Beyond the Noise: Redefining Enterprise Risk in the Era of AI Vulnerability Discovery

OPTAS™ by origina.

Table of Contents

Executive summary

While artificial intelligence (AI) has accelerated and reduced the cost of discovering software vulnerabilities, it has not necessarily enhanced enterprise safety. This paper addresses the rising flood of findings, alerts, and theoretical exposures that consumes security capacity without enhancing judgment about what truly matters.

Noise represents not only an operational challenge but also a significant business cost. Every alert requires triage. Every theoretical exposure prompts scrutiny across security, operations, audit, risk, compliance and executive teams. Over time, this leads to backlogs, increases pressure to act swiftly, and risks conflating activity with assurance.

Origina's stance is clear: while AI-driven discovery is crucial, it alone does not yield the expected outcomes for organizations. The real differentiator is the ability to convert intelligence into prioritized, actionable responses rather than simply reacting to a growing volume of alerts. That is the role Origina has built for: taking the noise out, so leaders can act on genuine threats.

Traditional vulnerability management models rely heavily on public disclosures, OEM advisories and standard severity scores. While valuable, these approaches are often reactive and can be misaligned with true business risk. Organizations require a proactive model that filters signals from noise, assesses genuine exposure – even in the absence of patches – and facilitates confident, informed decision-making.

At the same time, regulatory frameworks such as NIS2, DORA and PCI DSS are driving a shift from compliance as a checklist to resilience as a discipline. Responding to vulnerabilities is no longer enough: organizations must demonstrate understanding of critical dependencies, assess exposure, prioritize effectively, manage third-party risk and maintain evidence of due diligence.

In this environment, leaders need independent, context-driven intelligence. The key questions are no longer about volume, how many vulnerabilities exist, but about relevance and impact: Which risks truly matter in this environment? Which are exploitable? What action is required? What mitigation is practical? What evidence supports the decision?

Origina's Proactive Threat Assurance Service (OPTAS), offered on Origina-supported products, is designed to meet this need. By combining predictive intelligence with expert validation and actionable guidance, OPTAS enables organizations to achieve earlier visibility, sharper prioritization and context-specific action, supporting a shift from reactive response to proactive risk management.

For CIOs, the next steps are immediate and practical. Before the next board or risk cycle: identify which mission-critical platforms depend entirely on OEM patch cycles for their security posture; ask what proportion of last quarter's alerts led to validated, evidenced action rather than triage effort; and, for Origina-supported products, request an OPTAS briefing to see what proactive threat assurance looks like in your own environment. The window matters. Organizations that act now will meet AI-accelerated discovery with control and evidence. Those that wait will meet it with backlog.

The rapid evolution of the AI landscape

AI is rapidly evolving across various sectors, particularly in cybersecurity, where it accelerates vulnerability discovery. While this can improve security research, organizations must adapt their decision-making processes to cope with the influx of information. The challenge lies in distinguishing actionable insights from irrelevant noise, especially in complex enterprise environments.

AI-assisted tools can analyze large codebases, detect structural patterns, infer weakness classes and surface possible vulnerabilities at a scale that would have been difficult to achieve through traditional manual analysis alone. This changes the rhythm of vulnerability management. Discovery is no longer constrained by the availability of highly specialized researchers working through a limited number of targets. Instead, AI enables continuous, large-scale discovery across software ecosystems.

This is not inherently negative. Used responsibly, AI has the potential to improve security research, identify weaknesses earlier and support better defensive planning. However, the strategic impact is broader than the technical capability. When discovery accelerates, downstream decision-making must also evolve. Organizations need to know whether a finding is real, reachable, exploitable, applicable to their environment, and whether action is required. AI may increase the volume of possible answers, but it does not automatically provide enterprise-grade judgement. Nor is the acceleration one-sided: the same techniques are available to adversaries.

Elena Donea, managing director of WhatsExposed, a penetration-testing firm and Origina partner, has watched the change from the front line.

“AI has changed the maths for attackers,” she says. “Finding and weaponizing a vulnerability used to take a skilled researcher and time. Much of that can now be automated, so the window between a weakness existing and someone finding it is shrinking fast.”

That compression breaks any security model built around the calendar rather than the threat.

“If you test once a year and a flaw appears the week after, it sits open for the rest of the year while attackers scan continuously,” Donea says. “The answer is not more AI on its own. AI is very good at surfacing known patterns at scale, but exploiting business logic and chaining flaws together still takes human judgement.

“What works is expert-led testing run continuously, with AI doing the heavy lifting in the background. That is how you keep pace with a threat that no longer waits for your renewal date.”

The challenge becomes particularly acute in complex enterprise estates. Large organizations often rely on mature platforms, customized deployments, layered dependencies, open-source components and integrations that have evolved over many years. Even if the AI tooling is provisioned with full access, including to core codebases, these environments cannot be understood through generic discovery alone.

A vulnerability in a product does not automatically translate into business exposure. Equally, a lack of OEM patch availability does not automatically mean a customer has no practical mitigation path.

This is the point at which AI changes from a technology conversation to a governance conversation. Senior leaders must ask whether their organization has the operating model to absorb faster discovery, make better decisions and avoid being pulled into reactive cycles dictated by external, non-contextualized narratives.

The noise problem

The real challenge in AI-driven vulnerability discovery is filtering out noise from signals. Security teams face overwhelming volumes of alerts, which can create the illusion of progress while increasing operational strain. This dynamic can lead to vendor-driven narratives that complicate risk management.

Noise has a real business cost. Every alert must be triaged. Every theoretical exposure creates questions from security, operations, audit, risk, compliance and executive stakeholders. Every unresolved item adds to backlog and creates pressure to act quickly, even when the appropriate action is not clear.

The scale of the problem is now institutional: in April 2026, NIST moved its National Vulnerability Database to a prioritized enrichment model after CVE submissions grew 263% between 2020 and 2025, conceding that not every disclosed vulnerability can even be analyzed centrally, let alone actioned. Over time, organizations can become trapped in a cycle where activity is mistaken for assurance.[1]

This dynamic also creates space for vendor-led narratives, and for headline claims that outrun validated facts. When reports emerged in June 2026 that an AI model had identified vulnerabilities in classified US government systems within hours, officials were quick to stress that identifying a weakness is not the same as exploiting it. [2] If a provider can surface a large number of potential risks and then imply that the safest response is to upgrade, migrate or return to the OEM, security becomes entangled with commercial dependency. The customer is no longer simply managing risk; it is responding to a risk narrative shaped by the party that benefits from the proposed remediation path.

Origina's view is that this model is unsustainable. Enterprise leaders need independent intelligence that helps them understand what is genuinely relevant. The question is not, "How many vulnerabilities can be found?" The more important questions are: "Which risks matter to this environment? Which are exploitable? Which require action? What practical mitigation is available? What evidence can support the decision?"

In the AI era, the competitive advantage is not more noise. It is trusted interpretation.

[1] NIST Updates NVD Operations to Address Record CVE Growth
[2] Anthropic's Mythos model found vulnerabilities in classified US government systems, official says

The new risk reality

AI-driven discovery is contributing to a new risk reality defined by speed, complexity and accountability.

Speed is increasing because vulnerability discovery, analysis and exploit development can move faster than traditional enterprise response processes. Many organizations still operate vulnerability management through scheduled review cycles, patch windows, change advisory processes and operational dependencies. These mechanisms are necessary for stability, but they can become a constraint when the external risk environment accelerates.

Complexity is increasing because the enterprise attack surface is no longer limited to core vendor code. Modern platforms depend on libraries, frameworks, integrations, configurations, plug-ins and open-source components. A weakness may exist outside the narrow boundary of what an OEM considers its responsibility, yet still create practical exposure for the customer. For mature or End-of-Support systems, the challenge can be sharper: the platform remains operationally important, but vendor investment, research and patch availability may be reduced or absent.

Accountability is increasing because boards and regulators expect more than technical activity. They expect evidence that risk is understood, proportionate action has been taken and decisions are defensible. A generic severity score does not satisfy that requirement on its own. Nor does a patch-only model if the patch is unavailable, operationally unsuitable, ineffective or unrelated to the actual exposure in the customer's environment. AI-driven discovery sharpens this gap: it can surface credible weaknesses long before any OEM fix exists, yet boards will still expect a defensible answer in the interim.

This exposes the limitation of conventional OEM-dependent models. OEM advisories and patches remain important, but they are not a complete security strategy. They are often reactive, product-centric and aligned to the vendor's lifecycle priorities. They may not account for customer-specific configurations, compensating controls, business-critical constraints or practical mitigation alternatives.

Origina's position is that enterprise security must move beyond patch dependency. Patching may be one valid control, but it is not the only control and could be limited in specific contexts. Risk can also be reduced through hardening, configuration changes, isolation, monitoring, compensating controls and operational process improvements. The right answer depends on context.

Regulatory and governance pressures

The changing risk landscape is being matched by rising regulatory and governance expectations. Across sectors and jurisdictions, organizations are being asked to demonstrate stronger control over technology risk, operational resilience, third-party dependency and software governance.

Regulatory frameworks such as NIS2, DORA and PCI DSS reflect a broader shift from compliance as a checklist to risk management and resilience as an operating discipline. The expectation is not simply that organizations react when vulnerabilities are disclosed. They must demonstrate understanding of critical dependencies, assess exposure, prioritize action, manage third-party risk and maintain evidence of due diligence.

This has important implications for executive teams. Cybersecurity decisions are no longer isolated within technical teams. They affect business continuity, financial exposure, customer trust, regulatory posture and strategic transformation plans. A decision to patch, defer, mitigate, isolate, monitor or upgrade must be explainable. It must be connected to risk appetite, business impact and operational reality.

In this environment, the quality of evidence becomes critical. Organizations need documentation that shows what was assessed, why a risk was prioritized, what action was taken, what mitigations were implemented and what residual risk remains. This is where pure alerting falls short. An alert may tell a team that something might be wrong. It does not by itself create an audit-ready risk decision.

Origina's approach is aligned to this governance reality. The objective is to help customers make better, more defensible decisions by translating vulnerability intelligence into contextual risk assessment and practical mitigation guidance. This matters particularly for organizations running mature enterprise platforms where operational stability and regulatory assurance, even beyond those specific cybersecurity regulations, must be balanced carefully.

Origina's perspective

Origina champions independence and control in vulnerability management. By decoupling software maintenance from vendor roadmaps, Origina empowers organizations to make informed decisions based on validated risk rather than external pressures.

For many enterprise customers, software is not simply a set of products. It is the foundation for critical business processes, regulated services, customer operations and operational continuity. These systems often remain valuable and critical to operations long after an OEM has shifted its commercial roadmap elsewhere. The fact that a vendor wants to move a customer to a new platform does not mean the existing platform is no longer fit for purpose. It means the customer needs an independent basis for making that decision.

Origina exists to give organizations that independence. By decoupling software maintenance and support from OEM-controlled roadmaps, Origina helps customers extend the life of mission-critical systems, reduce unnecessary cost and maintain control over the timing and shape of technology change.

In security, independence is equally important. If the organization defining the risk also benefits from the upgrade path, the customer must be able to test the narrative. The right security decision may be to patch. It may be to harden. It may be to isolate a service, disable a feature, apply a compensating control, increase monitoring or plan a controlled upgrade. The point is that the customer should make that decision based on validated risk, operational and business context, not fear.

Origina's philosophy is therefore practical rather than ideological. It does not reject OEM patches. It rejects the idea that patching is the only credible security response or that OEM timelines should dictate enterprise risk priorities. The focus is on what protects the customer, what preserves resilience and what enables the organization to act with confidence.

This approach is built around four core principles: independence from vendor-driven narratives; contextual assessment of real exposure; risk-based prioritization; and practical mitigation that supports operational continuity. Together, these principles create an alternative to reactive, patch-centric vulnerability management.

"He who is first in the field and awaits the coming of the enemy will be fresh for the fight; he who is second in the field and has to hasten to battle will arrive exhausted."

The Art of War by Sun Tzu
(Tactical Dispositions)

The OPTAS solution

OPTAS (Origina's Proactive Threat Assurance Service) is Origina's strategic response to the new AI-driven risk landscape. OPTAS is designed to help organizations identify risks earlier and take proactive action. It combines predictive intelligence with human validation, emphasizing a model that supports anticipation over reaction.

OPTAS is an evolution, not a departure. It builds on the foundations Origina customers already rely on: a contextual, risk-based, defense-in-depth security model and an established Vulnerability Advisory service. OPTAS extends those capabilities into predictive, pre-disclosure territory, so organizations that already trust Origina with the health of their software gain earlier sight of the risks forming around it.

At its core, OPTAS combines predictive vulnerability intelligence with expert human validation and actionable mitigation guidance. This combination is essential. Predictive analysis can identify signals of emerging risk, but human expertise is required to determine relevance, exploitability, environmental context and appropriate response. OPTAS is therefore not simply an AI tool. It is a human-validated security service designed to convert early indicators into defensible action.

The model can be understood through four executive-level stages: predict, validate, prioritize and mitigate.

Predict: OPTAS uses data-driven analysis to identify indicators of emerging vulnerability risk. The objective is to create earlier visibility into areas of concern before traditional disclosure cycles force customers into compressed response windows.

Validate: Origina security experts review findings, assess confidence and determine whether the risk is relevant in the context of the supported software and customer environment. This step is central because it separates theoretical risk from practical exposure.

Prioritize: With contextual application, prioritize constrained resources in an informed manner, focusing on the risks that present the most immediate threat.

Mitigate: Where action is required, Origina provides practical guidance designed to reduce exposure. This may include hardening guidance, compensating controls, configuration recommendations, monitoring considerations or other risk-based mitigation steps. The aim is to help customers act sooner and with greater clarity, rather than waiting for an OEM patch or responding to generic alerts.

Behind that guidance sits Origina's independent engineering capability. Where mitigation requires more than configuration change or system hardening, Origina's engineers can develop independent, code-level remediations: fixes built

from scratch that address the exposure without modifying or infringing the OEM's intellectual property. It is the difference between advising on risk and engineering it down, and it widens the range of practical responses available when no vendor patch exists.

OPTAS takes a threat-actor perspective, focusing on how a system would be exploited rather than identifying individual vulnerabilities. This approach also means OPTAS does not require direct access to environments or proprietary code to identify new, novel vulnerabilities, further reducing risk to the service user.

The strategic value of OPTAS lies in closing the gap between timing, information and action. Most vulnerability models begin when a CVE is published, or a patch becomes available.[3] By then, the customer is already reacting.

OPTAS is designed to give organizations more time, better context and stronger control; the opportunity to act on their own terms.

For C-level leaders, this creates five benefits. First, earlier visibility into emerging risk supports proactive planning. Second, human validation reduces noise and helps teams focus on what matters. Third, contextual mitigation supports operational continuity and avoids unnecessary disruption. Fourth, independent guidance reduces reliance on OEM-controlled narratives. Fifth, evidence-based outputs support audit readiness and regulatory accountability.

This is the practical difference between discovery and assurance. Discovery identifies possibilities. Assurance helps an organization decide and act. OPTAS is built for control, independence and assurance.

FIGURE 1 - OPTAS SUMMARY

[3] Under responsible disclosure, OEMs are given the first opportunity to validate a newly identified vulnerability and develop a patch. If they act, public disclosure can be delayed for several months to ensure a fix is available – based on the assumption that OEMs are best placed to secure their own products. During this time, however, customers may remain exposed to a known risk.

Conclusion and strategic takeaways

AI-driven vulnerability discovery is reshaping enterprise security, accelerating both the pace and volume of findings, while increasing pressure on already stretched teams.

Security is about foresight; identifying risk early and acting with confidence. By focusing on anticipation rather than reaction, organizations gain clarity and early visibility, enabling them to address vulnerabilities in their own context before they escalate. OPTAS advances this approach through predictive insight, expert validation and pre-disclosure mitigation design, helping teams stay ahead of emerging threats. The result is a stronger, more resilient security posture with less disruption, fewer urgent interventions and reduced reliance on patch cycles. While disclosure provides visibility, anticipation delivers the advantage, giving organizations true control over their risk.

Executives should avoid equating more discovery with more protection. In the absence of context, prioritization, validation and mitigation guidance, more discovery can increase uncertainty. It can lead to reactive decisions, unnecessary upgrades, operational disruption and misplaced investment.[4]

The required shift is from vulnerability volume to risk relevance. Enterprises need to know what is exploitable, what is material, what can be mitigated and what evidence supports the decision. They also need independence from commercial narratives that frame every security concern as a reason to return to the OEM roadmap.

Origina's perspective is that the future of enterprise vulnerability management must be intelligence-led, context-driven and independent by design.

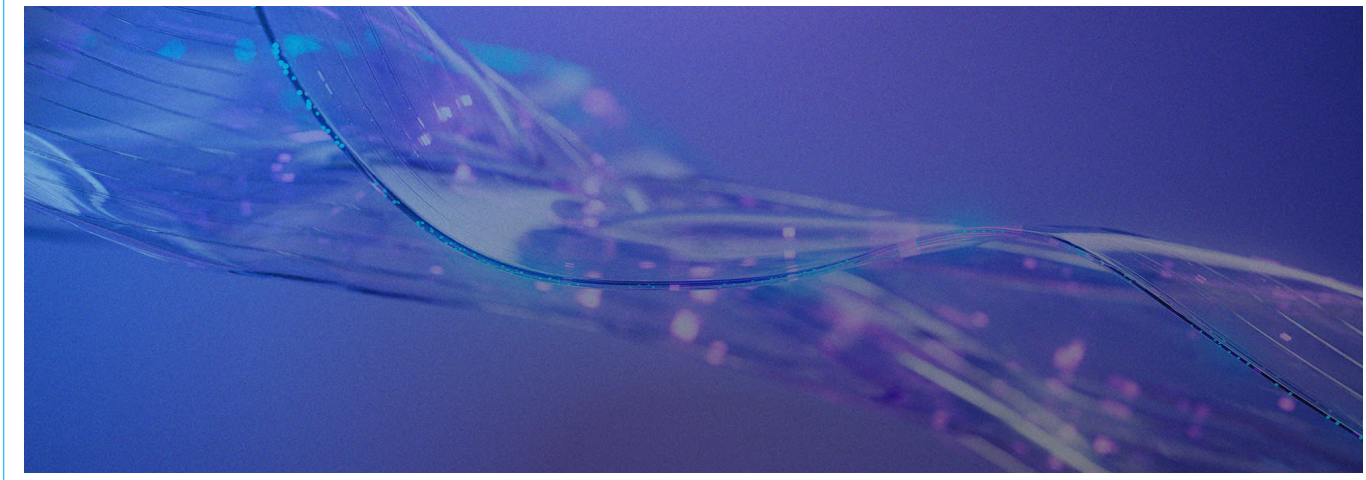

[4]Which is worse: 100 hours of downtime from a self-inflicted outage or 100 hours from a cyberattack? Both result in service disruption, reputational damage, and potential regulatory fines.

Security should give organizations more control over their technology estate, not less. It should support resilience, governance, business continuity and strategic growth, not force reactive change under pressure.

OPTAS provides a practical path towards that future. By combining predictive insight, human validation and actionable mitigation, it enables organizations to move from reactive patch dependency to proactive risk control.

For senior leaders, the strategic takeaways are clear: challenge noise; validate exposure; prioritize based on business risk; preserve control over technology decisions; and build defensible evidence for security governance. In the AI era, the enterprises that succeed will not be those that chase every alert. They will be those that apply foresight and act on the right intelligence with clarity and confidence.

FIGURE 2 - SPEED IS NOT THE PROBLEM

Executive actions: a working checklist

AI-driven vulnerability discovery is reshaping security faster than most operating models can keep pace with. The five actions below focus effort where it counts: relevant risks, independent assessment and evidence-based decisions. Each is expanded into a working checklist in the appendix.

1. Treat AI-driven discovery as a formal governance issue, not simply a technical input. AI-generated vulnerability intelligence needs clear ownership, defined validation standards and escalation paths, and cross-functional decision-making, so that findings are interpreted consistently, prioritized appropriately and managed in line with business risk, regulatory expectations and customer impact. In practice, that means legal, risk, security, engineering, product and customer-facing teams aligned on response expectations before a contentious finding arrives, not after; standing decision-making forums for high-impact or high-uncertainty findings; and defined thresholds for when an AI-generated finding requires executive visibility.

Fold AI-driven discovery into existing vulnerability management, risk management and incident response frameworks rather than building a parallel process. Above all, governance should distinguish between five states that are too often collapsed into one: theoretical exposure, suspected exploitability, validated exploit paths, active exploitation and business-impacting risk. Most of the noise this paper describes comes from treating the first state as though it were the last.

2. Challenge narratives that equate theoretical exposure with validated risk. AI tools can rapidly identify potential weaknesses, but potential exposure does not automatically equal exploitable or business-critical risk. The pressure to overreact is real, especially where media attention, vendor claims, or customer concerns inflate the perceived severity of an unvalidated finding, and resisting it requires evidence, not instinct.

Separate “could be vulnerable” from “is vulnerable in this environment”. Before accepting a finding as material, establish whether the affected component is present, reachable, enabled, configured in a vulnerable way and exposed to a realistic threat path. Challenge exploitability assumptions hardest where findings rest on code pattern matching, AI inference or generic CVE mapping, because these methods surface resemblance, not exposure. Communicate the distinctions to stakeholders in the same terms governance uses internally and avoid binary vulnerable-or-not framing where the honest answer depends on configuration, compensating controls, deployment context or operational usage.

3. Prioritize independent assessments for critical platforms. Automated findings, vendor guidance and theoretical exposure may not provide sufficient assurance for systems whose compromise would carry significant business, operational, regulatory or customer impact. This applies with particular force to legacy, highly customized or operationally sensitive estates, where patching may be complex, unavailable, delayed or disruptive, and where vendor statements tell you least about your actual position.

Independent validation matters most for platforms that are internet-facing, mission-critical, outside standard support, heavily customized, difficult to patch, subject to regulatory scrutiny or dependent on complex third-party components. A worthwhile assessment covers architecture, configuration, real-world exploitability, the effectiveness of existing controls and operational impact, and it produces actionable recommendations rather than generic vulnerability listings. Reassess critical platforms periodically, and specifically after major changes, new AI-discovered findings or shifts in threat activity, because an assessment describes a moment, not a state.

4. Strengthen evidence trails for vulnerability decisions. As vulnerability management becomes more complex, organizations need to be able to explain why decisions were made, and the explanation matters most when the decision is not to patch immediately: relying on compensating controls, accepting a risk or challenging the validity of a reported finding. Document the technical, operational and business rationale behind each material decision, record who reviewed and approved it, and capture the reasoning for prioritization, deferral, mitigation or acceptance.

The evidence set should cover asset ownership, affected versions, exposure status, exploitability assessment, business impact, compensating controls, testing performed, remediation options and residual risk, held in a form suitable for audit, customer assurance, regulatory review and internal governance. Link vulnerability decisions to enterprise risk registers where appropriate, review deferred or accepted risks at a periodicity aligned with the organization's risk appetite, and confirm on each review that the original assumptions still hold. A strong evidence trail is not bureaucracy; it is what disciplined, risk-based decision-making looks like from the outside.

5. Move towards proactive risk control. Shift from vulnerability identification to vulnerability control. Organizations should move beyond reactive patch-or-panic models: validate whether risk is real before escalating every finding as urgent, prioritize on business impact and deploy proportionate controls where direct remediation is not immediately possible or proportionate.

Controls that reduce likelihood or impact include network segmentation, access restriction, configuration hardening, monitoring, virtual patching, logging and detection, service isolation, privilege reduction and attack surface reduction. Use compensating controls when vendor patches are unavailable, delayed, risky or operationally impractical, and confirm through testing or independent validation that controls actually work, rather than assuming deployment equals protection. Track residual risk after controls are applied, and feed lessons learned back into baselines, monitoring and architecture. The destination is continuous risk reduction rather than point-in-time remediation: a model resilient to the pace of AI-driven discovery, not merely reactive to it.

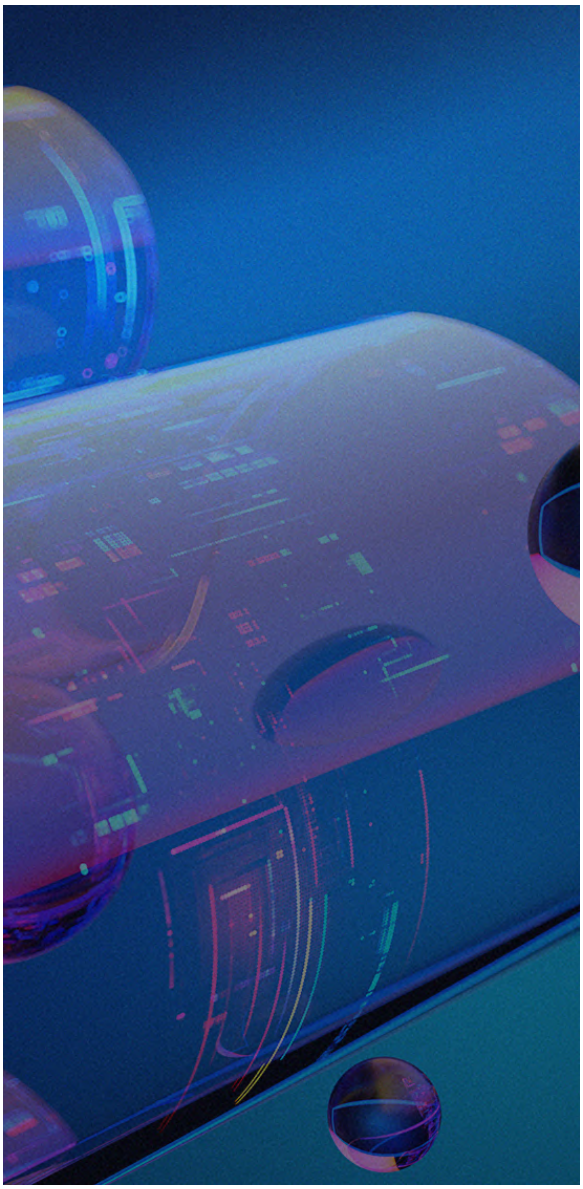

Appendix

Comparisons

Capability	Mythos/ Glasswing	OEM Patching	Generic MSSP	Origina OPTAS
Forecasts emerging risk	Yes	No	No	Yes
Covers EOS enterprise software	No	No	Partial	Yes
Validates risk in customer environment	No	No	No	Yes
Develops independent mitigations	No	Partial (generic)	No	Yes
Context-aware (customer usage)	No	No	No	Yes
Works without OEM involvement	Yes	No	Partial	Yes

TABLE 1 - COMPARATIVE OF AI TOOLS, TO PATCHING, MSSPS AND OPTAS

Mythos / Anthropic Notes – the reality

Mythos has faced restrictions from the U.S. government, and its outputs lack contextual consideration.

Anthropic disclosures:

- In May 2026, Anthropic reported 1,596 vulnerabilities.
- Only 88 had assigned CVEs at that time.
- Current tracking shows: 114 CVEs globally attributed to Anthropic; just 4 linked to the Glasswing project; none impacting Broadcom/ VMware technologies.

Broadcom joined the Glasswing project on April 7, 2026. Since then, three VMware Security Advisories (VMSAs) have been issued (as of late June 2026), none crediting Anthropic:

- **VMSA-2026-0003 (May 14, 2026):** Credit to Mathieu Farrell (@coiffeur0x90)
- **VMSA-2026-0004 (June 8, 2026):** Credit to Alexis Bernazzani (Visa Inc.)

Notes

[1] Under responsible disclosure, OEMs are given the first opportunity to validate a newly identified vulnerability and develop a patch. If they act, public disclosure can be delayed for several months to ensure a fix is available, based on the assumption that OEMs are best placed to secure their own products. During this time, however, customers may remain exposed to a known risk.

[2] Which is worse: 100 hours of downtime from a self-inflicted outage or 100 hours from a cyberattack? Both result in service disruption, reputational damage and potential regulatory fines.

Since 2012, Origina has helped the world's largest organisations take back control of their software.

We provide independent, enterprise-grade software maintenance that removes costly vendor-drive pressure - keeping critical systems stable, secure, and compliant, so organisations can operate by their own rules.

origina.

Ireland, Dublin
+353 1 524 0012

United States, Dallas
+1 888 206 4862

United Kingdom, London
+44 20 3318 3790

France, Paris
+33 1 59 03 03 60

Sydney, Australia
+61 7 3053 5831

Approved JB 072026