause and legal liability are the same thing?

Those questions do not have to lead to a new software category. But if enterprises increasingly need the same answers across cybersecurity, compliance, procurement, resilience and insurance, the market may eventually create one.

How does your organisation measure up?

Research note

This paper synthesises published research, standards, government material and established industry evidence. Where sources differ, the paper favours primary or methodologically transparent evidence, distinguishes observed findings from estimates, and avoids treating vendor claims as independently verified facts.

References & Further Reading

The references below are a consolidated, deduplicated selection of the strongest sources supporting the paper. Primary government, standards and original research sources are preferred over secondary summaries. Vendor sources are included only where they provide original market data or document a relevant technical capability.

1. **IANS Research & Artico Search.** [2024 Security Budget Benchmark Report](#) (2024). Original CISO benchmark; 755 respondents. Reports 13.2% average security budget as a share of IT spend.
2. **IANS Research & Artico Search.** [2025 Security Budget Benchmark Report](#) (2025). Original CISO benchmark; 587 respondents. Reports security budget share of IT spend falling to 10.9%.
3. **National Institute of Standards and Technology (NIST).** [SP 800-55 Vol. 1: Measurement Guide for Information Security — Identifying and Selecting Measures](#) (2024). Primary measurement guidance supporting the use of security measures and risk-oriented performance assessment rather than a universal protection percentage.
4. **SANS Institute.** [The State of Automation in Security Operations: A SANS Survey](#) (2024). Survey evidence on multi-tool SOC environments and the use of MTDD/MTTR as common performance measures.
5. **Kyndryl.** [Cyber Gauge 2024: Navigating the Complex Cybersecurity Landscape](#) (2024). Original survey evidence on security-tool/vendor complexity and interoperability challenges.
6. **Ponemon Institute / Sullivan Privacy Report.** [2025 Cybersecurity Threat and Risk Management Report](#) (2025). Survey reporting an average of 55 separate cybersecurity technologies among respondents; useful as directional evidence rather than a universal enterprise average.
7. **Cyber Safety Review Board / CISA.** [Review of the Summer 2023 Microsoft Exchange Online Intrusion](#) (2024). Primary incident review illustrating that substantial control and governance findings can coexist with incomplete causal evidence.
8. **Marsh.** [Cyber Hygiene Controls Critical as Cyber Threats Intensify](#) (2021). Broker evidence on underwriting scrutiny and commonly emphasised controls including MFA, EDR and tested backups.
9. **Marsh.** [Q4 2024 Update on the US Cyber Insurance Market](#) (2024). Current market evidence on underwriting scrutiny, cyber hygiene and third-party contract considerations.
10. **Marsh.** [Cyber Insurance and Its Role When Tech Outages Occur](#) (2025). Explains first-party and third-party coverage, including contingent business interruption, while reinforcing that coverage is policy-specific.
11. **Verizon Business.** [2025 Data Breach Investigations Report](#) (2025). Large incident dataset; 2025 findings reported third-party involvement in 30% of breaches.
12. **UK Government.** [Cyber Security Breaches Survey 2025/2026](#) (2026). Primary UK evidence on supplier-risk review: 48% of large businesses reviewed immediate suppliers and 24% reviewed their wider supply chain.
13. **CISA.** [Minimum Elements for a Software Bill of Materials \(SBOM\)](#) (2025). Primary guidance establishing SBOMs as a baseline mechanism for software and supply-chain transparency.
14. **SPDX / Linux Foundation.** [SPDX — System Package Data Exchange](#) (current). International open standard (ISO/IEC 5962:2021) for representing software components and related supply-chain information.
15. **CycloneDX.** [Authoritative Guide to SBOM / Software Dependencies](#) (current). Technical standard and guidance for representing direct and transitive software dependency graphs.
16. **OpenSSF.** [GUAC — Graph for Understanding Artifact Composition](#) (current). Open-source example of aggregating SBOMs, attestations and other software-security metadata into a relationship graph.

Source selection note

Secondary benchmark aggregators, generic statistics sites, search-result summaries and duplicative vendor commentary reviewed during the research were excluded where a stronger primary or original source was available. The reference list is intentionally selective: it supports the principal factual propositions in the paper without presenting every webpage consulted during the research.