Latio.

Security Operations Market Report

2026

Table of Contents

Executive Summary

Security Operations Center (SOC) tooling is in the middle of a capability and expectation disruption. As enterprises invest in modernizing their SOC program, there is a growing demand for platforms that improve analyst productivity, accelerate investigations, and simplify operations without introducing noise or complexity. In this report, we explore 5 key themes teams will consider as they prepare to modernize their SOC platforms:

- "AI SOC" tools fit well within traditional security operations categories, are different from one another, and should be assessed on which business outcomes they enable.
- Most "AI SOC" platforms are really Security Orchestration, Automation, and Response (SOAR) tools, but direct access to underlying data drastically improves performance, context, and investigative efficiency. In this report we highlight the providers breaking out of the SOAR mold.
- The SIEM market is undergoing a major architectural transformation, enabling more scalable, data-centric platforms and improved user experiences across the SOC. There has never been more flexibility in purchasing a SIEM.
- Users are open to upgrading their SIEM, but need help making the migration process as simple as possible. Many teams end up with data sprawled between multiple providers.
- Leading solutions are attempting to evolve into larger platforms covering data pipelines, detection engineering, threat hunting, and traditional SIEM capabilities.

The SOC market is ripe for disruption due to the collision of agentic workflows with new data architectures, in this report we deliver the latest developments across both categories.

Survey Results

To better understand how organizations are evaluating, adopting, and using SOC platforms, we surveyed security practitioners and leaders across a variety of sectors and industry sizes. The results highlighted several shaping the future of security operations.

Below are the key takeaways:

- Most teams (64%) consume alerts via an MDR managing an EDR and SIEM
- A majority of users are unexcited by their current SIEM capabilities, but felt migrating to another tool isn't worth the investment required (72%)
- AI investments are focused on incident response, with many teams preferring to build tools in house (80%)
- Most teams are prioritizing the optimization of their incident response processes, followed by evolving data architecture
- Larger platforms are appealing but come at a cost: security operations platforms are the most commonly used tools, but are also the ones practitioners expressed more comparative dissatisfaction with.

Most Teams Use an MDR Managing SIEM and EDR

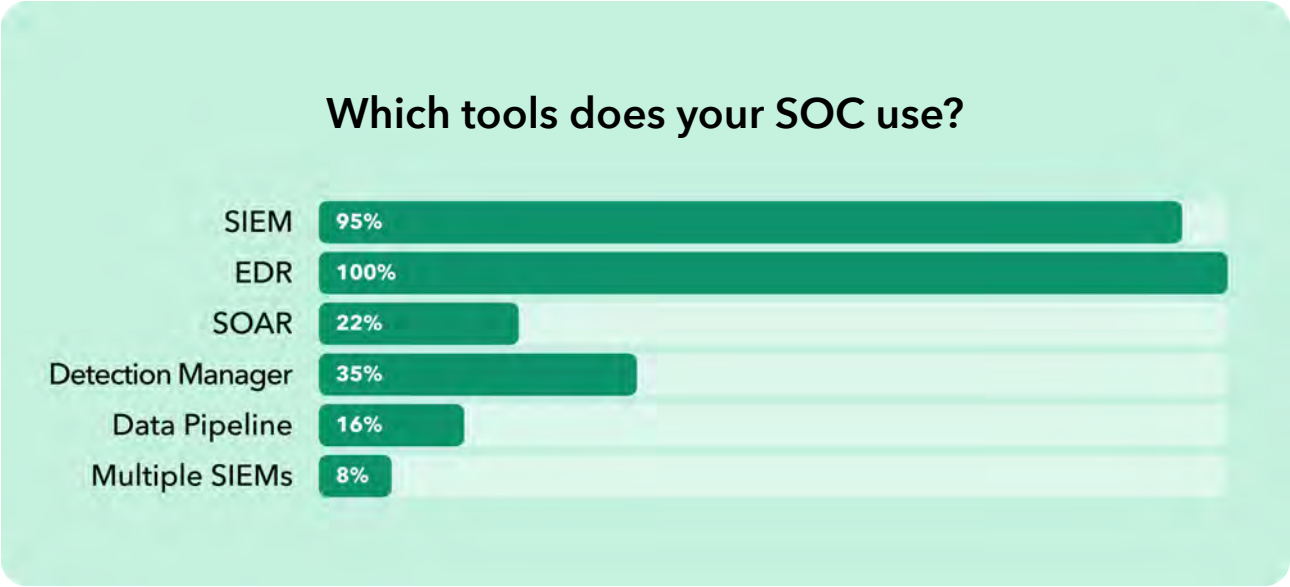

At the core of most SOC programs is the combination of a managed SIEM and EDR, typically delivered either through an MDR provider or directly by the vendor. A growing number of cloud-native companies have started to move away from dedicated SOC teams, instead shifting operational security responsibilities to product security teams.

Regardless of the hype around emerging solutions, SIEM and EDR technologies remain the keystone of modern security operations, while other tools see much lower adoption metrics. This is, in part, a result of MDRs providing services that often aren't incentivized to customize data pipelines and SOARs to their customer's environments. Other tools have benefits, but require an in-house SOC to operationalize.

Users Want to Modernize Their SIEM, but Migration Comes at a Cost

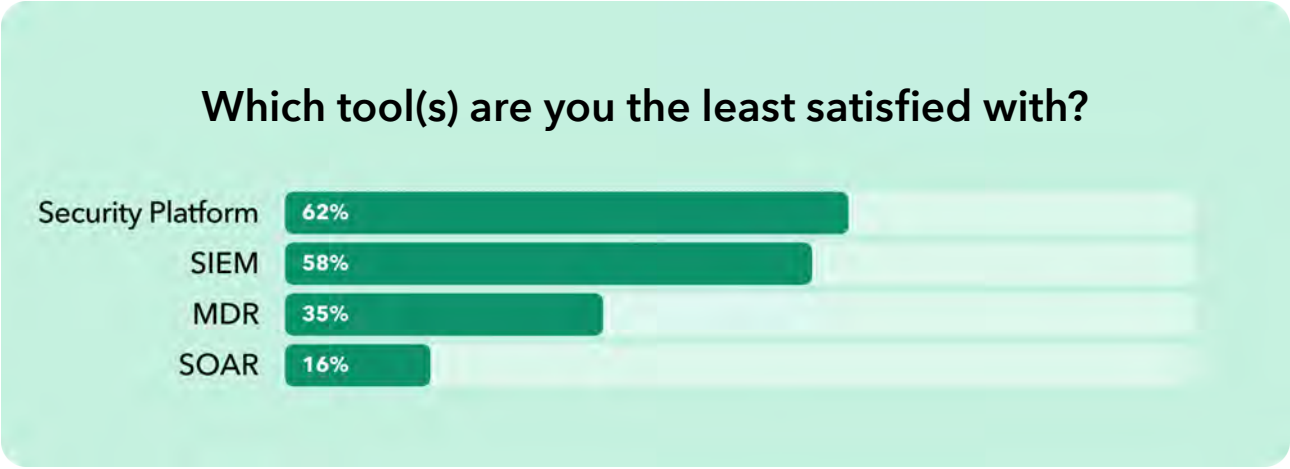

The majority of users surveyed found their SIEM to be underperformant, but not enough to justify the resources required for migration. This was reflected by responder's preference to stay on their current SIEM, despite finding their larger platforms to be the tool they're the most dissatisfied with. This is why there has been an emergence of startups focused on improving the usability of these tools by moving features like detection engineering and data pipelines outside of their SIEMs.

Teams are Investing in AI to Improve Core Metrics

Is AI SOC part of your 2026-2027 budget?

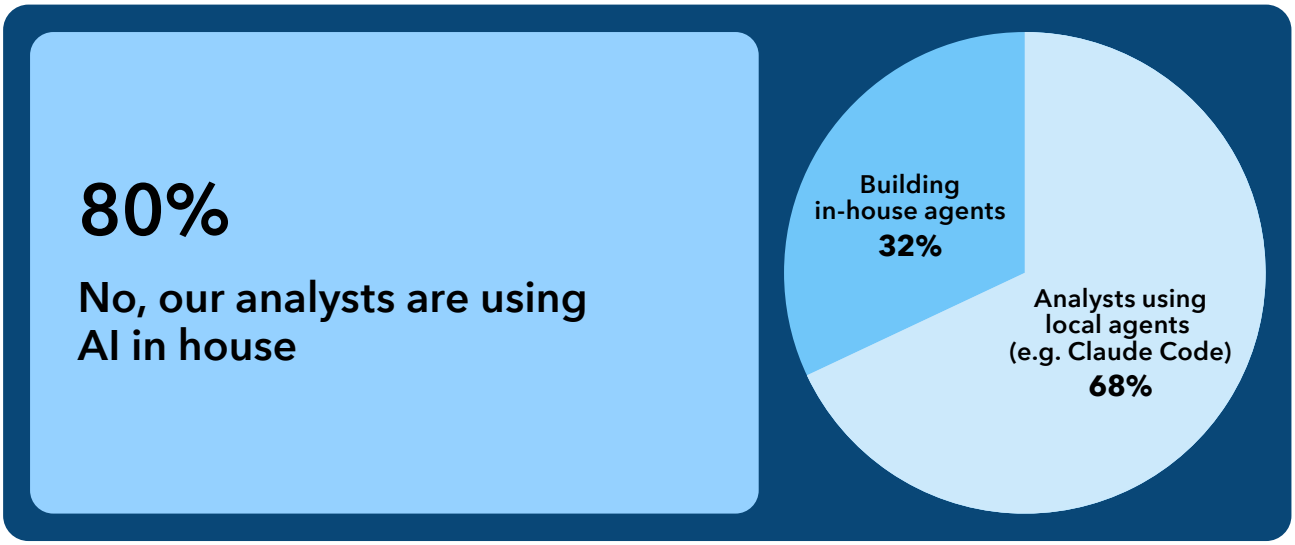

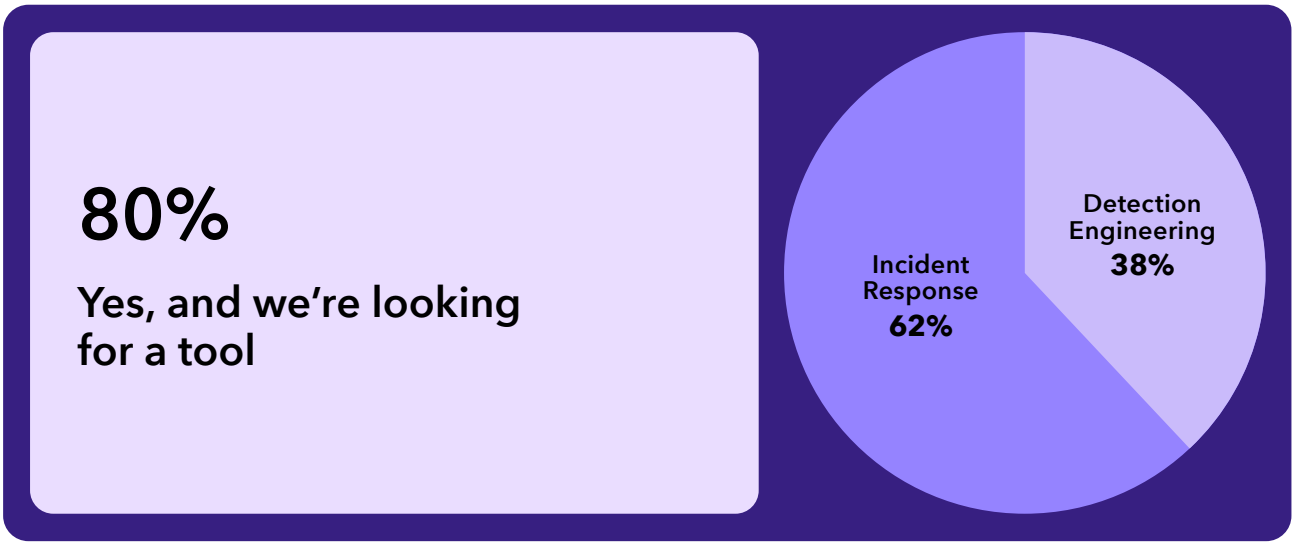

Security operations teams are approaching AI adoption as an operational enhancement layer rather than a rip and replace decision. For more advanced teams, they are enabling their analysts to leverage tools such as Claude Code into existing workflows, while enterprise organizations are focused on building custom agentic workflows internally. As executive leadership increasingly prioritizes AI initiatives, security teams are gaining dedicated modernization budgets, with the majority of investment focused on improving incident response operations. 80% of teams are choosing to use AI in house, while 20% were looking to bring in a dedicated vendor.

Improving Incident Response and Optimizing Data Emerge as Top Priorities

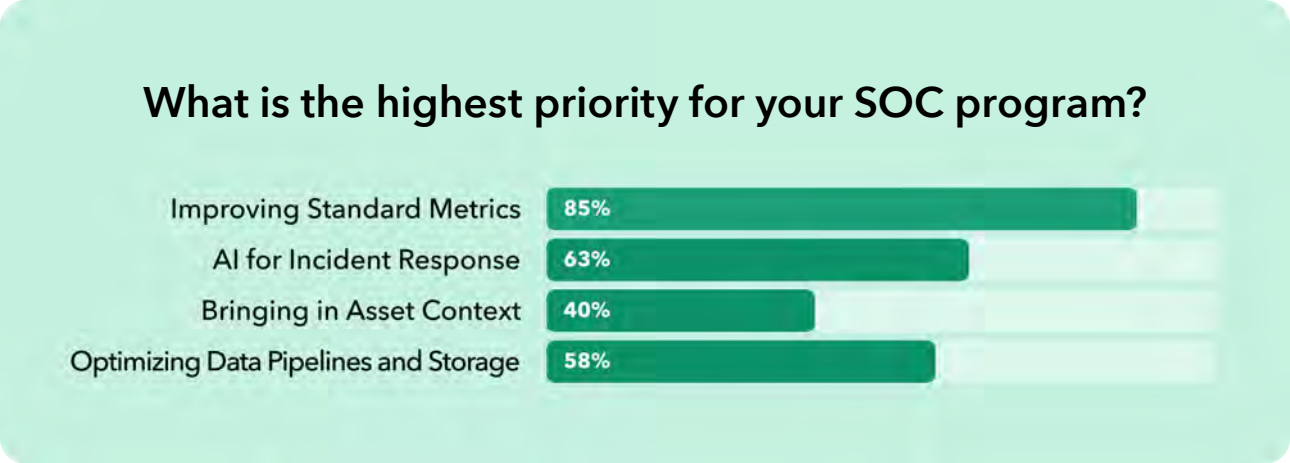

The highest priorities for teams centered around improving day to day metrics like mean time to respond and mean time to tune a detection rule. Organizations are viewing AI as a mechanism to improve these core metrics and reduce analyst busy work. Closely behind incident response modernization was the need to optimize underlying data architectures. For some, this meant adopting data pipeline tooling to optimize their log ingestion and processing efficiency, while for others it meant shifting to more flexible data architectures outside of the SIEM. Together, these trends point to the continued decentralization of SIEM, as organizations prioritize flexible, performance, and cost efficient storage.

The Evolution of Security Operations

What Is the Goal of a SOC

The security industry is broadly split into two disciplines: proactive and reactive. Proactive security helps to design and implement secure systems; while reactive security exists to respond to ongoing attacks. Security operations centers (SOCs) are oftentimes forced to do both, but their focus has always been on reactive security. The primary goal of a SOC is to detect and respond to security incidents.

Traditionally, reactive security has been the domain of security operations through five core components:

- Maintaining an inventory of log sources to ensure visibility across an organization
- Getting relevant logs into a place where they can be searched and stored (Data Pipelines)
- Creating detection and correlation rules around those log sources to detect unauthorized activity (Threat Detection).
- Responding to potential security incidents (Incident Response).
- Searching for evidence of attacks in historical data (Threat Hunting).

Security operations began with a focus on investigation capabilities around endpoint and network telemetry; the scope of security operations has now expanded to include operations in identity, SaaS tools, and vulnerability management.

The SOC is the bedrock of security practices, where teams detect and respond to live incidents in an environment. Their work often expands to several tertiary disciplines, from vulnerability to identity management.

How the SOC Has Evolved

Security Operations Timeline

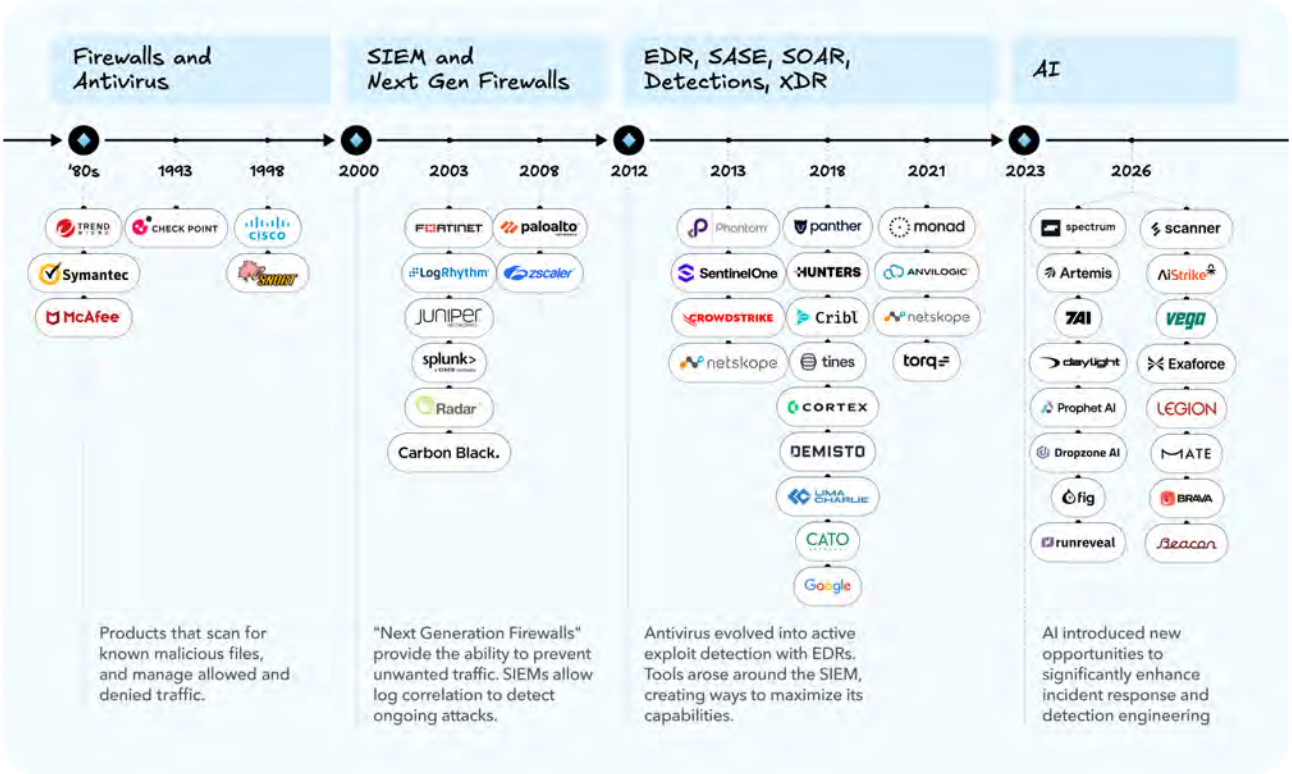

From file hashing to AI innovation, this section covers the overall evolution of the SOC. Addressing the investments over time helps to understand what challenges various companies were founded to solve, and where they fit in the overall ecosystem of priorities. Overall, we've seen core technologies (file scanning and log correlation) continue to underlie daily operations.

Early SOC platforms were designed to be firewalls and antivirus tools that worked by creating signatures of known malware and detecting them. For as many new anomaly detection and correlation engines exist, signature based detections continue to be the one of the most important functions for day to day operations. Most operators continue to use YARA rules for detecting malicious files, nearly 20 years after their invention, and they're typically the first line of defense for emerging threats.

Antivirus and firewalls evolved into EDRs and NGFW technologies respectively. EDRs went beyond signature based detections by looking for malicious patterns in real time, allowing them to spot new threats. NGFW brought intrusion detection and prevention into basic IP routing capabilities, enabling threat detection at the network layer.

SIEM tools emerged as a way for teams to build their own detections across a variety of different sources. These would quickly become the bedrock of security operations, as several types of operators were required to make them work. First, SIEM engineers build and maintain SIEMs, which require a lot of operational overhead as self-hosted systems. Second, threat detection engineers build the correlation rules that generate the potential security incident. Finally, security analysts review the detection alerts, validating if a real security incident was taking place.

Over time, EDR tools become more SIEM-like, leading to the rise of "XDR" functionality. The market never really decided what to do with XDR, with some treating it as an EDR + SIEM functionality, and others focusing on response capabilities as the differentiator. Today, the terminology has been largely dropped in favor of separating out the SIEM elements of a platform from the EDR elements. Yesterday's XDR providers now prefer to go-to-market with SIEM and EDR capabilities that complement each other.

Finally, SOAR tools emerged as workflow generators for day to day operations. These tools have always struggled to move beyond niche functionalities, as many security response actions are not as straightforwardly repeatable as it might seem. Many teams have built robust catalogues of automations with SOAR, but technical debt is a persistent challenge in operationalizing the capabilities.

Until recently, most of these technologies uneventfully consolidated into a few key players, with most organization's stacks consisting of only one to three vendors - EDR, SIEM, and Firewall. Advancements in data engineering and AI have made these straightforward choices once again more open, as organizations have access to cheaper storage, and AI drives more automation potential. This has led to a renewed interest in SOAR capabilities, and a re-evaluation of if SIEMs are still necessary - or more realistically, what the future of SIEM architecture looks like.

Saying "the SIEM is dead" has been dead since 2011. While the SIEM will never die, its backend architecture is evolving, creating new opportunities for teams ready to invest in the benefits of data lakes - from faster detections to more flexible storage.

The recent explosion of "AI SOC" tools are quickly expanding to become larger platforms. Early tools in this category focused only on automating incident response with LLMs - an evolution of the SOAR platforms that preceded them. These vendors have quickly realized that analysts sit downstream of other capabilities that need improvement to function properly - data ingestion, storage, and detection engineering. These other capabilities are forcing AI SOC tools to evolve outside of incident response automation to become data platforms in their own right. Later in this report we'll break down the current landscape of that evolution.

The Emergence of SOC Platforms

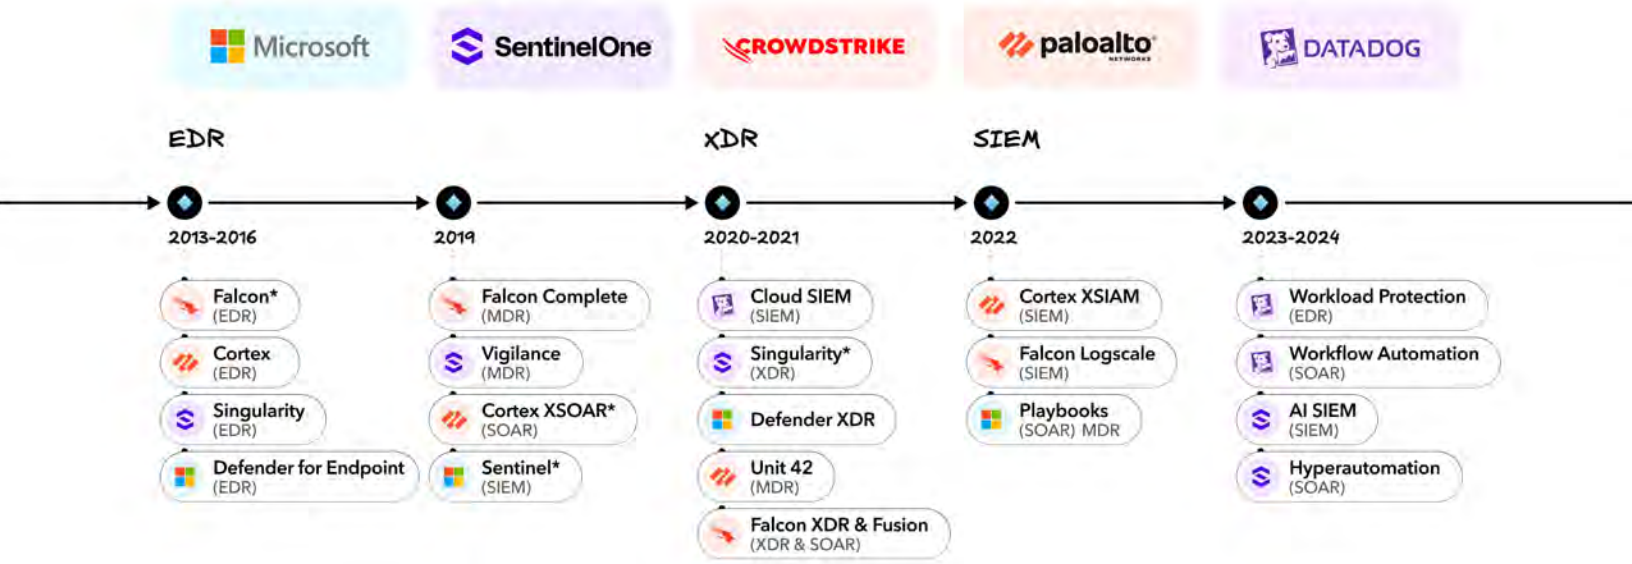

Several SOC tools followed a familiar growth into platforms: moving from endpoint protection into broader log correlation, with eXtended Detection Response (XDR) being the awkward middle. Modern SOC platforms provide several key components:

1. EDR is the core component of the SOC, monitoring endpoints for active exploitation
2. EDR platforms tend to expand into XDR platforms, as they add in workflows and broader log ingestion capabilities
3. Eventually, these platform's log ingestion capabilities expand into SIEM replacement territory, where they can function as end to end security tools.

While these platforms compete and overlap on paper, many enterprises end up with some combination of solutions due to cost and migration challenges over time. Typical enterprises utilize multiple log ingestion solutions, optimizing for cost and queryability of the underlying log source.

Modern Security Operations

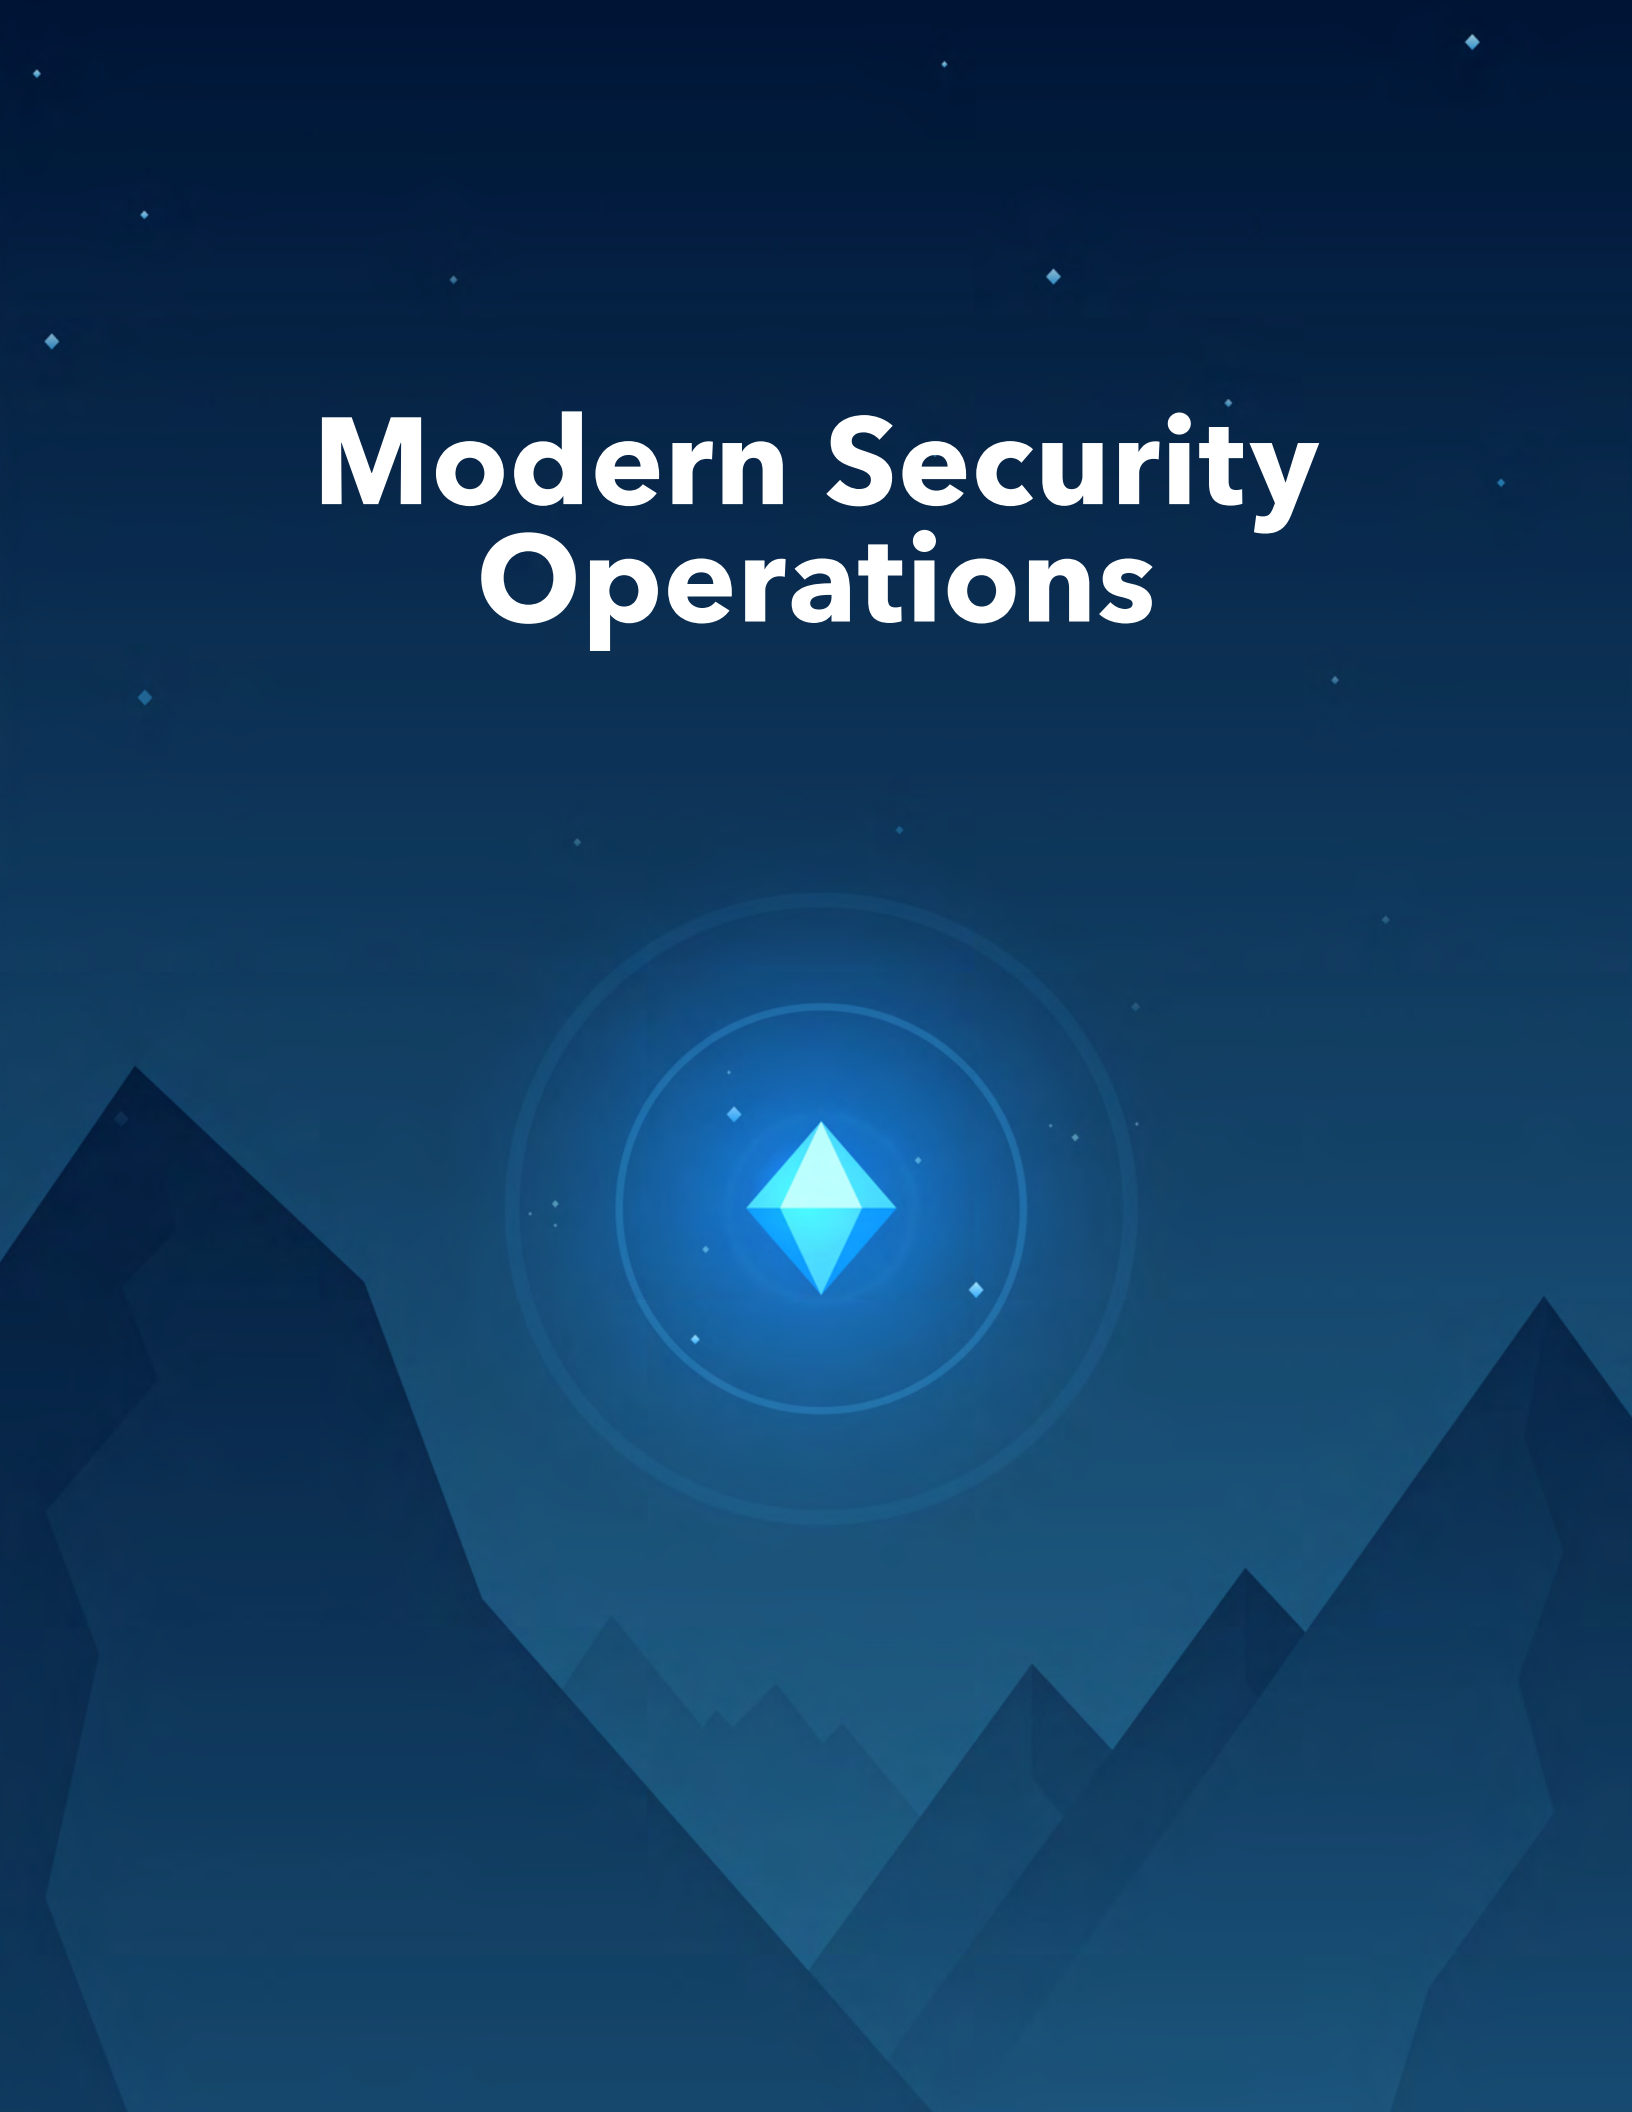

SOC tools have expanded beyond traditional ELK (Elastic, Logstash, Kubana) architectures. Teams now have access to a growing variety of emerging technologies built to improve their detection and response capabilities. This section covers the key components of a modern SOC architecture, and what tools are available for teams based on their current architectures and operational maturity.

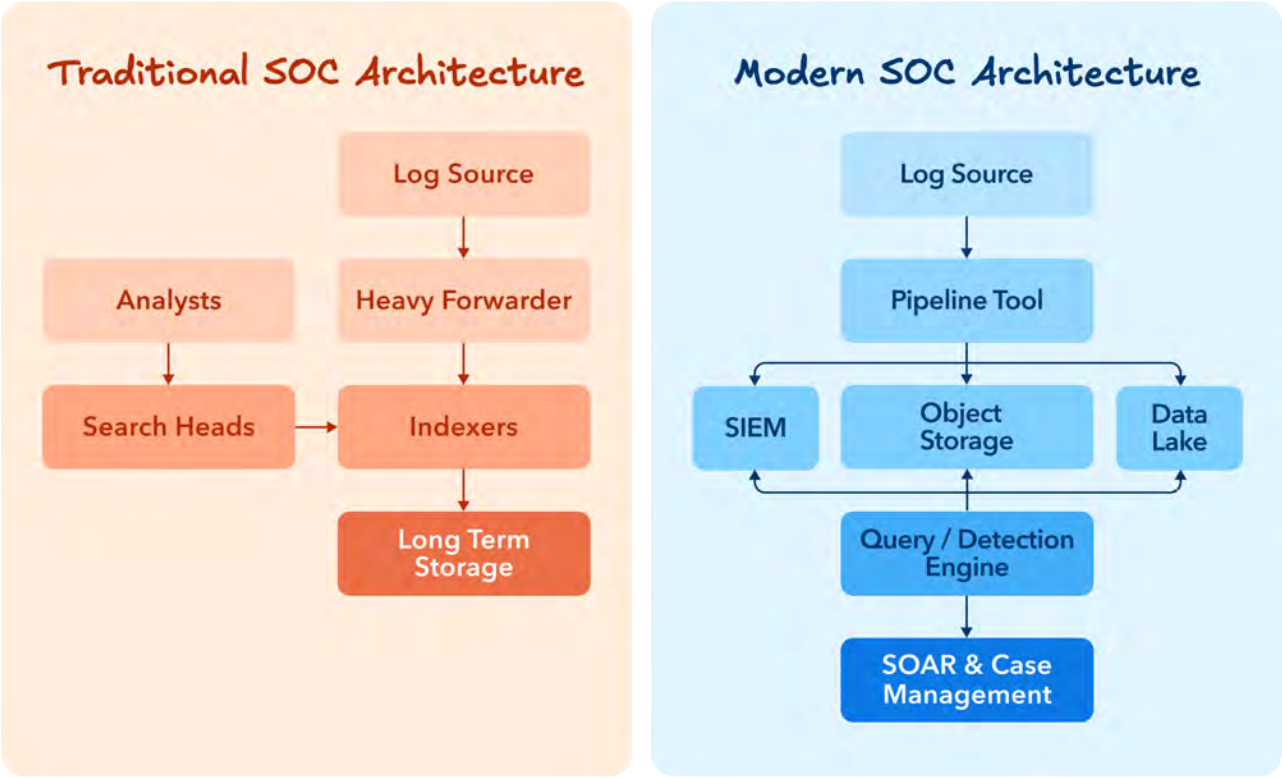

At the core of the modern SOC market is the decentralization of the underlying components of the SIEM. Traditional SIEM solutions consolidated data ingestion, indexing, searching, and long term storage into a single solution. That approach has become misaligned from how teams, and agents, store and access data.

Modern SOC teams optimize for accessing data that is more distributed than ever before. Data is generated, enriched, and retained across multiple domains. Teams outside of security including data engineering, IT, and developers all have their own data tools. In our Modern SOC Architecture diagram above, we outline the 5 layers that comprise the new SOC:

- Log Sources
- Data Pipelines
- Log Querying and Storage (SIEM, Object Storage, and Data Lake)
- Threat Detection
- SOAR

Instead of choosing between single platforms, teams are equipped with more flexibility to create a data architecture that works best for their needs and environment. In the rest of this report, we'll enumerate the components of a modern SOC architecture to see how the capabilities are evolving.

Log Sources

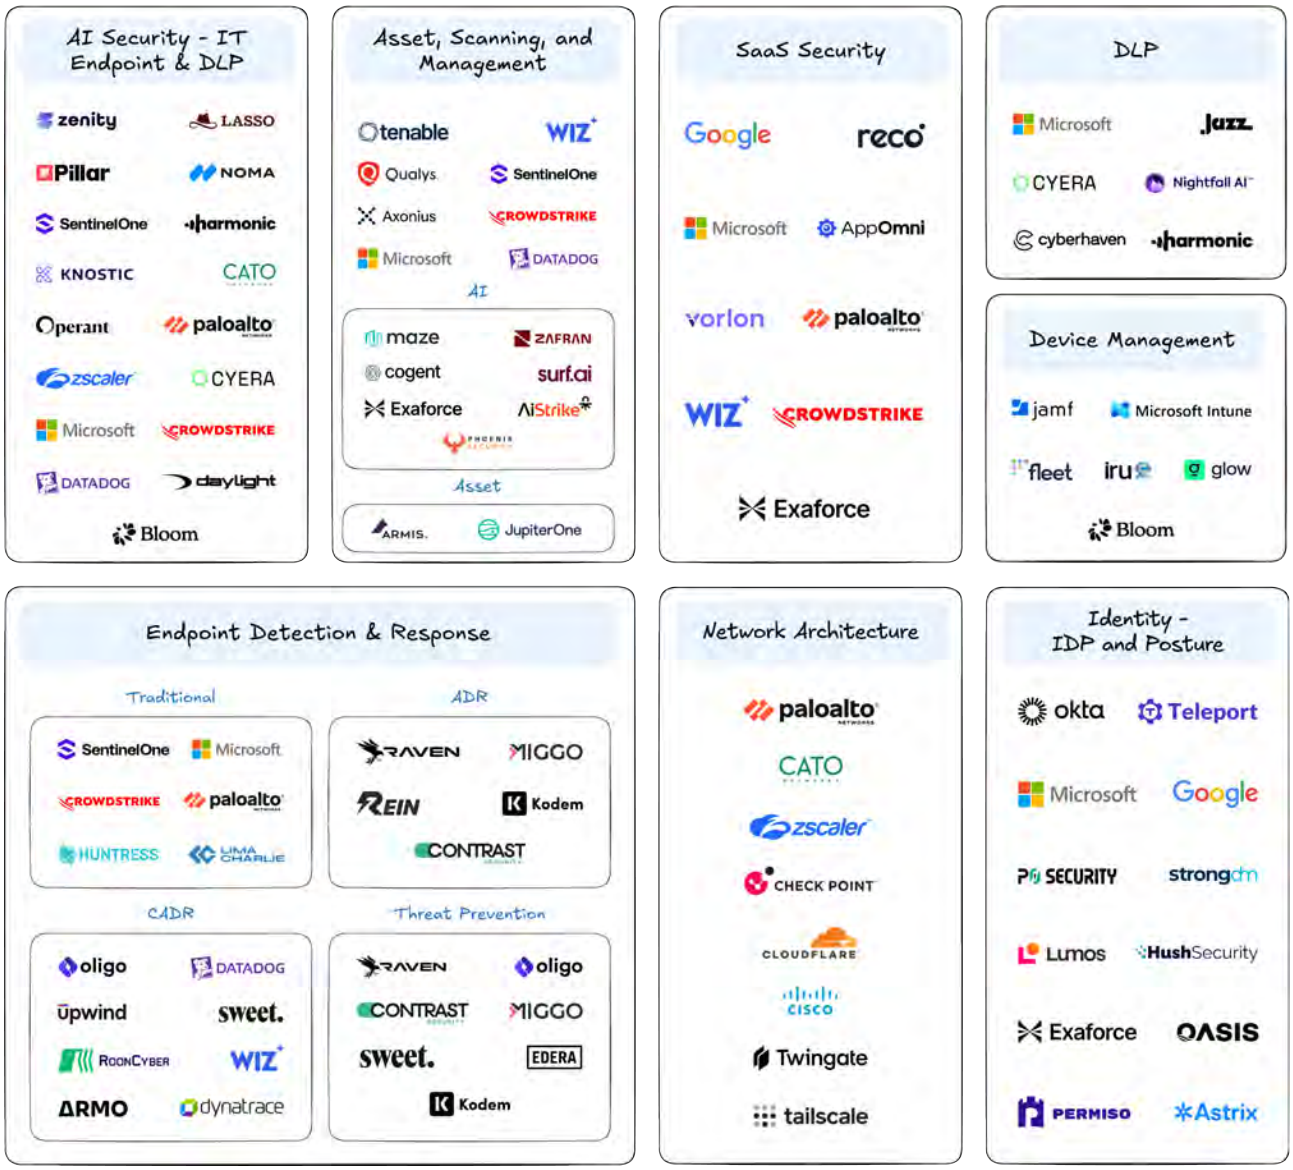

** Representative Vendors*

Security operations programs are built upon having complete and accurate log source coverage - ensuring teams have the visibility they need to detect emerging threats. When onboarding new vendors, SOC teams need to quickly understand their log sources: ingestion types, quality, and data structure.

Log sources tend to have one of two challenges: providing overly robust logs that explode ingest volume, or barely providing any logs, undermining detection capabilities.

While several log sources have remained the same, such as firewalls, EDRs, and DLP tools, there are several trends that are changing the market are worth highlighting:

- The importance and challenge of building meaningful SaaS detections
- The addition of application level telemetry to security investigation needs (ADR)
- Emerging AI security use cases which range from application detections to DLP
- Bringing in asset context to alerts via vulnerability management tools
- Uniting identity telemetry across various tools

Each of these five trends within log sources are quickly evolving to impact SOC practitioners and their daily work.

SaaS Detections

SaaS specific detection rules have been an aspirational goal for many teams, but several tools are available for improving these capabilities. On the one end are dedicated SaaS security management tools, which can take either a posture or runtime centric approach to detecting potentially malicious user activities. The other end are newer all-in-one security providers that bundle SaaS posture and detection events as part of their core offerings. Cloud security platforms like Wiz are also evolving to include more integrations as AI is expanding outside of traditional providers to other hosting solutions like Vercel and Railway.

Application Detection and Response

Even more than SaaS detections, application layer detections are challenging for security operations teams to operationalize. This is due to both the high level of domain specific expertise required to ingest application logs, and how specific application detections are to a specific application.

Thankfully, out of the box application detections and telemetry have become more accessible than ever to security teams - whether via standards like open telemetry, or via ADR and CADR solutions, which provide robust application layer protection and prevention for application workloads. These tools are becoming even more essential as AI makes first party applications easier than ever to attack, and time to exploit plummets.

AI Security and Telemetry

AI security boils down to two key use cases: protecting first party agents, and protecting employee usage of AI tooling. For protecting first party agents, ADR capabilities provide the core functionality - namely getting runtime insight into application logs and behaviors. Like any other applications, AI tools generate logs which can be correlated, detected, and blocked.

However, most teams are first focusing on protecting end-user adoption of AI, which often works on proxy based log collection, browser plugins, or most recently hooks. These tools give teams DLP insights into what data is being passed through AI systems. One core consideration for teams looking to operationalize these findings is that most resolutions come down to data access and identity controls, as AI will more quickly expose incorrectly scoped permissions.

Asset Context

Following in the path of CNAPP, large security operations providers are beginning to further unify asset context and blast radius information with detection and response capabilities. SOC practitioners are able to more quickly understand the overall context of an asset and what steps can be taken to reduce the scope of an ongoing incident. This unification is especially important for SOCs working on both vulnerabilities and incidents instead of separating those responsibilities. AI is especially useful for deducing asset ownership, and finding the most likely person who can address an issue.

Unified Identity Telemetry

Human SOC operators use a combination of assumptions and learned context to figure out the root cause of an incident. It's normal for operators to send out messages to entire teams in order to figure out the root user of an action, or if it was an approved operation. AI tools don't have this contextual knowledge organically.

In order to make AI agents more effective, some vendors are unifying identities across various providers. By creating an identity relationship graph, teams can more effectively baseline and respond to potential alerts by finding the responsible user.

Once teams have logs to search for potential security incidents, data pipeline tools are used to ingest, route, and enrich individual logs before they're used for security analysis.

Data Pipelines

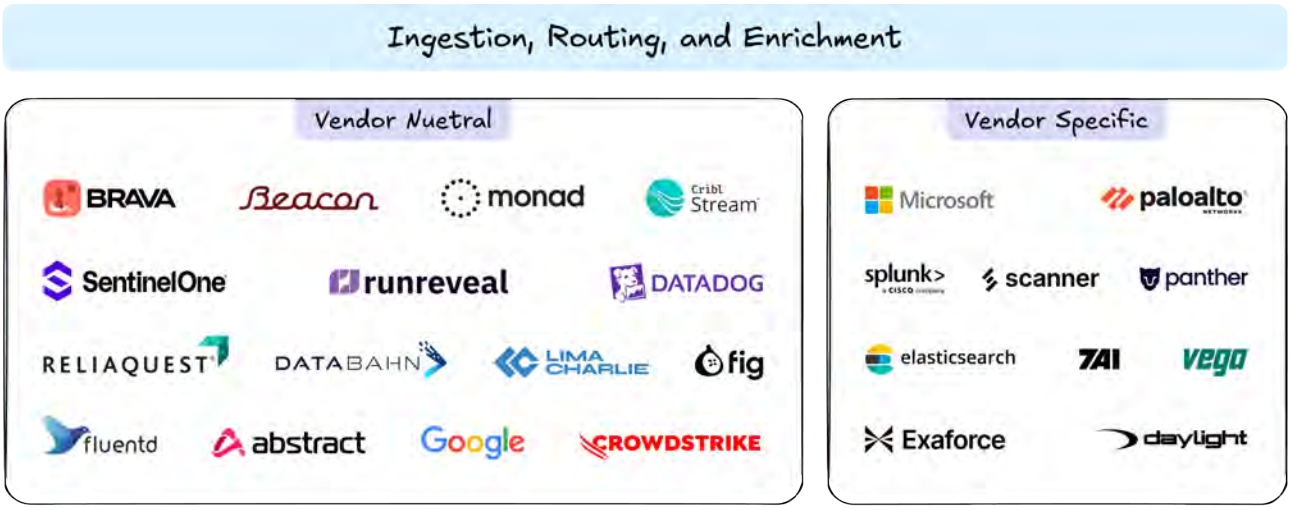

The transition to cloud became the foundation of the vendor neutral data pipeline category, which was defined by Cribl and later expanded into many SIEM providers who began building their own ingestion mechanisms. For these platforms, basic features are ingestion and routing, but the best can do real time enrichment. While routing tends to be the primary use case, these tools offer several additional benefits:

- **Normalize similar log data.** Most vendors have their own log values for similar concepts. Data pipeline tools can be used to standardize this data into a format like OCSF.
- **Test for log health** and that the correct fields are getting ingested. Silent log failures can lead to detections that stop firing, creating blind spots over time.
- **Lookup additional context** that may help with an investigation. Several tools are able to enrich individual logs with context that would otherwise need to fire post investigation, some examples include looking up additional user or geolocation data and appending it to the relevant log.
- **Optimize the cost** of short term and long term storage. This continues to be the most immediate value of data pipeline tools, optimizing what data is ingested and sending it to the cheapest storage location for its purpose.

There are three overall trends in the data pipeline space. First, many providers are offering their own data pipeline tools for routing data into their SIEM as a first party service. Second, several providers are running detections at log ingest rather than through retroactive searches. Finally, these tools are the bedrock of a SIEM migration strategy, as they allow teams to more independently migrate between backend providers. Finally, many practitioners we interviewed were frustrated with larger platforms expanding into broader data use cases and no longer prioritizing security use cases.

Teams should assess tools in this category by two primary factors: their level of commitment to a single vendor, and the amount of customization that can be enforced. Many of these tools allow vendor neutral routing as an afterthought, or as a result of an acquisition. Leaders in this category were selected for their flexibility and ability to function as a standalone solution.

Log Querying and Storage

SIEM Interoperability Graph

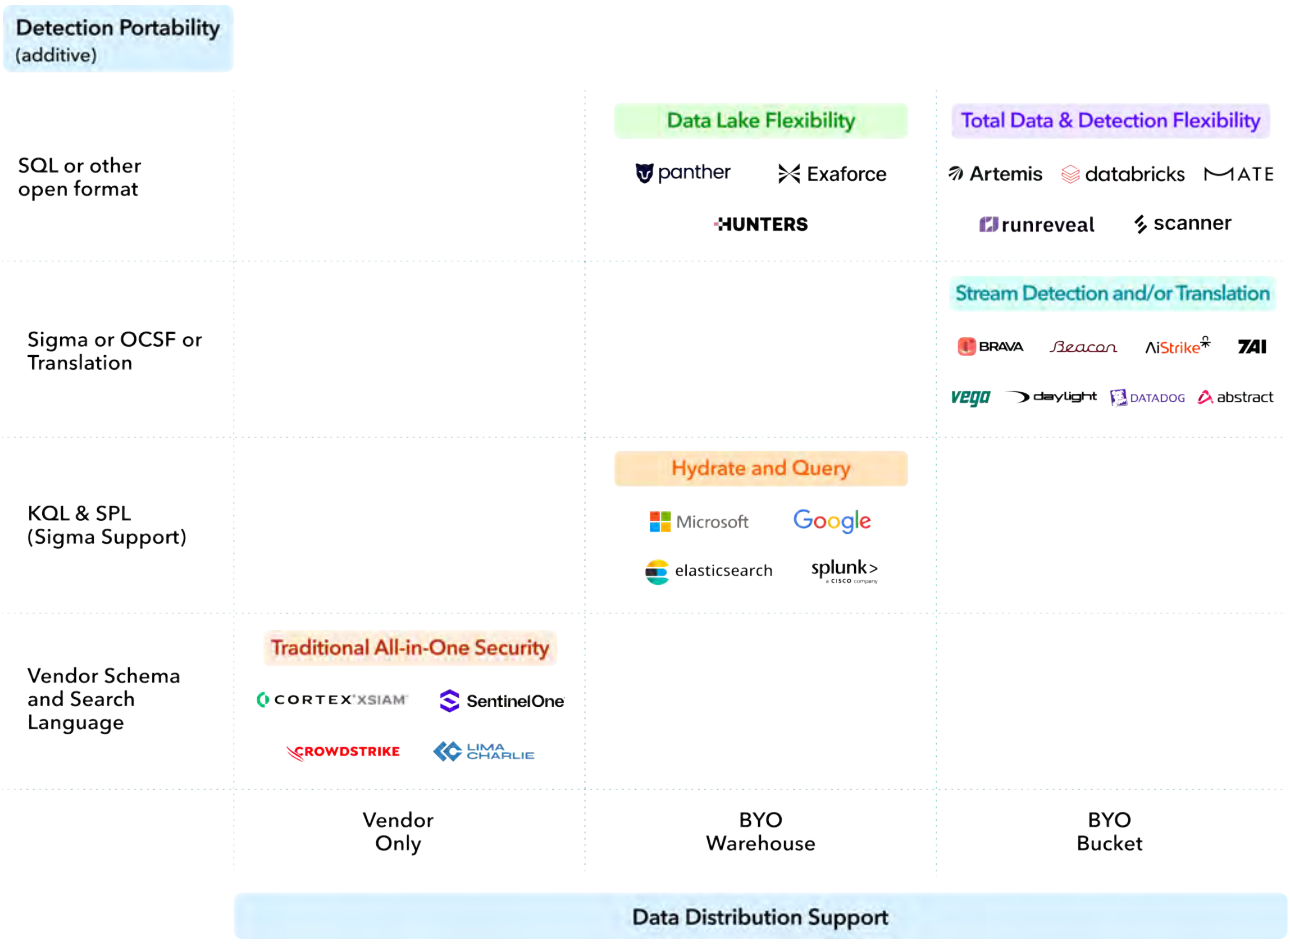

Data pipeline tools typically route most of their data into a Security Information and Event Management (SIEM) tool. Of all components of the SOC, SIEM architectures are changing most rapidly. The above graph depicts SIEMs by how open they are - how easy it is to migrate on and off of them, and how open their data methods are to being used across different teams. We've categorized SIEM type platforms into five categories, **each with their own pros and cons:**

Traditional All-in-One Security

Data is Ingested directly into a vendor SIEM, and users write detections with that vendor's language. The primary assumption of these tools is that your security data will exist almost entirely in their stack. Typically these are vendors unifying EDR data alongside other aspects of your environment. The benefit is having a unified data lake, but at the cost of being locked into that vendor's set of tools.

Hydrate and Query with Federated Search

These systems primarily function as traditional SIEMs, but have more flexibility and maturity in their detection ecosystems. They've prioritized having open access to robust data lakes. While it's possible to ingest other data into their systems, it's typically harder to leave and the query engine is assumed to be in their system. These tools have mature sets of detection capabilities and skillsets, but lack some of the evolving capabilities of startups.

Stream Detection and/or Translation

Newer entries to the SIEM ecosystem tend to build around observability functionalities while allowing query translations between other systems. These systems usually detect on log ingestion time, and then build indexes of data while it's ingested in order to speed up querying no matter where the data lives. These tools can empower new outcomes and be used to manage multiple tools, but can add another management layer to an existing architecture.

Data Lake Flexibility

These platforms tend to offer the most robust detection capabilities, prioritizing robust and fast querying of vast amounts of data. They typically don't directly support bringing your own

bucket models because of how poor search performance becomes when querying this data without a proper index; they do however have ways of ingesting S3 data into their data lakes. These tools allow you to have data sovereignty, but require a high level of data architecture or detection sophistication to take advantage of.

Total Data and Detection Flexibility

These systems combine the flexibility of SQL based detection rules with object level data. Sometimes this flexibility comes from building an index at log stream time, other times it comes from unique intellectual property on building indexes from historical data. These tools are the easiest to migrate on and off of, but can work less well alongside existing systems depending on their translation and federated search capabilities.

Beyond go to market strategy, these developments are occurring due to a transformation in data technologies, moving from Elasticsearch, Logstash, and Kibana (ELK) architectures to those powered by technologies like Apache Iceberg, Clickhouse, DuckDB and object storage. To put it simply, ELK is out, and data portability is in.

These new backend architectures are capable of drastically reduced costs, easier querying of data, and cleaner architectures. When not managed carefully, they can also cause half completed migrations, delayed queries, and unoptimized data. This is why transitioning off of traditional SIEMs should be done in a specific order: optimizing your data pipeline, decoupling your detection engine, then migrating the data.

Opportunities and Risks of SIEM Migrations

Opportunity	Risk
Consolidate Storage	Incomplete Migrations
Ingest cost savings	More money and resources spent on maintaining routing pipelines
Consistent detections	Broken detections when federated search doesn't work as expected
Flexible querying	Long query times on unoptimized data

Threat Detection

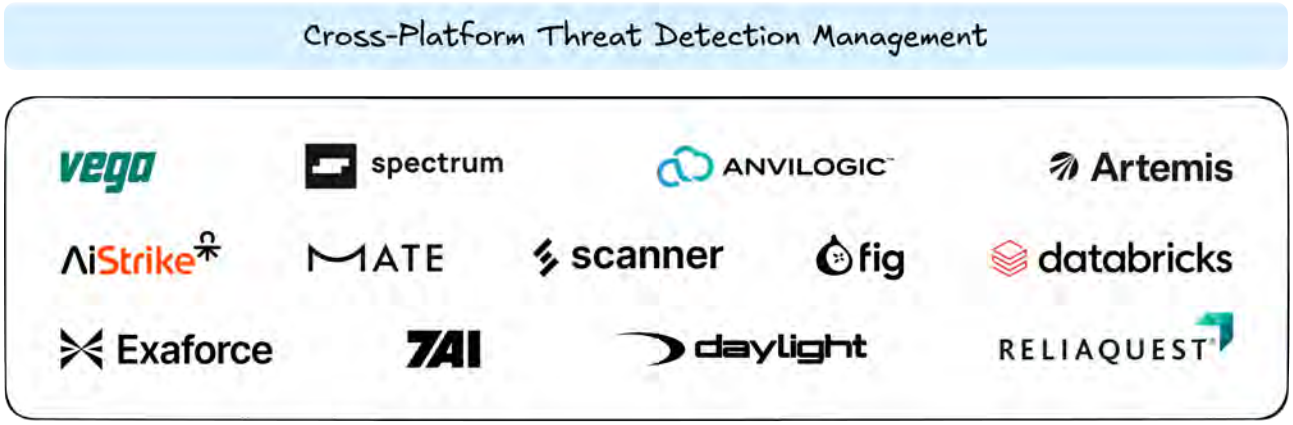

As teams move towards distributed SIEM architectures, a market has grown for tools that consolidate and standardize your detection engines. Threat detection tools excel at creating standardized queries for disparate log sources.

There are three major approaches to watch for in this category:

- **Creating an index based on ingestion or historical data.** These capabilities allow teams to achieve reliable searches on historical data without compromising on speed.
- **Integrating directly with third party SIEMs to create a unified detection library, and orchestrating detections where they already live.** These tools also tend to use LLMs to translate detections between various vendor languages. This is the most common approach, with tools working as a visibility, monitoring, and suggestion layer on top of your SIEM(s).
- **Running federated searches.** Most providers say they support federated searching, but the user experience varies largely between vendors. This can lead to long query times and inconsistent results. Provider's true capabilities here can be radically different and must be carefully assessed.

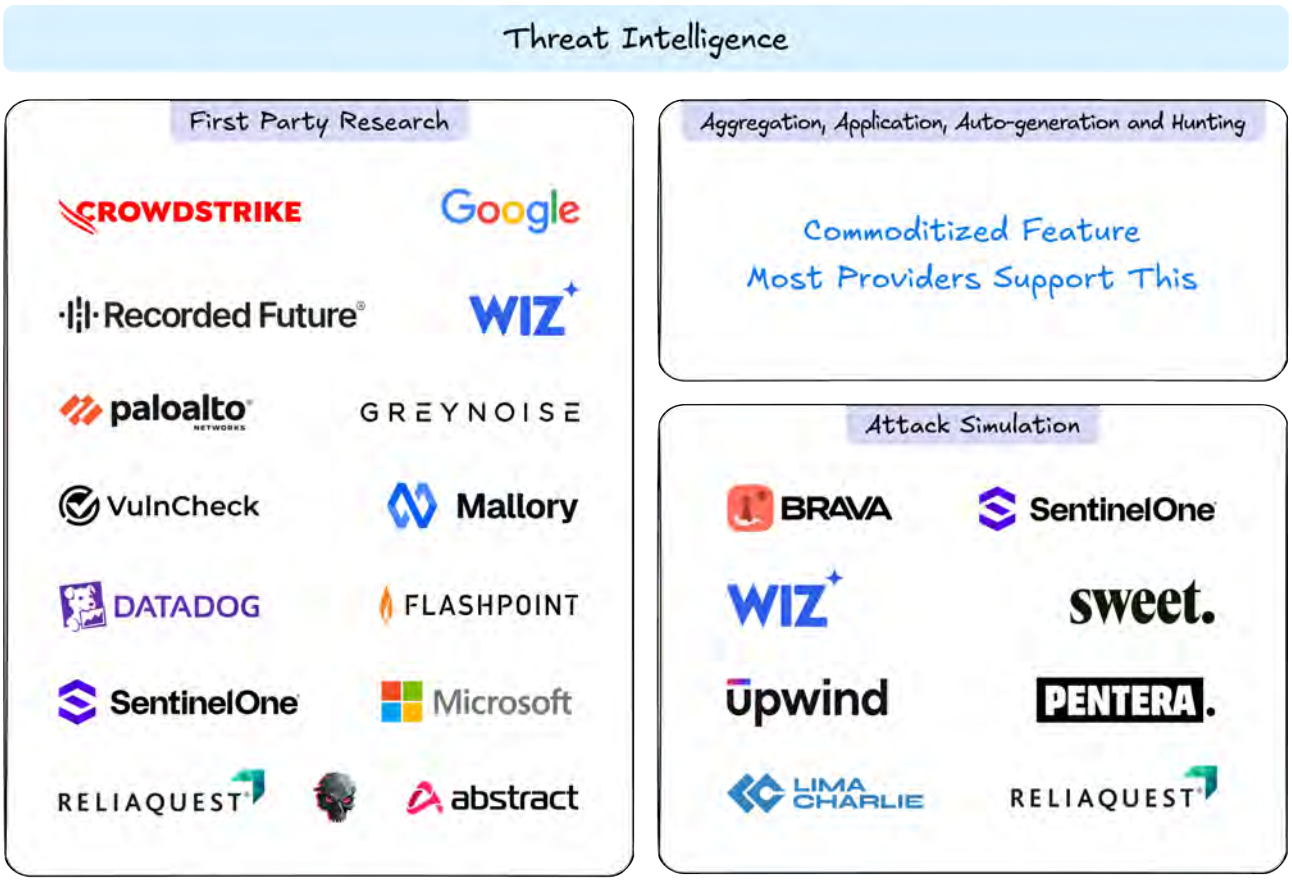

Detection engineering tools are also strong at ingesting threat intelligence feeds and converting them into detections across your infrastructure. Tools can also run threat hunts based off of the latest threat reports. Tools ingest third party news and IoCs, find the relevant logs in your stack, and then deploy searches across your infrastructure. This is a useful capability, allowing teams to run hunts and detections on emerging threats in hours instead of days or weeks.

Other tools are also innovating by bringing continuous attack simulation to the SOC, and using findings to improve prioritizations and detections. These tools tend to integrate into your cloud environment in order to understand the attack surface, before running continuous testing (of various depth) in order to check for what's exploitable.

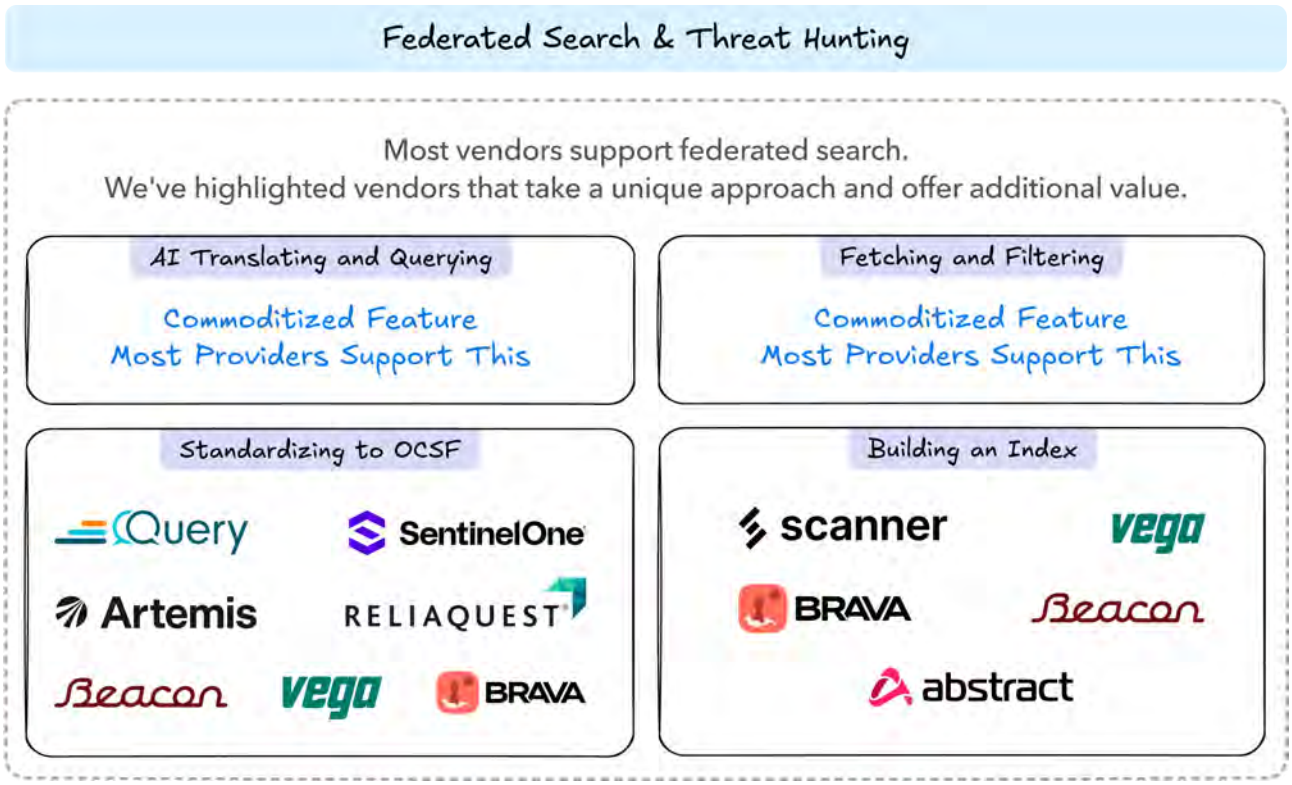

Enabling Flexible Querying of Underlying Data Sources

In order to enable AI agents to access underlying data, **federated search capabilities** have become essential for SIEM and AI SOC providers alike. Most of these tools work by either translating between vendor specific search languages and using APIs to search, or by bulk fetching the data and then filtering it on their side.

Both of these approaches work for most use cases, but have caveats teams should consider:

1. Hitting rate or data size limits on underlying providers
2. Inconsistent search speed due to throttling by the underlying data architecture
3. Slow or inability to query improperly structured data
4. Agent hallucination when data does not exist, or it's unknown if it does exist.

There are two unique approaches to solving this problem that can bring benefits to some teams: standardizing underlying data to OCSF, and building an index on ingest time. Standardization to OCSF allows agents to guarantee whether or not a specific type of data exists; however, OCSF support is notoriously limited and challenging to implement, so it may not be supported in all use cases.

Building an index on ingestion time, or in Scanner.dev's unique case, for historical data, allows teams to gain robust querying capabilities on otherwise unstructured data. This enables agents to query underlying data with more flexibility, making these solutions ideal for building robust searches.

Especially as teams look to use AI for running continuous threat hunts and investigations, the ability to quickly and reliably search data at scale has become more important than ever. Without a clean data architecture, teams will be unable to take advantage of the agentic capabilities outlined in the next section.

Security Orchestration, Automation, and Response (SOAR)

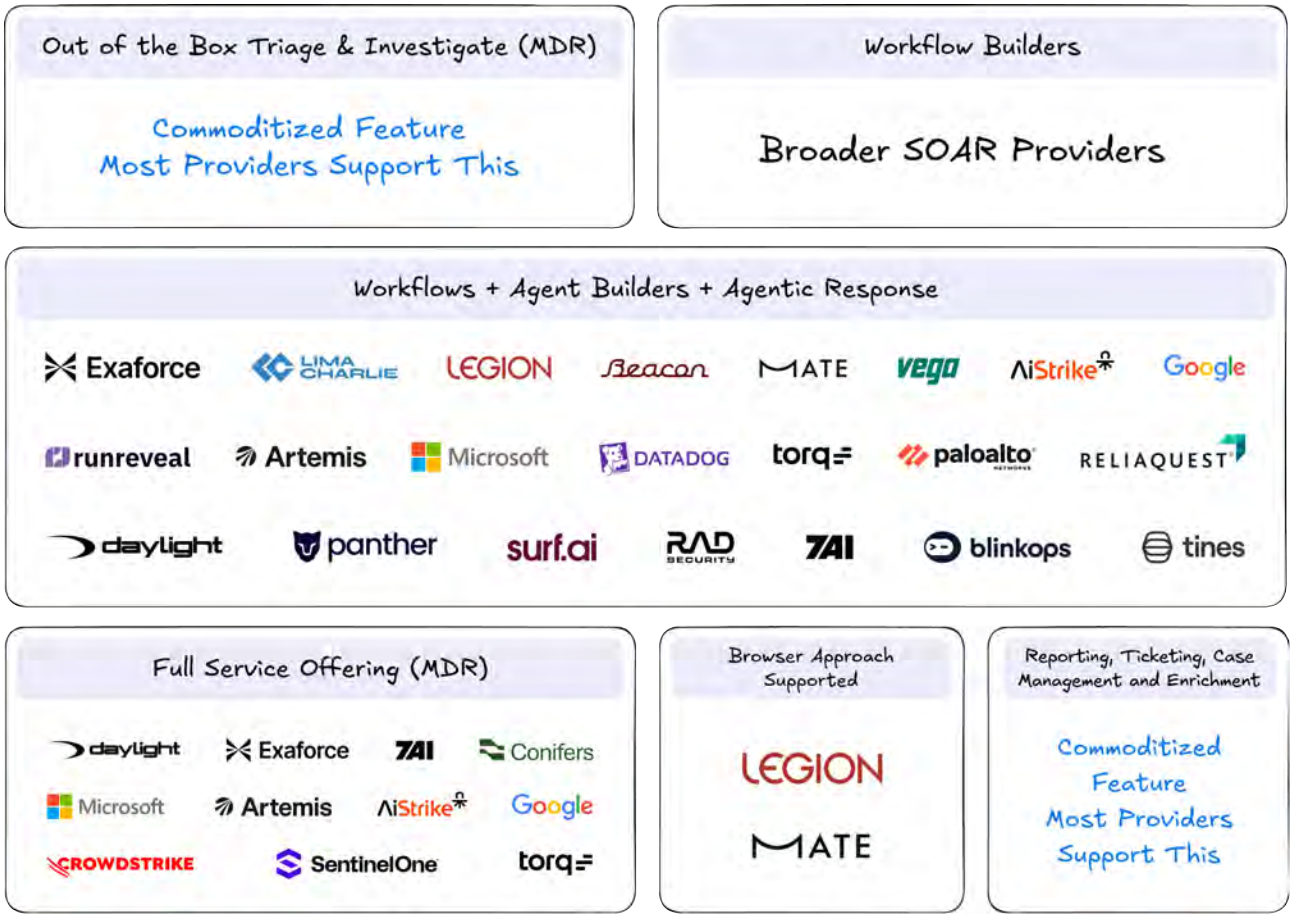

SOAR platforms have most radically evolved to incorporate AI, in no small amount due to the success of AI workflows in general. Tools like n8n and Claude Cowork have shown the value of orchestrating complex workflows with AI picking up properly scoped tasks in between.

The goal of SOAR is simple: make incident response faster and more effective. For the purposes of this report, we've chosen to bucket several disparate categories into SOAR, because they're quickly merging together to support incident response work:

- Traditional SOAR as workflow builders, allowing complete flexibility at the cost of maintenance overhead.
- Case management systems, traditionally a niche field in the SOC growing in importance as numerous tools compete to be the new analyst dashboard - most providers have case management systems built into their platform.
- Managed Detection Response (MDR) has evolved to incorporate software systems which aim to automate their in house prioritization and response work. The line between these tools and what's traditionally referred to as "AI SOC" is usually only 24/7 staffing.
- AI SOC tools, which run agent workflows under the hood and operate almost indistinguishably from SOAR, but benefit from the added flexibility of AI, oftentimes at the cost of a loss of control.

The story of SOAR has always fundamentally been about flexibility tradeoffs. On the one hand, total customization engines provide the ability to accomplish almost any action, but at the cost of maintenance, custom scripting, and deep integrations. On the other hand, many AI tools let AI take the wheel, but at the cost of auditability and repeatability that enterprises require.

The best "AI SOC" tools are those that mold to a user's business context out of the box, while enabling complete customization after the initial integration. Many AI SOC tools fail to scale well, as they can show an impressive demo, but fall over when tested in complex use cases, or when the proper data isn't available to them.

The next year of SOAR's development will be shaped by incumbent vendors adding AI features in order to incorporate the flexibility, and AI SOC vendors adding customization in order to support better repeatability.

The Future of Security Operations

The Future of AI in SOC

Start-Up AI SOC Market Map

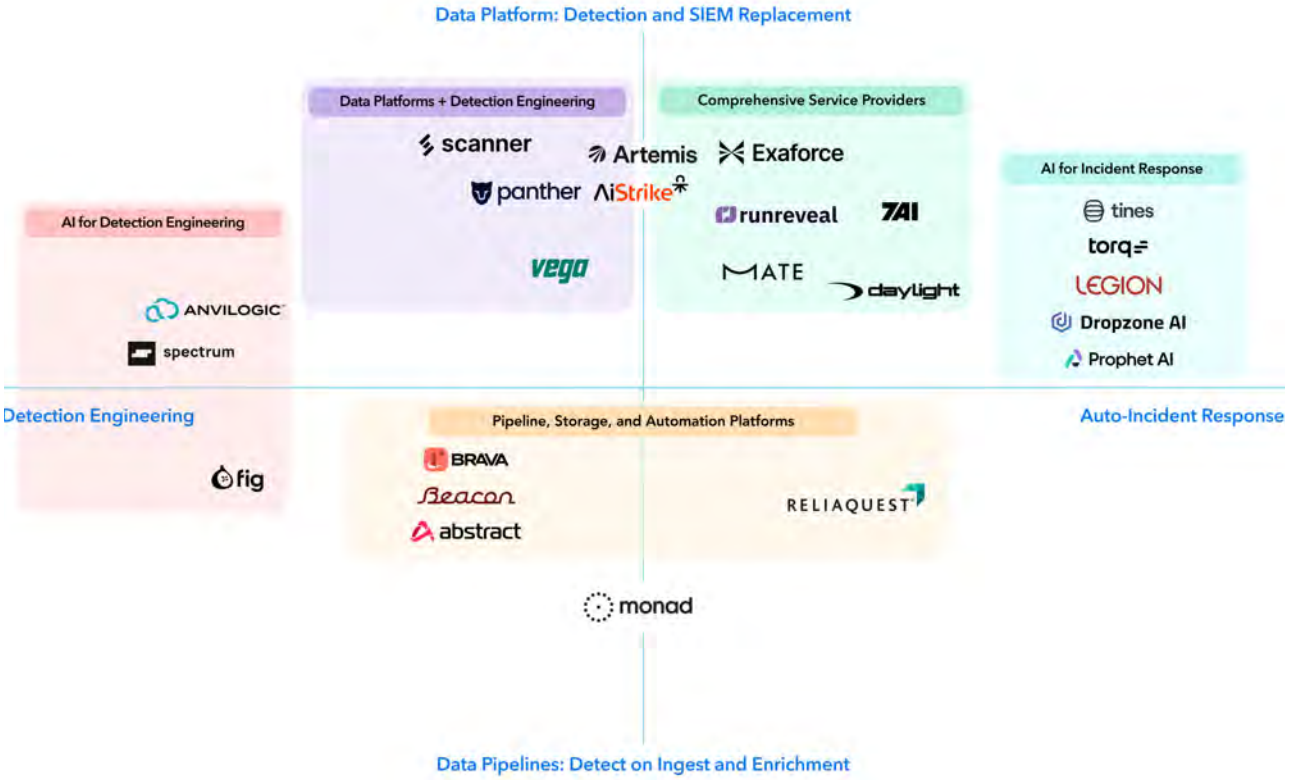

Unfortunately in 2026, "AI SOC" has become how every vendor positions themselves, leading to more confusion than ever. We've created a map that breaks down the market by answering two questions:

- Is this more of a data platform, or a data pipeline provider?
- Is this product focusing more on managing detections, or incident response?

In the long run, there is a convergence happening in this market, with leading platforms offering several capabilities:

- AI based incident response prioritization and investigation
- Detection management, engineering, and health improvement across your existing tools.
- Direct log ingestion, and support for third party data via bring your own bucket models
- Log routing and enrichment, creating ways for customers to more easily get their data into the platforms.

Companies looking to improve their incident response processes and looking for a pure "MDR as alert prioritization and 24/7 SOC" would benefit most from those in the top right quadrant, as these providers focus on automating incident response, and sometimes provide the ability to directly ingest data.

Enterprises looking to solve complex data problems, or orchestrate multiple SIEMs would benefit most of the companies in the top left quadrant. These tools have ways to deal with data from multiple SIEMs, and can function as a detection orchestration tool across the top of those existing solutions.

Finally, there are a wide range of benefits of the tools in the bottom quadrant, which focus on data ingestion and enrichment. These tools help monitor the health of incoming logs and route them to the right place. These tend to be built for larger enterprises or MDRs in order to provide optimized routing to reduce storage costs, but provide other benefits as well.

SIEM Expectations are Changing

SIEM Portability Graph

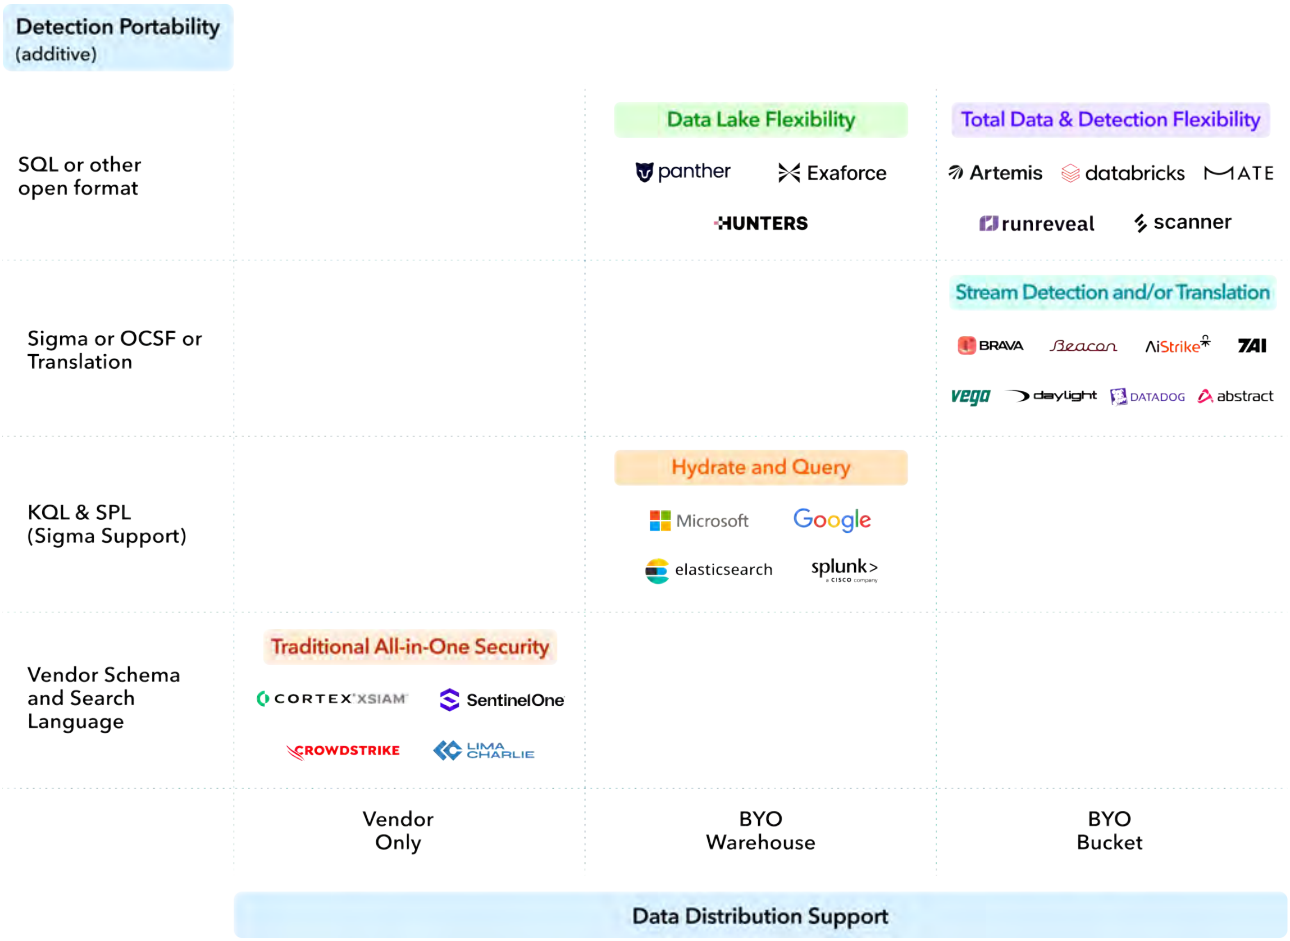

The major trend of the next two years will be the underlying data architecture of SIEMs continuing to evolve. Elasticsearch style tools were the backbone of the last ten years, building and deploying complex infrastructure for rapid log indexing and searching, with long term storage being added in as an afterthought. Long term storage searching has always been painful, as nodes need to be rehydrated with historical data in order to enable faster search.

Now, several emerging players are incorporating new data models, or building their own, in order to bring the couple the speed of the elastic stack with the data flexibility modern enterprises demand. Alongside flexibility in the underlying data stack, teams are valuing the flexibility and portability of their underlying detections. As teams build vast libraries of carefully tuned detection libraries, they don't want to be locked into vendors, usually without easy ways to maintain detection as code.

On the data side, vendors are allowing greater flexibility in where data is stored. Most vendors now offer ways to bring your own data warehouse instead of requiring you to use theirs. The real innovation is happening in creating ways for companies to search even their unoptimized historical logs in S3, which require creative ways of indexing unreliable data.

On the detection side, KQL has become a default query language for most of the industry when they're not ready to standardize into sigma rules, which can convert into any vendor specific search language. Beyond that, many newer data platforms are instead built for simple SQL queries, enabling teams to have completely flexible detections on a more common language across teams.

Innovative Capabilities

Detection on Stream

As data pipeline tooling becomes commonplace, many vendors choose to operate most of their detection capabilities as data enters the platform. Data pipeline tools are becoming more robust by adding in detection and enrichment capabilities, alongside building an index for data flexibility.

Building an index on stream can become important for a successful SIEM migration by enabling teams to store their data in cheaper locations. Detections can primarily fire on log ingestion, or the index can speed up search performance on object based storage.

Identity Baselining

Tracing a user's identity has become more important than ever for security operations, as teams need a reliable way to discern which groups of users are authorized for which actions, and what roles they have access to. Several vendors are baselining user behavior from their IDP to through to the endpoint in order to better contextualize alerting. This identity correlation improves most capabilities of the SOC - from creating log enrichment to allowing agents to investigate alerts more clearly.

Co-Pilots

While most AI SOC solutions are architected with an API first model, several providers also provide a browser extension to better empower existing analysts with workflows. These tools can record common patterns of response, and suggest in depth analysis or response actions without needing to pivot into an entirely separate tool. This approach is a great way to speed up incident response capabilities without spinning up an entirely separate platform.

Compression and Search

Many SOC tools are only as powerful as the underlying data, and the largest challenge for security teams can be making sense of historical data that hasn't been properly indexed. Scanner is the only solution in this category that's built for optimizing the queryability of unindexed historical data. This allows teams to show up with historical data and make it directly usable by the overall infrastructure.

Detection and Pipeline Health

While there are plenty of tools for configuring data pipelines and detection rules, these systems are typically disjointed. This creates a challenge when managing drift between systems - understanding what detection rules are still configured properly, what log sources are still sending, and if they have the correct underlying data to detect attacks.

Fig stands out for correlating this data together by sampling your ingested logs against your detection rules in order to reveal blind spots and improve the overall health of security operations teams. These capabilities focus on improving the health of team's existing tools.

Attack Simulation

Attack simulation has always been an underrated feature for SOC tools - from those like Pentera, which run full internal penetration tests, to emerging AI Pentesting tools which can attempt to exploit underlying infrastructure from outside. Brava has taken a unique approach, using AI to generate continuous exploit simulations, and using discoveries to optimize log ingestion, storage, and detection. This approach brings many potential benefits, from reducing hot storage of unnecessary data, to discovering blind spots in your detection logic.

Buyer's Guide

The Guide to Modernizing a SOC Program

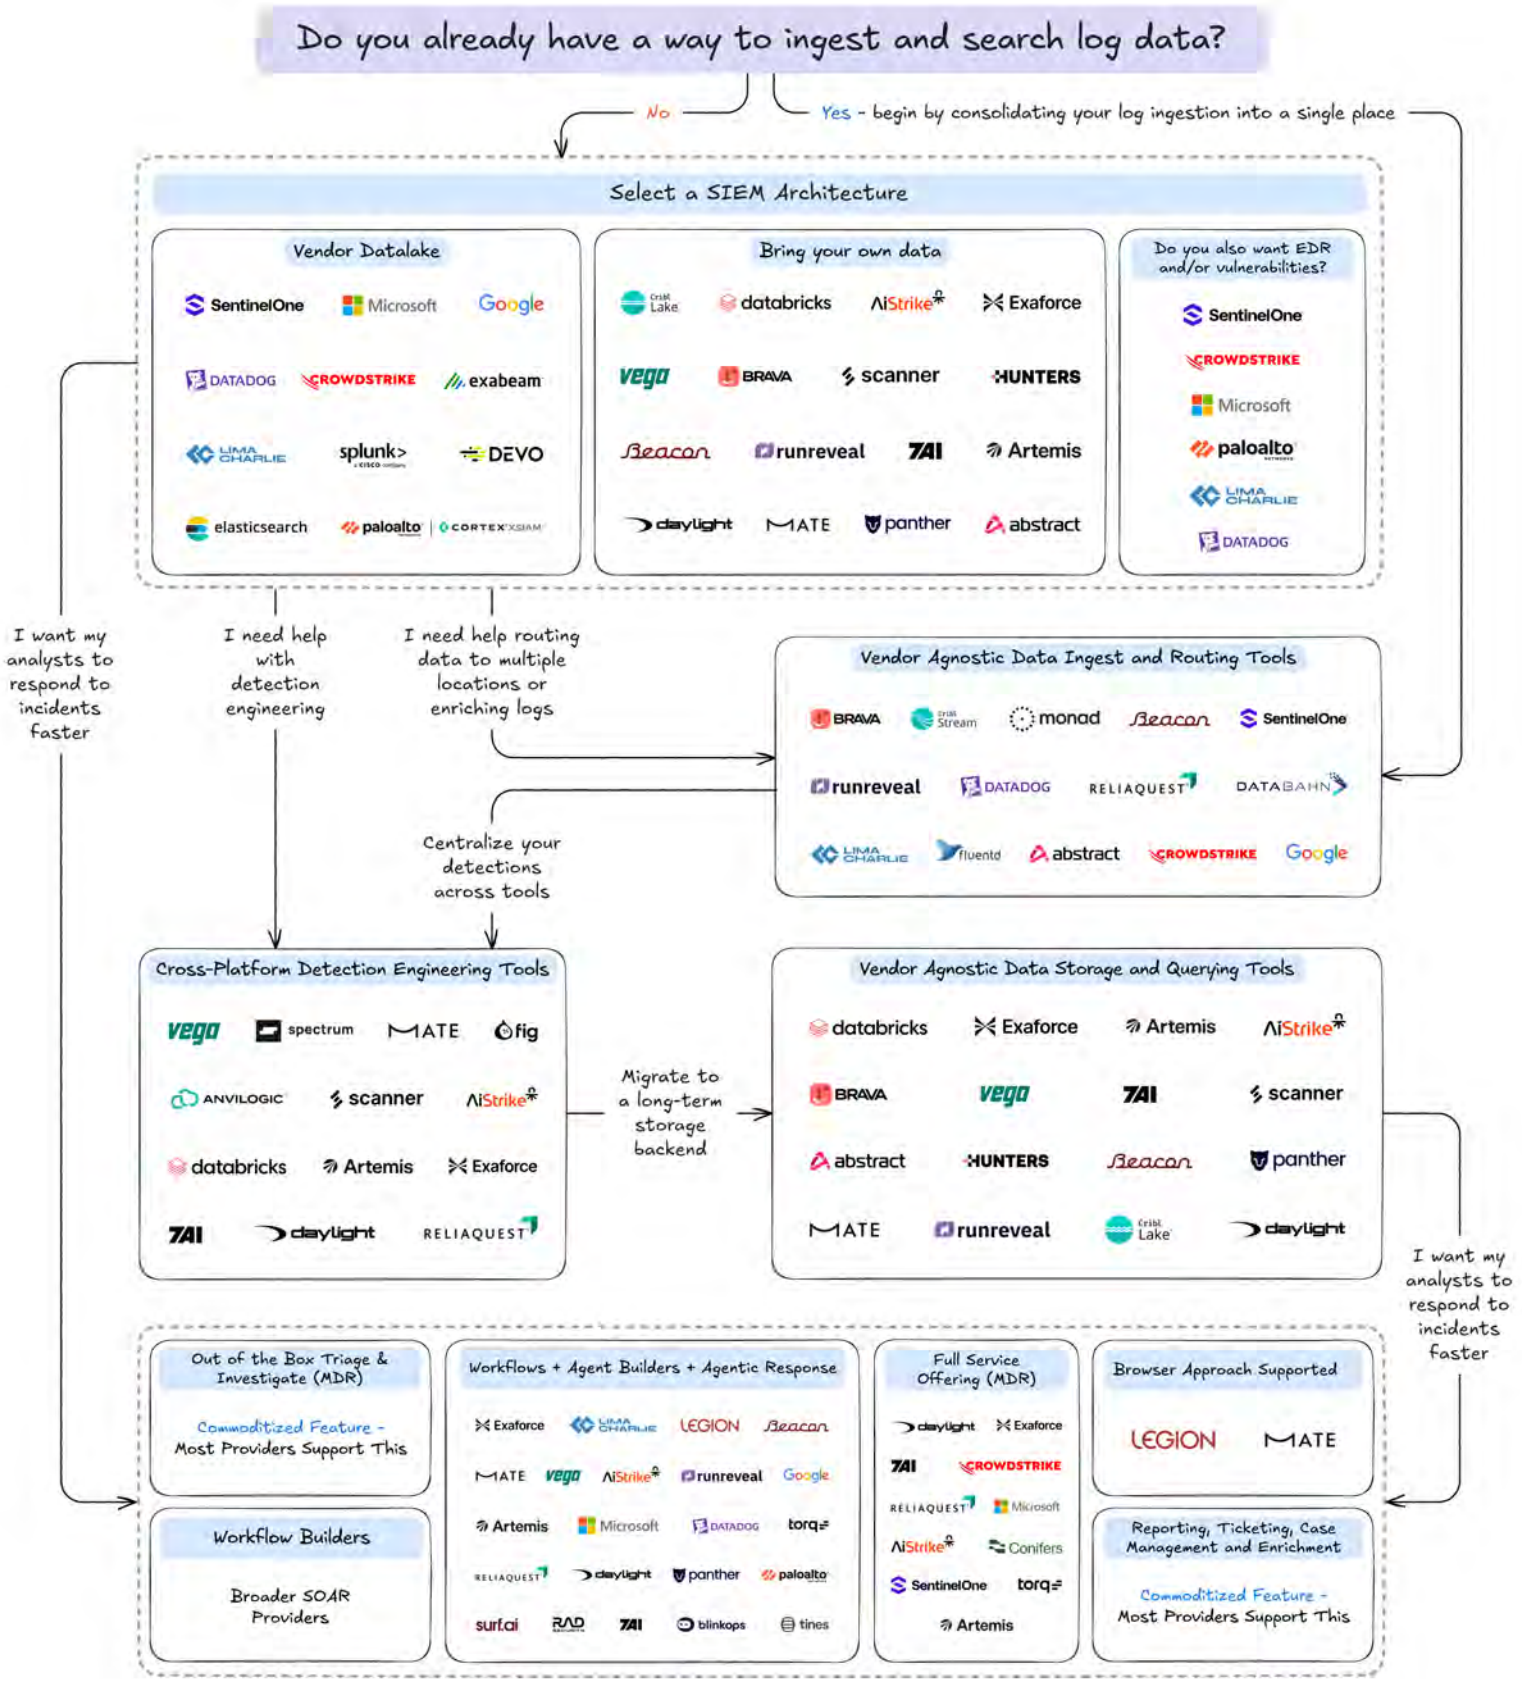

When to Buy a SIEM

The first step to building a security operations program is having a place to store and query logs. It's worth noting that many cloud native organizations start without this functionality, instead investing in a combination of vulnerability and runtime protection features (CNAPPs). However, as organizations grow, the need to correlate and search logs in order to respond to security incidents across various tools becomes essential. The need to purchase and integrate an EDR for employee endpoints can also drive considerations for a SOC tool.

When making a SIEM buying decision in 2026, the first question to be answered is "how much do you value flexibility versus single platforms?" Vendor specific data lakes provide an out of the box experience that tends to work particularly well for companies that already have large commitments to the underlying platforms. For example, Microsoft Sentinel and SentinelOne each provide obvious value as SIEMs depending on a company's underlying commitment to those vendors.

Next, companies need to decide if they value their SIEM including an EDR offering. In my opinion, these should be evaluated separately, as EDR technologies don't immediately benefit from SIEM architectures; however, many teams value having a single overall security vendor.

A final consideration is if you plan on consuming security operations as a managed service. When it comes to managed service approaches, providers generally either support a bring your own stack model, or managed versions of their services. Bring your own stack services offer greater flexibility in providers and pricing, but often at the cost of reduced support capabilities or extensive pass through costs.

The Start-Up AI SOC Market

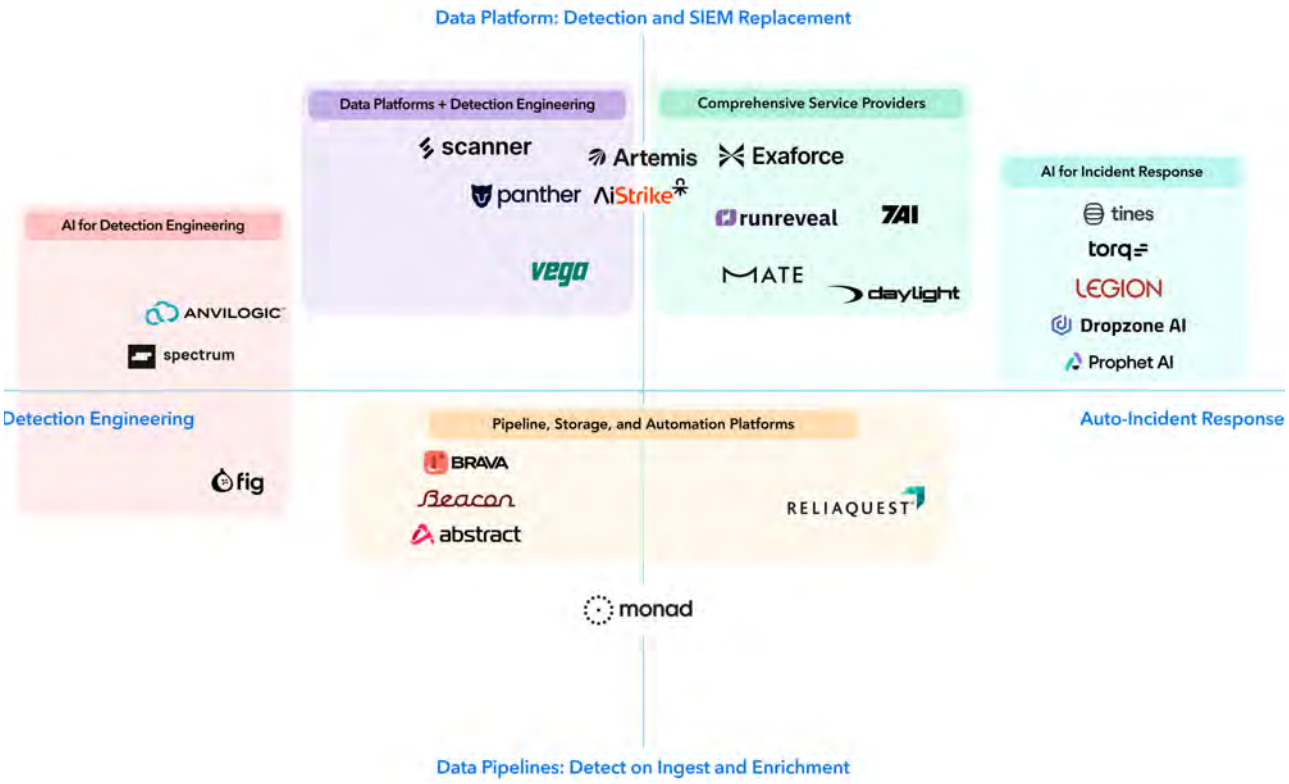

The managed service approach you choose ends up guiding your buying decision on "AI SOC." If you're looking for last mile support in 24/7 monitoring, prioritization, and incident response expertise, providers on the right side of the spectrum make the most sense. However, if you're looking for a "SOC in a box" which includes all of the functionalities needed from a single vendor alongside the prioritization and incident response capabilities, the data platforms will be better options. Teams looking to solve long detection tuning times or blind spots will benefit from vendors in the bottom left quadrants.

Modernizing Your Data Architecture

As data platforms and pricing evolve, many enterprises are slowly moving towards a distributed SIEM architecture. For most enterprises, multiple teams have an interest in their data lake architecture - from data teams, to business analytics, to developers, and security. Trends in data management therefore trickle down into security, as key data storage architectures impact their access to necessary telemetry.

While migration plans are always an exciting endeavor, too often they're left half completed. This has led many organizations to have multiple SIEM strategies - sometimes querying historical business data in one location, separate from their core security logs. This creates a nightmare version of the distributed SOC - where no one knows where the data lives, if it's formatted correctly, and how to access it. Rather than fixing this problem, AI makes it worse, as agents need consistent access to reliable data in order to reason properly about it, lest they make wrong assumptions about the data.

In order to mitigate this, we recommend a three step process for a successful migration:

- Gain complete visibility into your security telemetry and where it flows through a data pipeline provider, or manual tracking. Standardizing into a format like OCSF can be great for making the underlying data queryable.
- Consolidate your detection logic into a single authoritative location, whether a detection engineering solution or manual tracking.
- Migrate data to the new destination, preferring those that support open architectures.

Finally, the most challenging decision for many security teams will be if they should invest in an "AI SOC" solution. This depends on an organization's expectations. On the one hand, many solutions can make analyst life easier and more repeatable, helping to reduce alert fatigue and the time to tune detection alerts. On the other hand, these tools can be a black box, driving up costs with endless queries on hallucinated data. In our opinion, teams should focus on either purchasing a modern MDR, or partnering with vendors who can fix their underlying data architecture, before haphazardly pursuing AI querying - which can be an expensive side project.

Conclusion

The underlying data architecture of the SOC is being rebuilt at the same time that the response layer is being automated. These two developments are dependent on each other, agents need fast and reliable access to data.

If there's one piece of practitioner advice to take from this report, it's this: start with a plan for your data architecture. Every other improvement to the SOC depends on having enriched, properly routed, and well-formatted logs. No AI analyst automation can solve underlying data problems.

Teams can start this process with several smaller initiatives: investing in better detection management, migrating some logs to long term storage, or getting AI tooling for their SOC (or using it to replace an MDR). Buying an AI SOC tool hoping it fixes underlying data or detection problems is how teams end up with an expensive workflow engine that hallucinates over incomplete logs - driving up costs and wasting time in the process.

The good news is that the ecosystem has never been more buyer-friendly. Detection engineering tools can speed up detection efforts. Data pipeline tools have made routing logs to multiple destinations a normal procurement decision instead of a six-month engineering project. SIEM backends are competing on architecture again, not just on ingest pricing. And the AI SOC tools, despite the marketing noise around them, are evolving into true MDR replacement tools.

The teams that come out of the next two years in the strongest position will be the ones that treat their SOC as a data architecture and detection engineering problem first, and a response automation problem second. The vendors who win will be the ones who recognize that analysts sit downstream of everything else, and that fixing the analyst experience requires fixing the layers underneath it.

The SIEM isn't dead. But the way we've built SIEMs for the last twenty years is finally getting the refresh it's needed.

Vendor Spotlights

scanner

Scanner.dev is one of the few tools enabling AI through innovations in data architecture: they keep data in customer-owned object storage buckets and build the indexing and ingestion layers that makes it consistently queryable. Critically, this works on data already sitting in object storage, not just on new logs flowing in. Teams can point Scanner at years of historical CloudTrail or other logs they had given up on searching and make them queryable without re-ingesting anywhere.

The compute model is also extremely flexible and built for agentic querying. Scanner spins up Lambda workers in parallel at query time rather than running always-on search clusters, with a near real-time cache holding reference data for detections. The platform supports modern detection as code workflows and has a highly flexible query language familiar to SPL users.

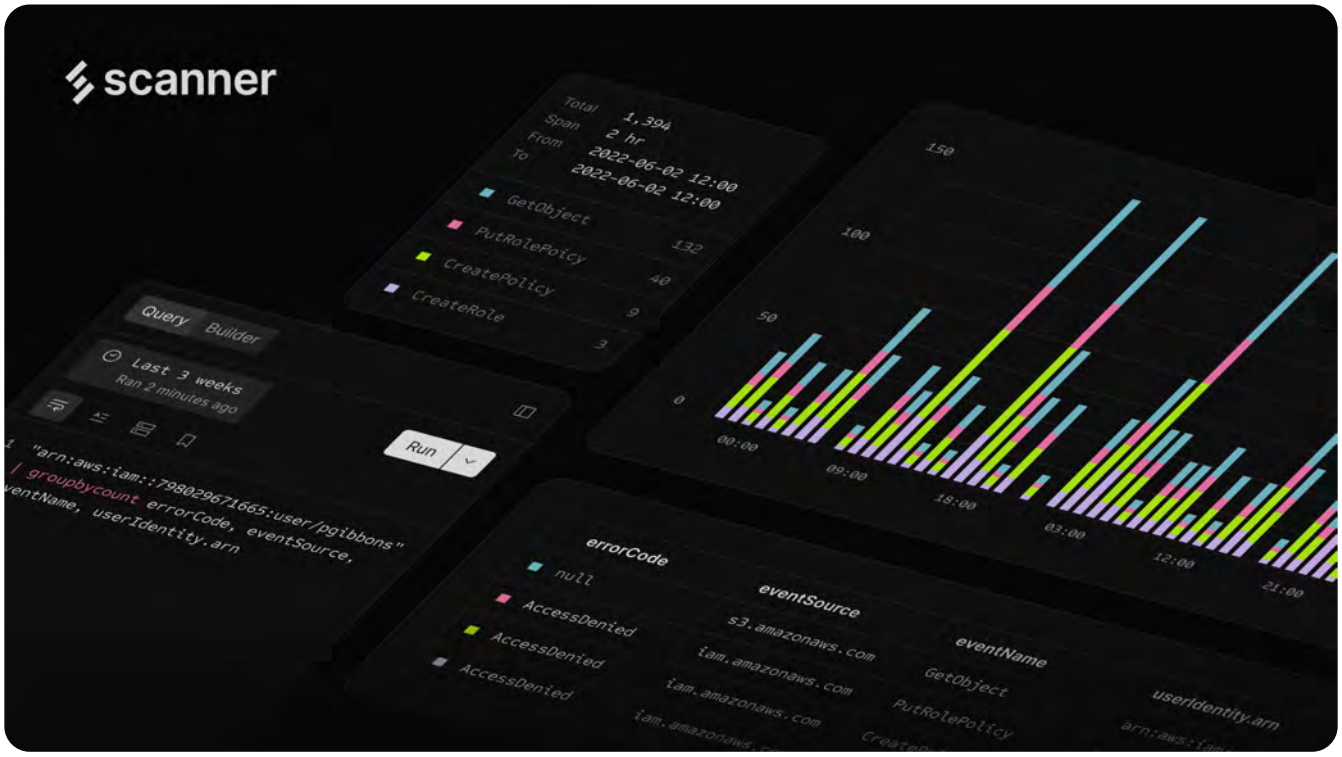

Threat intel logic layered on top of the index lets analysts pivot on indicators across years of historical data during a live investigation. Scanner also ships an MCP server that gives AI agents direct access to the data lake for investigations and detection engineering, which is one of the more practical implementations of agent-accessible security data on the market. Scanner also provides agent skills for SOC and detection engineering work that teams can use to power their own agents or Claude Code sessions.

Scanner is Best For

Teams looking to modernize their SIEM into a more flexible architecture that better enables AI queries, operationalizes large volumes of historical logs in object storage, and reduces storage costs without giving up the ability to investigate older incidents.

Exaforce

Exaforce has built everything a mid-market business needs for security into a single platform - an agentic operations tool that unites asset, vulnerability and runtime operational data into a single place. More than handling log ingestion, SIEM, and MDR services, Exaforce builds a security knowledge graph at ingestion time, linking events, identities, configurations, and cloud activity as the data lands rather than reconstructing context every time an alert fires.

Exaforce was one of the first tools to build AI SOC as more than an investigation automation tool. By giving agents direct access to underlying data, modernizing data architectures, and building complete workflow engines, they're able to deliver an end-to-end solution that offers a true one stop shop for companies looking to make their AI SOC or MDR journey simple.

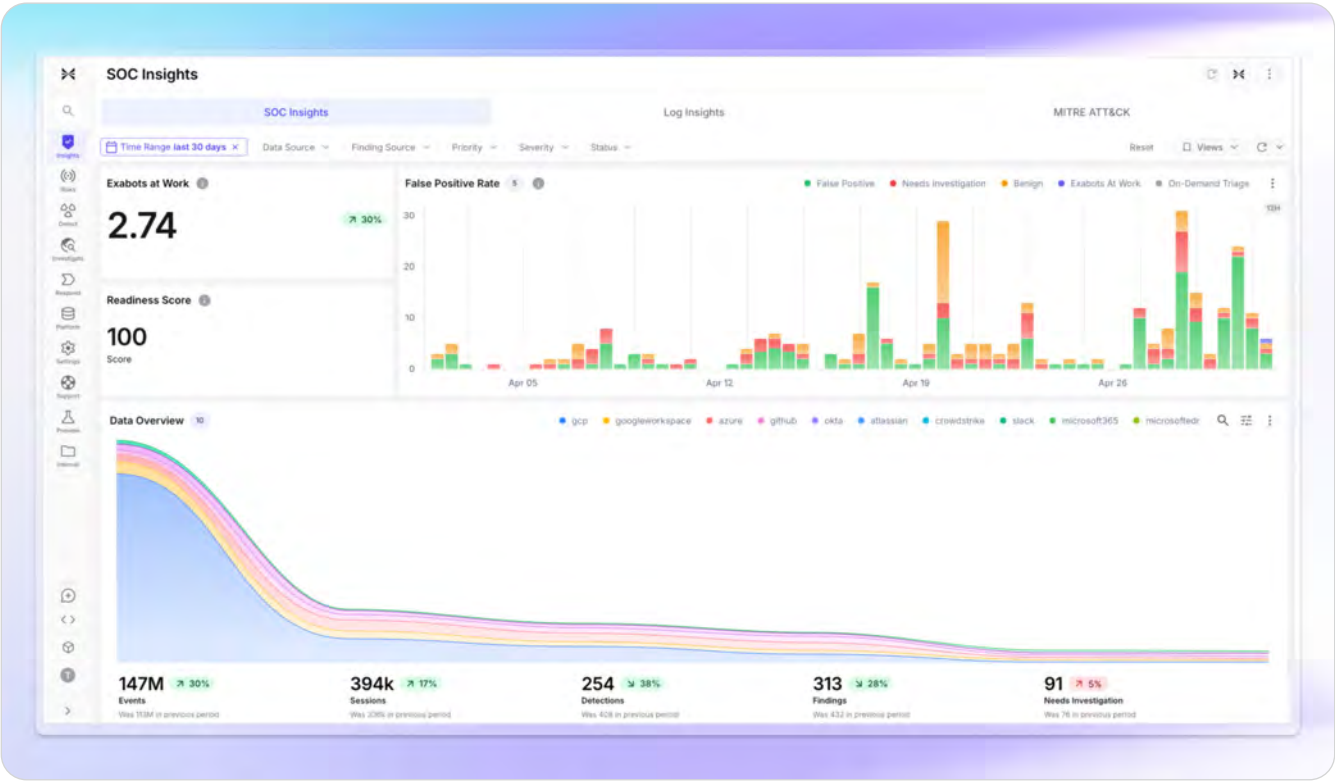

For teams that aren't ready to fully migrate, Exaforce ingests rules from existing SIEMs and runs queries directly against tools like CrowdStrike using CQL natively, with tuning suggestions surfaced as part of the platform's detection engineering layer. The platform's multi-model AI - blending semantic understanding, behavioral baselining, and LLM-based reasoning - is what drives the higher-tier work like investigation and response, and Exaforce's MDR offering wraps the platform in a managed service for teams who want the outcomes without running it themselves.

Exaforce is Best For

Teams looking for an AI SOC platform that can stitch detection, triage, investigation, and response across cloud, identity, and endpoint data without forcing a full SIEM migration upfront, and teams wanting an MDR service backed by a modern agentic platform.

Brava maximizes the value of every log through their use of attack simulation AI agents, which validate your present and future threat detection capabilities. First, their agents identify blind spots across your ecosystem by simulating various attacks and testing your detection logic. They then utilize the findings to optimize your log ingestion by reducing unnecessary logs. Finally, they provide an optimized data storage and searching layer to empower agentic incident response.

The Data Pipelines product handles ingestion, routing, indexing, and long-term storage, with support for forwarding to existing destinations like Splunk and Cribl Stream alongside Brava's own Aquarium backend - which indexes and compresses data for cost-effective retention. The Coverage product runs continuous attack simulations against your telemetry, scoring readiness against the MITRE matrix and producing concrete recommendations on what to enable, tune, or ingest to close gaps.

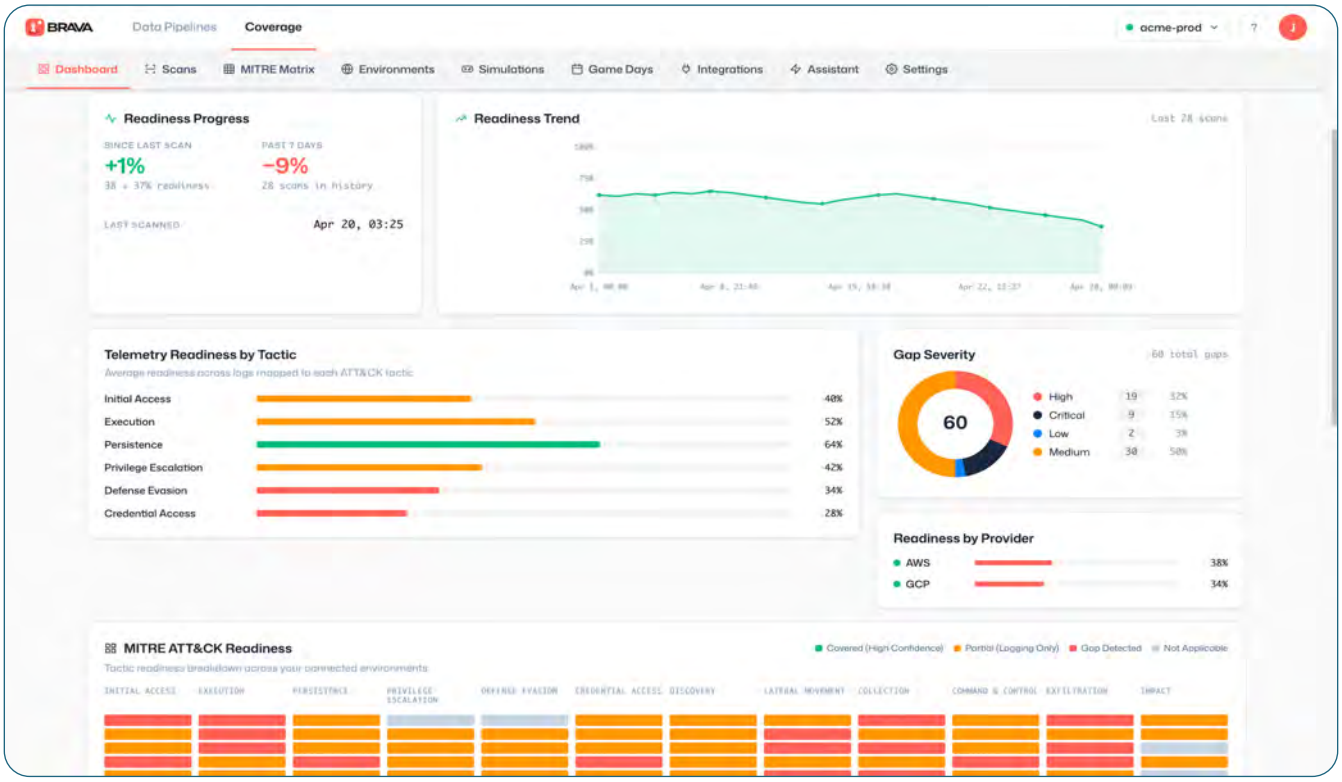

Threat hunts and investigations can run as free-text across all data sources or as goal-driven AI investigations using prebuilt templates that automatically surface the most relevant evidence across the indexes. Teams that want to keep Splunk as the analyst interface can do so with Aquarium as the underlying data layer; teams that want a single platform can run end-to-end inside Brava without giving up federated search across the rest of their stack.

Brava is Best For

Teams looking to improve detection coverage while reducing SIEM spend by pruning low-value telemetry, especially organizations that want attack simulation, data pipeline management, and detection engineering tied together rather than purchased as separate point solutions.

Daylight offers AI-native managed security operations services powered by a data platform, but with a commitment to the human aspects of being a service provider. Daylight's primary differentiation goes beyond AI or data lookup capabilities to focus on the quality of the human analysts and threat hunters behind the product. Daylight uses AI to execute detection and investigation workflows, with human experts focused on higher-order judgment, threat hunting, detection engineering, and customer-specific adaptation.

Our report survey results indicated that many teams are dissatisfied with their MDR providers, typically due to long wait times, irrelevant case escalations, or black box solutions that don't provide insights. Daylight provides customers with a direct line to their service providers, giving them confidence in tuning the tool and services directly to their needs.

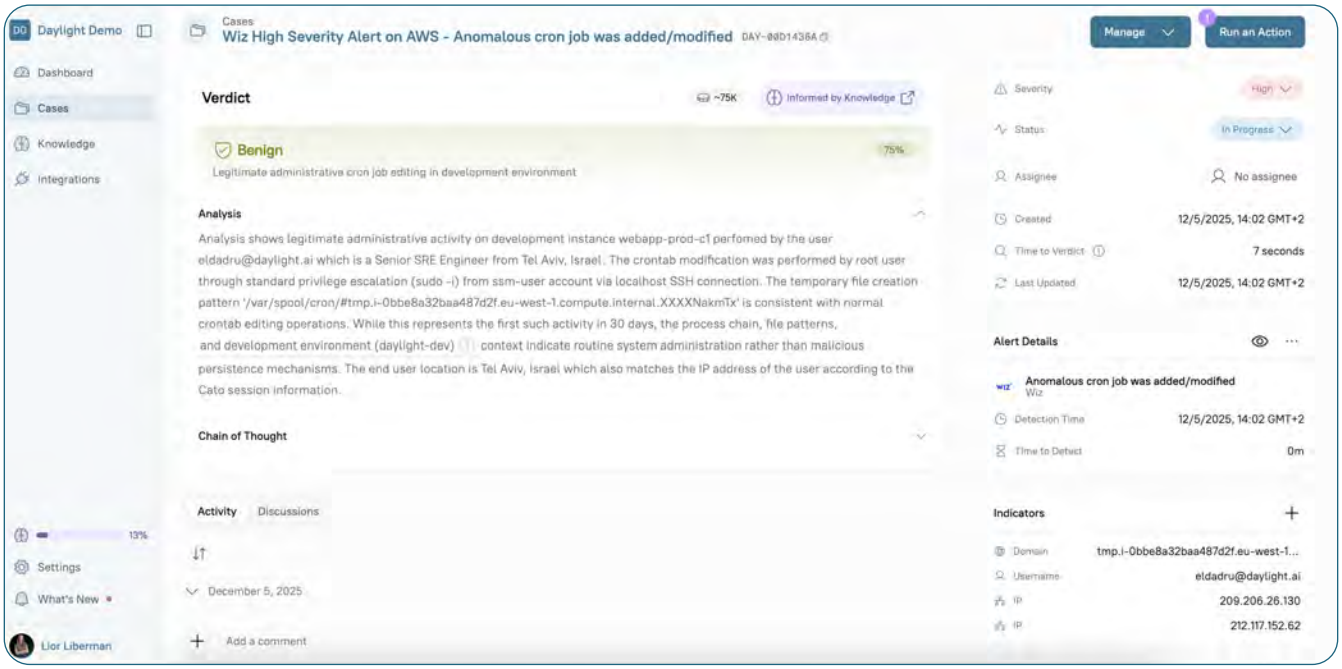

The platform is built around a flexible data layer for giving agents access to customer's security data. Customers can run detection on stream against data Daylight ingests directly, plug Daylight's agents into an existing SIEM, or use Daylight's own managed security data lake. The Integration Layer covers 140+ sources spanning EDR, identity, cloud, email, SIEM, and business tools like Slack, GitHub, and Notion, and agent permissions are configured per integration so customers keep tight control over what each automated workflow is allowed to read or change. Detection rules and continuous threat hunts feed the platform's AI investigator, which handles enrichment, evidence collection, and verdict before anything escalates to a human analyst.

Daylight is Best For

Teams looking for an MDR provider that pairs continuous agentic investigation and response with human expertise, especially organizations that want the option to bring their own SIEM or data lake without giving up the velocity of a fully managed, automation-first service.

AiStrike has built one of the most fully functional AI SOC platforms, spanning from detection engineering optimization to complete SIEM and MDR replacement. By combining robust data capabilities alongside evolving AI SOC use cases - like detection engineering, investigation and incident response - AiStrike enables teams to either augment their existing infrastructure or replace it completely.

First, AiStrike builds an asset relationship database alongside traditional SIEM data ingestion capabilities. This means that agents can know what the affected entity is, what it's connected to, and what other signals across endpoints, identities, and infrastructure are touching the same blast radius. Beyond assets, AiStrike also baselines identity and asset behavior, enabling in depth detection capabilities.

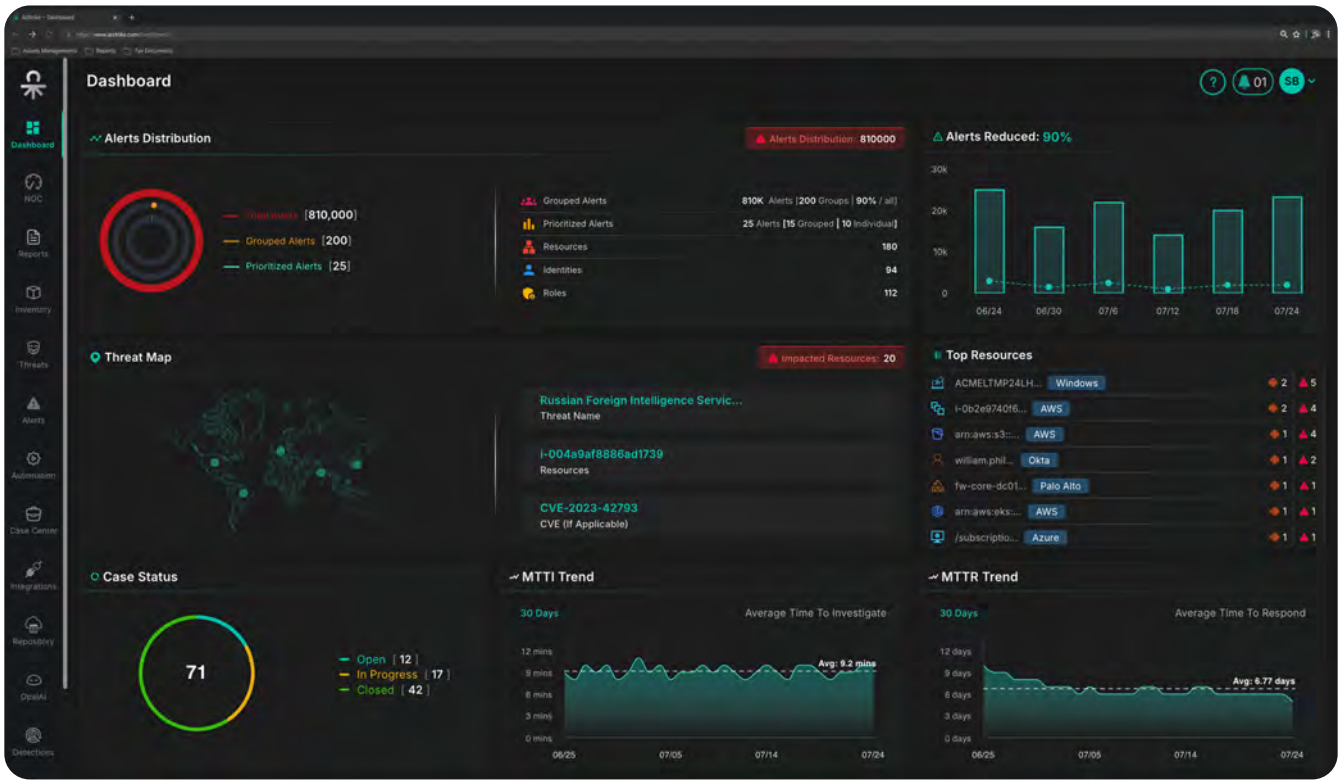

AiStrike provides several ways to improve detection capabilities. The platform grades the customer's detection program across feed quality, detection quality, MITRE coverage, threat exposure, and efficacy, then surfaces specific tuning recommendations, including AI rule tuning that flags noisy detectors and generates an optimized rule for review before deployment. Threat intelligence is integrated rather than bolted on, with automatic feed updates, threat actor tracking, and CVE prioritization driving both alert enrichment and ongoing detection consolidation as the threat landscape shifts. The result is a platform that goes beyond incident response optimization to include threat hunting and detection engineering as well.

AiStrike is Best For

Teams looking for a complete SIEM replacement, or teams who want a single layer that consolidates detections, enriches alerts with asset relationship context for investigations, and continuously grades and tunes the detection engineering and threat hunting workflows against current threat intelligence.

Artemis

Artemis is emerging from stealth with a hybrid architecture offering a complete SOC platform - able to operate as a SIEM replacement, detection optimizer, and AI incident responder. The platform makes per-query decisions on whether to ingest data directly or run a federated search against the SIEM or data lake where it already lives, optimizing for cost rather than forcing customers into a single storage commitment. For AI investigations that need to span sources outside Artemis, the platform manages the rules centrally and pushes queries out to the underlying SIEM, avoiding the typical federated search problem of slow, inconsistent results against data of unknown shape.

The detection engineering layer is where Artemis most excels. The platform supports modern detection engineering features, generates detections tuned to the customer environment, and offers a structured path for migrating rules without rewriting from scratch. Identity is baselined into user profiles that get enriched into a reference set over time, giving response agents contextual data without additional lookups.

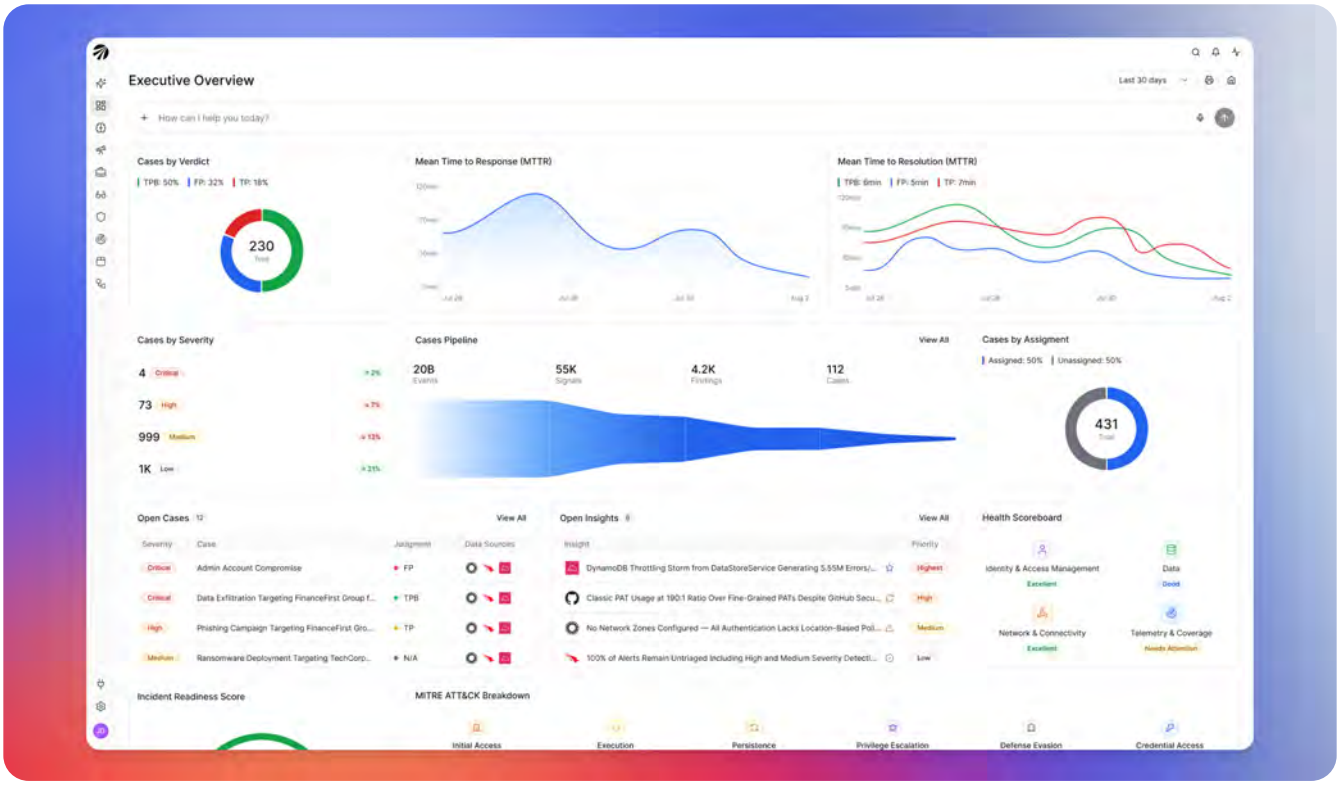

For day-to-day analyst work, Artemis prioritizes and contextualizes cases from its own detections and a variety of third party sources. An integrated intelligence feed drives ongoing threat hunts, and Artemis offers an MDR service that runs threat hunting on customers' behalf and escalates cases back into the platform for analyst review. To help the analyst respond quickly, Artemis dynamically generates response actions, reducing MTTR. These features combine to make Artemis one of the most complete tools on the market, providing immediate triage and investigation data alongside long term data architecture benefits.

Artemis is Best For

Teams looking for a single place to manage SIEM capabilities while running AI powered investigations and threat hunting across multiple log sources, especially organizations migrating off legacy SIEMs while balancing cost-aware data storage.

SentinelOne

Latio. SECURITY OPERATIONS PLATFORM LEADER 2026

SentinelOne was one of the first security providers to normalize first and third-party data into OCSF and store in a searchable data lake architecture, and the benefits are more clear than ever: giving customers a single place to manage all of their security programs, from EDR to CNAPP to AI Security. Consolidating this data enables teams to run cross-domain investigations without stitching integrations between separate products, and gives Purple AI (their agentic analyst) a unified data plane to reason across instead of guessing at relationships between siloed tools. Customers can pair the platform with Wayfinder Threat Detection & Response, which provides 24/7 managed detection and response with threat hunting directly on the Singularity Platform.

SentinelOne's capabilities continue to expand outside of EDR as they've opened the platform to provide complete SIEM, SOAR, AI Analyst and AI security capabilities. The overall platform has several meaningful differentiators from its competitors, namely vendor-agnostic data pipelines, attack simulation capabilities, and a robust AI security solution from the Prompt Security acquisition.

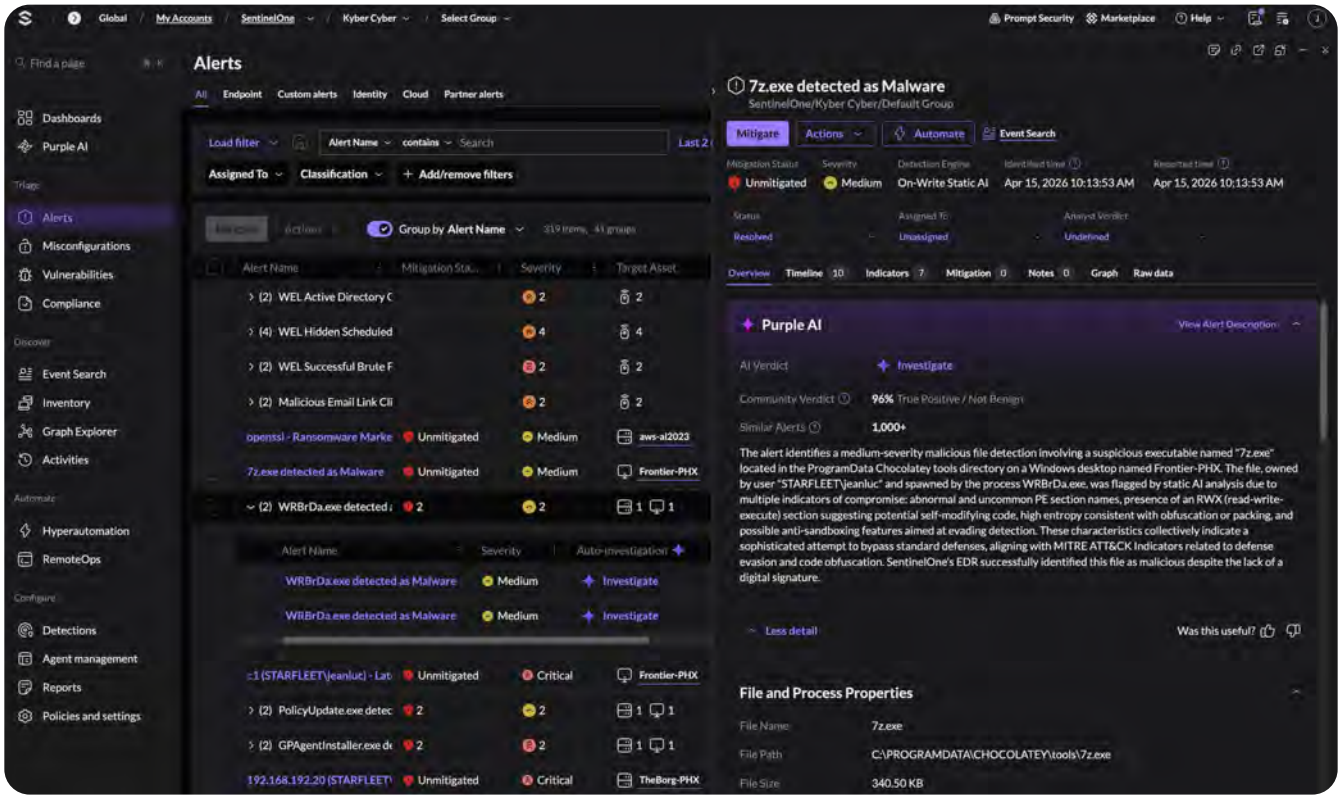

The Observo AI acquisition has been the most strategically interesting move in the lineup. Most SIEM vendors treat data pipelines as a feature bolt-on; SentinelOne is bundling pre-ingestion filtering, normalization, and enrichment directly into Singularity AI SIEM via Singularity AI Data Pipelines (formerly Observo AI). On the AI security side, the Prompt Security integration goes deeper than most platform AI security plays. SentinelOne now offers one of the most comprehensive AI Security solutions on the market, covering everything from AI Red Teaming to MCP Proxies in a single place.

SentinelOne is Best For

SentinelOne is a strong fit for organizations either looking for a complete security platform, or those with endpoint coverage looking to extend into SIEM, AI SOC, and AI security without adding vendors.

DATADOG

Latio. SECURITY OPERATIONS PLATFORM LEADER 2026

Datadog has evolved into a holistic security platform by covering developers, security operations, and workloads in a single place - consolidating customer logs, metrics, traces, and runtime telemetry. This makes it one of the few tools that can legitimately run cross-domain investigations spanning a Kubernetes pod, an Okta login, an S3 bucket, and a CloudTrail event in one tool. Datadog's unified context engine and workflow tools are also what give Bits AI Security Analyst (Datadog's agentic analyst) the flexibility to be effective, building directly on the investment Datadog has put into AI SRE workflows.

Cloud SIEM has matured into a complete SecOps platform with detection-as-code, a strong out-of-the-box rule library tuned for modern cloud architectures, and 100+ Content Packs that ship detection rules, dashboards, parsers, and SOAR workflows. The platform has several meaningful differentiators, from vendor-agnostic Observability Pipelines to robust workflow engines, to detection on ingest in addition to historical searches.

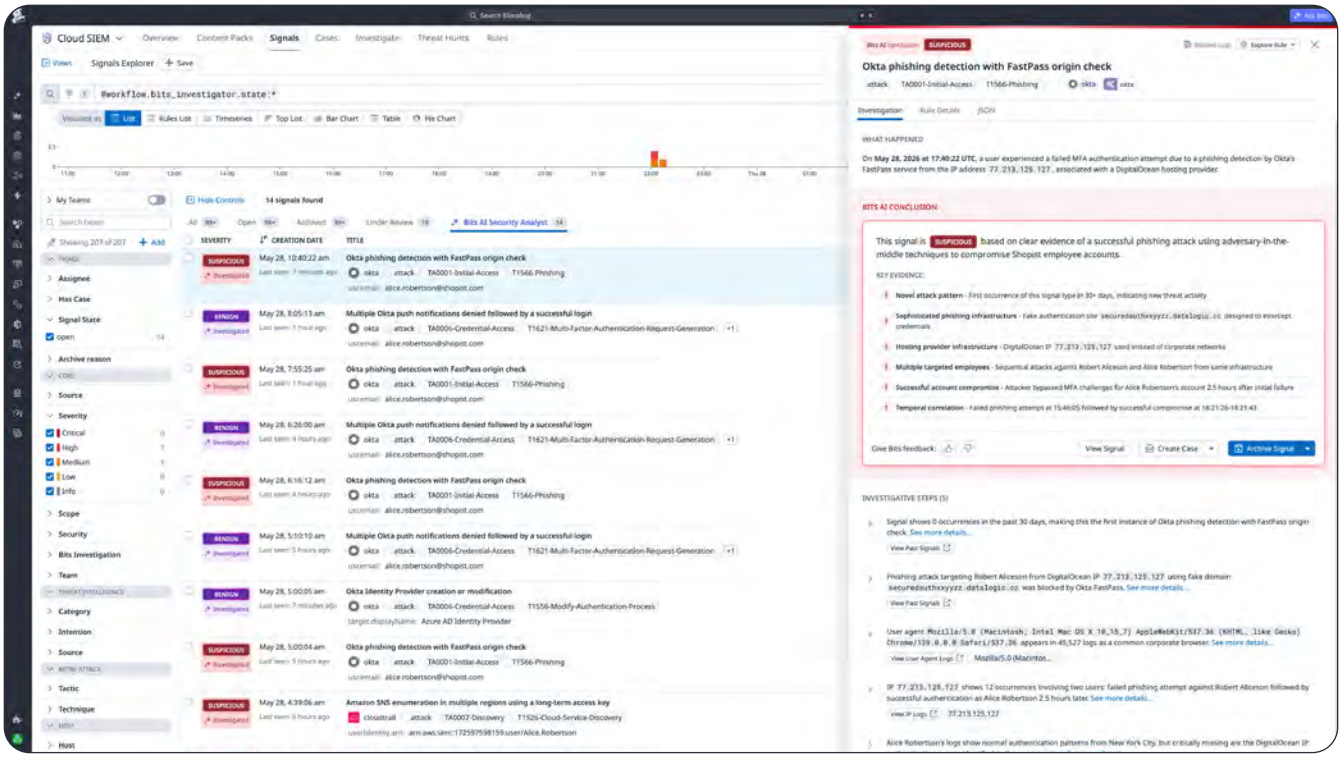

Bits AI Security Analyst and Datadog's strong out of the box tools and resources for detecting cloud native threats are the most robust aspects of its lineup. Datadog has long made strong investments in their detection engineering team, and provide several leading industry resources for cloud native detection and response. The team also offers protection capabilities that go beyond the check the box of EDR for tools like monitoring for malicious packages, VSCode extensions, and more.

Datadog is Best For

Teams already running Datadog for observability who want to consolidate developer security, cloud security, and security operations onto a single platform, and more broadly, teams looking for a holistic security offering with strong cloud-native detection coverage, flexible cost tiering, and one of the most capable agentic analysts on the market.

MATE

Mate is one of the few AI SOC platforms with a differentiated underlying architecture: their knowledge-graph is the actual foundation of the product, building a Security Context Graph from one simple integration. The graph captures context about an organization - from asset ownership to available tools - and becomes the lookup layer that every agent reasons against before taking an action. Consolidating context into a single graph lets investigation, response, and detection engineering share the same source of truth rather than each agent rebuilding context from raw signals every time an alert fires.

Mate's capabilities are organized around two tightly-coupled disciplines: agentic investigation and federated detection engineering. Agents run in a plan mode with explicit workflow checklists per investigation type, giving detection engineers fine control over how the agent reasons without forcing them to overspecify every step. Mate also supports a self-hosted deployment that runs alongside customer-owned data lakes.

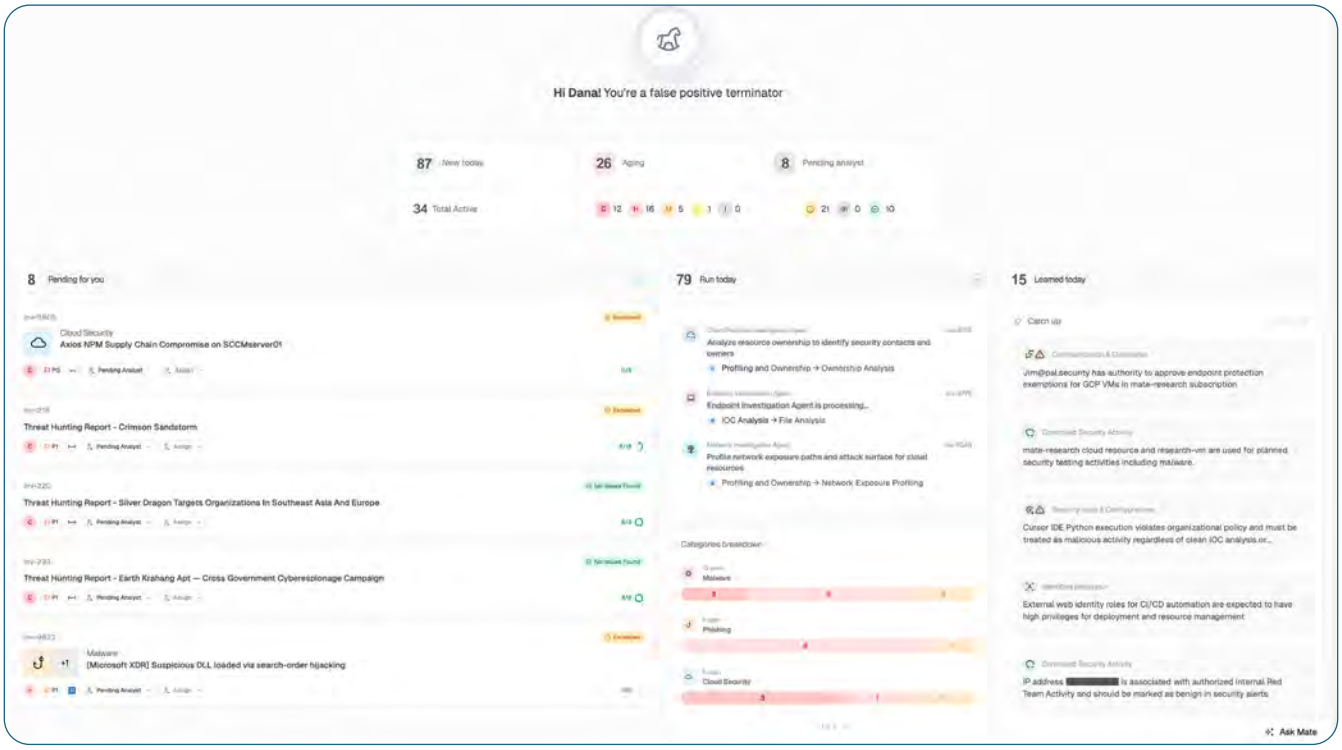

Mate's concept of continuous detection, continuous response (CD/CR) creates a feedback loop between detections and AI analysis. Threat intel can be quickly transformed into a hunt across various log sources. Detections are exported to Sigma so the logic remains portable, and the CD/CR loop feeds analyst resolutions back into the detection layer so rules tighten over time without an explicit detection engineering cycle.

Mate is Best For

Teams looking for a contextual AI SOC platform that accelerates core incident response metrics through a rich coupling of detection and analysis capabilities.

7AI has evolved more rapidly than almost any other player in the agentic SOC space, expanding from incident response automation into a broader security data platform that now shows up in nearly every section of this report. The original product was focused on agentic triage and investigation, but has grown to cover federated detection rule management, threat hunting, response, and a data lake. Creating focused response agents is one of the more practical design choices in the space - endpoint investigations and email/phishing investigations have fundamentally different evidence shapes, and giving each its own agent (with its own playbook depth, query library, and verdict criteria) is what makes the platform feel like it's actually doing full analyst work rather than just summarizing alerts.

7AI ships with a complete set of out-of-the-box connectors for standing up a SOC, such as threat intel tools, EDR, identity, email security, and SIEM. This makes 7AI an especially strong out of the box MDR competitor, able to help teams create a more full fledged security operations platform no matter the current state of their security program.

7AI is continuing to expand beyond what traditional MDR services offer with their Harden and Data Lake capabilities. Harden moves 7AI beyond pure response into proactive capabilities, from identifying missing telemetry to finding exploitable weaknesses to active threat hunting. The Data Lake module brings optimized storage, federated search across local and SaaS sources, and a knowledge graph that turns prior cases and unstructured context into a queryable layer for the agents.

7AI is Best For

Teams looking to improve their core security operations metrics through flexible investigation capabilities, or teams looking to upgrade their MDR with an all-in-one platform that can expand their capabilities.

Microsoft

DATADOG

SentinelOne

CROWDSTRIKE

paloalto NETWORKS

Artemis

AiStrike

vega

spectrum

BRAVA

monad

Beacon

MATE

daylight

LEGION

Exaforce

ZAI

scanner

Exaforce

Artemis

AiStrike

BRAVA

vega

Beacon

LIMA CHARLIE

ANVILOGIC™

torq

Microsoft

Query

Exaforce

monad

AiStrike

runreveal

vega

Latio.

Ever wonder: Am I using the right security tools for my business, or am I building the right product for the market?

Everyday companies are making decisions based on the information that is available to them, which is often incomplete and based on vibes rather than usage.

That's where Latio comes in.

Founded in 2023 by James Berthoty, Latio was built to solve a critical problem James was facing: there was no reliable, credible way to evaluate a vendor's capabilities until after an agreement was signed. Latio exists to make the buying and building processes better by getting accurate information to the most relevant teams.

We focus on the product, the practitioner, and the market rather than slides and hype cycles. We believe the greatest predictor of a great security tool and program is finding the right product fit for both vendors and buyers.

We are creating a future where every decision is based on tests, market insights, experience, and hard work, where it's easy to find the right product you're looking for.

Our mission is to help every team find the right security product. So we test every product, to make it easier for you to pick the right one.

A special thank you to everyone who has supported this mission, without you, none of this would be possible.

Learn more:

- latio.com
- Schedule a product briefing
- Schedule a security program sync
- Follow us

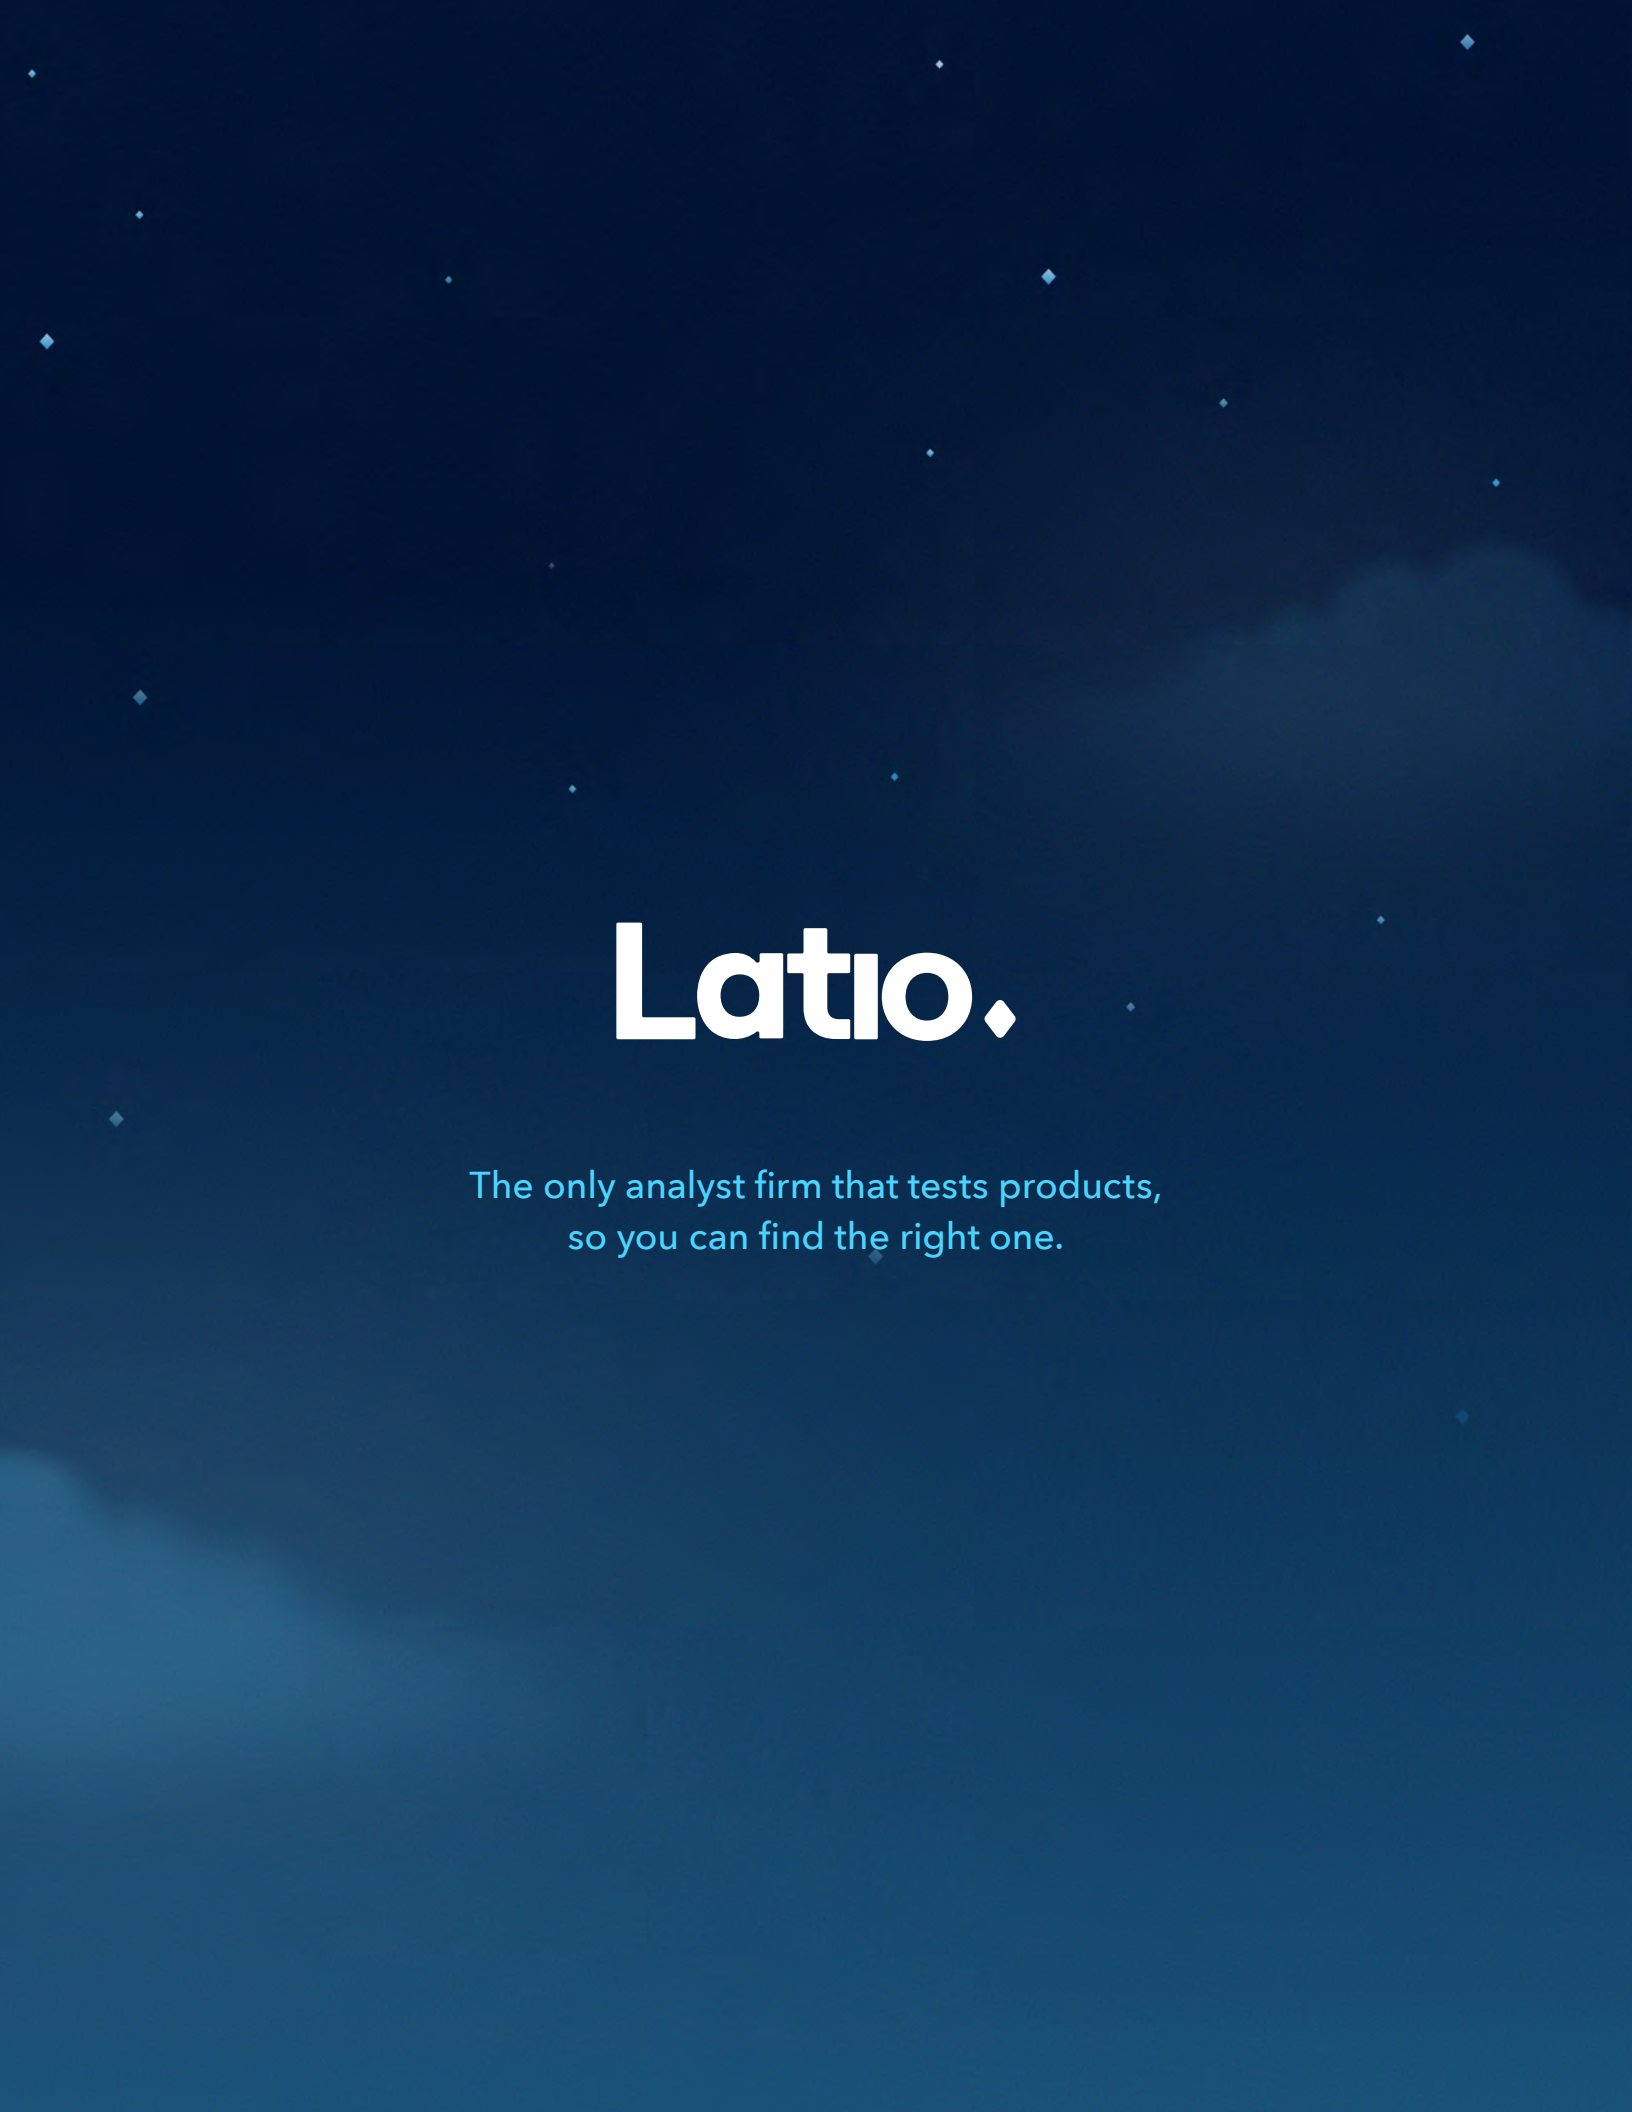
Latio.
The only analyst firm that tests products,
so you can find the right one.