GIGAOM

GigaOm Perspective Brief



# Five Alerts, One Attack:

## The Cost of a Fragmented SOC

Commissioned by CybrHawk



**Chris Ray, Field CTO, GigaOm**



## The One-Incident Question

If you are a SOC leader right now, you probably own more detection than you can use. A SIEM collects the logs, an EDR watches the endpoints, an identity platform flags risky sign-ins, a cloud tool watches the control plane, and an NDR sensor or two sits on the network. Each one fires alerts, and each alert is accurate on its own terms. But the attacker is running one operation, and your analysts are working five queues.

Here is the question that should be on your whiteboard: **when something happens across my environment, does my stack hand an analyst one incident with the whole story, or five alerts that someone has to stitch together by hand?**

Most current stacks cannot answer that well, and the reason is structural. Each tool keeps its own data model, its own idea of what an asset is, and its own severity scale. The SIEM was supposed to bring those views together, but it often became a search layer analysts query after the fact. The stitching still happens in someone's head.

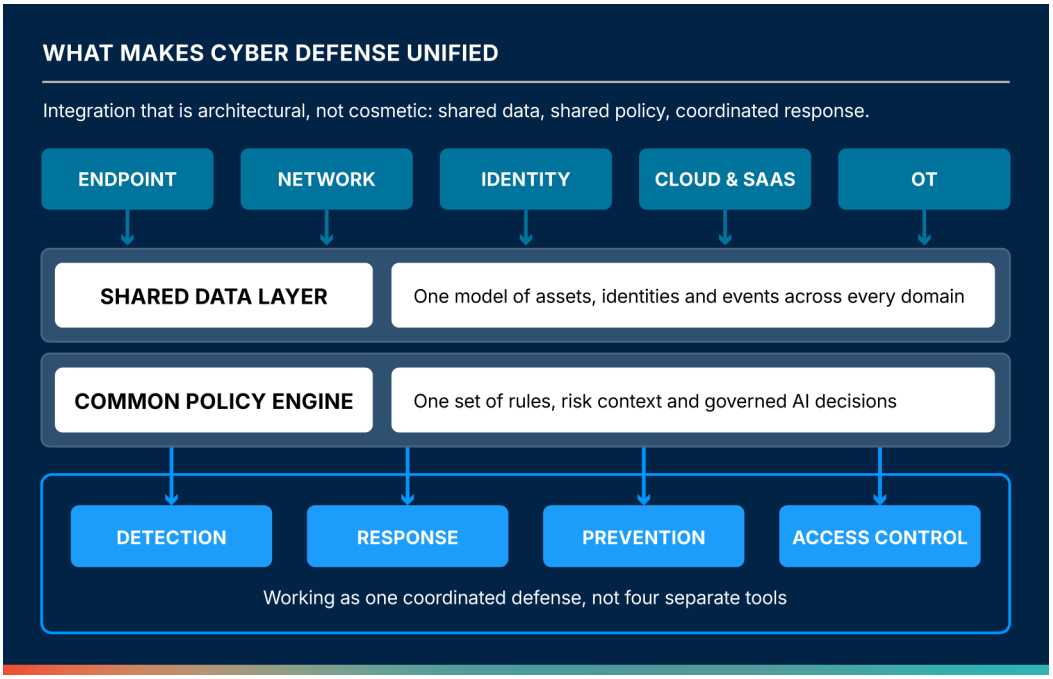
Consider a common sequence. A user approves a push notification they did not initiate. Twenty minutes later, an encoded PowerShell command runs on their laptop. An hour after that, the same account touches a file server it has never accessed, and a cloud bucket shows an unusual download. Your identity tool, EDR, network sensor, and cloud tool each raise a medium-severity alert in a different console. None of them is wrong, but none tells the analyst this is one intrusion moving toward data.

This brief offers a vendor-neutral way to evaluate platforms that promise to close that gap, starting with a definition, because "unified" is now applied to almost everything in a sales deck.

## What GigaOm Means by Unified Cyber Defense

Unified cyber defense is not yet a standardized product category. It is a strategic security posture, and GigaOm defines it as **an architecture in which detection, response, prevention, and identity and access control across endpoint, network, identity, cloud, and operational technology (OT) environments share a common data layer and a common policy engine, and act together through coordinated response.**

In practice, that means a detection in one domain automatically informs posture in the others, threat intelligence updates detection logic rather than only populating a dashboard, AI-driven triage and response are explainable and governed, and compliance evidence comes from the same controls that defend the environment.




Unified cyber defense is defined as much by what it is as what it is not. The test is whether integration is architectural or cosmetic. A single pane of glass over disconnected tools is not unified cyber defense, and neither is a rebranded bundle of acquired point products that still keep separate data and separate policy. A single vendor's platform, an open layer that works across the tools you already run, or a managed service can deliver a unified posture, but in every case, the shared data, shared policy, and coordinated response must be demonstrable.

## What You Are Actually Facing

Based on conversations with SOC leaders and security operations teams across mid-market and enterprise organizations, these patterns are consistent:



**The same incident lands in several queues.** Analysts start every investigation by pivoting between consoles to answer basic questions: which user, which device, what else did it touch? That pivot time rarely shows up in metrics, but it is often the largest share of triage.



**Your SIEM is essential and unloved.** The SANS 2026 SOC Survey found SIEM proficiency is the most requested skill in new hires, at nearly double the demand for EDR skills, yet EDR, not SIEM, earned the highest satisfaction rating of any SOC technology.



**Context lives in someone's head.** Which servers hold regulated data, which accounts belong to executives, which vulnerabilities are being exploited: this knowledge often sits with a few senior analysts. When they are off shift, severity reverts to the tool's default.



**Identity and OT are someone else's problem until they are not.** In the same survey, only 16% of respondents said their SOC fully covers OT and other at-risk smart systems, fewer than half monitor them at all, and identity signals often sit in a separate team's tool. Attackers do not respect those boundaries.

**AI is everywhere and in no workflow.** SANS found 79% of SOCs use AI or machine learning tools, but only 36% have built them into a defined workflow, which makes results hard to trust or measure.

**The integration tax never stops.** Parsers drift when vendors change log formats, and correlation rules break quietly. Only 24% of SANS respondents reported fully integrated dashboards and shared workflows between their SOC and network operations teams.

For the **CISO**, the same problem is a budget and accountability question. Detection spending keeps rising, but the board still asks how long it took to understand the last incident. A CISO evaluating consolidation needs to know whether a platform reduces the work of understanding an incident, or just the number of invoices.

### FIVE QUESTIONS EVERY SOC LEADER SHOULD BE ASKING RIGHT NOW

Use these as a scorecard in every vendor conversation. A strong answer is specific and demonstrable in your environment.

1. When my identity, endpoint, network, and cloud tools each fire on the same attack, how many cases does your platform create, and what does the analyst see first?
2. Which of my existing sensors can you ingest and correlate on day one without replacing them, and how long until that data is parsed and usable?
3. How does your severity score change when you know which assets and accounts matter most to my business?
4. For every AI-generated verdict or automated action, can my analyst see the evidence behind it, override it, and know whether a rule or a model produced it?
5. Where does your software end and your service team begin, and what happens to detection quality if I run the platform myself?

## What Your Shortlist Should Require

Based on GigaOm's coverage of SIEM, XDR, NDR, and OT security, what separates an adequate unified platform from an effective one is the following. Each requirement maps to a stage of the unified defense loop shown below.

### THE UNIFIED DEFENSE LOOP

From five alerts to one incident and one coordinated response, across every domain.




| | |
|---|---|
| 1 COLLECT | Telemetry from endpoint, network, identity, cloud and OT. |
| 2 NORMALIZE | Your existing sensors, one data model. |
| 3 CORRELATE | Resolve users and devices, build one case. |
| 4 PRIORITIZE | Rank by asset, privilege and threat context. |
| 5 RESPOND | Coordinated, governed action across every domain. |

1. **Cover your real blind spots *(reach)*.**

What matters is depth in the domains where you are weakest today; often identity, SaaS, and OT.

**TEST IT:**
Map the vendor's supported data sources and protocols against your own asset inventory, and for OT, confirm protocol-level visibility for the specific industrial protocols on your network.

2. **Work with the sensors you already own *(openness)*.**

A unified layer should ingest your EDR, firewalls, and identity provider by syslog, API, or webhook and normalize them into one data model without forcing a migration.

**TEST IT:**
Connect your actual EDR, firewall, and identity provider during the trial, and measure the time from connection to parsed, correlated data.

3. **Correlate across domains into one case *(coordination)*.**

The core promise is fewer, richer cases. That depends on resolving users, devices, and cloud identities to the same entity across every data source, then building a single timeline.

**TEST IT:**
Run a multistage scenario in the proof of concept, such as a suspicious sign-in followed by endpoint script execution and an unusual cloud download, and count how many cases the platform opens and how much of the story the first case shows.

4. **Let exposure and business context set priority *(context)*.**

A medium-severity alert on a domain controller or a CFO's laptop is not a medium-severity event. The platform should take asset criticality, identity privilege, current threat intelligence, and whether a vulnerability is being actively exploited, and use them to rank work.

**TEST IT:**
Give the vendor a short list of your crown-jewel assets and accounts, then replay the same alerts and confirm the ranking changes in ways you agree with.

5. **Respond as one, with AI you can audit *(governance)*.**

Response should act on every affected domain at once, and AI triage and automated actions shorten investigations only if analysts can see why a verdict was reached, correct it, and control which actions run on their own.

**TEST IT:**
Have the platform triage 20 to 30 historical alerts with known outcomes, record how often your analysts agree with each verdict and whether it shows its evidence, and confirm high-impact actions can require approval.

## The Architecture Decision

The choice in front of most SOC leaders is not whether to buy another tool. It is whether to keep stitching domains together in people and playbooks, or to move the stitching into a shared data and policy layer.

| | STITCHED POINT TOOLS | UNIFIED CYBER DEFENSE |
|---|---|---|
| **Where correlation happens** | In analysts' heads and SOAR playbooks | In a shared data layer and policy engine |
| **Unit of work** | Alert | Case with an entity timeline |
| **Severity** | Each tool's own scale | Asset criticality, privilege and exploitation status |
| **Adding a data source** | New console, new runbook | New feed into the same model |
| **AI role** | Ad hoc assistant beside the workflow | Governed triage step inside the workflow |
| **OT and identity** | Separate teams and tools | Same case, same timeline |

**Unify the data before the console**

*(principle: the data model comes first).*

A single pane of glass over separate data stores still leaves correlation to the analyst. Ask how the platform represents an asset and an identity, and whether every source maps to it.

**Keep the sensors you trust**

*(principle: compose, don't force replacement).*

A single-vendor platform can deliver unified defense, but most enterprises will run a mixed estate for years. Favor platforms whose correlation and response work as well on third-party telemetry as on their own agents.

**Gate autonomy by consequence**

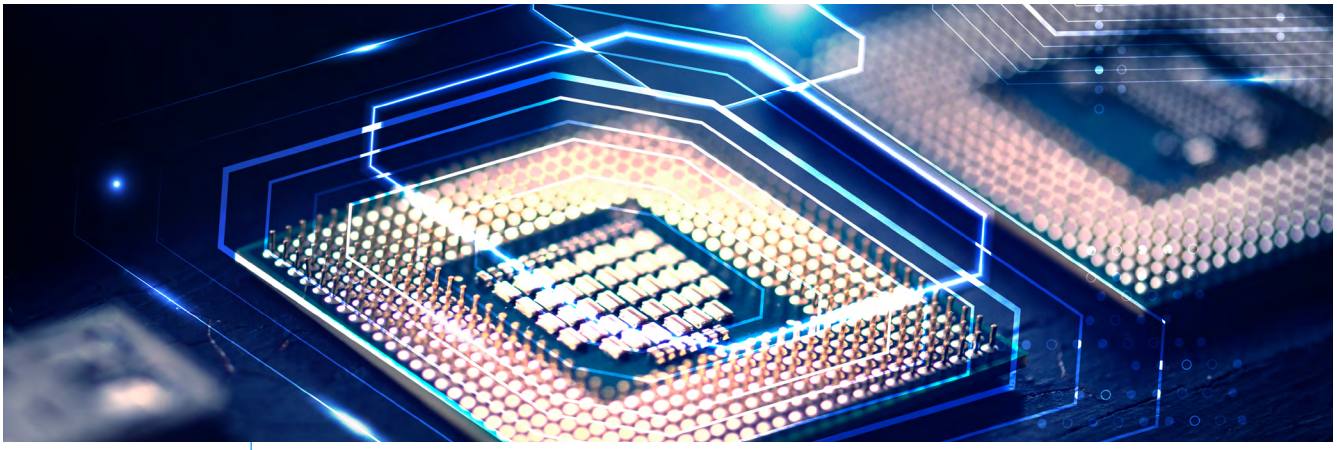
*(principle: human approval scales with impact).*

Enrichment and closing known false positives can be automated early. Isolating a production server or disabling an executive's account should require approval until the platform earns trust on your data.

**Make the service boundary explicit**

*(principle: know who is doing the work).*

Many platforms here are sold with a managed SOC. Know which outcomes come from software and which from the vendor's analysts, because that determines what you keep if the arrangement changes.



## Where CybrHawk Fits

CybrHawk is one of the vendors addressing the framework above. This section is based on vendor briefings, product demonstrations in August 2026, and product documentation, rather than a new scored evaluation. GigaOm has not yet evaluated unified cyber defense as a category in a GigaOm Radar.

CybrHawk builds its platform on a multitenant SIEM core that keeps each tenant's data fully isolated, which suits managed security service providers (MSSPs). Ingestion is vendor-agnostic, using syslog, API, and webhooks for firewalls, EDR products, and identity systems, and the platform ships with approximately 1,900 base alert types. A cloud module covers Microsoft 365, Google Workspace, AWS, and Azure; identity threat detection and response (ITDR) monitors privileged access; and HawkINT, the platform's threat intelligence module, enriches alerts with external indicators. Customers include MSSPs and mid-market and enterprise organizations in government, retail, and manufacturing.

Two areas deserve the most attention from a SOC leader: how CybrHawk puts AI into the triage workflow, and how it brings OT into the same data plane as IT.



### AI Inside the Workflow

Detection logic is roughly 80% rule-based and 20% machine learning; the AI is concentrated in what happens after alerts fire. Its AI Alert Escalator addresses the one-incident question directly: every 10 minutes it groups related alerts from all sources by entity and attack stage, and escalates any group of three or more that shows a logical attack progression against the same entity as a single scored incident. The AI Triage Assistant analyzes an individual alert across endpoint, identity, and network sources and proposes a verdict that an analyst accepts or rejects. The AI Analyst runs preapproved response playbooks, such as containing ransomware or a compromised account, and writes incident summaries with root-cause analysis.

This is a sensible division of labor: correlation and triage are automated, while response autonomy is bounded by preapproved playbooks.



### OT in the Same Data Plane

CybrHawk's network sensor passively monitors a mirrored copy of traffic, so it never sits inline with industrial processes. In OT environments it covers PLCs, SCADA systems, HMIs, and engineering workstations, flags deviations from established operational patterns, and supports industrial protocols such as Modbus. Because OT detections land in the same platform as identity, endpoint, and cloud alerts, an intrusion that crosses from the corporate network into a plant network can surface as one incident rather than two investigations by two teams.



**What Stands Out, and What Does Not**

Three things are more distinctive than breadth alone: OT, network, and identity telemetry in the same correlation engine as endpoint and cloud; multitenancy with full data isolation, which makes the platform a practical base for MSSPs serving customers with different EDR and firewall vendors; and compliance readiness views for PCI DSS, NIST, GDPR, HIPAA, and CMMC built on the same telemetry, so a team can see which controls are passing or failing before an audit.

SIEM-based correlation, an AI chat assistant, and compliance dashboards are now common across this market, and buyers should not treat any of them as a differentiator on its own. The Alert Escalator's 10-minute collection cycle suits multistage intrusions that unfold over hours, but fast-moving attacks will rely on individual high-severity detections and playbooks rather than correlated incidents. In the demonstrations GigaOm saw, the analyst interface was functional but less refined than those of larger, longer-established platforms. And because CybrHawk sells both software and a SOC-as-a-service offering, some outcomes buyers hear about may depend on CybrHawk's own analysts.

Measured against GigaOm's definition, CybrHawk is strongest on the shared data layer and coordinated detection. Prevention and access control largely stay with the firewalls, EDR, and identity providers the customer already runs, so the shared-policy and coordinated-response parts of unified defense depend on how far CybrHawk's playbooks can act through those third-party controls.

**Diligence items for a proof of concept:**

- Run the five tests above on your own data, including a multi-domain scenario and historical alerts with known outcomes.
- Measure alert reduction from the AI Alert Escalator on your own data rather than relying on published figures, and check what it filters out.
- Review every preapproved response playbook before enabling it, and confirm you can require approval for high-impact actions.
- For OT, confirm protocol coverage on your network and the deployment model for the network sensor.
- Ask for sustained ingestion rates, retention periods, and search performance at your expected volume.
- Have your analysts work the console for a full shift and judge usability for themselves.
- If you are considering the managed service, get the division of responsibility between your team and CybrHawk's in writing.

## The Gap Between Access and Impact Has Collapsed

The case for unifying detection is not new, but the time available to do the stitching by hand has shrunk sharply. Verizon's 2026 Data Breach Investigations Report (DBIR) found that exploitation of vulnerabilities is now the most common way breaches begin, at 31%, overtaking credential abuse, while the median time to fully remediate critical vulnerabilities on CISA's Known Exploited Vulnerabilities list was 43 days. A third party was involved in 48% of breaches, and ransomware in 48%.

When an intruder can move from a vulnerable edge device to a stolen identity to a cloud workload before the patch is applied, a SOC that correlates by hand across consoles is working on yesterday's timeline.

Meanwhile, AI is already inside most SOCs, but mostly outside their workflows. With 2027 budgets being set now, this is the right moment to decide where AI triage should sit, what evidence it must show, and which actions it is allowed to take, before habits form around tools no one governs.

# Analyst's Take

CybrHawk is solving the right problem at the right layer. The hard part of security operations today is not seeing an attack in any single domain; most organizations already own good sensors. The hard part is quickly understanding that several signals are one attack and deciding what to do first.

The main way this category fails is by becoming another console: if correlation does not produce fewer and better cases, a unified platform adds one more place to look. The second failure mode is trust, because AI triage that cannot show its evidence will be ignored by experienced analysts and overtrusted by inexperienced ones.

CybrHawk's strongest arguments are the two areas most stitched stacks handle worst: AI that correlates alerts into incidents before an analyst sees them, and OT telemetry that lands in the same incident as identity and endpoint activity. For organizations with plant, utility, or facilities networks, that second point alone can justify a proof of concept.

For CybrHawk specifically, the limiting factor is maturity rather than scope. The platform's breadth is wide for a company of its size, but buyers should test depth and usability in their own environment.

Success also depends on shared ownership. The SOC owns detection quality, but identity, cloud, and OT teams own the data sources and the asset context that make unified correlation work. Bring them into the proof of concept from the start, begin with correlation and context before autonomy, and expand automated response one action at a time.

**Whatever platform you choose, measure it by one outcome: how many cases your analysts have to open to understand one attack.**

## Report Methodology

**THE GIGAOM PERSPECTIVE BRIEF** provides an independent analyst view of an emerging technology decision. This brief was commissioned by CybrHawk; the definition, evaluation framework, and analysis are GigaOm's own. It was developed from CybrHawk briefings in August and October 2026, product demonstrations in August 2026, CybrHawk's published product documentation, and GigaOm market coverage. GigaOm has not yet published a Radar evaluating unified cyber defense as a category. Third-party statistics are drawn from the SANS 2026 SOC Survey and the Verizon 2026 Data Breach Investigations Report. GigaOm did not independently test the platform; buyers should validate capabilities and performance on their own data.



**ABOUT THE AUTHOR**

**Chris Ray** is Field CTO for Security and Risk at GigaOm, where he leads coverage of security operations, network security, SASE, OT security, and related markets.

Copyright© Knowingly, Inc. 2026 "*GigaOm Perspective Brief: Five Alerts, One Attack*" is a trademark of Knowingly, Inc.
For permission to reproduce this report, please contact sales@gigaom.com.



GIGAOM

GigaOm provides technical, operational, and business advice for IT's strategic digital enterprise and business initiatives. Enterprise business leaders, CIOs, and technology organizations partner with GigaOm for practical, actionable, strategic, and visionary advice for modernizing and transforming their business. GigaOm's advice empowers enterprises to successfully compete in an increasingly complicated business atmosphere that requires a solid understanding of constantly changing customer demands.

GigaOm works directly with enterprises both inside and outside of the IT organization to apply proven research and methodologies designed to avoid pitfalls and roadblocks while balancing risk and innovation. Research methodologies include but are not limited to adoption and benchmarking surveys, use cases, interviews, ROI/TCO, market landscapes, strategic trends, and technical benchmarks. Our analysts possess 20+ years of experience advising a spectrum of clients from early adopters to mainstream enterprises.

GigaOm's perspective is that of the unbiased enterprise practitioner. Through this perspective, GigaOm connects with engaged and loyal subscribers on a deep and meaningful level.