

Beyond the AI SOC label

Where ten platforms differ in data, context, reporting and action


EXECUTIVE SUMMARY

AI-driven investigation is the shared proposition. The more consequential buying decisions concern the evidence a platform can retain, the environment it understands, the reports it delivers and the actions it can execute.

Under this paper's stated breadth rubric, **Fluency and Exaforce share the highest tier for SIEM foundation, raw telemetry history and configuration intelligence**. Fluency's supported distinction is its customer assessment suite. Exaforce, 7AI and Torq describe broader response execution. These are differences in documented scope—not a laboratory ranking of security effectiveness.

Overall capability profile

Fluency Exaforce



Shown: Fluency, Exaforce.

Reading the chart: 1 = narrow or externally dependent; 2 = substantial; 3 = broad within the dimension-specific rubric. Translucent fills show overlap. Unknown or preview-only dimensions remain gaps; incomplete profiles are not filled. Area is not a product-quality score.

What defines the AI SOC proposition?



Evidence → context → a reviewable decision · Concept illustration

An AI SOC applies AI-supported workflows to security investigation and operations. For this paper, the shared proposition is the set of functions below. This is a working comparison boundary, not a formal certification or a claim that every vendor implements each function equally well.

Common capability	What it is intended to accomplish
Automated triage and investigation	Turn alerts into an investigation plan, collect evidence, evaluate explanations and produce a disposition.
Cross-source evidence gathering	Search connected security systems and bring related activity into a coherent investigation.
Environmental context	Interpret activity using relevant users, assets, permissions, prior behavior and organizational procedures.
Explainable findings	Show supporting evidence, reasoning and unresolved questions so an analyst can review the result.
Response support and human control	Recommend, route or execute the next steps within the customer's operating model. Autonomous containment is not required for every product.
Reduced analyst workload	Reduce repeated evidence gathering, investigation preparation and documentation. The amount saved requires measurement.

The AI SOC label is the baseline—not the differentiator. Scoring every platform for claiming AI investigation would contribute little to the decision. This paper instead examines the data and operational capabilities around that shared proposition.

Shared positioning does not prove equal results

“Every alert investigated,” “autonomous” and “reduced MTTR” can describe different scopes. A verdict, case closure, human escalation, containment and verified recovery are separate outcomes. Comparisons require the same starting point, evidence requirements, alert mix and stopping condition.

Likewise, a preserved case summary is not a full raw-event archive; an organizational memory is not necessarily a configuration inventory; and a generated incident narrative is not a complete compliance assessment. These distinctions determine the six dimensions that follow.

Different operating models can be appropriate

An MSSP may want findings and recommendations that its own team implements. An enterprise may prioritize cross-tool containment. A customer with an established SIEM may prefer an investigation overlay, while another may need a replacement data foundation. Broader integration within one product is relevant to this comparison, but does not automatically make it the better choice for every buyer.

A published-material assessment, with explicit evidence boundaries

The comparison primarily uses vendor-published product pages, technical descriptions and workflow documentation available at the research cutoff. Fluency additionally supplied owner confirmations, internal implementation documentation and one customer investigation artifact. **This is an informed comparison prepared for Fluency—not an independent product certification or a controlled benchmark.**

How a capability earns credit

- A specific first-party description can support a provisional scope score. The same rule applies to all ten companies.
- Reviewed artifacts and implementation documents are identified separately from a vendor’s statement.
- Scores express the scope supported by the material reviewed, not a proven ceiling on the product.
- Unknown means evidence was insufficient. Preview-only features receive no generally available product credit.

Evidence labels

V	Vendor-published description
O	Owner-confirmed capability
D	Artifact or implementation/workflow documentation reviewed
S	Managed-service or underlying-platform dependency
?	Unresolved scope
P	Preview; excluded from scoring

What this paper does not establish

No common live workload was run across the products. The study therefore does not rank detection recall, false-negative rates, investigation correctness, search completeness, containment success, MTTR, scalability or total cost. Vendor percentages, case studies and speed claims are not converted into comparative performance scores.

Public documentation is uneven. A missing description is not proof that a capability is absent. Fluency’s additional access improves specificity but creates an evidence asymmetry; not every Fluency capability or report was executed for this review. No independent field-by-field configuration comparison was completed.

Retention duration is separate from included volume, ingestion cost, search cost and licensing. Fluency’s 12-month default is owner-confirmed; the included volume was not established here. Exaforce’s published long-term retention description is not treated as a verified commercial allowance. Unknown pricing or capacity is never labeled free or unlimited.

The six ordinal scales answer different questions. A “3” means the broadest defined scope on that axis, not three times the effectiveness of “1.” No weighted total or radar-area winner is asserted. Polygon area changes with axis order and can conceal gaps.

The shortlist follows the prior ten-company scope, including Fluency. It is not exhaustive or independently ranked by market share or funding. Web pages and offerings may change after the cutoff. Customer-identifying information from the Fluency artifact is excluded.

Different strengths, different dependencies

The matrix captures supported breadth using the definitions in the following sections. These are provisional analytical judgments, not vendor-certified measurements.

Company	SIEM	Raw history	Configuration	Reporting	Hygiene / exposure	Response
Fluency	3	3	3	3	2	1
Torq	1	1	3	2	2	3
TENEX	1	1	?	1	1	2
Exaforce	3	3	3	2	2	3
7AI	3	2	2	1	2	3
Dropzone	1	1	1	1	1	2
Intezer	1	1	2	1	?	2
Qevlar	1	1	1	1	P	?
Prophet	1	1	1	1	?	2
Legion	1	1	1	1	?	2

? = unresolved; P = preview, unscored. The component sections retain the evidence label and explanation for every cell.

Data foundation and history

Fluency and Exaforce combine their own data foundation with approximately a year or more of historical telemetry. 7AI documents both federated access and optional native hot storage; its long-term included storage scope remains unclear. The other reviewed offerings principally operate over existing security infrastructure. That dependency matters when underlying telemetry expires.

Configuration context is broader than a graph label

Fluency's property tables and event data, Exaforce's configuration/identity model, and Torq's described historical context graph all support substantial environmental context. Their exact field coverage, refresh behavior and point-in-time reconstruction were not tested side by side. No superior architecture is inferred merely from using the word "graph."


Customer deliverables and response ownership separate the offers

Fluency's supported suite spans incident, hygiene, compliance and vulnerability/gap assessment outputs. Exaforce and Torq document additional custom or workflow-generated reports, while several others substantiate investigation reporting more clearly than a broader assessment suite. On response, Fluency deliberately provides findings and recommendations for MSSPs to implement; the planned Nobody Security response model is excluded from the current-product comparison.

A fair selection asks which gap the customer needs to close: a data platform, an investigation layer, an assessment/reporting service, response execution—or a combination. The same radar can support different buying decisions.

SIEM foundation

Who owns the detection and data foundation? An integrated SIEM can collect, normalize, search and detect on the data itself. An investigation overlay can be valuable without replacing that foundation. The distinction affects architecture, dependencies and which system must remain available—not the inherent quality of AI reasoning.




Scoring definition

- 1: works with an external SIEM/detection stack.
- 2: own search/detection layer, with collection or storage not established.
- 3: own collection, normalization, search, detection and storage capability, including optional native storage. Federation itself is not a penalty. This evaluates the data foundation, not investigative quality.

Company	Score / evidence	Assessment and boundary
Fluency	3 O+D	Owner confirms full SIEM with real-time ingestion and behavioral analytics; supplied investigation demonstrates raw-data searches. [Fluency evidence]
Torq	1 V	Connected SIEM and source integrations underpin orchestration. Alert ingestion and case management do not establish a full raw-log SIEM replacement. [2]
TENEX	1 V/S	Managed Google SecOps/Sentinel deployments; underlying SIEM is third-party. A managed service is not an independently built SIEM. [4]
Exaforce	3 V	Documents its own ingestion, normalized data platform, storage and detection; this is a vendor claim, not a hands-on validation. [6]
7AI	3 V	Documents collection, streaming detection, OCSF query normalization, federated queries and optional 7AI hot storage. Native storage remains an optional deployment choice. [10]
Dropzone	1 V	Investigates through existing security tools; no independent raw-log SIEM replacement established. [14]
Intezer	1 V	Investigates alerts and collects forensic artifacts alongside the existing stack; forensic storage is not a complete SIEM. [16]
Qevlar	1 V	Product workflow begins with an alert from the customer's SIEM or EDR. [17]
Prophet	1 V	Investigations, hunting and detection engineering over connected security tools; own full SIEM not established. [20]
Legion	1 V	Executes analyst workflows through existing tools and browsers; own full SIEM not established. [21]

Raw telemetry history

Can an investigation still reach the evidence later? Delayed discovery requires preserved raw events. This axis distinguishes an owned historical archive from reliance on a connected source's retention. Saved case narratives and contextual memory may help investigations but do not preserve every underlying event.




Scoring definition

- 1: investigation history depends on external raw-event stores; saved cases/context do not count as a full raw archive.
 - 2: native raw storage documented but long-term duration not established.
 - 3: searchable raw history of approximately a year or more explicitly supported.
- No score implies free/unlimited volume, equivalent search latency or price.

Company	Score / evidence	Assessment and boundary
Fluency	3 O+D	Owner confirms 12 months of historical data by default and configurable behavioral windows including 90 days. The supplied investigation selected a 30-day lookback; included volume not specified. [Fluency evidence]
Torq	1 V	Raw-event archive depends on connected sources. Torq separately documents historical context and case memory; this score does not mean it can only investigate today's incidents. [1]
TENEX	1 V/S	Historical data lives in the underlying SIEM service; the TENEX pages reviewed do not establish an independent native archive allowance. [4]
Exaforce	3 V	Claims 90 days of correlated in-memory context and over a year of queryable raw data. Default included volume, commercial inclusion and cold-query performance unverified. [6]
7AI	2 V	Optional native hot storage is documented. Included duration/volume and long-term native raw archive remain unspecified. [10]
Dropzone	1 V	Federated searches depend on retained source telemetry. Context memory and saved investigations are different from preserving all raw events. [15]
Intezer	1 V	Stored forensic artifacts and investigation results do not establish a full source-telemetry archive; historical reach depends on connected data. [16]
Qevlar	1 V	Source queries and past investigations do not establish an independent full raw-event archive. [17]
Prophet	1 V	Historical hunting and backtesting depend on connected data; independent raw archive not established. [20]
Legion	1 V	Retains workflow knowledge and uses existing tools; no independent full raw-event archive established. [22]

Configuration intelligence

Can the system explain activity in the context of the environment? Identity, application, device, permission and posture information can change the meaning of the same event. This axis measures the breadth of maintained context used with activity. Tables and graphs are implementation choices; field coverage, freshness and reconstructability determine practical value.




Scoring definition

- 1: environment context/enrichment supported.
- 2: concrete configuration, posture or forensic state used for a bounded domain.
- 3: maintained multi-domain resource/configuration context linked to activity. A graph is not inherently superior to tables. Exact historical state reconstruction requires a separate test for every vendor.

Company	Score / evidence	Assessment and boundary
Fluency	3 O+D	Property tables for identities, users, applications, Defender machines/vulnerabilities/AV health/recommendations/Secure Score, plus audit and firewall changes. August resource validation supports concrete collection. Exact cross-vendor field coverage has not been tested. [Fluency evidence]
Torq	3 V	June 2026 technical description claims maintained entity/permission relationships, streamed or reconciled changes, and bitemporal historical context. Coverage is broader than device workflows alone; connector completeness and actual replay not tested. [1]
TENEX	? ?	The reviewed service pages do not specify property-table collection or independent configuration-history coverage; unknown rather than absent. [5]
Exaforce	3 V	Claims identity/config snapshots and correlated relationships. No field-level evidence shows better collection than Fluency; historical reconstruction not independently tested. [7]
7AI	2 V	Ongoing environment posture checks and control mapping are documented. Systematic configuration inventory and versioned state coverage remain unclear. [11]
Dropzone	1 V	Client context memory and investigation enrichment are documented; not equivalent to demonstrated broad property-table collection. [14]
Intezer	2 V	Concrete files, processes, registry changes and memory collection supports forensic state. It does not establish a broad configuration-history inventory. [16]
Qevlar	1 V	Live environment graph/context claims support enrichment. Public pages reviewed do not enumerate a maintained configuration dataset. [17]
Prophet	1 V	Environment context and policies inform investigations; property-table coverage and version history not established. [20]
Legion	1 V	Learns procedures, analyst activity and organizational context. Browser access alone is not proof of systematic configuration collection. [21]

Customer reporting

What can the customer actually receive? MSSPs and security leaders need deliverables beyond incident narratives: hygiene, control posture, vulnerability findings and gaps. This axis credits identifiable report workflows and assessment families, rather than treating a certification, dashboard or general-purpose prompt as a completed report suite.




Scoring definition

- 1: incident/hunt/operational reports supported.
- 2: explicit additional posture/compliance report workflow or custom evidence-report generation.
- 3: identifiable multi-family customer assessment suite spanning incident, hygiene, compliance and vulnerability/gap outputs. Dashboards, certifications, generic prompts and posture findings alone do not qualify as report families. Scores describe supported scope, not a proven maximum.

Company	Score / evidence	Assessment and boundary
Fluency	3 O+D	Customer closure artifact reviewed; owner-confirmed assessment suite corroborated by CISO reporting documentation, executable hygiene reports and compliance renderer. Not every report family was run on a live tenant for this comparison. [Fluency evidence]
Torq	2 V/D	Actual Intune hygiene reporting template plus customizable compliance/audit reporting. Full framework coverage and quality of customer deliverables not validated. [3]
TENEX	1 V/S	Service explicitly includes dashboards and post-incident reports. A standardized multi-family customer assessment suite is not established. [5]
Exaforce	2 V	Exabot Respond explicitly generates custom compliance, executive and incident reports. This supports more than incident summaries, but not a verified predefined assessment suite. [8]
7AI	1 V	Investigation outputs are supported. ISO/SOC 2 posture scores and ranked gaps are not, by themselves, proof of exported customer assessment reports. [13]
Dropzone	1 V	Investigation and hunt reports supported, including exposure findings. This does not establish separate hygiene/compliance assessment families. [15]
Intezer	1 V	Investigation records and SOC operational reporting supported; no broad customer assessment suite established. [16]
Qevlar	1 V	Comprehensive investigation reports documented; broader recurring customer assessment suite not established. [17]
Prophet	1 V	Investigation outputs and detection-coverage analysis supported. An ATT&CK matrix is not proof of a compliance/hygiene customer report suite. [20]
Legion	1 V	Workflow and incident reporting supported. Generic board-report examples do not establish a standardized customer assessment suite. [21]

Hygiene and exposure

Does the platform support proactive posture work? An investigation may reveal a vulnerable host or a weak setting. Dedicated posture workflows go further by repeatedly identifying and prioritizing exposure. A complete remediation lifecycle also needs ownership, tracking and verification—none of which should be inferred from a findings list alone.




Scoring definition

- 1: exposure analysis inside investigations/hunts or a stated assessment service.
- 2: dedicated available posture/hygiene/vulnerability workflows.
- 3: a specified end-to-end workflow including discovery, prioritization, remediation tracking and verified closure. No company qualifies for 3 on reviewed evidence. Preview-only scope stays unscored.

Company	Score / evidence	Assessment and boundary
Fluency	2 O+D	Dedicated identity/endpoint hygiene and vulnerability posture capabilities. Collection and reporting do not alone prove closed-loop remediation. [Fluency evidence]
Torq	2 V/D	Concrete device-compliance and misconfiguration workflows. End-to-end vulnerability lifecycle not established by those examples. [3]
TENEX	1 V/S	Vulnerability assessment and mitigation services advertised; native software lifecycle and deliverable detail remain unclear. [5]
Exaforce	2 V	Dedicated SaaS hygiene checks and exposure findings. Complete vulnerability remediation lifecycle not established. [9]
7AI	2 V	Ongoing posture sweeps, framework mappings and prioritized fixes. Pending attestation/GRC controls are explicitly distinguished. [11]
Dropzone	1 V	Threat hunts can surface vulnerabilities, misconfigurations and visibility gaps; a dedicated lifecycle is not established. [15]
Intezer	? ?	Forensic depth does not establish a dedicated hygiene or vulnerability-management workflow. Insufficient evidence to rate this category. [16]
Qevlar	P P	SOC/vulnerability agents are explicitly labeled preview on the current page. Do not credit as generally available without confirmation. [18]
Prophet	? ?	Detection engineering and coverage gaps are not vulnerability management. Dedicated hygiene/exposure workflow not established. [20]
Legion	? ?	Generic exposure questions and process automation do not establish a dedicated hygiene/vulnerability workflow. [22]

Response automation

Who carries out the response, and under what control? Recommendations, configurable actions and governed multi-domain workflows represent different operating models. The relevant distinction is documented execution scope. A high score does not establish safer actions, faster recovery or a better fit for an MSSP that intentionally keeps implementation with its own staff.



Scoring definition

- 1: recommendations/case routing only are explicitly the offered scope.
- 2: configurable execution or containment documented, with limited cross-tool workflow detail.
- 3: multi-domain response actions, workflow composition, approval controls and audit documented. A service promise alone does not earn 3. A single escalated case does not cap a platform at 1.

Company	Score / evidence	Assessment and boundary
Fluency	1 O+D	Owner confirms current Fluency scope is findings and recommendations for MSSPs to implement. Supplied case corroborates escalation. Planned low-impact responses in Nobody Security are a separate future offering and receive no current-product credit. [Fluency evidence]
Torq	3 V	Concrete cross-tool actions, configurable workflows, approval paths and logging documented. Execution success, reversibility and accuracy were not tested. [2]
TENEX	2 V/S	MDR tier explicitly includes automated response playbooks across endpoint/identity/network. Service and third-party platform dependencies prevent treating this as demonstrated standalone automation breadth. [5]
Exaforce	3 V	Specific endpoint, identity, cloud and email actions, workflow composition, approval gates and audit trails documented. Vendor claims; end-to-end execution not tested. [8]
7AI	3 V	Specific isolate/revoke/reset/quarantine actions, custom workflows, approval controls and case audit documented. Vendor claims; end-to-end execution not tested. [12]
Dropzone	2 V	Concrete account disablement/IP blocking and auto-containment documented. Scope depends on integration permissions; broad custom orchestration not established. [14]
Intezer	2 V	Embedded user blocking and SOAR/workflow triggering documented. Outcome effectiveness not independently tested. [16]
Qevlar	? ?	Homepage describes containment, while the AI SOC workflow says analysts take suggested next steps. Native execution scope remains unresolved. [19]
Prophet	2 V	Scoped Agent Actions can run autonomously or with approval. Broad cross-domain catalog and verified execution outcomes not established. [20]
Legion	2 V	Product explicitly executes workflows through browsers with oversight and staged autonomy. Specific containment catalog and reliability not established. [21]

Evidence behind the comparison

Vendor pages were reviewed for the comparison as of 1 October 2026. Links identify the material used; they are not immutable archived copies. Statements in these sources remain vendor statements unless separately identified.

1. **Torq: context graph and historical reconstruction**
<https://torq.io/blog/ai-soc-context-graph/>
2. **Torq: response and reporting scope**
<https://torq.io/use-case/automated-soc-incident-response/>
3. **Torq: Intune reporting workflow**
<https://kb.torq.io/en/articles/9350010-workflow-template-compliance-generate-report-on-non-compliant-devices-intune>
4. **TENEX: platform and deployment model**
<https://tenex.ai/>
5. **TENEX: service tiers and threat management**
<https://tenex.ai/services/threat-management/>
6. **Exaforce: data platform and retention**
<https://www.exaforce.com/platform/data-platform>
7. **Exaforce: investigation and configuration context**
<https://www.exaforce.com/platform/exabot-investigate>
8. **Exaforce: response and custom reporting**
<https://www.exaforce.com/platform/exabot-respond>
9. **Exaforce: SaaS posture and hygiene**
<https://www.exaforce.com/solutions/saas-attack-surface>
10. **7AI: federated SIEM and optional native storage**
<https://7ai.com/platform/federated-siem>
11. **7AI: posture and framework mapping**
<https://7ai.com/platform/security-posture>
12. **7AI: response actions and controls**
<https://7ai.com/platform/response>
13. **7AI: investigations**
<https://7ai.com/platform/investigations>
14. **Dropzone: investigation, memory and containment**
<https://www.dropzone.ai/ai-soc-analyst>
15. **Dropzone: MSSP operations and hunt reports**
<https://www.dropzone.ai/for-mssps>
16. **Intezer: forensic collection and response**
<https://intezer.com/enterprise>
17. **Qevlar: investigation workflow and reporting**
<https://www.qevlar.com/ai-soc>
18. **Qevlar: vulnerability agents preview**
<https://www.qevlar.com/soc-and-vulnerability>
19. **Qevlar: broader response positioning**
<https://www.qevlar.com/>
20. **Prophet: investigation, detection engineering and actions**
<https://www.prophetsecurity.ai/>
21. **Legion: workflow execution and operating model**
<https://www.legionsecurity.ai/>
22. **Legion: AI SOC implementation approach**
<https://www.legionsecurity.ai/blog-posts/implementing-ai-soc>

Additional Fluency evidence

Owner confirmations in the research discussion establish 12 months of default historical data, configurable behavioral windows including 90 days, configuration/property collection, and the MSSP findings-and-recommendations response model. Internal materials reviewed include the CISO Pack, Microsoft Defender Resource Analysis (validated in August 2026), customer reporting implementations and a supplied investigation report.

The internal documents support specific implementation scope; they are not public third-party certifications. The investigation artifact demonstrates historical querying, recovery from a wrong-table query, contextual analysis, customer-facing findings and escalation. One case cannot establish fleet-wide accuracy or a general performance benchmark.

Prepared for Fluency · Research synthesis and presentation · Version 1.0 · All product names belong to their respective owners.