sumo logic

SIEM evaluation guide

Safeguarding your future

SIEM evaluation guide

The Security Information and Event Management (SIEM) market is shifting. As organizations grow more complex and adversaries adopt new techniques, the tools we rely on for threat detection, investigation, and response must evolve as well. This guide helps you objectively assess whether your current SIEM remains the right fit for your organizational needs or if recent innovations, particularly in artificial intelligence and automation, offer a path toward a more efficient security posture.

Evaluating a SIEM shouldn't be a fire drill. The right evaluation process builds the foundation your team needs to detect faster, respond smarter, and stay ahead of what's next. This roadmap will help you gain a clear understanding of your security environment. By taking a proactive look at your capabilities now, you can build a SOC for long-term resilience rather than just meeting today's baseline requirements. In other words, going from an effective SIEM to an exceptional one.

The first step in this journey is recognizing the common operational shifts that typically prompt a closer look at a SIEM solution.

Common triggers for evaluation

Most SIEM evaluations are prompted by one of a handful of situations. If any of these sound familiar, it's worth taking a closer look.

- **Recent security incidents and pen tests** should be a wake-up call. They expose glaring vulnerabilities and gaps that organizations can no longer ignore and highlight critical capabilities you might need going forward in your solution.
- **The advancement of AI is reshaping the threat landscape** at an unprecedented pace as adversaries are leveraging AI to craft more sophisticated attacks, launch faster campaigns, and evade traditional detection methods. Your SIEM must match that reality to stay ahead of threats.
- **Regulatory requirements and changes** aren't just bureaucratic hurdles. They're crucial mandates that require robust security measures to avoid crippling penalties.
- **Growth initiatives**, whether expanding into new markets or adopting cutting-edge technologies, demand a security solution that can scale and adapt.
- **Budget cycles offer a prime opportunity** to invest wisely in security rather than continuing with outdated or inadequate solutions.
- **Industry mergers and acquisitions** force you to confront whether your SIEM can withstand the shockwaves of market consolidation and still provide robust protection.

How to use this guide

This guide draws from analyst-defined SIEM critical capabilities, Sumo Logic customer insights, and our expert analysis. This multifaceted approach ensures you comprehensively understand the SIEM market's current and future trends.

Five steps to evaluate your SIEM:

1. Are you collecting the right logs?
2. How is data transformed in your SIEM?
3. Does your SIEM offer advanced analytics?
4. Does your SIEM offer effective investigation?
5. Does your SIEM facilitate response?

What steps matter most to you based on your role

CISO

Steps one, two, and five are most relevant to you. They're key to optimizing budget and security tooling with a SIEM solution while delivering confidence, visibility, and robust protection in order to safeguard your organization's future.

SOC MANAGER

Steps three, four, and five outline the necessary advanced analytics, investigation and collaboration capabilities to efficiently detect threats and enable you and your team to continually secure your organization.

SECURITY ANALYST

Steps four and five take a closer look at how ease of use and rapid, accurate insights from your SIEM solution can speed investigations and remediation.

What generation of SIEM do you have?

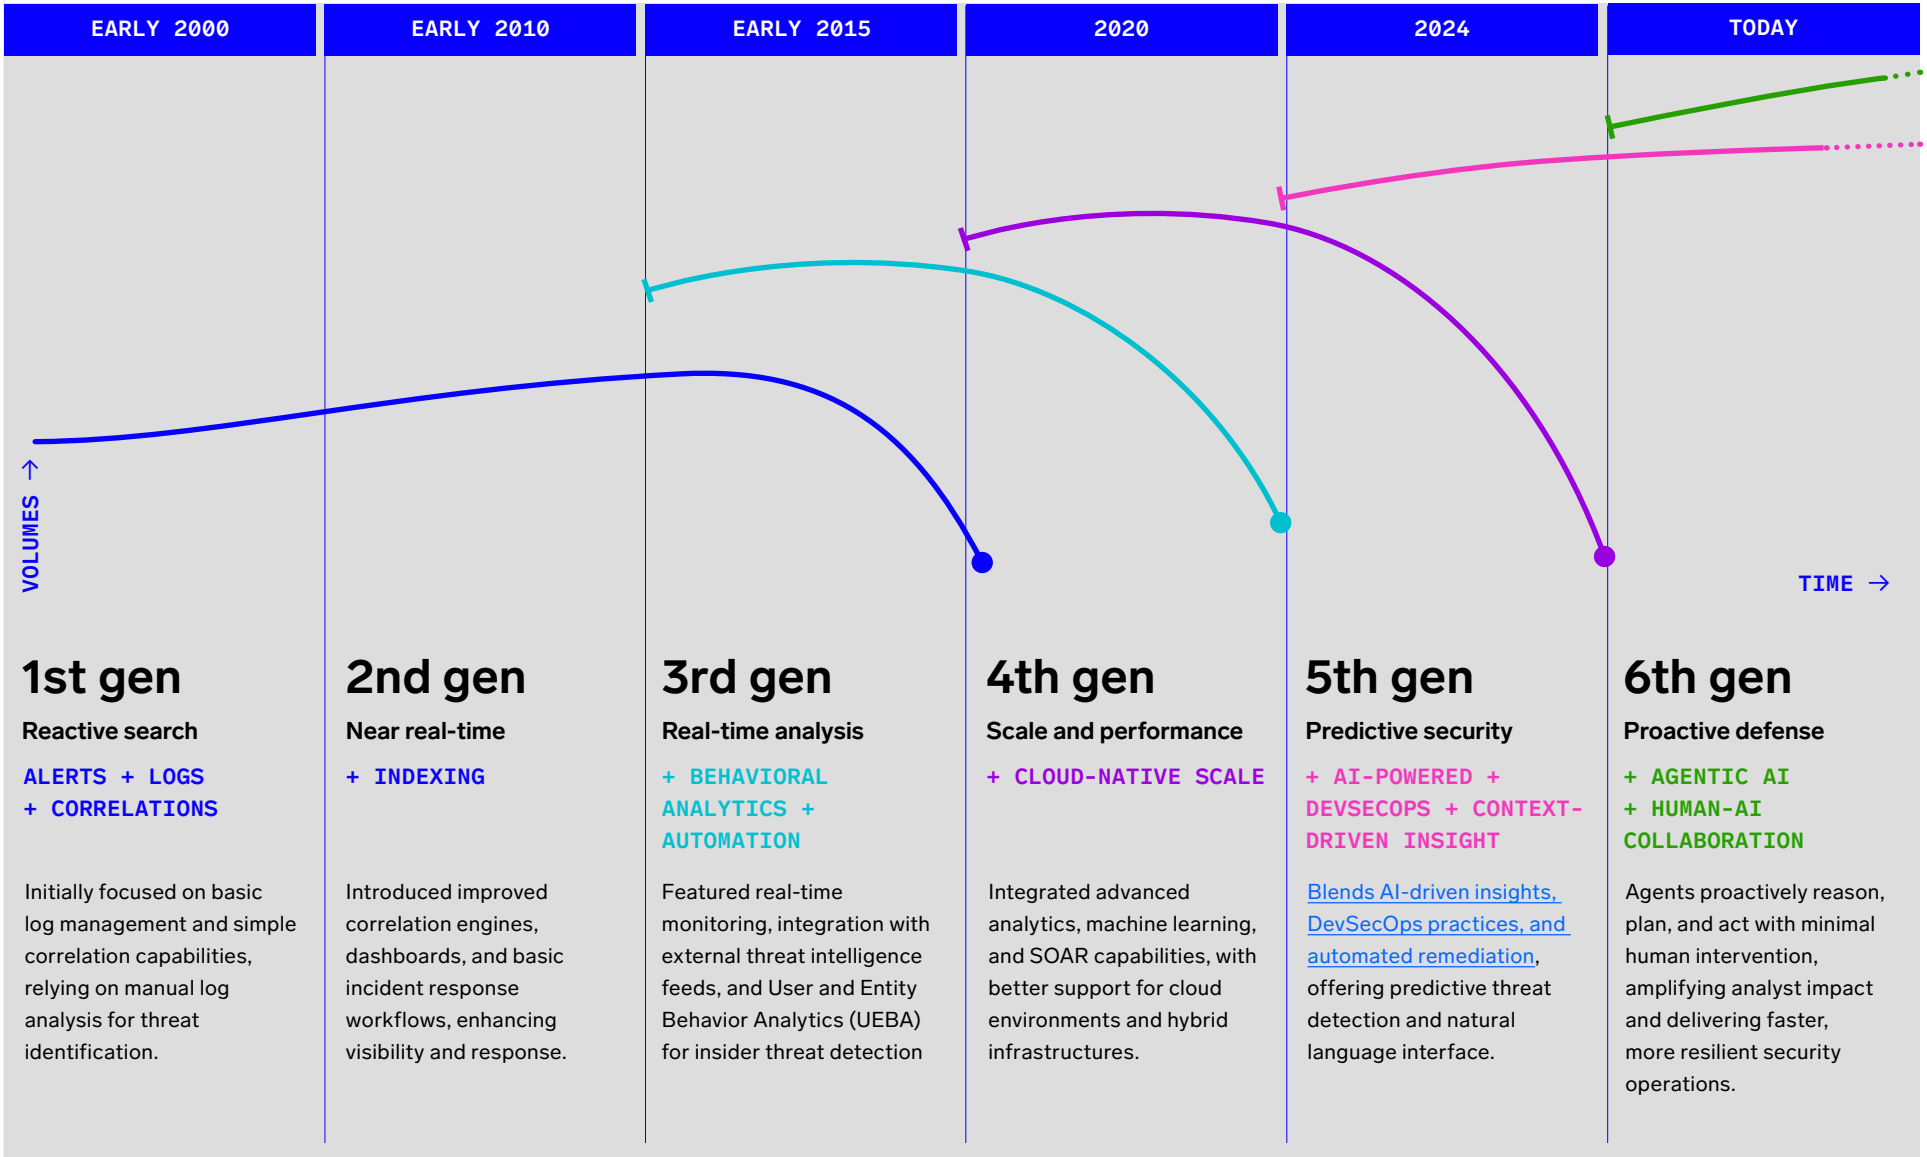

1st gen

Reactive search

ALERTS + LOGS + CORRELATIONS

Initially focused on basic log management and simple correlation capabilities, relying on manual log analysis for threat identification.

2nd gen

Near real-time

+ INDEXING

Introduced improved correlation engines, dashboards, and basic incident response workflows, enhancing visibility and response.

3rd gen

Real-time analysis

+ BEHAVIORAL ANALYTICS + AUTOMATION

Featured real-time monitoring, integration with external threat intelligence feeds, and User and Entity Behavior Analytics (UEBA) for insider threat detection

4th gen

Scale and performance

+ CLOUD-NATIVE SCALE

Integrated advanced analytics, machine learning, and SOAR capabilities, with better support for cloud environments and hybrid infrastructures.

5th gen

Predictive security

+ AI-POWERED + DEVSECOPS + CONTEXT-DRIVEN INSIGHT

Blends AI-driven insights, DevSecOps practices, and automated remediation, offering predictive threat detection and natural language interface.

6th gen

Proactive defense

+ AGENTIC AI + HUMAN-AI COLLABORATION

Agents proactively reason, plan, and act with minimal human intervention, amplifying analyst impact and delivering faster, more resilient security operations.

Sixth generation SIEM

Currently, the sixth-generation SIEM represents a leap from AI-assisted security into agentic operations. Where fifth-generation solutions provided AI-driven recommendations and predictive insights, sixth-generation SIEMs deploy purpose-built AI agents that automatically reason, plan, and act — proactively hunting threats, investigating incidents, and executing response actions with minimal human intervention.

AI continuously correlates signals across the entire attack surface and adapts to emerging attack patterns in real-time. Agents close low-risk cases autonomously, escalating only those that require human judgment. It introduces human-AI collaboration, where analysts are no longer burdened by repetitive triage but can instead focus on high-value decisions, such as proactive threat hunting and building a better security posture. The result is a reduction in alert fatigue, mean time to respond (MTTR), and operational strain caused by the ongoing security talent shortage. If your SIEM is sixth-gen ready, excellent; otherwise, you may want to seriously reevaluate whether your current solution can meet the demands of today's threat landscape.

With over ten years of log analytics leadership and over 1,000 security customers, we have created this evaluation guide as a resource to help you objectively evaluate your current SIEM solution. Using the provided criteria and associated scorecard, you will gain the comprehensive understanding needed to drive the decision process for replacing your existing SIEM.

5 steps to evaluate your SIEM

At the core of evaluating your SIEM is ensuring that it can address threat detection, investigation, and response (TDIR) challenges.

1 Are you collecting the right logs?

Logs are the most fundamental artifact of digital computing.

Effective log collection is the cornerstone of a robust SIEM solution. When investigating an incident, all critical logs need to be online, available, and part of your analysis. In complex IT environments, data are generated from an expanding range of sources: on-premises systems, cloud services, SaaS applications, network devices, your security and intrusion prevention stack, endpoints, and increasingly AI systems themselves.

LLM applications, AI agents, and model API gateways are now active participants in your environment, generating telemetry that is as security-relevant as any traditional log source. Aggregating this diverse and growing data landscape accurately and efficiently is one of the most pressing challenges in modern security operations, and where your SIEM's scalability, ingestion speed, and licensing model become decisive factors.

A modern SIEM must seamlessly integrate with varied data sources, support real-time ingestion, and handle multiple formats, including vendor-agnostic, open-source technologies such as OpenTelemetry. In an AI-first security environment, this breadth of collection is the foundation for every layer of AI-driven detection, behavioral analytics, and agentic response.

Gaps in your data pipeline create blind spots and directly degrade the intelligence and autonomous capabilities your security operations depend on. Comprehensive data collection is essential to ensure that no critical data are missed and your security team has the information to detect and respond to threats promptly and effectively.

LOG COLLECTION EVALUATION CRITERIA

Comprehensive source integration

Collects data from all relevant sources, including on-premises, cloud, and hybrid environments. This includes logs, network flows, endpoint data, data from various SaaS applications, and AI systems, including LLM applications, AI agents, and model API gateways.

Real-time data ingestion

Ingests data in real time for timely threat detection. Real-time ingestion is equally critical for AI models and agents acting on live data as delayed inputs directly degrade the accuracy of AI-driven detections and automated response actions.

Support for diverse data types

Supports a wide range of data types, including logs, events, metrics, and other relevant data formats, to provide a holistic view of your organization's security posture. Broader data diversity also strengthens AI model accuracy, enabling more precise behavioral baselines and anomaly detections across your entire environment.

Log storage and data retention

Securely stores security log data with AES 256 encryption at rest and TLS encryption in transit, and retained for up to seven years or pursuant to regulatory requirements in your industry. Long-term retention gives AI models the historical data depth needed to identify slow-moving threats, establish accurate behavioral baselines, and continuously improve detection efficacy over time.

2 How is data transformed in your SIEM?

Once collected, SIEM transforms the data into a format that is AI-ready, ensuring accurate analysis and action.

Data transformation involves normalization, enrichment, and correlation processes that convert raw data into meaningful insights. Analysts often face the daunting task of making sense of vast amounts of data from diverse sources, and in an AI-first environment, the quality of that transformation directly determines the effectiveness of every AI-driven capability built on top of it. Without proper transformation, data remains fragmented and difficult to analyze, and AI models operating on poorly normalized or inconsistently enriched data will produce unreliable detections, false positives, and incorrect agentic decisions.

Normalization is paramount for understanding network activities. It transforms disparate data from various sources into a unified format or schema that simplifies analysis. Despite many solutions touting normalization capabilities, their implementations often fail in effectiveness and ease of use. Normalization is the backbone of detection engineering, threat hunting, and security operations and increasingly, the backbone of AI model training and inference.

Converting raw messages into standardized records ensures that machine learning models and AI agents operate on consistent, high-quality inputs rather than noisy, inconsistent data that undermines their accuracy. Data transformation must occur at ingestion speed, as AI agents operating autonomously cannot afford to act on stale or incomplete transformed data. Any lag between collection and transformation directly degrades the quality of agentic decisions and automated response actions.

Not all schemas are created equal. The strength of normalization lies in the use of parsers and mappers to handle both structured and unstructured data. Parsers decode and extract crucial information from raw data, converting it into a readable, structured format. Mappers then align this data with a predefined schema, ensuring uniformity across diverse data sources. This process is especially critical when dealing with unstructured data, which can vary wildly in format and content yet contains some of the most critical, custom application insights, including telemetry from LLM applications and AI agents that do not conform to traditional log structures.

Modern SIEMs should leverage AI to automate parser creation for new and unfamiliar log sources, reducing the manual overhead of keeping pace with an ever-expanding data landscape and ensuring coverage gaps don't emerge as new AI systems are deployed across the environment.

Enrichment enhances normalized data with contextual information, such as threat intelligence feeds and asset data, helping analysts understand its significance and making it easier to detect and respond to threats. In an agentic SOC, enrichment becomes even more important as AI agents rely on richly contextualized insights to make accurate autonomous decisions, prioritize incidents correctly, and execute response actions with confidence. Beyond traditional threat intelligence enrichment, modern SIEMs should apply AI-driven enrichment to score and contextualize records based on behavioral baselines and real-time threat context.

DATA TRANSFORMATION CRITERIA

Effectiveness of normalization

Ensure consistent schema/field mapping, including for emerging sources like AI application telemetry and agent logs.

Parser and mapper quality

Check parser accuracy and update ease, and whether AI can auto-generate parsers for unfamiliar log sources.

Performance and scalability

Evaluate processing speed and capacity for large, unstructured, and AI-generated telemetry at ingestion speed.

Integration capabilities

Assess compatibility with data sources, APIs, and plugins — including AI platforms and LLM APIs.

Contextual data integration

Confirm robust, AI-enriched integration of threat intel and asset data, and assess the precision and relevance of the context it adds.

Ease of use

Evaluate UI intuitiveness, configuration flexibility, and support for automation and investigation triggers.

Impact on detection and response

Ensure detection accuracy and investigation efficiency can be tuned to your organization, and check how normalization/enrichment feed AI detection models, behavioral analytics, and agentic response.

3 Does your SIEM offer advanced analytics?

Data analytics powers the detection of sophisticated cyber threats, and in an AI SOC, it is the engine that transforms data into proactive defense.

Advanced capabilities have evolved past signature-based detection and basic correlation. including applying machine learning, behavioral analytics, and AI-driven threat detection to analyze large data volumes, and identify patterns that traditional methods might miss. SIEM solutions incorporating machine learning models continuously adapt to new threat patterns, improving over time. These models feed directly into a SOC's ability to move from passive detection to proactive defense, surfacing threats earlier, with greater precision, and less reliance on analyst intervention.

User and entity behavior analytics is crucial in modern SIEM solutions as it builds detailed profiles of normal behavior across users and entities. UEBA identifies deviations that may indicate insider threats that signature-based detection often misses. These behavioral baselines feed directly into AI models, enriching detections with context to make autonomous decisions more accurate and reliable.

Entity-centric detection, correlation and pattern recognition are key components of effective threat detection. By linking data from various sources, SIEM solutions can detect complex attack patterns that individual data points might not reveal. SaaS-based deployment is advantageous here as it makes it easy to continuously update detection content curated and customized for your environment by expert threat research teams, so your analytics stay current against emerging attack techniques without manual maintenance overhead.

Incorporating these AI-driven features ensures a SIEM solution can handle the dynamic nature of modern cybersecurity threats, providing security teams with the necessary tools to protect their organizations.

ADVANCED ANALYTICS EVALUATION CRITERIA

Anomaly detection
Verify the SIEM flags deviations from normal behavior using AI models that continuously adapt baselines, not static rules.

Behavioral analysi
Assess how well the SIEM monitors user/entity activity to surface unknown and insider threats.

Holistic view
UEBA should unify activity across users and entities, catching correlations single-point monitoring misses, and feeding directly into autonomous AI agents.

Enhanced contextual analysis
Check that AI dynamically adjusts risk scoring using real-time context, improving prioritization and cutting false positives.

Continuous updates, reduced maintenance
Confirm it's a true SaaS solution, with detection rules and AI models auto-updated against the latest threats.

AI-infused capabilities
Check that AI runs throughout the analytics layer: log clustering, noise reduction, alert fatigue reduction.

Scalability and flexibility
Assess how well the solution scales resources with demand, especially as AI workloads grow.

Transparency and ease of customization
Ensure easy rule-building and explainability. Can it show why it flagged an event?

Predefined rules and templates
Appraise the rule/template library's quality and how often the vendor's threat research team updates it.

Rule testing and simulation
Review the ability to test rules pre-deployment and validate AI model behavior against historical data.

4 Does your SIEM offer effective investigation?

Identifying potential threats is only the beginning.

Effective investigation capabilities determine how quickly and accurately your team can understand an incident, from origin to response. Security investigation capabilities have been fundamentally transformed by AI.

Where analysts once manually correlated events across data sources, modern SIEMs apply AI to automatically surface related signals, detailed timelines, and identify root causes in a fraction of the time. AI no longer just assists investigations, it conducts them by autonomously looking at threats, assembling evidence, and presenting analysts with a complete investigative context rather than a queue of individual alerts to process. The result is a decrease in the mean-time-to-understand, the window between detection and effective response. Natural language processing further enhances investigation and threat hunting capabilities, enabling a broader set of analysts, and even DevOps teams, to understand security data without deep query language knowledge.

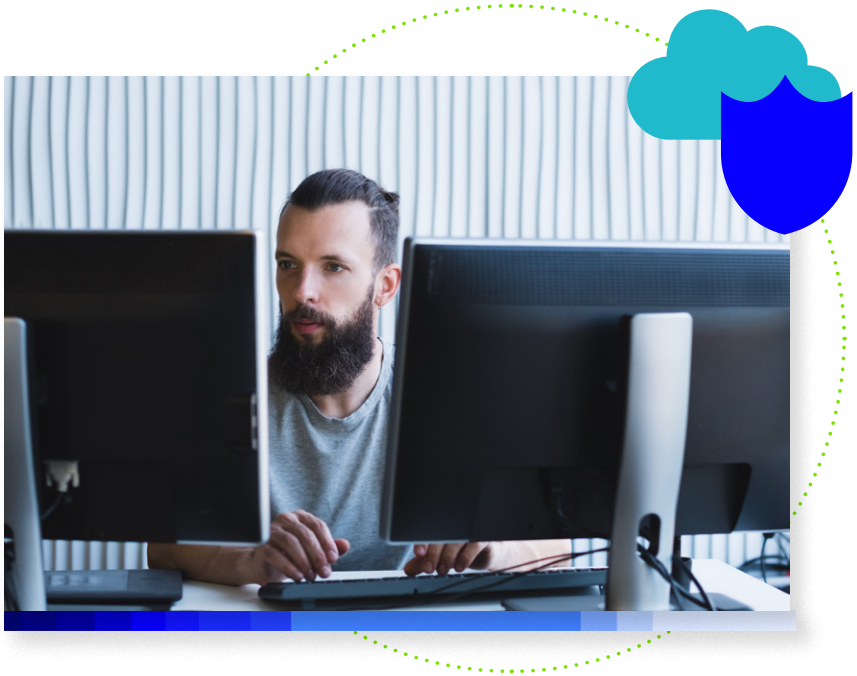

INVESTIGATION EVALUATION CRITERIA

Identifying initial compromise

Ensure the SIEM traces incidents to the initial point of compromise, and check whether AI builds the timeline automatically rather than through manual correlation.

Tracking threat propagation

Assess tools for tracking lateral movement and attacker methods, and whether AI can autonomously follow threat paths across the environment without analyst-driven queries.

Determining impact

Verify the SIEM helps quantify full impact — data exfiltration, downtime, regulatory exposure.

Alert triage and prioritization

Confirm alerts are classified and prioritized by severity, risk, and evolving context using ML and predefined rules, rather than static scores. Evaluate whether AI agents can fully resolve low-risk cases and escalate complex ones with full investigative context.

Contextual information

Determine if alerts include affected assets, users, and threat intelligence and how much context AI assembles automatically at alert time.

AI-driven threat hunting

Assess whether the SIEM proactively surfaces suspicious patterns by correlating signals, without analyst-initiated queries.

False positive reduction

Validate that detection rules improve via analyst feedback, reducing noise over time.

Rapid root cause identification

Ensure powerful analytics provide full visibility across on-prem and cloud, including natural language search, structured/unstructured querying, and ML-based log deduplication.

Natural language interface

Assess whether analysts can query data, build detections, and generate summaries in natural language, without deep query-language expertise.

AI-assisted investigation workflows

Evaluate AI-generated summaries, recommended next steps, and automated evidence assembly, turning investigation from a skills-intensive manual process into an AI-augmented workflow accessible to more of the team.

5 Does your SIEM facilitate response?

Collaboration is essential for effectively managing and responding to security incidents in cybersecurity.

Effective response is where SOCs ultimately succeed. Incidents and breaches occur at the application level across Kubernetes containers, GitHub repositories, and cloud infrastructure components that security teams don't typically own or monitor, meaning response must extend beyond the security team itself and across the full DevSecOps organization. With resource-constrained teams, the ability to automate and act quickly is a baseline requirement. A SIEM solution that facilitates teamwork and communication can significantly enhance an organization's team efficiency and effectiveness through features like customizable dashboards, automated reporting, compliance tracking, and streamlined response processes.

In addition, a modern SIEM must also perform autonomous response. AI agents that execute predefined playbooks, close low-risk cases without human intervention and escalate complex incidents with full investigative context are defining the capability of proactive defense.

Analysts are free from repetitive triage and routine case management to focus on high-level decisions and continuous improvement of the security program. Underneath this autonomy are the collaboration, reporting, and compliance capabilities to stakeholders across the entire organization. Together, these capabilities form a response framework that is faster and more intelligent than any previous generation of SIEM could deliver.

COLLABORATION EVALUATION CRITERIA

Customizable dashboards

Review users' ability to build role-relevant dashboards with real-time metrics and alerts, with AI-driven insights surfaced proactively.

Collaboration tools

Validate support for shared workspaces across teams and real-time case management, with AI-generated summaries surfaced to speed decisions.

Role-based access controls and views

Ensure permissions and views can be tailored by role for relevance and usability.

Automated and custom reporting

Check that reporting is consistent and time-saving, with custom options for compliance or execs, and whether AI can generate narrative summaries.

Regulatory compliance

Assess support for data collection, retention, and reporting under GDPR, HIPAA, and PCI DSS, including specialized dashboards and automatic audit-ready reports with full audit trails.

Incident response workflows

Ensure customizable workflows guide analysts and integrate with AI-driven triage. Look for MCP server support so AI clients connect directly to tools, ticketing, and LLMs.

Automated playbooks

Evaluate the playbook library's depth, ease of customization, and support for human-in-the-loop escalation on high-risk incidents.

Agentic response automation

Assess whether AI agents can autonomously close low-risk cases and escalate complex ones with full context.

Post-incident reviews

Ensure the solution supports analyzing response effectiveness and updating playbooks for continuous improvement.

Rating system based on evaluation criteria

How to interpret your results:

Score 0 - 30	Consider a new solution
Score 31 - 40	Adequate coverage
Score 41 - 50	Your current SIEM provides adequate security coverage

SIEM Scorecard

CAPABILITY	IDEAL SIEM CAPABILILITIES	CURRENT SIEM CAPABILITIES		POINTS
DATA COLLECTION	• Cloud-native SIEM • Collects from on-prem, cloud, hybrid, SaaS, endpoints, and network devices • AI system & LLM telemetry ingestion • Out-of-the-box connectors • Support for structured and unstructured data	Limited data sources, manual integration required; significant visibility gaps	0-3 points	
		Supports multiple data sources with some automation; ingestion gaps remain for cloud or AI telemetry.	4-7 points	
		Comprehensive, seamless data collection across all sources including AI systems; compliant long-term retention	8-10 points	
DATA TRANSFORMATION	• Automatic alert triage and threat correlation • Multiple threat intelligence feeds & AI-driven contextual enrichment • 100's of app integrations, parsing, and normalization for new & unfamiliar log sources • Automated & playbook enrichment • Raw data storage for search	Basic normalization, minimal enrichment; parsers require significant manual effort	0-3 points	
		Effective normalization, good enrichment capabilities; enrichment not fully integrated into AI workflows	4-7 points	
		Advanced normalization and enrichment, robust correlation; enrichment feeds AI detection and agentic response	8-10 points	
ADVANCED ANALYTICS	• User and Entity Behavior Analytics with adaptive baselines • Pre-built and customizable rules • AI-powered anomaly detection, log clustering, & noise reduction • Dynamic risk scoring adjusted by real-time environmental context • Peer benchmarking • Continuously updated threat research • AI-detection explainability	Basic analytics, signature-based detection capabilities; limited behavioral detection	0-3 points	
		Advanced analytics with good behavioral analysis. UEBA present but baselines may be static; some AI features	4-7 points	
		AI-driven analytics throughout, adaptive UEBA; dynamic risk scoring; continuously updated rules; explainable detections	8-10 points	
INVESTIGATION	• Dynamic alert prioritization; automated triage and false positive reduction • Natural language querying • AI-assisted root cause identification and timeline analysis • AI-generated investigation summaries, evidence assembly, and next steps	Basic investigation tools, manual correlation, limited root cause analysis, no natural language interface.	0-3 points	
		Good investigation tools and effective root cause analysis, some automation, natural language query limited or partial	4-7 points	
		Advanced investigation capabilities, AI-driven timelines, triage, and robust root cause analysis, natural language interface; strong forensics.	8-10 points	
RESPONSE	• Predefined and customizable playbooks with human-in-the-loop escalation • SOAR integration • Agentic AI closes low-risk cases autonomously; escalates complex incidents • Automated & playbook enrichment • Automated and AI-narrative reporting • Full DevSecOps collaboration	Basic case management, minimal collaboration tools; limited playbooks; no agentic response.	0-3 points	
		Good case management and some collaboration features; playbooks available but limited; no agentic response.	4-7 points	
		Comprehensive case management, strong integration with SOAR, agentic response automation, deep playbook library, excellent collaboration tools.	8-10 points	
			TOTAL SCORE	

Consider a new solution (0-30)
A low SIEM evaluation score indicates significant underperformance in critical areas, exposing your organization to substantial security risks, operational inefficiencies, and potential regulatory non-compliance. This score is particularly dangerous as adversaries are using AI to launch faster, more sophisticated attacks, and a solution that lacks AI-powered detection, agentic response, and automated triage cannot keep pace.

Immediate action is required, including reconfiguring the existing system, investing in additional training for your security team, or evaluating more capable alternatives that offer AI-driven analytics, autonomous threat hunting, and agentic response capabilities to better protect your organization and enhance your overall security posture.

Adequate coverage (31-40)
A medium score indicates that the SIEM solution is functional but has notable limitations that affect optimal security operations. It adequately covers basic requirements but lacks the advanced AI and agentic capabilities needed for comprehensive threat detection, analysis, and response. While the solution may perform reasonably well in data collection and basic analytics, it likely falls short in AI-driven investigation, autonomous triage, and seamless integration with modern security tooling.

As adversaries increasingly leverage AI to evade detection and accelerate attacks, a SIEM without adaptive behavioral baselines, natural language querying, and agentic case resolution creates a growing capability gap. To address these limitations, consider incremental improvements, integrating supplementary tools, or gradually transitioning to a more robust solution built for the demands of AI-era threats.

Great security foundation (41-50)
A high score in your SIEM evaluation indicates excellent performance across key functionalities, reflecting a security posture built for the modern threat landscape. Your SIEM demonstrates strong AI-driven analytics, adaptive UEBA, automated triage, and agentic response capabilities that allow your team to move from reactive defense to proactive threat hunting.

Analysts are freed from repetitive alert triage to focus on high-value decisions, while AI agents handle low-risk case resolution autonomously and escalate complex incidents with full investigative context. To maintain and enhance this performance, consider leveraging your high-performing SIEM for a comprehensive DevSecOps strategy that ensures continuous security monitoring, rapid threat response, and tight collaboration between development, security, and operations teams as you work toward the full human-AI collaboration model that defines sixth-generation security operations.

Ready to evaluate Sumo Logic SIEM?

Book a demo

About Sumo Logic

Sumo Logic, Inc. helps make the digital world secure, fast, and reliable by unifying critical security and operational data through its intelligent platform. Built to address the increasing complexity of modern cybersecurity and cloud operations challenges, we empower digital teams to move from reaction to readiness—combining agentic AI-powered SIEM and log analytics into a single platform to detect, investigate, and resolve modern challenges. Customers around the world rely on Sumo Logic for trusted insights to protect against security threats, ensure reliability, and gain powerful insights into their digital environments. **For more information, visit** www.sumologic.com**.**

sumo logic