



Cyber Phenomenon Series

The 2026 SIEM Market Landscape

Executive Assessment of Coverage Capability Operating Effort and Cost

Scott Foote

Last Updated: September 11, 2026

Phenomenati Consulting
www.phenomenati.com

6 Liberty Square, #2736
Boston, MA 02109
(508) 709-7990 (office)

CONFIDENTIALITY NOTICE: The contents of this document, including any attachments, are intended solely for stakeholders of Phenomenati Consulting, may contain confidential and/or privileged information, and are legally protected from disclosure.

Contents

1	Executive Judgment	1-1
2	How the Marketplace is Changing	2-2
2.1	SIEM is Becoming a Wider Operating Platform	2-2
2.2	Storage and Detection are Being Priced Separately	2-2
2.3	Names and Product Boundaries Need Checking	2-2
2.4	AI Claims Need Operational Evidence	2-2
3	Assessment Method and Score Definitions	3-3
4	Provider Scorecard	4-4
5	Commercial Comparison of Platform Providers	5-5
6	Commercial Comparison of Additional Providers	6-6
7	Provider Assessments 1 to 4	7-7
7.1	Splunk Enterprise Security	7-7
7.2	Microsoft Sentinel	7-7
7.3	Google Security Operations	7-7
7.4	CrowdStrike Falcon Next Gen SIEM	7-7
8	Provider Assessments 5 to 8	8-8
8.1	Palo Alto Networks Cortex XSIAM	8-8
8.2	Exabeam New Scale Fusion	8-8
8.3	Securonix Unified Defense SIEM	8-8
8.4	Elastic Security	8-8
9	Provider Assessments 9 to 12	9-9
9.1	Devo Intelligent SIEM	9-9
9.2	Sumo Logic SIEM	9-9
9.3	Rapid7 Incident Command	9-9
9.4	Panther	9-9
10	Provider Assessments 13 to 16	10-10
10.1	Hunters Next Gen SIEM	10-10
10.2	Gurukul Next Gen SIEM	10-10
10.3	Fortinet FortiSIEM	10-10
10.4	OpenText Enterprise Security Manager	10-10
11	Provider Assessments 17 to 20	11-11
11.1	IBM QRadar SIEM	11-11
11.2	ManageEngine Log360	11-11
11.3	Graylog Security	11-11

11.4	Wazuh	11-11
12	Provider Assessments 21 to 23	12-12
12.1	Datadog Cloud SIEM	12-12
12.2	SentinelOne Singularity AI SIEM.....	12-12
12.3	NetWitness SIEM	12-12
13	The Work that Remains After Deployment.....	13-13
13.1	Planning Labor Without Inventing a Benchmark	13-13
13.2	MDR is a Different Commercial Decision.....	13-13
14	Reference Workload and Total Cost Model	14-14
14.1	Calculation.....	14-14
15	Three Year Ownership Scenarios.....	15-15
15.1	What the Arithmetic Tells the Executive Team	15-15
16	Published Prices and Cost Sensitivity.....	16-16
16.1	Examples That Can Be Calculated From Public Evidence	16-16
16.2	Sensitivity That Belongs in the Investment Memo	16-16
17	Selecting the Right Finalists.....	17-17
17.1	Treat the Incumbent as a Real Alternative	17-17
18	Proof of Value and Procurement Requirements	18-18
18.1	Require One Comparable Commercial Schedule	18-18
19	Governance and Transition Priorities	19-19
19.1	Use the SIEM as a Controlled Evidence System	19-19
19.2	Govern AI and Automated Response as Privileged Systems	19-19
19.3	Phase Migration Around Demonstrated Coverage	19-19
20	References	20-20

1 Executive Judgment

A SIEM purchase should begin with the security outcomes the organization must sustain and the people available to sustain them. The broadest platform can be the wrong investment when its data engineering and detection workload exceed the team's capacity. A low license fee can also be expensive once infrastructure, integration, tuning and recovery are counted.

For a heterogeneous enterprise, begin with Splunk, Microsoft Sentinel and Google Security Operations, then add one platform aligned with the existing endpoint estate and one credible challenger. For a lean SOC, give more weight to supported integrations, maintained detection content and repeatable investigation workflows. These are advisory shortlists, not findings from a controlled product benchmark. [1][3][5][6][7][17][20]

Decision Priority	Initial Evaluation Candidates
Breadth and bespoke detection	Splunk; Sentinel; Google SecOps; Elastic
Consolidation around existing security tools	CrowdStrike; Cortex XSIAM; SentinelOne
Behavioral and identity investigation	Exabeam; Securonix; Gurucul
Small operating team	Rapid7; Hunters; Sumo Logic; an established stack provider
Control over data and deployment	FortiSIEM; QRadar; OpenText ESM; NetWitness; self-managed Elastic
License cost discipline	Graylog; Wazuh; Log360, subject to a complete quote

The most reliable commercial finding is that published prices are rarely comparable. They may cover a small entry tier, an extreme volume discount, one data tier or only part of a SOC platform. Most providers require a scoped quote. The scorecard preserves that uncertainty instead of manufacturing transaction prices.

Use the report to select three finalists, issue one common workload and retention schedule, and measure both detection results and customer labor during a proof of value. Authorize a purchase only after the bill of materials, three-year cost exposure and exit terms are reviewable.

2 How the Marketplace is Changing

2.1 SIEM is Becoming a Wider Operating Platform

Large providers increasingly combine security information and event management with user and entity behavior analytics, security orchestration and response, threat intelligence, data pipelines and AI-assisted investigation. Splunk, Cortex XSIAM, Exabeam and Securonix all describe this wider scope. Procurement must identify which capabilities are included in the selected edition and which require additional products. [1][7][8][10]

2.2 Storage and Detection are Being Priced Separately

Microsoft separates analytics and data lake tiers. Elastic exposes ingestion and retention meters. Graylog distinguishes active processed data from data routed to a lake. This can improve economics, but inexpensive retained data is not automatically available to every continuous detection or fast investigation workflow. [3][11][29]

2.3 Names and Product Boundaries Need Checking

Provider or Label	Current Procurement Implication
Rapid7 InsightIDR	The current package page presents Incident Command Essential, Advanced and Ultimate. Compare those entitlements rather than an old InsightIDR price. [17]
Exabeam and LogRhythm	The current portfolio contains cloud New-Scale Fusion and self-hosted LogRhythm SIEM. They should not be treated as identical deployment options. [8][9]
IBM QRadar	This report scores the on-premises hardware and virtual appliance offering marketed by IBM. Obtain a written roadmap for the exact SKU. [24][25]
ArcSight ESM	The current OpenText product name is Enterprise Security Manager. Scope its dependent data and operating components. [23]
Elastic Security	Hosted, serverless and self-managed deployments differ. Current serverless pricing also includes optional workflow and AI usage meters. [11][12][13]

2.4 AI Claims Need Operational Evidence

Vendor pages often promise dramatic improvements in search, alert quality, investigation time or staffing productivity. Their examples use different populations, baselines and configurations. None of those percentages is treated here as an independently verified comparative result. The practical test is whether the platform reduces analyst effort while preserving evidence quality, missed-incident rates and human control.

This is a representative provider landscape, not a market-share census. Additional candidates include Guardsix for a sovereignty-focused evaluation and Anvilogic for an alternative SecOps architecture. Their public positioning warrants consideration when those requirements dominate, but they are outside the scored set. [37][38]

3 Assessment Method and Score Definitions

The reference buyer is a midsize enterprise with a mixed cloud and on-premises estate, 1,500 monitored assets, an existing SOC, 30 source families and some custom application logs. The reference load is 250 GB/day in year one, growing 20% annually. This is an explicit comparison assumption, not an estimate of your environment.

Scores are editorial judgments from current primary product, pricing and technical documentation. No hands-on comparative testing or negotiated vendor quotations were available. Product pages establish offered scope; they do not establish production effectiveness. Feature and operating grades should be treated as provisional within approximately one point, with wider uncertainty for unusual environments.

Dimension	What the Score Assesses	Meaning of 1 and 5
B Breadth	Source diversity; cloud, identity, endpoint, network and application coverage; investigation and enterprise needs.	1 narrow coverage; 5 broad enterprise coverage with substantial extensibility.
D Differentiation	Distinctive detection, investigation, automation, behavioral, engineering or data capabilities.	1 basic or weakly differentiated; 5 a strong, relevant functional distinction.
C Configure	Expected customer work to collect, normalize, map identities, migrate rules and reach useful production coverage.	1 very high effort; 5 very low effort for the stated deployment.
M Maintain	Expected customer work for platform upkeep, content, source drift, capacity, access and billing.	1 very high effort; 5 very low ongoing effort.
P Purchase	Annualized software or SaaS expense for the reference footprint, where evidence permits a calculation.	5 up to \$25k; 4 above \$25k to \$75k; 3 above \$75k to \$200k; 2 above \$200k to \$400k; 1 above \$400k.
T TCO	Provisional ownership burden from operating architecture, dependencies, labor and known cost exposure.	1 very high; 2 high; 3 moderate or unresolved; 4 relatively favorable; 5 exceptionally favorable with strong evidence.

All numerical scales point in the same direction: higher is more favorable. C and M are shown separately because easy deployment can coexist with costly upkeep. The T grade is a qualitative operating assessment, not a vendor-specific dollar forecast and not a measure of risk reduction.

P marked 3* is a neutral placeholder where a comparable price is unestablished. It does not mean the vendor actually costs \$75k to \$200k, and these entries cannot be ranked commercially. P marked 1L uses published list pricing under the stated assumptions; discounts and scope changes can alter it. Wazuh's 5 applies only to self-managed software licensing.

The report intentionally avoids an aggregate league table. Averaging an unknown purchase price into an overall rank would conceal the most consequential uncertainty. Resolve mandatory requirements first, then replace provisional grades with workload evidence and written quotes.

4 Provider Scorecard

Higher scores are more favorable. B breadth; D differentiation; C configure; M maintain; P purchase; T ownership burden. P 3* means unpriced, not average cost; 1L is the reference list-price calculation. T grades carry low commercial confidence until quotes are received.

Provider	B	D	C	M	P	T
Splunk Enterprise Security [1][2]	5	5	2	2	3*	2
Microsoft Sentinel [3][4]	5	4	3	3	3*	3
Google Security Operations [5]	5	5	3	4	3*	4
CrowdStrike Falcon Next Gen SIEM [6]	4	5	4	4	3*	4
Palo Alto Networks Cortex XSIAM [7]	5	5	3	4	3*	3
Exabeam New Scale Fusion [8][9]	4	5	3	3	3*	3
Securonix Unified Defense SIEM [10]	4	5	3	3	3*	3
Elastic Security [11][12][13]	5	4	3	3	3*	3
Devo Intelligent SIEM [14][15]	4	4	3	3	3*	3
Sumo Logic SIEM [16]	4	4	4	4	3*	4
Rapid7 Incident Command [17]	4	4	4	4	3*	4
Panther [18][19]	4	5	3	3	3*	3
Hunters Next Gen SIEM [20]	4	4	4	4	3*	4
Gurukul Next Gen SIEM [21]	4	5	3	3	3*	3
Fortinet FortiSIEM [22]	4	4	3	2	3*	3
OpenText Enterprise Security Manager [23]	4	4	2	2	3*	2
IBM QRadar SIEM [24][25]	4	4	2	2	3*	2
ManageEngine Log360 [26][27]	3	3	3	3	3*	3
Graylog Security [28][29]	3	3	3	3	3*	3
Wazuh [30][31]	3	3	2	2	5	2
Datadog Cloud SIEM [32][33]	4	4	4	4	1L	2
SentinelOne Singularity AI SIEM [34]	4	4	4	4	3*	4
NetWitness SIEM [35][36]	4	4	2	2	3*	2

Configuration assumptions are stated in each profile. CrowdStrike and SentinelOne operating grades assume an existing compatible endpoint estate. Google uses an Enterprise SaaS scenario. Elastic uses Serverless Complete. Self-managed alternatives generally require more infrastructure labor than their cloud counterparts.

5 Commercial Comparison of Platform Providers

Published commercial evidence as of the research date. A request for a quote is a material pricing limitation, not a negative product finding. The questions in the last column define what must be priced; they do not assert that every item is separately charged.

Provider	Pricing Structure or Scope	Published Evidence and Quotation Requirements
Splunk	Activity, workload or ingest; ES edition and deployment matter.	Quote required. Price platform, ES, retention and optional services together. [2]
Microsoft	Analytics ingestion with PAYG or commitment; separate lake meters.	Regional calculator and agreement required. Include processing, queries, retention and playbooks. [3][4]
Google	Ingestion-based Standard, Enterprise or Enterprise Plus.	Quote required; one-year hot retention included. Confirm rules, intelligence and pipeline limits. [5]
CrowdStrike	Falcon SIEM package and data entitlements.	Quote required. Identify third-party volume, first-party allowances, retention and pipeline scope. [6]
Palo Alto Networks	XSIAM platform and scoped capabilities.	Quote required. Price all data, endpoint and response requirements plus transition services. [7]
Exabeam	New-Scale Fusion or separately scoped LogRhythm deployment.	Quote required. Distinguish analytics augmentation from a full replacement. [8][9]
Securonix	Unified Defense SIEM package.	Quote required. Confirm volume meter, hot retention, content, AI and automation inclusions. [10]
Elastic	Serverless ingest plus retained GB-month and optional services.	Complete advertises as low as \$0.11/GB ingest and \$0.019/GB-month retained. Volume tiers apply. [11][12]
Devo	Intelligent SIEM Starter or full Intelligent SIEM.	Quote required. Starter lists two behavioral models and two automation playbooks. [15]
Sumo Logic	Essentials or Enterprise Suite; product activation and flex configuration.	Quote required. SIEM and some AI services require explicit activation. [16]
Rapid7	Asset-based Incident Command tiers.	Quote required. Advanced lists 180-day logs; 13-month alert/audit retention is different. [17]
Panther	Platform subscription and chosen architecture.	Quote required. Confirm ingest, query, lake, AI, support and managed-content entitlements. [18][19]

6 Commercial Comparison of Additional Providers

Provider	Pricing Structure or Scope	Published Evidence and Quotation Requirements
Hunters	Platform with own or managed data lake.	Quote required. Clarify which party pays lake storage, compute and data movement. [20]
Gurukul	Modular platform and deployment choice.	Quote required. Separate data-lake infrastructure from platform and support charges. [21]
Fortinet	FortiSIEM deployment and entitlement schedule.	Quote required. Confirm capacity, collectors, HA and bundled automation scope. [22]
OpenText	ESM correlation platform and associated components.	Quote required. Include retention, infrastructure, connectors, support and recovery. [23]
IBM	EPS/FPM or MVS; subscription or perpetual.	Quote required. Include appliances or VMs, storage and recurring support. [25]
ManageEngine	Log360 component and monitored-resource configuration.	Quote required. Price the complete source inventory and deployment, not one component. [27]
Graylog	Active processed volume per day or annual consumption units.	Security starts at \$18,000/year, from 10 GB/day or 100 GCUs; larger deployment needs a quote. [29]
Wazuh	Self-managed software; optional commercial services or cloud.	\$0 self-managed software license. Infrastructure, labor, support and cloud are additional. [30][31]
Datadog	12-month analysis and retention package per GB analyzed.	Published \$5/GB starting rate billed annually; \$7.50 on demand. Workflows billed separately. [33]
SentinelOne	AI SIEM platform and selected entitlements.	Quote required. Confirm data pipeline, retention, response and AI consumption. [34]
NetWitness	SIEM deployment and optional platform modules.	Quote required. Scope network forensics and packet storage separately. [35][36]

Do not scale a small starting tier linearly into a large enterprise price. Do not multiply a lowest marginal unit rate by all usage. A volume commitment, regional rate, capability tier, support level and data retention schedule must be attached to every meaningful commercial comparison.

7 Provider Assessments 1 to 4

Product descriptions are sourced facts. Fit, operating burden and score explanations are advisory judgments for the reference scenario.

7.1 Splunk Enterprise Security

Broad heterogeneous telemetry and a unified investigation environment make Splunk a strong candidate where unusual applications and bespoke detection logic matter. Current ES packaging incorporates SIEM, UEBA, SOAR and detection engineering, with edition differences that need explicit scoping. [1][2]

The 5/5 capability grades recognize flexibility and functional depth. The 2/2 operating grades reflect the expected work of maintaining normalization, detection content, permissions and search performance even on Splunk Cloud. Retaining skilled staff can make an existing deployment more economic than a replacement. Require separate prices for the platform, ES edition, retention and automation.

7.2 Microsoft Sentinel

Sentinel combines broad ingestion, behavior analytics and orchestration with an analytics tier and a separate data lake. It is a natural evaluation candidate for Microsoft-centered organizations, while its collection scope extends beyond Microsoft products. [3][4]

Breadth scores 5; functionality scores 4 because the assembled workflow and entitlements require validation across services. Configuration and maintenance score 3: connectors, KQL content, identity mapping, playbooks and billing controls remain customer responsibilities. Challenge the quote with non-Microsoft logs, delayed events and retrieval from the lake. Avoid counting classic Log Analytics and simplified Sentinel ingestion charges twice.

7.3 Google Security Operations

Google combines SIEM, SOAR, curated detections and threat intelligence. Its published packages include one year of hot telemetry retention; rule capacity, curated content, Gemini and threat intelligence differ by edition. [5]

Breadth and differentiation score 5 for the combined data, detection and intelligence offering. Configuration scores 3 because normalization and rule migration still require expertise; ongoing maintenance scores 4 under managed service delivery. The TCO grade of 4 assumes that included retention displaces a real requirement. Prove parser field quality, rule limits and the exact intelligence package before committing.

7.4 CrowdStrike Falcon Next Gen SIEM

Falcon connects third-party security data with the broader Falcon environment and combines detection, investigation and automation. Its current positioning includes data pipeline management through Falcon Onum. [6]

The 5 for differentiation reflects the opportunity to connect endpoint evidence directly to cross-domain investigations. The 4/4 operating grades assume an established Falcon environment and standard sources; an organization without that foundation should lower both grades by one. Obtain a SIEM-specific quote and compare first-party and third-party allowances, retention and automation. Endpoint bundle pricing is not a full SIEM price.

8 Provider Assessments 5 to 8

Product descriptions are sourced facts. Fit, operating burden and score explanations are advisory judgments for the reference scenario.

8.1 Palo Alto Networks Cortex XSIAM

XSIAM presents a combined SOC platform spanning SIEM, SOAR, endpoint, network and cloud detection capabilities. Its value proposition is coordinated investigation and response across this wider platform. [7]

Breadth and differentiation score 5 for integrated scope, without accepting the vendor's percentage improvement claims as benchmarks. Configuration scores 3 because deployment changes workflows, data models and response permissions; maintenance scores 4 once stable. TCO remains 3 until actual retiring-tool savings are established. Compare a complete XSIAM program against the complete incumbent stack and reserve budget for migration and dual operation.

8.2 Exabeam New Scale Fusion

New-Scale Fusion combines SIEM, behavioral analytics, risk scoring and automation, including analytics for human and non-human entities. Exabeam also markets LogRhythm SIEM as a self-hosted product. These are distinct deployment choices. [8][9]

Differentiation scores 5 for behavioral investigation and the option to augment an existing SIEM. Operating grades of 3 reflect the work of getting identities, peer groups and use-case baselines right. The scorecard applies to the cloud Fusion deployment; do not apply its maintenance estimate to self-hosted LogRhythm. Test service accounts and shared identities, and price analytics-only augmentation separately from replacement.

8.3 Securonix Unified Defense SIEM

Securonix brings SIEM, UEBA, SOAR and threat intelligence into a cloud platform and advertises 365 days of searchable hot data. Threat content delivery and behavior-based investigation are central to its positioning. [10]

The 5 for differentiation reflects behavioral analysis and integrated investigation. Configuration and maintenance score 3 because source normalization, identity context and anomaly tuning still require ownership. Confirm which retention, content and automation entitlements are contractual. Its TCO grade remains 3 until detection value, retained data access and the full commercial package are demonstrated against the reference workload.

8.4 Elastic Security

Elastic offers hosted, serverless and self-managed deployments. Serverless Security Analytics Complete adds advanced analytics and workflow capabilities, while billing separates ingestion, retention and optional services. These deployment models do not have identical capability or cost profiles. [11][12][13]

Breadth scores 5 and differentiation 4 for flexible data analytics and security integration. Serverless removes cluster administration but leaves schema, detection and integration engineering, supporting 3/3 operating grades. Self-managed Elastic would lower maintenance to 2. Model actual marginal tiers, support, egress and AI executions; advertised lowest unit rates cannot be applied to every ingested GB.

9 Provider Assessments 9 to 12

Product descriptions are sourced facts. Fit, operating burden and score explanations are advisory judgments for the reference scenario.

9.1 Devo Intelligent SIEM

Devo pairs a security data platform with SIEM, SOAR and ThreatLink case management. Its published Starter package limits behavioral models and automation playbooks, whereas the fuller package removes those stated functional limits. [14][15]

Breadth and differentiation score 4 for data analytics and the integrated SOC workflow. Operating grades of 3 recognize that collection, detection design and reporting still need engineering. TCO remains neutral until retention, service levels and all processing charges are priced. Test high-volume joins and concurrent hunts with the same parsers, time windows and analyst workload used for competitors.

9.2 Sumo Logic SIEM

Sumo Logic combines log analytics and security operations. Its Enterprise Suite includes the SIEM-oriented capability set, but the current pricing page requires product activation and identifies separate activation for some AI functions. [16]

The 4/4 capability and operating grades reflect a managed cloud service with entity-based detection and a shared operational environment. They do not assume that a base logging subscription buys the full security workflow. Obtain a configuration schedule covering SIEM, log tiers, automation, support and AI. Validate complex custom application sources before assigning a lower staffing budget.

9.3 Rapid7 Incident Command

Rapid7 now presents its SIEM packages as Incident Command Essential, Advanced and Ultimate with asset-based pricing. Published log retention differs by tier, and the page distinguishes alert and audit retention from raw log retention. [17]

Capability grades of 4 and operating grades of 4 make it a credible candidate for a lean SOC. Predictable asset metering improves budgeting, but it does not prove a lower purchase price. Confirm the contractual asset definition, cloud counting, eligible event sources, retention extensions and response functions. Price managed detection separately; product support is not an outsourced incident response team.

9.4 Panther

Panther offers an AI SOC platform with a detection engineering heritage. Its documentation supports Python rules, a CLI workflow, a no-code builder and vendor-managed detections with continued updates. [18][19]

Differentiation scores 5 for code-based detection control and engineering workflow. Configuration and maintenance score 3: serverless delivery does not remove rule testing, enrichment and software development work. A team with strong Python and CI/CD skills may achieve better economics than a team that must acquire them. Test code review, regression tests, runtime failures, no-code usability and the complete storage and compute bill.

10 Provider Assessments 13 to 16

Product descriptions are sourced facts. Fit, operating burden and score explanations are advisory judgments for the reference scenario.

10.1 Hunters Next Gen SIEM

Hunters emphasizes vendor-maintained detection content, automated investigation and a security data lake that can be vendor-managed or customer-owned. Its product page describes OCSF normalization. [20]

Grades of 4 reflect an appealing operating model for a small SOC. They do not endorse claims of zero ongoing management or no need for detection engineering. Validate coverage of nonstandard applications and how customers request, inspect and override detections. The scorecard assumes a managed lake; owning the lake adds capacity, access, query and cost responsibilities. Separate managed content from any human monitoring service.

10.2 Gurukul Next Gen SIEM

Gurukul emphasizes behavior analytics, identity context, SOAR and a bring-your-own-data-lake architecture. It advertises SaaS, cloud, on-premises and hybrid deployment choices. [21]

Differentiation scores 5 for behavioral and identity analysis combined with storage flexibility. The 3/3 operating grades apply to SaaS and recognize the effort of model calibration and identity quality. Bring-your-own storage may reduce duplication but increases engineering and accountability unless managed explicitly. Require tenant-scale performance evidence, customer-controlled thresholds and a priced split between analytics, infrastructure and support.

10.3 Fortinet FortiSIEM

FortiSIEM combines security event analysis with a configuration management database, IT/OT context, UEBA and native automation. The current product page describes integration of FortiSOAR capabilities within the SIEM offering. [22]

Breadth and differentiation score 4 for the combination of operational context and security analytics. Configuration scores 3 and self-managed maintenance 2 because collection, discovery credentials, capacity, upgrades and recovery stay in the operating plan. Request itemized entitlements for all analytics and automation functions. Validate OT-specific detections against the actual industrial protocols and data sources rather than accepting IT/OT branding as coverage evidence.

10.4 OpenText Enterprise Security Manager

The former ArcSight ESM product is now presented as OpenText Enterprise Security Manager. Its offering centers on real-time correlation, data collection, native threat intelligence and SOAR. [23]

Capability grades of 4 recognize a mature correlation-oriented approach; operating grades of 2 reflect specialist content, integration and platform work in the self-managed scenario. Existing skills and working detection content can materially improve renewal economics. For a new deployment, require the full architecture, support and storage bill. Do not treat a correlation engine license as the complete logging and investigation solution.

11 Provider Assessments 17 to 20

Product descriptions are sourced facts. Fit, operating burden and score explanations are advisory judgments for the reference scenario.

11.1 IBM QRadar SIEM

IBM's current product and pricing pages market on-premises hardware and virtual appliances. Licensing can use EPS/FPM or Managed Virtual Servers, with subscription or perpetual choices. The product includes behavior analytics and supports Sigma content. [24][25]

Capability grades of 4 recognize established correlation and network-oriented investigation. Configuration and maintenance score 2 for appliance capacity, rule tuning, upgrades and recovery. This assessment covers the currently marketed on-premises offering, not a presumed equivalent cloud SKU. Require a written support horizon, product roadmap and migration terms. Evaluate renewal against verified incremental improvement, not only against the cost of a new platform.

11.2 ManageEngine Log360

Log360 combines SIEM with Active Directory-oriented auditing and a wider set of data protection functions. Its current offering includes visual playbook building, detection tuning and an incident workbench. [26][27]

Grades of 3 indicate a practical candidate for conventional enterprise IT and audit use cases, with more proof required for complex global environments. On-premises operation still requires service, storage and component maintenance. Request a quote listing each component, source and capacity entitlement. Demonstrate cloud identity, custom SaaS and workload coverage independently; integrated product labels do not establish feature parity with a larger SOC platform.

11.3 Graylog Security

Graylog Security adds SIEM detection, risk-based investigation and guided workflows to the Graylog logging foundation. It is distinct from both Graylog Open and Graylog Enterprise; the current Security entry price is \$18,000 per year at a small starting tier. [28][29]

Grades of 3 reflect a focused and approachable proposition with enterprise-scale proof still required. The scorecard assumes self-management; cloud operation shifts infrastructure effort into the contract. The entry price does not establish the price at 250 GB/day. Validate maintained parsers, Sigma behavior, identity context and recovery. Data routed away from active analytics must remain visible in the detection coverage plan.

11.4 Wazuh

Wazuh provides an open-source security platform with agent-based monitoring, file integrity, configuration assessment, vulnerability detection and log analysis. The self-managed software has no license fee; a separately managed cloud offering exists. [30][31]

Purchase scores 5 specifically for zero software license cost. Breadth and differentiation score 3; enterprise-grade cross-source behavior analysis and complex orchestration need proof or additional tools. Operating grades of 2 reflect infrastructure, scaling, upgrades, integration and custom content work. TCO therefore scores 2 for the reference enterprise scenario. For narrower monitoring requirements and an experienced team, that assessment can improve substantially.

12 Provider Assessments 21 to 23

Product descriptions are sourced facts. Fit, operating burden and score explanations are advisory judgments for the reference scenario.

12.1 Datadog Cloud SIEM

Datadog ties security investigation to application and infrastructure observability, with entity analytics and AI-assisted triage. Its current pricing page presents a 12-month analysis and retention offering at a published starting rate of \$5 per GB analyzed, billed annually. [32][33]

Capabilities and operating ease score 4 where shared observability data and workflows are valuable. Purchase scores 1L under the reference annualized list-price calculation; TCO scores 2 pending discounts and tool consolidation. Those grades do not apply to a smaller selected-log footprint. Request exact metering definitions and additional workflow, AI and logging costs; old event-based or shorter-retention prices are not interchangeable.

12.2 SentinelOne Singularity AI SIEM

SentinelOne combines a security data pipeline with cross-source investigation and native response in Singularity AI SIEM. Its product page describes endpoint, cloud, identity and third-party telemetry within that workflow. [34]

Capability and operating grades of 4 recognize the potential to consolidate collection and investigation, especially for an existing SentinelOne customer. A mixed endpoint estate should reduce operating confidence until tested. Prove third-party parser quality, custom detections, access controls and analyst explainability. Obtain a SIEM-specific commercial schedule; an endpoint subscription is not evidence that all SIEM data and retention are included.

12.3 NetWitness SIEM

NetWitness combines log collection and metadata enrichment with a broader platform that includes network and packet investigation. Its SIEM page describes regional log processing with centralized visibility and support for heterogeneous sources. [35][36]

Capability grades of 4 recognize the investigative value of logs alongside network evidence. The 2/2 operating grades reflect the engineering and capacity burden of a self-managed deployment. Full packet capture is a separately scoped architecture and is excluded from the baseline cost comparison. Include the associated capture, storage, bandwidth and privacy controls when network forensics is a deciding requirement.

13 The Work that Remains After Deployment

Cloud hosting transfers responsibility for the service infrastructure. It does not transfer accountability for which data arrives, whether it is interpreted correctly or whether anyone acts on the resulting alert. A credible operating model gives each of those responsibilities an owner.

Workstream	Configuration Work	Recurring Ownership
Source integration	Enable audit events; provision credentials; deploy collectors; validate fields and timestamps.	Source owner and SIEM engineer monitor source health, schema changes and collection gaps.
Detection content	Map threats to data; enable and test rules; identify blind spots.	Detection engineer reviews changes, tunes suppressions and retests after releases.
Identity and assets	Resolve users, service accounts, workloads and asset criticality.	IAM and asset owners maintain mappings, lifecycle events and enrichment quality.
Response	Design approvals, credentials, playbooks and emergency stops.	SOC and incident commanders exercise response and review privileged actions.
Service reliability	Size collection, buffering, storage and recovery.	SRE or provider manages service health; customer verifies recovery and evidence continuity.
Cost and retention	Classify data, retention and entitlement requirements.	FinOps and security owners review bills, growth, query load and retention exceptions.
Assurance	Set access, segregation, exports and audit trails.	Security and privacy teams test access, retention, vendor changes and evidence integrity.

13.1 Planning Labor Without Inventing a Benchmark

For budgeting, test 1.5 to 2.0 incremental full-time equivalents for a managed platform with conventional sources, and 3.0 to 3.5 for a self-managed or heavily customized platform. These are scenario assumptions for this report, not measured industry staffing ratios or vendor staffing promises. Existing SOC analysts are budgeted separately in the model.

The primary drivers are source diversity, custom detection requirements, data quality, uptime expectations and analyst turnover. In a narrow 50 GB/day estate, a fraction of that staffing could be sufficient; a global regulated environment can require substantially more. Replace these assumptions with recorded customer hours in the proof of value.

13.2 MDR is a Different Commercial Decision

Managed detection and response can change the staffing model, but an MDR contract must specify who monitors, investigates, contains and reports incidents. Compare the combined SIEM and human service contract with the combined internal platform and staffing cost. Do not count both eliminated analyst positions and an unchanged internal SOC budget as savings.

14 Reference Workload and Total Cost Model

The following model illustrates how purchasing and operating choices interact. All dollar inputs in the architecture model are analyst-selected assumptions in USD, not vendor quotations, market averages or promises that a product meets the workload at that price.

Assumption	Reference Value
Monitored environment	1,500 assets; 30 source families; mixed cloud and on-premises; existing endpoint controls.
Daily telemetry	Year 1 250 GB; year 2 300 GB; year 3 360 GB. Decimal GB; 365 days per year.
Annual telemetry	91,250 GB; 109,500 GB; 131,400 GB. Three-year total 332,150 GB.
Required retention	90 days available to interactive investigation; 365 days total retrievable evidence. Verify actual continuous detection access by tier.
Operating scenario	One primary deployment; disaster recovery and collector buffering included in the infrastructure assumption; no packet capture.
Staff cost	Fully loaded \$180,000 per FTE per year; held constant to isolate workload effects.
Shared SOC labor	Four existing investigation FTEs, \$720,000 per year. This is not a promise of continuous 24/7 staffing.
Platform labor	Incremental SIEM engineering, content and administration. Separate from shared investigation labor.
Variable spend	Subscription/support and infrastructure allowances grow 20% yearly in this scenario; actual contracts may scale differently.
Accounting	Undiscounted three-year cash-equivalent cost; one-time implementation and terminal exit allowance included.

14.1 Calculation

Three-year TCO = implementation and migration + subscription and support + infrastructure and data services + platform labor + investigation labor + exit cost.

For a year-one annual cost that grows 20%, the three-year multiplier is $1 + 1.20 + 1.44 = 3.64$. Fixed annual labor uses a multiplier of 3. The model excludes taxes, financing, breach losses, capital depreciation and unchanged enterprise tools outside the scoped SIEM program. Additional MDR, packet capture, new endpoint subscriptions or dedicated multi-region deployments must be added when required.

A single hypothetical budget cannot normalize feature differences. All four architectures must still pass the same coverage and response requirements. A cheaper platform that misses a mandatory detection is not an economic equivalent.

15 Three Year Ownership Scenarios

Illustrative budget inputs in USD thousands. These are four hypothetical architectures; they are not price estimates for named providers. Subscription and infrastructure figures are budget allowances, not derived vendor unit rates.

Budget Input	A Managed Focused	B Broad Managed	C Commercial Self Managed	D Open Source Self Managed
Year 1 subscription and support	150	260	90	20 support; \$0 license
Year 1 infrastructure and data	15	25	65	80
Platform FTEs	1.5	2.0	3.0	3.5
Year 1 platform labor	270	360	540	630
One-time implementation	100	200	180	220
Exit allowance	25	40	35	40
Annual shared investigation labor	720	720	720	720

Scenario	Platform and Transition	Shared SOC Labor	Three Year TCO	Annual Average
A Managed focused	1,535.6	2,160.0	3,695.6	1,231.9
B Broad managed	2,357.4	2,160.0	4,517.4	1,505.8
C Commercial self managed	2,399.2	2,160.0	4,559.2	1,519.7
D Open source self managed	2,514.0	2,160.0	4,674.0	1,558.0

15.1 What the Arithmetic Tells the Executive Team

The focused managed scenario totals \$3.696 million over three years, while the open-source scenario totals \$4.674 million under these assumptions. The difference is driven largely by assigned labor, not proof that commercial software is inherently cheaper. Reduce the open-source platform team from 3.5 to 1.5 FTEs and its three-year TCO falls by \$1.08 million to \$3.594 million, reversing the result.

Likewise, one additional platform FTE costs \$540,000 over three years in this model. That can erase a substantial license discount. Budget discussions should therefore expose both the commercial offer and the staff hours required to sustain the promised coverage.

The model keeps shared investigation labor equal because no comparable alert-quality or analyst-productivity results were established. After a controlled pilot, substitute measured alert volumes, handling times and required coverage hours. Vendor-reported productivity percentages are not a basis for removing positions from a funded staffing plan.

16 Published Prices and Cost Sensitivity

16.1 Examples That Can Be Calculated From Public Evidence

Public Evidence	Reference Calculation and Limitation
Wazuh self-managed license	Published no-license-cost software yields \$0 license expenditure. This does not remove infrastructure, support, implementation or labor. [30]
Graylog Security entry tier	Starting at \$18,000/year gives a three-year entry subscription floor of \$54,000 before growth and other costs. The published starting footprint is much smaller than the reference 250 GB/day; no enterprise extrapolation is justified. [29]
Elastic Complete minimum rates	At \$0.11/GB, 91,250 GB produces a \$10,037.50 year-one ingest-only mathematical floor. Tiered rates, retained GB-months, support and optional services must be added. This is not an attainable full-year quote at the reference volume. [11][12]
Datadog twelve-month package	At the displayed \$5/GB starting rate, $250 \times 365 \times \$5 = \$456,250$ for year-one analyzed volume. The page describes 12-month analysis and retention; workflows are extra. Confirm the monthly billing meter, analyzed volume and negotiated discounts. [33]

These examples have different retention and entitlement scopes. They are useful price anchors, not equivalent full-system offers. Datadog's broader included historical analysis should be valued where needed; Elastic's ingest-only floor omits the principal storage and service elements.

16.2 Sensitivity That Belongs in the Investment Memo

Change in Assumption	Effect on Three Year Cost or Comparison
One more platform FTE	\$540,000 at the assumed fully loaded rate.
Ingest effective price rises \$1/GB	\$332,150 across the reference three-year volume, before retention or taxes.
Growth rises from 20% to 40%	Three-year volume rises to 397,850 GB, 19.8% above the base. Fixed labor does not automatically grow at the same rate.
20% lower recurring variable spend	Scenario B saves \$207,480 over three years on its subscription and infrastructure allowances; labor is unchanged.
Ninety days of duplicate subscription	A \$260,000 annual incumbent adds about \$64,110 at 90/365 of annual cost, plus any duplicate operating work.
Longer searchable retention	Calculate retained GB-months and search workload explicitly; a storage discount does not establish fast query or detection equivalence.

Retention should be modeled month by month when building from an empty repository. A steady-state full-year storage quantity overstates the early ramp; a model that ignores historical backlog understates migration cost.

17 Selecting the Right Finalists

Organizational Context	Practical Shortlist	Decision Condition
Microsoft-centered enterprise	Sentinel; Google; Splunk	Test non-Microsoft sources and the full cost of the chosen data tiers.
Established Falcon estate	CrowdStrike; Sentinel or Google; one independent challenger	Price the incremental SIEM and the full renewal stack, and prove third-party coverage.
Palo Alto platform strategy	Cortex XSIAM; Splunk or Google; incumbent	Validate retiring-tool savings and the complexity of response migration.
SentinelOne estate	SentinelOne; Google or Sentinel; Panther or Hunters	Prove cross-source correlation and the completeness of the commercial package.
Small SOC with standard sources	Rapid7; Hunters; Sumo Logic	Run a customer-led onboarding and operations trial with limited vendor assistance.
Engineering-oriented cloud company	Panther; Elastic; Datadog	Measure detection-as-code workflow, custom telemetry support and unit cost at growth.
Insider and identity priorities	Exabeam; Securonix; Gurucul	Use real identity ambiguity and peer-group examples; measure false escalations.
Local processing or complex networks	FortiSIEM; QRadar; OpenText; NetWitness	Verify all support, AI, recovery and telemetry flows against the deployment requirement.
Tight license budget	Graylog; Wazuh; Log360	Price engineering time and test every mandatory use case before accepting the apparent saving.

These are suggested evaluation sets. They do not imply exclusivity: a vendor can be viable in several categories. A regional or sector-specific service provider can also change implementation quality and operating cost materially.

17.1 Treat the Incumbent as a Real Alternative

For an existing SIEM, compare three paths: improve the current configuration, add a targeted analytics capability, or replace the platform. Include rule re-engineering, evidence migration, training, parallel monitoring and lost institutional knowledge. Exabeam and Gurucul explicitly support augmentation approaches that make this a practical procurement option. [8][21]

Select a winner only after the finalist has passed mandatory technical and governance gates. Among qualifying products, compare measured staff effort and priced three-year ownership. A small numerical difference in an advisory score is not a sufficient reason to replace a working system.

18 Proof of Value and Procurement Requirements

Run a four-to-six-week evaluation using the same source inventory, data period and attack scenarios for every finalist. Track vendor assistance separately from customer work. The following acceptance criteria are proposed defaults to tailor before the trial; they are not industry benchmarks.

Test	Evidence to Require
Coverage and parsing	All mandatory source families connected; critical fields, actor identity and time normalized. Record parsing failures and data loss, not connector counts alone.
Threat detection	Replay at least 20 relevant attack and misuse scenarios plus realistic benign activity. Record misses, duplicates, false escalations and time to actionable evidence.
Historical investigation	Search a 90-day window under concurrency and retrieve a 365-day example. Record elapsed time and all query or rehydration charges.
Configuration labor	Customer configures five new source families, one unusual application and one cross-source rule. Record hours, skills and vendor intervention.
Maintenance and resilience	Change a schema, rotate credentials, pause a collector and replay a backlog. Prove health alerts, buffering, recovery and rule regression controls.
Response safety	Exercise containment with approval, rollback and scoped credentials. Verify the audit trail and behavior when a dependency fails.
AI reliability	Require evidence-linked summaries; inject misleading log text; test access boundaries and escalation to a human. Measure incorrect conclusions.
Cost and exit	Reconcile measured usage to a sample invoice. Export raw events, cases, rules and audit records into usable formats and price the transfer.

18.1 Require One Comparable Commercial Schedule

Ask each finalist to quote years one through three for 250, 300 and 360 GB/day; the same 1,500 assets; 30 source families; 90-day interactive and 365-day total retention; required response functions; support; recovery; and implementation. Include alternative prices for 500 GB/day and a seven-day incident surge at twice the normal rate.

The contract should define billable units before and after filtering, normalization, enrichment, duplication and compression; included quotas; overruns; unused commitments; renewal ceilings; scope changes; data export and termination assistance. Obtain an explicit list of features that are generally available, preview, roadmap-only or separately licensed.

19 Governance and Transition Priorities

19.1 Use the SIEM as a Controlled Evidence System

Telemetry can contain credentials, identifiers, customer content and sensitive operational details. Define collection purpose, permitted fields, access roles, retention and deletion before onboarding sources. Preserve a restricted evidence path when broad analyst access requires redaction. A compliance dashboard is evidence support, not proof that a legal obligation has been fulfilled.

For cloud delivery, verify processing regions, support access, subprocessors, key custody, incident notification, export rights and deletion behavior in the actual service agreement. Local log storage alone does not establish local processing if AI assistants, support tools or integrations send data elsewhere. These are procurement controls to validate, not findings that any named vendor fails them.

19.2 Govern AI and Automated Response as Privileged Systems

Require distinct service identities, scoped credentials, evidence citations and complete audit trails. Keep destructive response actions subject to an approved policy and a tested emergency stop. Evaluate prompt injection in ingested content, cross-tenant information exposure, unsupported conclusions and model or subprocessor changes. Do not accept natural-language ease of use as a substitute for access control.

19.3 Phase Migration Around Demonstrated Coverage

Phase	Exit Condition
Inventory and baseline	Source owners, active rules, open cases, retention duties, existing spend and staff effort are documented.
Pilot and commercial validation	Mandatory scenarios pass; measured data volumes reconcile to a written price and operating model.
Parallel operation	Critical sources and detections run in both environments; gaps and contradictory results are resolved.
Controlled cutover	SOC approves coverage and response readiness; historical evidence remains retrievable and access-controlled.
Retire and optimize	Export and deletion obligations are satisfied; unused commitments end; workload, noise and unit cost are reviewed monthly.

The board-level decision is whether the proposed system provides sufficient evidence and response capability at a cost the organization can sustain. Require three accountable commitments: a documented coverage target, a funded operating owner and a priced exit. The provider scorecard narrows the field; demonstrated outcomes and contractual economics should decide the purchase.

20 References

All sources were accessed on 11 September 2026. These are primary vendor product, pricing or technical sources. Product claims remain vendor statements unless explicitly identified as calculations or advisory judgments. URLs are retained as visible text without external hyperlink relationships.

[1] Splunk Enterprise Security.

https://www.splunk.com/en_us/products/enterprise-security.html

[2] Splunk pricing.

https://www.splunk.com/en_us/products/pricing.html

[3] Microsoft Sentinel pricing.

<https://www.microsoft.com/en-us/security/pricing/microsoft-sentinel/>

[4] Microsoft Sentinel plan costs and understand pricing and billing.

<https://learn.microsoft.com/en-us/azure/sentinel/billing>

[5] Google Security Operations product and package comparison.

<https://cloud.google.com/security/products/security-operations>

[6] CrowdStrike Falcon Next Gen SIEM.

<https://www.crowdstrike.com/en-us/platform/next-gen-siem/>

[7] Palo Alto Networks Cortex XSIAM.

<https://www.paloaltonetworks.com/cortex/cortex-xsiam>

[8] Exabeam New Scale Fusion Security Operations Platform.

<https://www.exabeam.com/platform/exabeam-new-scale-fusion-security-operations-platform/>

[9] Exabeam LogRhythm SIEM.

<https://www.exabeam.com/platform/logrhythm-siem/>

[10] Securonix Unified Defense SIEM.

<https://www.securonix.com/products/siem-solutions/>

[11] Elastic Cloud Serverless pricing for Elastic Security.

<https://www.elastic.co/pricing/serverless-security>

[12] Elastic Security Serverless billing dimensions.

<https://www.elastic.co/docs/deploy-manage/cloud-organization/billing/security-billing-dimensions>

[13] Elastic pricing and deployment comparison.

<https://www.elastic.co/pricing>

[14] Devo Security Data Platform.

<https://www.devo.com/>

[15] Devo Platform Packaging.

<https://www.devo.com/packaging/>

[16] Sumo Logic pricing and package comparison.

<https://www.sumologic.com/pricing>

[17] Rapid7 Incident Command packages.

<https://www.rapid7.com/products/siem/packages/>

[18] Panther AI SOC Platform.

<https://panther.com/>

[19] Panther documentation Writing Python Detections.

<https://docs.panther.com/detections/rules/python>

- [20] Hunters Next Gen SIEM product overview.
<https://www.hunters.security/product>
- [21] Gurukul Next Gen SIEM.
<https://gurukul.com/products/next-gen-siem/>
- [22] Fortinet FortiSIEM.
<https://www.fortinet.com/products/siem/fortisiem>
- [23] OpenText Enterprise Security Manager.
<https://www.opentext.com/products/enterprise-security-manager>
- [24] IBM QRadar SIEM.
<https://www.ibm.com/products/qradar-siem>
- [25] IBM QRadar SIEM pricing.
<https://www.ibm.com/products/qradar-siem/pricing>
- [26] ManageEngine Log360.
<https://www.manageengine.com/log-management/>
- [27] ManageEngine Log360 quotation configuration.
<https://www.manageengine.com/log-management/get-quote.html?top-menu=>
- [28] Graylog Security.
<https://graylog.org/products/security/>
- [29] Graylog pricing plans and features.
<https://graylog.org/pricing/>
- [30] Wazuh open source security platform.
<https://wazuh.com/>
- [31] Wazuh Cloud.
<https://wazuh.com/cloud/>
- [32] Datadog Cloud SIEM.
<https://www.datadoghq.com/product/cloud-siem/>
- [33] Datadog pricing Cloud SIEM section.
<https://www.datadoghq.com/pricing/?product=cloud-siem>
- [34] SentinelOne Singularity AI SIEM.
<https://www.sentinelone.com/platform/ai-siem/>
- [35] NetWitness SIEM.
<https://www.netwitness.com/modules/security-information-event-management/>
- [36] NetWitness unified threat detection and response platform.
<https://www.netwitness.com/>
- [37] Guardsix sovereign security platform.
<https://guardsix.com/>
- [38] Anvilogic Agentic SecOps Platform.
<https://www.anvilogic.com/>