Beyond the Scan

The Modern Intelligent Service Provider's Approach to CTEM

Chad Fullerton
VP, Information Security, ECI

September 2026

Found Out By Accident

This paper is written for the leaders accountable for risk inside regulated organizations - chief executives, boards and audit committees, risk and compliance officers, general counsel, finance and operating leaders, and the CISOs and IT leaders who work alongside them. Technology exposure stopped being a technology-department problem some time ago. It now arrives as a disclosure obligation, an examiner question, or a line item in a board packet, and it lands on whoever signs.

This paper argues that Continuous Threat Exposure Management, done well, is not a monitoring or alerting exercise. It is the difference between deciding what to do about a problem on your own terms and having it surfaced for you by an attacker, an auditor, or a regulator.

On May 5, 2026, Community Bank - a Pennsylvania-based subsidiary of CB Financial Services (NASDAQ: CBFV) - discovered that an employee had been running non-public customer information through an artificial intelligence application nobody at the bank had approved. The data involved included customer names, Social Security numbers, and dates of birth. Two days later, the bank's leadership determined the incident was material. On May 11, it filed a Form 8-K with the SEC citing unauthorized use of AI as the cause - the only filing of its kind on the SEC's record, according to a review by American Banker.

Nothing else broke. No systems went down, no customer lost access to their account, no ransom was demanded. That is precisely what makes the filing worth stopping on. The exposure had been sitting there, unremarkable and undetected, until someone happened to look. Every organization reading this and assuming it couldn't happen internally should sit with the fact that Community Bank almost certainly assumed the same thing, right up until it didn't.

This is the argument of this paper in miniature: an organization can be diligent and well-intentioned and still find out about its own exposure by accident, because the tools built to manage that exposure were never designed to see the whole picture, and because most managed service providers only show clients what they've paid to see. If you're responsible for your organization's technology risk, the question underneath everything in this paper is simple: would you know before your regulator does?

I lead ECI's GRC service line, which means I spend as much time inside compliance frameworks as I do inside raw technical data. I've watched the gap between those two worlds cost clients more than any single vulnerability could. This paper is my case for closing it — and for what Continuous Threat Exposure Management (CTEM) has to become to actually do that.

From Alerts to Decisions

Most conversations about exposure management start with alerts: how many came in, how many were false, how many got missed. Those numbers are real and they matter, but they are the wrong starting point. A modern intelligent service provider's job was never to respond to alerts faster. It is to make sure fewer decisions ever have to get made under alert conditions in the first place, because the problem was already understood, already prioritized, and already being handled before it needed a response at all.

The alternative is a ticket factory: a provider that processes volume and calls it security. Prophet Security's second annual State of AI in the SOC survey, published in August 2026, found the average organization leaves 28% of its alerts uninvestigated. That figure should be uncomfortable on its own - more than one alert in four is never looked at by anyone. What follows makes it considerably worse.

Sixty percent of respondents said an alert they never investigated later turned out to be material - it risked customer data exposure, system downtime, or business disruption - and for 34%, that happened three or more times in a single year. Separately, 40% have switched off a detection rule, or expect to, because they did not have the capacity to work through what it produced.

Source: Prophet Security, State of AI in the SOC 2026 (250 security leaders and practitioners, fielded by ViB, August 2026).

28%

Uninvestigated Alerts

Over a quarter of all incoming threat indicators are ignored entirely due to severe analytical fatigue and resource constraints.

60%

Material Consequences

A majority of unexamined alerts later manifest as actual, material security breaches or operational business disruptions.

34%

Repeated Failures

Experienced material consequences from uninvestigated alerts three or more times within a single year.

What an Unread Queue Actually Costs

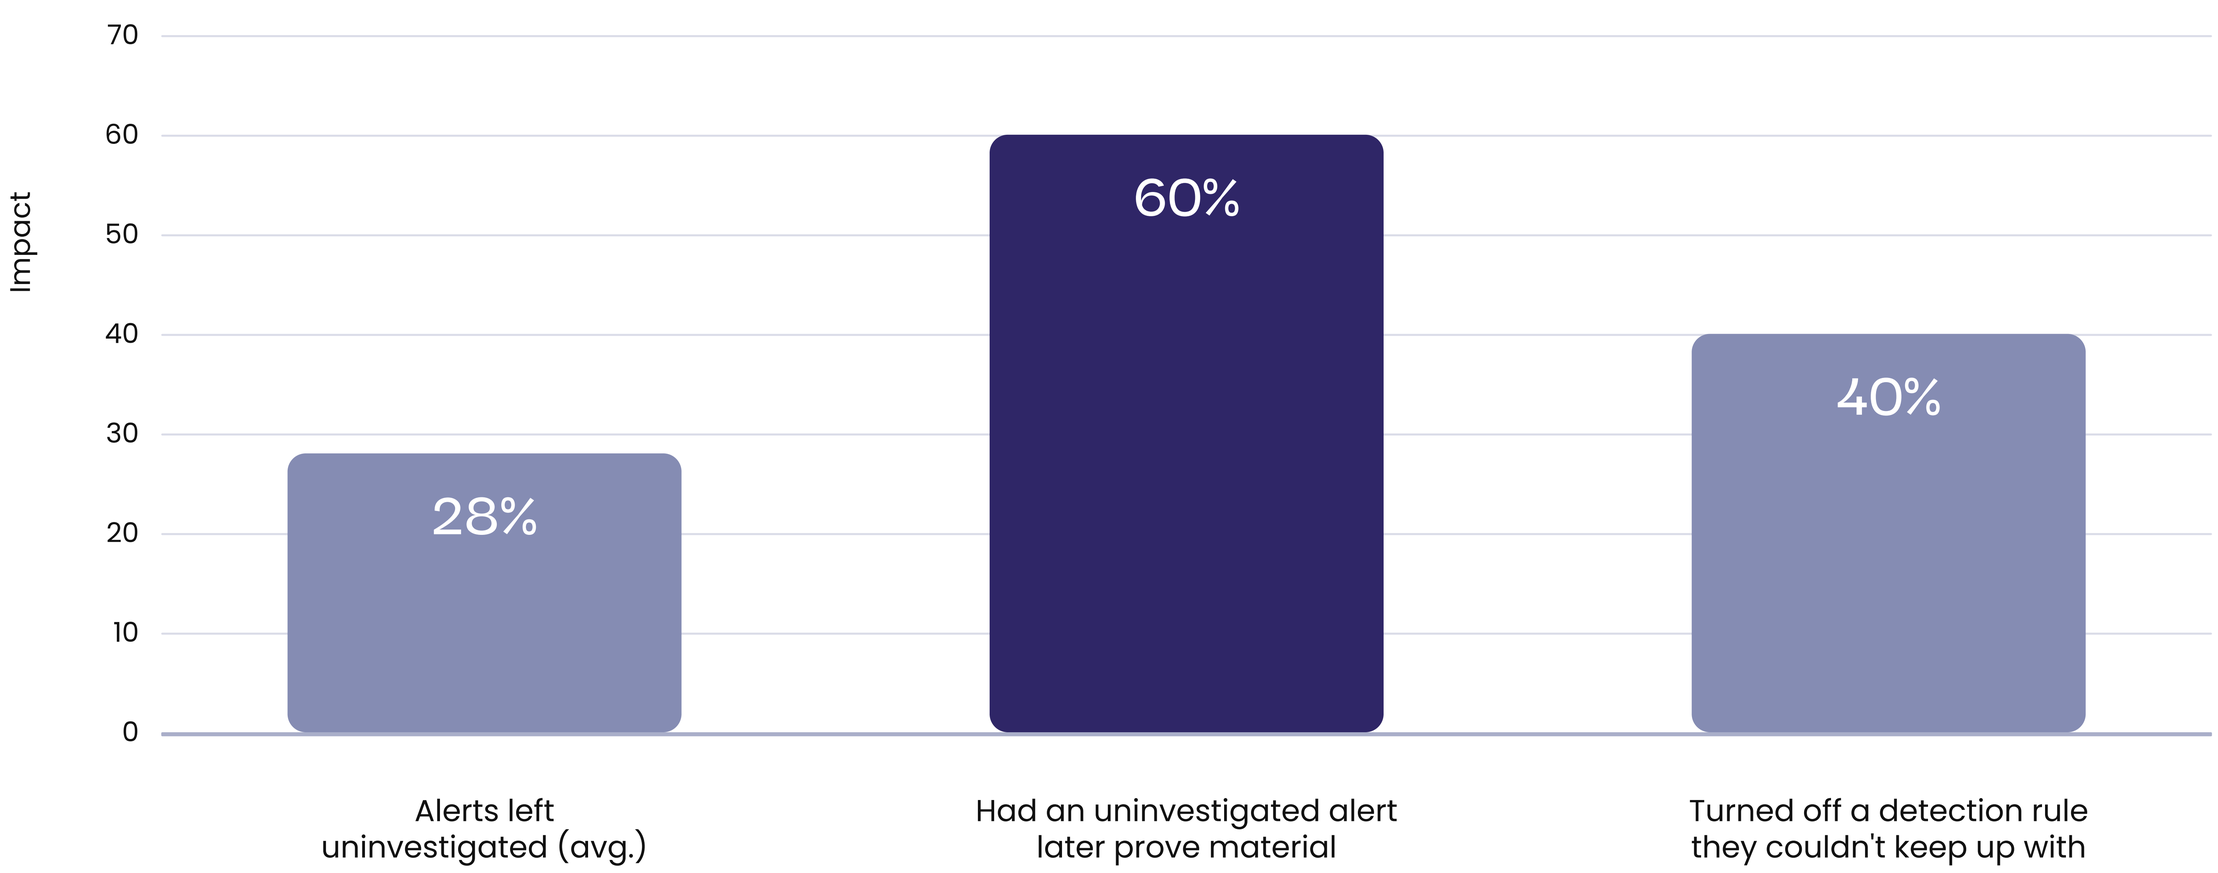

None of that is a call to triage harder.

Two of those three numbers describe consequences rather than workload, and the third describes an organization quietly reducing what it is able to see. Switching off a detection because nobody can work its output does not reduce exposure; it stops measuring it, which means the real figures run higher than the reported ones. Supporting research points the same direction - Microsoft and Omdia's State of the SOC 2026 put the false-positive rate at 46%, and the SANS Institute's 2025 Detection and Response Survey found 73% of security teams name false positives as their single biggest detection challenge.

Meanwhile the window for getting any of this right keeps narrowing. CrowdStrike's 2026 Global Threat Report puts average eCrime breakout time - the interval between initial compromise and lateral movement - at 29 minutes, with the fastest observed at 27 seconds. Most survey respondents need 30 minutes or more simply to investigate a single alert. You cannot out-hire or out-tune your way past that arithmetic. A model built around responding to alerts has already lost, structurally, regardless of how capable the analysts are.

The same pattern is compounding a layer up, in the exact places most exposure management programs don't look. IBM's 2026 Cost of a Data Breach Report found that shadow AI - AI tools and agents operating outside any approval process - factored into 43% of the breaches it studied, more than double the year before, even as the share of organizations requiring IT approval before deploying AI actually fell. Separately, an April 2026 Cloud Security Alliance survey found that 82% of organizations have AI agents running in their environment that nobody signed off on.

Average eCrime Breakout Time

29 Min

The window for proactive containment is shrinking rapidly. Defenders can no longer rely on manual triage pipelines.

46%

False-positive rate per Microsoft & Omdia State of the SOC 2026

82%

Of orgs have unsanctioned AI agents running in their environments

Shadow AI's Share of Breaches Studied

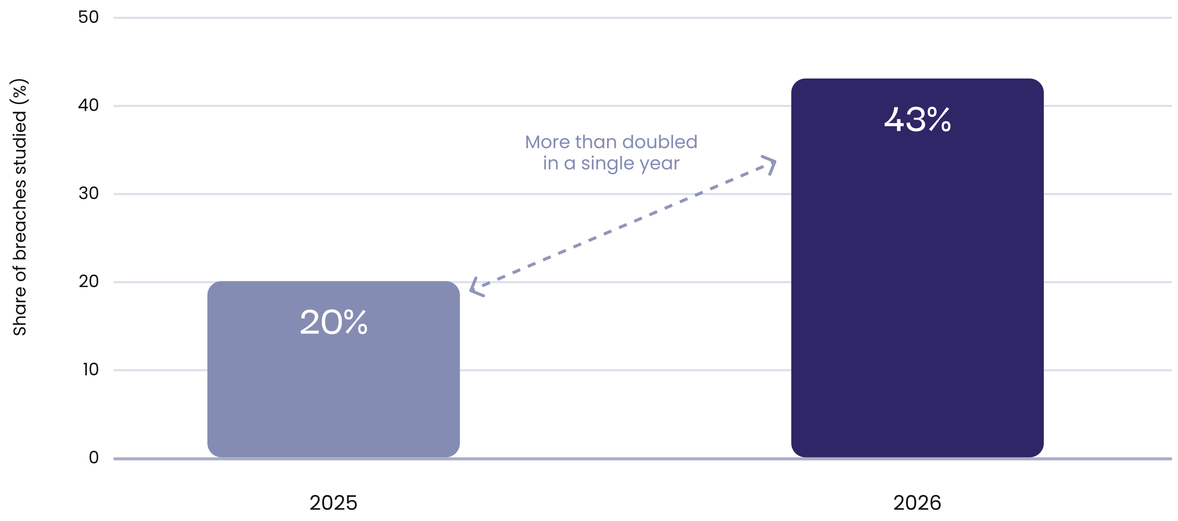

You cannot improve what does not get measured. But the answer to that is not more data feeds — most organizations already have more feeds than anyone has time to use. It is bringing the feeds that already exist into one place where an actual decision gets made: this matters, this doesn't, this gets fixed automatically, this gets escalated to a person, this becomes a question for your own risk appetite. That central decision-making platform, not another alert stream, is the real point of what a modern intelligent service provider has to build.

Source: IBM Cost of a Data Breach Report 2026 (Ponemon Institute), based on 602 breached organizations studied March 2025-February 2026.

What CTEM Was Built to Solve

That decision-making platform is what Continuous Threat Exposure Management, done right, should already be. It's worth being precise about where the term came from, and where common practice still falls short of it.

Gartner introduced Continuous Threat Exposure Management in 2022 as a corrective to a specific failure mode: point-in-time vulnerability scans, CVSS-driven backlogs, and siloed tools that could no longer keep pace with how attackers actually move through modern, cloud-native environments. The framework runs as a five-stage, repeatable cycle - scoping, discovery, prioritization, validation, and mobilization - designed to turn exposure data into measurable risk reduction rather than a static list of findings.

In practice, though, most CTEM programs still orbit a narrower set of tools than the framework's five stages suggest. Industry analyses of what actually powers a CTEM program point overwhelmingly to external attack surface management, cloud asset inventories, breach-and-attack simulation, and vulnerability prioritization - useful tools, but tools built almost entirely around infrastructure and code, not around the AI platforms, SaaS sprawl, and compliance posture that now make up an equally large share of real exposure.

That's the traditional view of CTEM: a genuinely better way to run vulnerability management and cloud security posture management, wrapped in a continuous operating model. It is not, on its own, a complete picture of organizational risk — and for a managed service provider trying to operate as more than a vulnerability shop, it can't be the whole strategy.

There is a further gap worth naming, because it shapes what this paper proposes. CTEM as originally framed was not built against any standard. It is a methodology — a sequence of stages describing how to run a program — and it deliberately stayed independent of the security and compliance frameworks organizations are actually measured against. That independence helped adoption, but it left every program to invent its own relationship between an exposure finding and the control language its regulators, auditors, and board understand. The result is two parallel vocabularies describing the same environment, reconciled by hand, usually under deadline. Aligning exposure management to recognized frameworks from the outset is not a decorative addition to the model. It is what allows a finding and an obligation to be expressed in terms that reconcile on their own.

The Modern Intelligent Service Provider's Approach to CTEM

A Modern Intelligent Service Provider – a MISP – has to treat CTEM as the operating model for an entire environment, not just its infrastructure.

A Modern Intelligent Service Provider — a MISP — has to treat CTEM as the operating model for an entire environment, not just its infrastructure. That means folding in domains most CTEM programs never touch: SaaS and shadow IT, artificial intelligence and shadow AI, data security, endpoint posture, cloud and network drift, and compliance — each pulling from real, live data in your own environment rather than a quarterly assessment, and each tied back to one throughline: why does this matter to this particular business.

Vulnerability management and cloud posture stay central to that picture — they don't disappear. What changes is that they become two domains among several, not the entire program.

SaaS and AI deserve particular attention, because this is where most CTEM programs currently have the least visibility and where the risk is compounding fastest. Unsanctioned SaaS tools, AI copilots operating outside SSO, and autonomous agents no one approved are not edge cases anymore — IBM's data on shadow AI's growing share of breaches makes that clear.

A MISP's exposure program needs to answer the same questions for an AI agent that it already answers for a laptop: what does it have access to, what account is it running under, when did it last authenticate, and who approved it in the first place. Most of that context already exists somewhere — in Entra ID, in Defender for Cloud Apps, in the ticket that recorded why the tool was requested — it just isn't being pulled together and read as a single exposure story.

Two misconceptions make this harder than it needs to be, and both are common in otherwise well-run organizations. The first is that a paid or enterprise-tier AI subscription is inherently safe — that because the organization is paying, its data is not being used to train anyone's model and the risk is handled. Price is not the control. What matters is the contractual data-handling terms, where prompts and attachments actually travel, and whether the tool is reachable through a corporate identity or a personal one. The same product can be low-risk on a governed corporate tenant and a reportable exposure on an employee's personal account, and most organizations have no view of which of the two is happening.

The second is treating conversational AI and autonomous agents as one risk instead of two. Organizations that have carefully governed which chat assistants employees may use often have no inventory of the agents running against their data, and organizations focused on agent governance frequently have no visibility into what is being pasted into a browser tab. These are different exposures with different shapes: one is about what leaves, the other is about what acts. A serious exposure program has to answer for both, and the Community Bank filing is a reminder that the first kind is entirely sufficient to produce a material disclosure on its own.

There is a third layer that is easier to miss still, because it never shows up as a discrete incident. Organizations frequently have no real visibility into how data actually flows through the AI systems they've already approved: what informed a given output, how that system's behavior has shifted as it kept learning from user interaction and feedback, or where the data behind a specific answer originally came from. As models keep adapting, control, traceability, and accountability get harder to maintain, not easier — which raises the stakes on data leakage, intellectual property exposure, regulatory exposure, and outcomes nobody specifically intended.

37%

rise in cloud-conscious intrusions in 2025 alone, with valid account abuse behind roughly a third of all cloud incidents. — CrowdStrike, 2026 Global Threat Report

Regulators have not caught up to the pace of this yet, but the direction is clear. NYDFS issued industry letters in May 2026 warning regulated financial institutions about the cybersecurity risks posed by frontier AI models, and FINRA's 2026 Annual Regulatory Oversight Report added a dedicated section on generative AI for the first time. Both are framed as guidance interpreting existing rules rather than new binding requirements — for now. In practice, most organizations do not prioritize a risk like this until it maps cleanly to a specific regulatory requirement, framework, or examiner expectation. The organizations that get ahead of that mapping before it is mandatory will spend far less catching up than the ones waiting for the requirement to arrive.

Cloud and network posture follow the same logic on a different data source. An environment that passes today's review can drift out of alignment within months as configurations change, new resources spin up, or requirements shift. The same pattern shows up on the network side, where posture and weaknesses shift quietly as changes accumulate. Both need to be watched as continuously as vulnerabilities are, not revisited once a year in an

Reduced to a single number, the scope gap is stark. Across the seven domains this paper argues belong in a modern exposure program, traditional CTEM tooling directly and fully addresses two. To be explicit about what the figure below is: seven of seven is the target state this paper is arguing for — a standard against which any provider, ECI included, should be measured. It is not a claim that anyone has finished building it.

Exposure Domains Covered

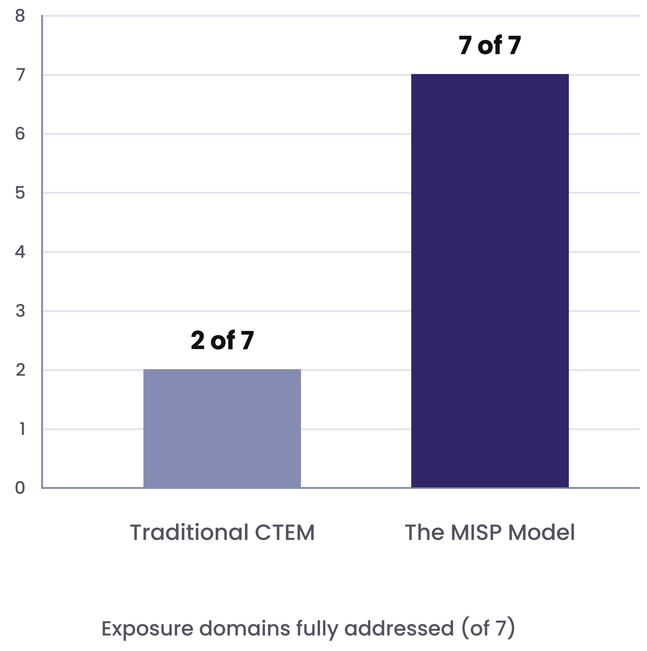

The MISP Model represents the target-state standard argued for in this paper, not current-state coverage by any single provider. 'Fully addressed' reflects direct, continuous coverage rather than partial or occasional visibility.

Domain	Traditional CTEM	The MISP's Expanded View
Vulnerability Mgmt.	Scored primarily by CVSS / EPSS	Scored by asset criticality, exploitability, and business context
Cloud Posture	Point-in-time CSPM findings	Continuous drift detection across M365, Azure, and AWS
SaaS & Identity	Rarely in scope	Shadow IT, SSO status, data permissions, oversharing
Artificial Intelligence	Not addressed	Shadow AI, LLM and agent risk, personal vs. corporate account use
Endpoint / XDR	Alert volume reported	Alerts received vs. auto-triaged, SOC engineering value shown
Network	Occasional pen-test findings	Continuous drift and weakness detection
Compliance	Framework mapping / audit checkbox	Contextualizing overlay grounded in individual business risk

Grounding Every Finding in Business Risk

Vulnerability management is the clearest example of what changes when context enters the picture, and it's worth going deeper here because it shows where the whole model is headed. Risk-based vulnerability management already moves past CVSS alone, layering in the CISA Known Exploited Vulnerabilities catalog and EPSS exploit-prediction scores.

A MISP's exposure program goes a step further: it adds the asset's own context — internal or external facing, what's actually running on it, who owns it, whether it sits inside a defined recovery time objective — so a moderate-severity finding on a crown-jewel system can outrank a critical-severity finding on an isolated test box. That's not a theoretical improvement. It's the difference between fixing the number and fixing the risk.

An engagement from my own portfolio makes the point better than the abstraction does. We took on a healthcare software client whose vulnerability counts looked alarming on first inspection. The organization ran on Macs at scale, with no central management platform in place to patch or configure them consistently. Measured by raw severity count, that was the emergency, and it would have been an easy one to sell against.

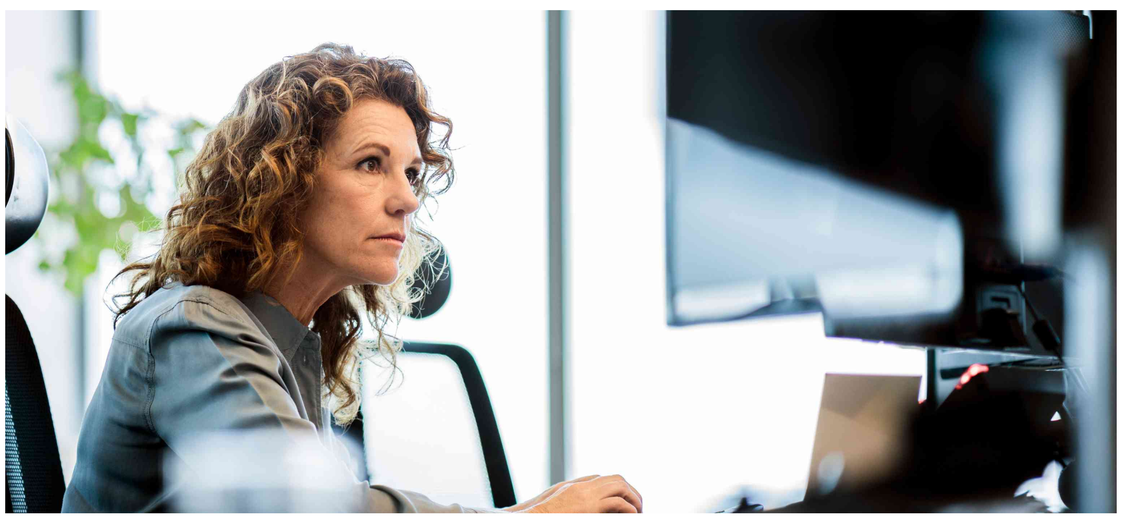

The context changed the conclusion. The client's actual product was a cloud-hosted platform running in Azure, and their development environment was cloud-operated rather than local. The endpoints were how people reached the work. They were not where the work lived, and they were not the path to the data that mattered. Once the compensating controls already present in that architecture were properly accounted for, the real risk carried by the unmanaged Mac fleet fell considerably — and the priority of the cloud environment, along with the SaaS product's own scope, rose to meet it.

That reprioritization is what made the rest possible. Instead of spending the engagement chasing endpoint findings, the effort went into continuous review of the scope that actually carried the risk, and stayed there as the client was taken through a HITRUST audit. They passed with zero findings.

The honest version of that story is that the loudest thing the scanner reported was never fixed. What happened instead was that its real weight was established, the controls that made it less dangerous than it appeared were documented, and the effort was spent where the exposure genuinely was. A severity-driven program cannot arrive at that conclusion, because it has no mechanism for asking what an asset is actually for.

Asset context is one dimension of that question. How an exposure would actually be used is another, and it is a deeper problem than this paper can properly treat. Attack-pattern libraries — MITRE's CAPEC catalog for adversary patterns, ATT&CK for observed behavior — let a program reason about whether a weakness is reachable through a route an adversary genuinely takes, rather than whether it cleared a scoring threshold. Modeling exposure that way is a discipline in its own right and deserves its own paper. The point here is narrower: severity scores and asset context together still stop short of the question that matters most, which is whether anything can actually be done with this. It is worth being blunt about the cost of getting that wrong in the other direction.

Detection and response tooling that produces findings faster than anyone can qualify them does not merely waste analyst time. It generates action on things that never warranted it, erodes confidence in the next finding that arrives, and in the worst cases drives remediation that damages an environment more than the original exposure ever would have. A program that cannot separate a real path from a scored number is not neutral in its effects. It carries costs of its own, and they are frequently invisible on the reports that describe it as working.

That shift in emphasis just became mandatory, whether providers are ready or not. In April 2026, NIST changed how it operates the National Vulnerability Database: going forward, NVD reliably enriches only CVEs tied to CISA's Known Exploited Vulnerabilities catalog, software used by the federal government, or software designated critical under Executive Order 14028. Everything else — the large majority of new CVEs — still gets listed, but not scored, not mapped to affected products, and not classified, for the foreseeable future. Any vulnerability program built on the assumption of a complete, centrally-scored CVE feed lost that assumption this year.

That makes remediation for the sake of remediation less viable, not more. Patching everything above a CVSS threshold was already an inefficient strategy; it is now an incomplete one, because a meaningful share of what matters will not have a NIST-issued score to sort by at all. What replaces it is exactly the model this paper is describing: remediation targeted at the assets that actually carry risk, real weight given to compensating controls where remediation is not immediately viable, and all of it surfaced through the same central decision-making platform that already knows which assets matter and why, rather than waiting on an external database to say what's urgent.

All of that rests on something unglamorous: an accurate, continuously maintained configuration management database. Asset context is the input that makes every other judgment in this model possible — without a reliable picture of what exists, what runs on it, who owns it, and what it connects to, risk scoring degrades back into sorting by severity number, and automation becomes something no responsible provider would turn on.

Organizations that treat the CMDB as an inventory chore rather than the foundation of their risk model tend to discover the gap at exactly the wrong moment.

> A program that cannot separate a real path from a scored number is not neutral in its effects.

The automation opportunity that follows is already precedented in the industry, not speculative.

Tanium's model for automated remediation, for example, tiers assets by business impact: lower-tier assets — test environments, low-impact workloads — become strong candidates for automated remediation within defined risk thresholds, while production and business-critical systems still route through change approval. Applied across a client base, that logic produces a straightforward set of expectations you should be able to hold a provider to. If a change is validated as low-impact and carries no meaningful chance of disruption, it should simply be made, and you should be told afterward. If it is moderate — a device reboot, an application restart — you should receive a maintenance window you can approve early or move. If it is significant — a network device, something that touches production — the decision stays yours, but it should arrive with a proposed schedule already attached, rather than as an open question that waits until someone chases it.

There is a failure mode on the automation side that deserves naming directly, because the case made above understates it. Controls that adapt to risk signals — conditional access that tightens when a session looks anomalous, automated isolation of a host, policy that reconfigures itself as posture shifts — are effective precisely because they act without waiting for a person. That same property is what makes them capable of breaking things that were working perfectly well. A tightened access policy locks a legitimate workflow out at the worst possible hour. An automated remediation pushes a configuration that an application quietly depended on. An isolation action takes a production system offline over a signal that turns out to be noise. Each of those is a security control performing exactly as designed and producing an operational incident anyway

That is not an argument against automation.

It is an argument for deploying automation the way any other production change gets deployed. Adaptive controls need performance and availability baselines captured before they go live, so degradation is measurable rather than anecdotal. They need staged rollout against a representative subset rather than an entire estate at once. They need a rollback path that does not depend on the person who built it being reachable. And they need monitoring that tracks user and system impact alongside security outcomes, because a control that improves one metric while quietly damaging the other is not actually working. A provider that can describe those safeguards for its own automation is running a service. One that cannot is running an experiment on your environment.

This is also where XDR reporting earns its place. Knowing how many alerts arrived, how many were automatically triaged and closed, and how many escalated into a genuine incident serves two purposes at once. Operationally, it is what tells a provider how to staff against rising volume instead of guessing at it. From where you sit, it is the difference between a provider asserting that it is working and a provider showing you the shape of that work — including the large share of it that never reaches you as a ticket. Work that never gets surfaced is, from your side of the relationship, indistinguishable from work that was never done.

There is a practical reason the automation question matters to you and not only to your provider. Any organization that has waited on a patch cycle behind a queue of someone else's tickets has felt the constraint: the team of one, or the group that has to add headcount every time it takes on more work. When low-impact remediation is genuinely automated — scanned, scoped, tested, and pushed without a person touching every ticket — that constraint stops governing how quickly your environment gets attention. Whether a provider has solved it determines whether you are quietly competing with their other commitments for capacity, particularly as AI raises the volume every provider has to absorb.

Measuring What Actually Matters

Measurement only helps if the measure is one a non-technical leader can act on. The top of that structure should be a single, plain-language standing: a health score, expressed as a grade or a clear rating, that tells an executive where the organization sits today and, more importantly, which direction it has moved since the last time they looked. That score should be anchored in a recognized framework so it means something outside the room, and then weighted by the organization's own risk profile so it means something inside it. Two firms with identical framework coverage do not carry identical risk, and a score that cannot reflect that difference is decoration.

Underneath the score sit the operational measures that explain it. Mean time to detect and mean time to respond remain the clearest indicators of whether a provider is actually shortening the window an adversary gets to work in. Alongside them: how quickly identified weaknesses are contained, what share of the estate is genuinely under coverage rather than assumed to be, how exposure counts are trending by domain, and how many actions were taken on the organization's behalf without anyone needing to be chased for a decision. Each of those is a number a board can follow quarter over quarter without a translator.

The context for why those numbers matter is not subtle. IBM's 2026 report found the mean time to identify and contain a breach rose to 247 days, reversing five consecutive years of improvement. That is roughly eight months between the beginning of an incident and the end of it — and it is getting longer, not shorter, at exactly the moment attacker tooling is getting faster.

247 days

average time to identify and contain a breach in 2026 — up from 241, reversing five straight years of decline. (IBM Cost of a Data Breach Report 2026)

The score itself is never the point. The point is that it moves, that the reason it moved is traceable to specific decisions and specific actions, and that everyone from the board to the operations team is reading the same underlying data when they ask why.

Compliance as Connective Tissue, Not the Finish Line

It helps to be precise about what “compliance” actually means in a given organization, because it usually means one of four quite different things. Some obligations are genuinely audited, with an examiner or assessor who will ask for evidence. Some are required by a regulator that does not routinely audit against them, where the obligation is real but the verification is episodic. Some are self-attested, where the organization asserts its own posture and carries the consequences if that assertion turns out to be wrong.

And some are simply good practice — CIS Controls, NIST CSF 2.0, the NIST Risk Management Framework and its AI counterpart — adopted because they are sound, not because anyone requires them. Each of those carries a different evidentiary burden and a different cost of being wrong, and a program that treats all four identically will over-invest in some and leave the organization exposed on others.

Security and compliance frameworks — NIST CSF 2.0, CIS Controls, SOC 2, the sector-specific regimes that apply to regulated industries — are genuinely useful for bringing all of this data together into something a leadership team can act on. But a framework alone was never meant to be the end state. Two organizations can map to the same framework and carry entirely different risk, depending on what's actually running in their environment, what data it touches, and what their board is willing to tolerate. A MISP's compliance overlay has to start from your own written policies, measure those policies against the framework's actual control language, and then map each control to the technology and data already being monitored across every other domain — so that whether you are meeting a given requirement is something the platform can answer directly, not something a consultant reconstructs by hand every audit cycle. None of that has to be manual anymore. Building it that way is where the most meaningful gains in the discipline are still available.

Visibility as the Starting Point

Visibility has historically arrived on a cycle rather than continuously, and the reason is more practical than commercial. Assembling a defensible picture of an environment has genuinely been hard work. It has meant people pulling exports from a dozen consoles, normalizing them into a common shape, reconciling what disagrees, chasing down what is missing, and then translating the result into something a leadership team can read. That effort is real, it takes expertise, and it is precisely why the output has typically been produced as a discrete piece of work with a beginning and an end.

The trouble is that the four compliance postures described earlier all need something slightly different from that picture, and none of them are well served by one assembled at a point in time. An audited obligation needs evidence retrievable on demand, not reconstructed the month before fieldwork.

A regulator-expected obligation needs demonstrable readiness in the long stretches between examinations. A self-attestation needs a defensible basis at the moment it is signed — which is a specific date, and rarely the date the last assessment closed. And a framework adopted as good practice needs a live benchmark to measure drift against, or it becomes a document rather than a control. In each case the underlying data exists somewhere in the environment. The constraint has been the cost of gathering and interpreting it.

That constraint is what changes when the domains in this paper feed one platform rather than a dozen disconnected ones. When SaaS, AI, endpoint, cloud, network, vulnerability, and compliance data are already aggregated, normalized, and mapped to control language, producing the picture stops being a project and becomes closer to a query. Information that was previously slow to obtain and slower still to digest becomes something an organization can look at when it wants to, in a form it can act on.

That does not remove the need for expertise. It moves expert time toward the part that actually requires judgment — interpreting what the data means for a specific business, deciding what to do about it, and doing it — rather than the mechanical work of assembling the picture in the first place. Heading into an environment that changes as quickly as this one now does, that is the shift a Modern Intelligent Service Provider should be working to deliver.

Reporting That Reaches the Right Audience

None of this matters if it only gets presented one way. The same underlying exposure data has to reach at least three audiences that need genuinely different views of it — not the same dashboard with more or fewer rows.

A board or audit committee needs the least and the most consequential: where the organization stands, whether that is better or worse than last quarter, what the material exposures are in plain language, and what has been formally accepted as risk. An executive needs a slightly tighter frame — what is being asked of them to approve, what it costs, and what happens if they do nothing, with a date attached. An operations team needs everything underneath: who authenticated an unsanctioned tool, when, what ticket justified standing it up in the first place, and what the remediation path looks like. That is the layer where the work actually happens, and it is the layer that makes the two above it credible.

Building all three views from the same underlying data, rather than assembling a separate report for each audience, is what keeps them consistent with one another — so the figure in the board packet and the figure in the operations queue have the same origin and survive being compared. It is also the only way the work you would otherwise never see actually reaches you: vulnerabilities closed quietly, incidents triaged before they became incidents, drift caught before it became exposure. None of that travels upward on its own. Someone has to build the path deliberately, in language suited to whoever is reading.

What This Looks Like for ECI

Everything in this paper points toward the same conclusion: a modern intelligent service provider needs one place where all of this comes together — SaaS and AI exposure, vulnerability and cloud posture, compliance context, and your own risk appetite — and turns it into decisions, not just more dashboards. That is what we are building with ELLA: bringing these domains into a single platform built to support that kind of decision-making, rather than adding another feed to an already crowded stack.

That is the standard we are holding ourselves to. Reaching all seven domains at full depth is a direction of travel rather than a finished claim, and we would rather publish the standard and be measured against it than describe a product that does not yet exist. The providers who get there first will be the ones organizations trust most as AI reshapes what both attackers and defenders can do. A modern intelligent service provider is not defined by which tools it licenses. It is defined by whether it can turn everything those tools already know into one honest, real-time picture of where you actually stand — and help you decide what to do about it before an attacker, an auditor, or a regulator decides for you.

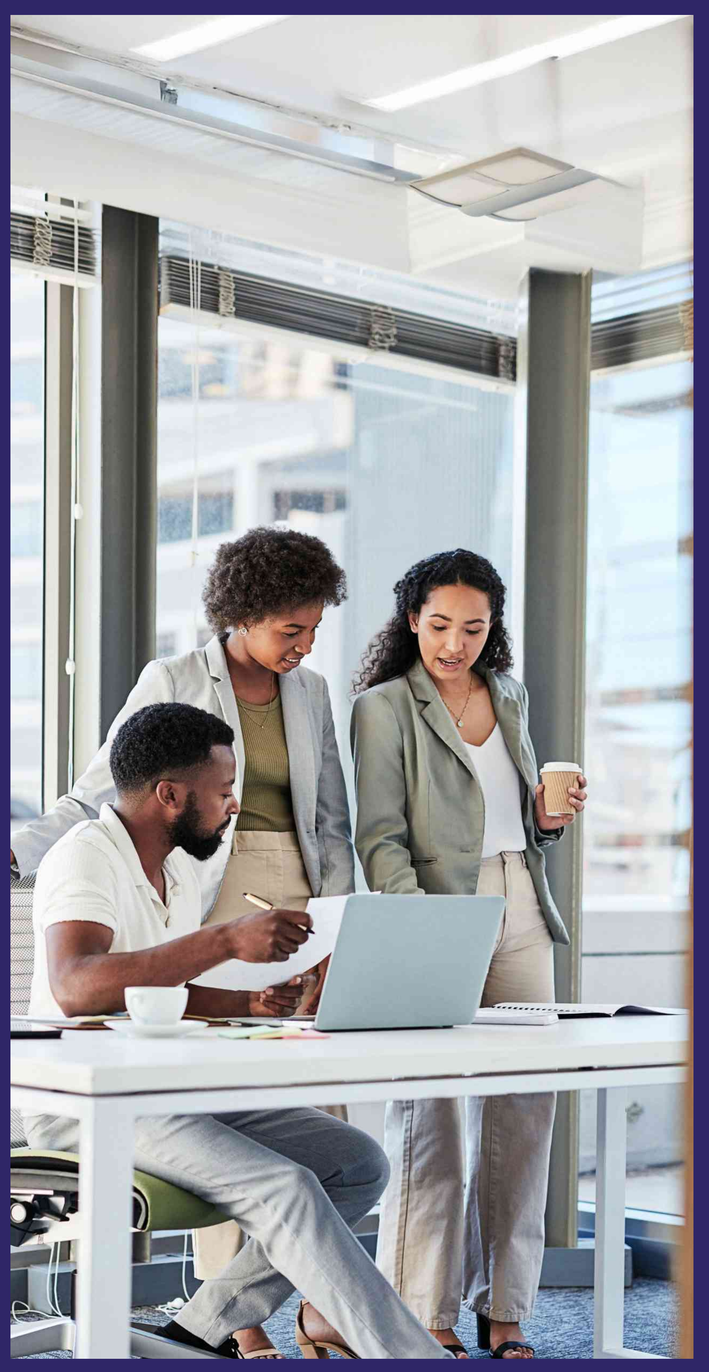

Where to Start

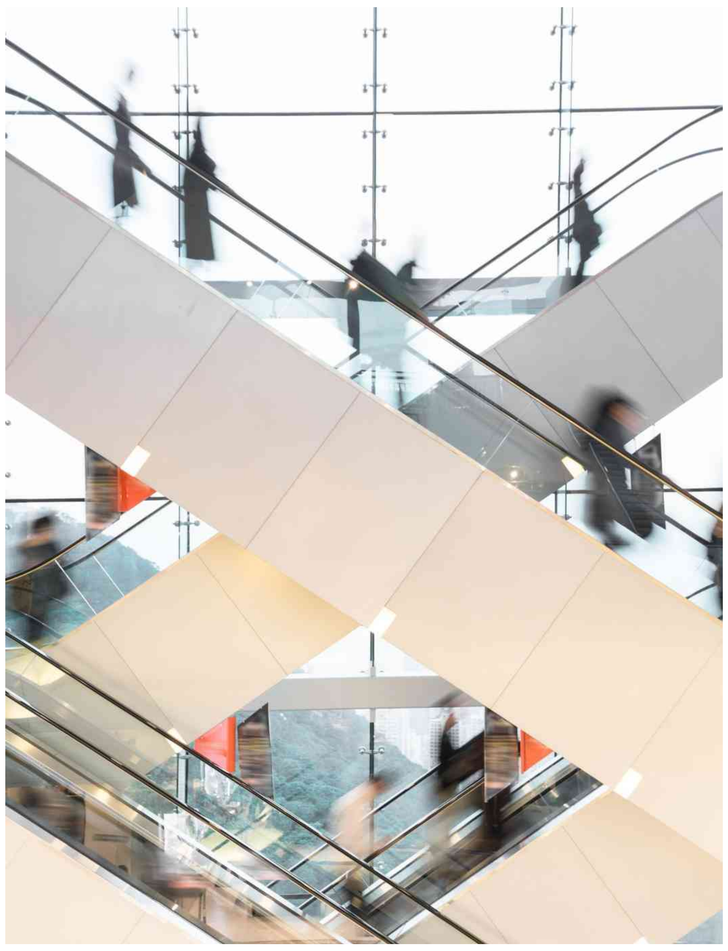

The organizations that work through these questions deliberately over the next few years will be the ones making decisions about their exposure. The ones that defer them will be responding to decisions made for them — by an adversary, an examiner, or a disclosure deadline. That difference is worth more than any single control, and it starts with a conversation rather than a contract.

None of this has to be solved at once, and none of it requires a purchasing decision. What it requires is putting a handful of questions on the agenda before circumstances put them there for you. A few worth raising internally, and with whoever supports your environment:

Is your asset picture good enough to prioritize on? Risk decisions degrade quietly when the underlying inventory does, and the degradation is rarely visible until it matters.

Do you have GRC expertise at the table alongside technical staff, rather than engaged only at audit time? The distance between knowing the correct control and knowing whether it fits your risk appetite is where most programs quietly stall.

Do you know how your obligations — audited, regulator-expected, self-attested, and adopted as good practice — map to what is actually running in your environment today? And how old that mapping is?

Is your asset picture good enough to prioritize on? Risk decisions degrade quietly when the underlying inventory does, and the degradation is rarely visible until it matters.

What will your position on AI governance be when it becomes an examiner question rather than a good idea — and who in your organization owns drafting it?

About the Author

Chad Fullerton is VP of Information Security at ECI, where he owns the firm's GRC service line and works across product strategy, client delivery, and business development. His background spans GRC advisory, CMMC, TISAX, NIST CSF 2.0, SOC 2, and SEC/FCA financial services compliance, alongside deep experience in MSSP delivery models.

Disclaimer

This paper reflects industry data and research available as of September 2026 and ECI's product direction as of the same date; both may change as the threat landscape and ECI's own platform continue to evolve. References to ECI capabilities describe the direction the company is building toward and may not reflect every capability available today.

US: +1 800 752 1382

Singapore: +65 6622 2345

UK: +44 207 071 6802

Hong Kong: +852 3189 0101

eci.com

This paper provides an independent assessment of enterprise AI platforms for alternative investment managers and boutique financial services firms. It is not a product pitch. It is a decision framework built on real-world deployment experience across 750+ firms managing complex regulatory and security requirements.