

B|ACKNOISE

OPERATIONAL SECURITY EFFECTIVENESS

DEEP PURPLE REPORT IV

V2.0 - Updated version August 2026

Published by ERIUM

contact@erium.fr

erium.fr/blacknoise-aev/

INTRODUCTION

For the fourth consecutive year, BlackNoise is publishing its annual report on operational security effectiveness.

This edition analyses more than 2,000 attack simulations carried out in 2025 and during the first half of 2026 against production infrastructures. It also incorporates regulatory developments and changes in the threat landscape through August. The campaigns from the first half of 2026 confirm the observed trends without altering them.

The major trend of 2026 can be summed up in one sentence:

The growing volume of published vulnerabilities, accelerated by AI-assisted offensive capabilities, makes comprehensive remediation unattainable.

Alongside “Patch First,” a complementary principle is emerging: “Response First.” Rather than waiting for widespread remediation, this approach ensures that the most exploitable vulnerabilities trigger a rapid, automated, and proportionate response. This response can be implemented within a few days, whereas a full patch cycle spans several months. But a response capability cannot be simply decreed - it must be verified: AEV is the tool that allows this chain to be tested under real-world conditions, measures the time between the execution of the attack and the blocking action, and ensures the principle remains effective over time as configurations and threats evolve.

This report is written for CISOs and CIOs who want to benchmark their trajectory against their peers and calibrate their defensive investments against field data.



Pierre TEXIER / CTO
pierre.texier@erium.fr




Arnaud LE MEN / CEO
arnaud.lemen@erium.fr

ABOUT BLACKNOISE AEV

Developed by Erium, BlackNoise is a fully agentless Adversarial Exposure Validation (AEV) platform, available in both SaaS and on-premise models.


Designed and built by security operations practitioners, it is used by public and private organisations to validate the operational effectiveness of their cyber investments under real-world conditions.

BlackNoise combines attack surface mapping (ASM), automated compromise testing (Automated Pentest) and an advanced engine for executing adversary tradecraft (BAS). This convergence addresses several use cases: validating detection and response capabilities, demonstrating compliance (DORA, NIS2, TIBER-EU) and training the cyber teams responsible for incident and crisis management.



EXECUTIVE SUMMARY — TL;DR

What more than 2,000 attack simulations run on production environments reveal.



Key figures for full-year 2025, BlackNoise AEV

Four findings

- **Detection is improving** (42/100, +11% vs 2024), but expansion into cloud, OT/ICS and the supply chain is shifting the blind spots.
- **Discovery, Exfiltration and Lateral Movement** remain the weakest areas, particularly against Living-off-the-Land techniques.
- **Response validation is taking hold**: organisations are timing MTTR and testing automated countermeasures before activation.
- The **frequency of simulations explains progress** far better than shifts between in-house, MSSP or hybrid SOC models.

The priority

Move from periodic control to demonstrated capability. AEV is becoming the continuous validation layer between open vulnerabilities, defensive tooling, expert pentesting, crisis management, compliance and team training.


The **five recommendations for 2026/2027** appear at the end of this report.

1. CONTEXT AND METHODOLOGY

The **Deep Purple Report** consolidates the technical results of simulations run by the BlackNoise AEV platform during full-year 2025 and the first half of 2026. All tests are mapped to the MITRE ATT&CK model and produce an effectiveness score based on MTTD (Mean Time To Detect), MTTA (Mean Time To Acknowledge) and MTTR (Mean Time To React). The multi-year series stop at 2025, the last complete year, in order to preserve year-on-year comparability; the H1 2026 campaigns feed the qualitative findings and trend analysis.

1.1 Scope of the organisations tested

The simulations run in 2025 and the first half of 2026 targeted the infrastructures of mid-caps and large groups across all sectors: public sector, finance and insurance, industry, telecommunications, aerospace, transport, retail and healthcare.



Breakdown of environments tested in 2025

Industrial OT/ICS environments doubled their share ($\times 2$ vs 2024), driven by the growing availability of dedicated detection tools. Cloud is accelerating sharply, with a first structured coverage of SaaS environments.


1.2 Methodological notes

- The results reflect the organisations and scopes tested by BlackNoise. They do not constitute a statistically representative sample of the market.
- MTTD, MTTA and MTTR indicators are calculated according to BlackNoise conventions and the availability of logs, alerts and operator feedback.
- The use cases describe observed or recommended practices. They constitute neither a guarantee of regulatory compliance nor a promise of cost savings or financial performance.

2. RESULTS: ASSESSING DEFENSIVE CAPABILITIES

2.1 Overall detection score

Detection capabilities improved for the fourth consecutive year. The pace is levelling off — a sign of maturity on established scopes, and of expansion into new ground.



Average detection score, 2022 to 2025

The score of 42/100 in 2025 (+11% vs 2024) reflects two realities. The most advanced organisations are approaching a plateau on the scopes they have mastered. Expansion into OT/ICS, cloud and the supply chain introduces new blind spots that weigh on the overall score. The gap between organisations that train regularly and those just starting out remains significant: regularity is the primary driver of progress, whatever the SOC model.

2.2 Tops and flops



Detection tops and flops by MITRE ATT&CK tactic, 2025

Persistence moves into first place among the tops: detection mechanisms have been refined against access-maintenance techniques. Discovery remains top of the flops for the third consecutive year, with the footprint of Living-off-the-Land (LotL) techniques still the main cause. Exfiltration climbs back into the flops: diversification towards alternative protocols and standard cloud services makes it harder to distinguish legitimate traffic from malicious exfiltration.


5 best-detected techniques

- T1003: OS Credential Dumping
- T1558: Steal or Forge Kerberos Tickets
- T1562: Impair Defenses
- T1548: Abuse Elevation Control Mechanism
- T1078: Valid Accounts, Default Accounts

5 least-detected techniques

- T1046: Network Service Discovery
- T1087: Account Discovery
- T1048: Exfiltration Over Alternative Protocol
- T1567: Exfiltration Over Web Service
- T1021: Remote Services, Lateral Movement

2.3 Focus on detection sources



Detection sources by tool, 2022–2025

EDR rebounds to 69%: systematically enriching its configurations from simulation results, particularly on PowerShell vectors and LotL techniques, confirms that endpoint visibility remains the cornerstone of modern detection. SIEM/XDR continues to gain ground at 49%: pairing EDR with SIEM/XDR is becoming the standard among mature organisations. NDR advances to 21%, driven by OT/ICS monitoring and lateral movement detection.

Recommendations

- Enrich EDR configuration from simulation results, prioritising PowerShell and LotL techniques.
- Consolidate the EDR + SIEM/XDR pairing as the backbone of the detection stack.
- Deploy NDR on OT/ICS and cloud scopes where EDR cannot be the only sensor.


KEY TAKEAWAY

Performance comes less from any single tool than from the way signals are correlated and responses chained together. At comparable levels of tooling, organisations that train regularly progress further than those concentrating their effort on a single annual test.

3. AEV IS SCALING ACROSS ORGANISATIONS

3.1 Observed use cases

The volume of simulations accumulated by BlackNoise this year makes it possible to broaden the analysis: in addition to the simulation results, the analysis now includes how organizations use them. The AEV not only measures exposure but also drives its reduction.



Validating detection capabilities — 90%

MTTD and MTTA have become steering KPIs presented to executive committees and written into MSSP service contracts. Simulation is no longer a one-off audit: it is a continuous barometer.

Validating response (MTTR) — 60%, rising

Organisations are no longer content to detect; they are timing their ability to react. Before any large-scale rollout, the countermeasures triggered by EDR — process kill, file deletion, endpoint isolation — are validated in real environments to confirm their effectiveness and the absence of side effects.

NIS2 / DORA compliance — 60%

AEV produces velocity indicators, cyber scores and TLPT-style deliverables. The degree of vulnerability exploitation, derived from Automated Pentest, provides a picture of the flaws that are genuinely exploitable.

Crisis exercise preparation — 50% and above

More than one organisation in two uses technical simulations to feed its crisis management exercises: realistic injects, testing of alert chains, verification that the right contacts can actually be reached.

Pentest targeting, supply chain and rule forging — 20%

Organisations draw on simulation data to identify genuinely exploitable attack paths before launching a pentest, extend their scope to partners and suppliers, and build rules calibrated on their own blind spots rather than applying generic ones. This is the direct application of the Response First doctrine.

3.2 Training teams: Human Readiness

Simulations offer a unique training framework: analysts face real scenarios, in their production environment, with their usual tools.

The logs, alerts and decisions from a campaign make it possible to isolate the sequence that went wrong (a signal ignored, an incorrect classification, a late escalation, an incomplete action or an unsuitable playbook, and so on).



As in Formula 1, the simulator makes it possible to rework the sector where performance dropped.


That sequence is turned into a short training or awareness path: technical replay, quiz, role-play, then a new simulation. Three gaps recur systematically: theoretical playbooks that diverge from their actual application, insufficient triage velocity, and single points of failure in the alert chain — an invalid on-call number, an email landing in spam, a notification drowned in false positives.

KEY TAKEAWAY

Embed AEV simulations in a structured training programme, at a minimum quarterly cadence on critical scenarios, with a systematic review of playbook-versus-reality gaps after every campaign.

4. A TWO-SPEED DEFENCE

On sensitive or transactional environments, a few seconds are enough to cause impact. Serious attacks also demand humans trained to take back control.



AEV puts both machine time and judgement time to the test.


On sensitive, industrial or high-impact transactional scopes, a coordinated attack can produce a business effect before any human triage. XDR platforms are beginning to contain these chains automatically from multi-source correlated signals: they reduce lateral movement and buy analyst's time.

The right boundary is not "human or machine". It separates repeatable, high-confidence actions from decisions that are ambiguous, high-impact or dependent on business context.

KEY TAKEAWAY

Enable automated blocking first on high-confidence, high-impact scenarios, after validation under realistic conditions, then replay after every significant change.

5. PENTEST vs. AEV: DEPTH AND CONTINUITY



AEV removes the need to wait for the next pentest to verify a fix or a rule.

What pentesting does better

Human creativity, exploration of complex chains with a human touch, understanding of context, targeted demonstration of impact, and expert insight into design flaws.

What AEV makes possible

Frequency, replay coverage, MTDD/MTTR measurement, non-regression control validation, API integration and large-scale data processing, remediation tracking and targeting of future expert tests.

KEY TAKEAWAY

Automate repetitive exploration and remediation validation; concentrate pentester time on the complex scenarios that warrant a human approach, on critical assets and on genuinely exploitable paths.

6. IMPLEMENTATION OBSERVED IN THE FIELD

Scenarios gain credibility when they produce real signals, real alerts and real decisions.

Enriching and adding credibility to cyber crisis exercises

Situation — The organisation wants to test alerting, triage, escalation and decision-making without being confined to a tabletop scenario.

AEV contribution — An attack chain is executed on a controlled scope; teams receive real signals and the crisis unit works from credible technical material.

Outcomes — Timeline, decisions, playbook-versus-reality gaps, MTTD/MTTR, remediation plan.

Preparing and running a TLPT programme

Situation — Formal testing is infrequent, demanding and heavily governed. Teams need to progress between two exercises, and to demonstrate continuity of evidence in the meantime.

AEV contribution — AEV prepares scenarios, tests controls, facilitates and automates purple teaming and replays remediations, without substituting for the prescribed red team test.

For designated financial entities, DORA requires a TLPT at least every three years; the TIBER-EU 2025 framework sets out the steps. NIS2 carries the same expectation of evidence and continuity: in France, transposition remained incomplete as of summer 2026, which does not diminish the need to prepare.

Outcomes — Evidence pack, MITRE coverage, exercise report, replay evidence, action tracking.

Validating a service provider's SLAs

Situation — Contractual SLAs describe response times, but the handling chain itself is rarely tested.

AEV contribution — Repeated campaigns across several time windows, with tracking of acknowledgements, escalations and remediations.

Outcomes — MTTD/MTTA/MTTR, proof of service, procedural gaps, improvement plan.

KEY TAKEAWAY

One campaign, several audiences. The SOC and CIRT work the logs; vulnerability management teams qualify exploitability; GRC collects the evidence; the provider is assessed on results; crisis management gains a credible scenario. The value comes from reusing the same data.

7. TRENDS

7.1 From BAS to AEV: a paradigm shift

For three years, BlackNoise reports documented the evolution of BAS as a measurement tool.

In 2025, this technology category completed its transformation: BAS has been absorbed into Adversarial Exposure Validation. BAS measured; AEV validates. It confirms, continuously and automatically, the real feasibility of an attack, its potential impact, and the organisation's ability to detect it, respond to it and recover from it. The convergence of these components into a unified platform marks the shift from a testing logic to continuous exposure management.

7.2 Offensive targeting: Automated Pentest and ASM

By delegating the systematic exploration of attack paths to automation, pentesters concentrate their expertise on complex scenarios: shorter, more targeted campaigns with a markedly better signal-to-noise ratio. Automated Pentest confirms actual compromise, not merely simulation, and pinpoints exploitable vulnerabilities.

Upstream, the Hacker View Map is becoming the starting point for campaigns: before simulating an attack, the question is where an attacker would come in. ASM maps internet-exposed assets — domains, IPs, services, certificates, technologies — and feeds scenarios directly. The most realistic paths start from real exposures, not hypothetical perimeters.

7.3 Patch First and Response First

The continued growth in the number of vulnerabilities, now accelerated by offensive AI, puts exhaustive remediation out of reach. Patch First is not going away: it is being complemented by Detect Rule First.

The answer is not to abandon patching, but to supplement it. “Response First” involves prioritizing the ability to respond to the most exploitable vulnerabilities within the organization's context: detecting the exploit, then blocking the process, isolating the host, or terminating the session before the attacker achieves their objective. The reasoning is pragmatic: for a critical vulnerability, an automated response provides protection within seconds, whereas deploying the patch across the entire infrastructure can take several hours, days, or even weeks.

The defensive posture is based on the assumption that some vulnerabilities will remain unpatched, and compensates for this with a robust response capability—provided, of course, that this capability is actually in place and not merely configured.

This is precisely what AEV puts into practice: reenacting attack paths that exploit high-priority vulnerabilities, verifying that the detection-response chain is actually triggered, measuring the time between execution and blocking, and correcting anything that fails to trigger—before an attacker takes the same path. The 20% of organizations that already use AEV to validate their rules and automated response mechanisms are the pioneers of a practice that is set to become widespread.

7.4 The R in xDR: automated blocking is accelerating

Validation through simulation accelerates the activation of the automated blocking functions carried by modern defensive tools — the “R” in EDR, MDR, NDR and XDR. As long as these capabilities remain unvalidated, organisations keep them in detection-only mode, wary of false positives and side effects. AEV provides evidence that blocking is reliable on the scenarios tested and removes the psychological barrier to deployment. Process kill, endpoint isolation and network flow blocking trigger in milliseconds, where human MTTR is counted in minutes.

7.5 OT/ICS: widening coverage

Industrial environments accounted for 12% of campaigns in 2025, double the 2024 figure. This growth is driven by the increasing availability of dedicated tools — industrial NDR, passive agents, honeypots — and by executive awareness, accelerated by NIS2. What BlackNoise observes is consistent with the early years of BAS on enterprise IT: low detection rates, numerous blind spots, and SOC teams unfamiliar with the TTPs specific to industrial protocols (Modbus, DNP3, OPC-UA).

7.6 NIS2, DORA and TIBER-EU compliance: AEV as evidence

Regulatory compliance accounted for 60% of AEV use cases in 2025. What regulators ask for, and what AEV can produce, is objective and continuous evidence: technical data demonstrating that the mechanisms work, at a documented frequency, with measurable velocity indicators. Cyber scores by scope, MTTD/MTTR trends, MITRE tactic coverage and exercise reports are now fully fledged components of compliance files.

7.7 AI in detection and remediation

AI is gradually taking its place in detection and response chains. Its value lies not in replacing analysts but in automating repetitive tasks: alert triage, log summarisation, query generation, script analysis. Human time shifts towards the decisions where judgement, context and experience count.


These benefits must be weighed against significant limitations. Models excel at the already-seen, far less at the emergent. They can hallucinate without any explicit error signal. They lack organisation-specific context. To this are added the risks inherent in AI itself: data leakage, prompt injection, atrophy of analyst skills. And attackers are adopting AI too: mass phishing, deepfakes, polymorphic code.

KEY TAKEAWAY

The more tools detect, correlate and block automatically, the more essential it becomes to test what they actually see, ignore or interrupt. This is the structural role of AEV in an automated defensive environment.

8. INVESTING WITHIN A FLAT BUDGET

The most realistic path often consists in reallocating time and spend already committed, rather than adding another layer of control.



Reallocation principle — to be adapted to the maturity level and operating model.

Four ROI levers within a flat budget

- **Automated evidence** — eliminate the manual collection of compliance material: results are replayable and directly usable by GRC.
- **Targeted pentesting** — focus the expert budget on high-value paths, instead of re-exploring the same scopes with every campaign.
- **Training on real gaps** — train analysts on their observed failures, rather than funding generic content and disconnected lab environments.
- **Validated automated response** — reduce the human time spent on repetitive containment actions, once automated blocking has been proven.

KEY TAKEAWAY

AEV makes visible the duplication, the low-yield controls, and the places where an additional euro genuinely produces progress.

9. FIVE RECOMMENDATIONS FOR 2026/27

1. Qualify open vulnerabilities through attack paths

Link unpatched vulnerabilities to critical assets, actual exploitability, and compensating controls by enhancing detection capabilities.

2. Validate automated response

Test isolation, blocking and deactivation before broad activation, then after every major change.

3. Train for taking back control

Select a few high-impact scenarios and replay the full chain (SOC, business lines, executives, crisis unit) to ensure decision processes are effective.

4. Turn every gap into training

Convert weak sequences into technical replay, targeted learning paths, quizzes or role-play.

5. Reuse the same evidence

Feed the SOC, CIRT, vulnerability management, GRC, the executive committee, compliance, providers and remediation plans from the same set of results.

IN CONCLUSION

This fourth edition confirms that operational security validation has come of age. Regularity outweighs exhaustiveness: the organisations that progress most are not those that test most broadly, but those that test most often. Blind spots move without disappearing, and remain architectural rather than technical.

AEV becomes unavoidable wherever constrained budgets must still demonstrate operational effectiveness, compliance, resilience and skills retention. Targeted investment, demonstrable resilience.

SOURCES AND REFERENCES

Unless otherwise stated, the figures are drawn from BlackNoise campaigns run in 2025 and the first half of 2026. The year-on-year series stop at 2025, the last complete year. The regulatory context and threat landscape are current as of summer 2026.

- 1 Gartner, Market Guide for Adversarial Exposure Validation, 11 March 2025.
- 2 Erium, Awards & Accolades, accessed August 2026: BlackNoise named a Representative Vendor in Gartner's Adversarial Exposure Validation coverage.
- 3 Google Cloud, Cloud Threat Horizons Report H1 2026: the window between disclosure of a vulnerability and active exploitation narrowed from weeks to a few days in the second half of 2025.
- 4 Google Cloud / Mandiant, AI Risk and Resilience, 9 March 2026: operationalisation of adversarial AI, adaptive tooling and agents operating with minimal human supervision.
- 5 Microsoft Learn, Automatic attack disruption in Microsoft Defender, updated 23 June 2026: real-time containment from correlated XDR signals.
- 6 European Central Bank, TIBER-EU Framework 2025: alignment with DORA, purple teaming guidance, reporting and remediation tracking.
- 7 Regulation (EU) 2022/2554 (DORA), Article 26, and Delegated Regulation (EU) 2025/1190: TLPT at least every three years for designated financial entities.
- 8 Gartner, Top Cybersecurity Trends CISOs Must Act on in 2026, 17 March 2026: the adoption of AI agents in the SOC calls for a new analyst training regime and human-in-the-loop controls.
- 9 ENISA, NIS Investments 2025, 8 December 2025: compliance is the leading investment driver (70%); 28% of organisations take more than three months to remediate critical vulnerabilities.
- 10 European Commission, 8 July 2026: France referred to the Court of Justice of the European Union for incomplete transposition of the NIS2 Directive.
- 11 Erium, CYBAS / France 2030 project announcement: interoperability between BlackNoise, HarfangLab, Gatewatcher, Snowpack and BreachHunt.

Gartner does not endorse any vendor, product or service depicted in its publications. Its research reflects the opinions of its research organisation and should not be construed as statements of fact.

B|ACKNOISE

DEEP PURPLE REPORT IV

Updated version august 2026

Rennes (BZH)

Paris (FR)

Bologna (IT)

Doha (QA)

Publication directors: Pierre Texier (CTO, Erium); Arnaud Le Men (CEO, Erium)

erium.fr/blacknoise-aev | contact@blacknoise.co

Published by the Erium Group