COMPARATIVE EVALUATION / EVIDENCE REFRESH

CMP-003

Enterprise SIEM Platforms

Security analytics, detection, investigation and response

CURRENT EVIDENCE SNAPSHOT

Official product sources refreshed through 2026-09-17

Controlled Mythos laboratory rerun: NOT ASSERTED

BENCHMARK DOCTRINE

Gates before scores. Evidence before certainty. Profile fit before universal ranking.

VERSION 1.2.0-candidate

CANDIDATE COMPARATIVE EVALUATION / PUBLIC

Enterprise SIEM Platforms

Security analytics, detection, investigation and response

DOCUMENT ID
CMP-003
VERSION
1.2.0-candidate
STATUS
Candidate Comparative Evaluation
PUBLICATION DATE
2026-09-17
SCHEDULED REVIEW
2026-12-16
CANONICAL ROUTE
/library/evaluations/cmp-003/

4 CANDIDATES

18 ATOMIC CRITERIA

9 MANDATORY GATES

3 WORKLOAD PROFILES

E2 CURRENT MAX EVIDENCE

CURRENT DECISION BOUNDARY

The July 0.1-point Sentinel/Splunk difference is not decision certainty. Current data-lake, graph, AI, SOAR, and endpoint changes are material. Reproduce ingestion, detection, investigation, response, retention, migration, and cost behavior before profile selection.

NO CURRENT UNIVERSAL WINNER IS PUBLISHED FROM THIS REFRESH.

July 24 baseline scores are preserved as lineage and sensitivity evidence, not silently reissued as September laboratory results.

Microsoft Sentinel

Cloud service / updates through 2026-08-24

REFRESH TRIGGER

Data lake, graph, MCP-assisted investigation, detections-as-code, and expanded UEBA materially change the platform boundary.

Splunk Enterprise Security

8.7.0 / 2026-09-02

REFRESH TRIGGER

8.7.0 adds AI SOC Analyst and synchronized CIM updates; known issues and upgrade compatibility remain decision inputs.

Google Security Operations

SIEM 2026-09-15 / SOAR 6.3.100

REFRESH TRIGGER

Parser updates plus preview reaction triggers and case playbooks expand response automation and lifecycle considerations.

Elastic Security

9.5.4 current docs

REFRESH TRIGGER

Current 9.x line continues rapid detection, endpoint, AI-assistant, and migration changes; exact fit depends on deployment model.

FINAL BENCHMARK SYSTEM / DECISION ARCHITECTURE

01 SCOPE

Freeze decision, candidates, versions and exclusions.

02 GATES

Test non-compensable trust, state, recovery, change and exit.

03 SCORES

Measure atomic capability with approved profile weights.

04 CONFIDENCE

Discount by source grade, freshness, repeatability and relevance.

05 PROOF

Run normal, failure, migration, recovery and operator scenarios.

06 DECISION

Record conditional fit, residual risk, owner and revalidation trigger.

EVIDENCE CONFIDENCE / MODIFIERS

E0	E1	E2	E3	E4	E5
0.00	0.38	0.64	0.78	0.91	1.00
UNSUPPORTED	VENDOR ASSERTION	OFFICIAL SOURCE	MYTHOS LAB	PRODUCTION	INDEPENDENT

THIS EDITION

Current official product and release sources have been refreshed to E2. No E3 controlled Mythos lab rerun is claimed. Therefore current profile rankings remain intentionally unissued until the proof plan is executed.

What changed since the July benchmark snapshot

Changes below are sourced from current official release documentation. They identify revalidation triggers; they are not translated into new numeric scores without controlled proof.

Microsoft Sentinel

Cloud service / updates through 2026-08-24

MATERIAL DELTA

Data lake, graph, MCP-assisted investigation, detections-as-code, and expanded UEBA materially change the platform boundary.

CRITERIA TOUCHED

C01 / C05 / C06 / C07 / C08 / C12 / C14

MODERATE REVALIDATION

Splunk Enterprise Security

8.7.0 / 2026-09-02

MATERIAL DELTA

8.7.0 adds AI SOC Analyst and synchronized CIM updates; known issues and upgrade compatibility remain decision inputs.

CRITERIA TOUCHED

C01 / C07 / C10 / C11 / C14

MODERATE REVALIDATION

Google Security Operations

SIEM 2026-09-15 / SOAR 6.3.100

MATERIAL DELTA

Parser updates plus preview reaction triggers and case playbooks expand response automation and lifecycle considerations.

CRITERIA TOUCHED

C01 / C05 / C06 / C07 / C11 / C14

MODERATE REVALIDATION

Elastic Security

9.5.4 current docs

MATERIAL DELTA

Current 9.x line continues rapid detection, endpoint, AI-assistant, and migration changes; exact fit depends on deployment model.

CRITERIA TOUCHED

C01 / C05 / C07 / C12 / C13 / C14

MODERATE REVALIDATION

Mandatory gates remain non-compensable

No current release is marked PASS solely from documentation. E2 means the official source supports the capability path; LAB and RETEST identify proof still required. HOLD blocks a current decision.

MANDATORY GATE	MICROSOFT SENTI	SPLUNK ENTERPRI	GOOGLE SECURITY	ELASTIC SECURIT
C01 Functional coverage	RETEST	RETEST	RETEST	RETEST
C02 Security / trust	E2	E2	E2	E2
C03 Governance / approval	E2	E2	E2	E2
C04 State integrity	E2	E2	E2	E2
C07 Observability / evidence	RETEST	RETEST	RETEST	RETEST
C09 Resilience / continuity	LAB	LAB	LAB	LAB
C11 Change safety / rollback	LAB	RETEST	RETEST	LAB
C16 Portability / exit	LAB	LAB	LAB	LAB
C18 Assurance confidence	HOLD	HOLD	HOLD	HOLD

LEGEND

E2 official-source support only

LAB controlled proof required

RETEST release delta touches this gate

HOLD decision gate remains closed

Historical profile sensitivity - preserved, not reissued

Values below are the 2026-07-24 evidence-adjusted baseline. They show sensitivity to operating-model weights. The September refresh does not recompute rank because E3 lab proof has not been reproduced.

MICROSOFT-CENTRIC SOC

BASELINE SPREAD 1.3 PTS

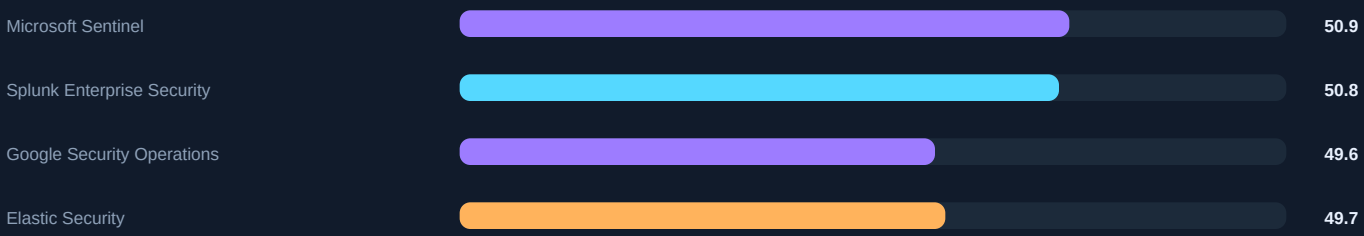

HIGH-SCALE SECURITY DATA PLATFORM

BASELINE SPREAD 1.3 PTS

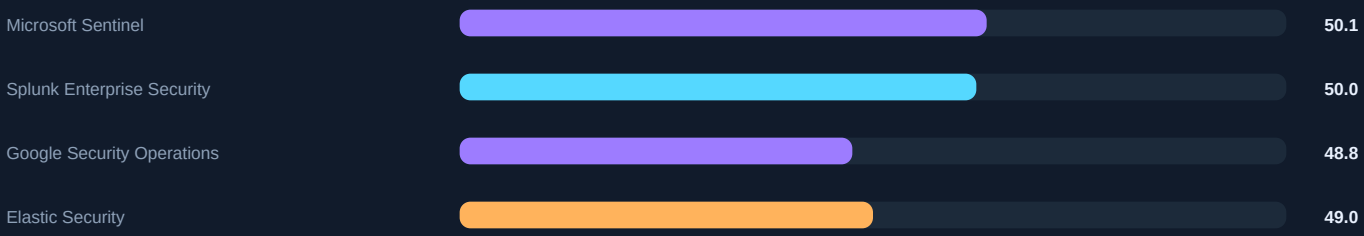

SOVEREIGN AND EXTENSIBLE SOC

BASELINE SPREAD 1.4 PTS

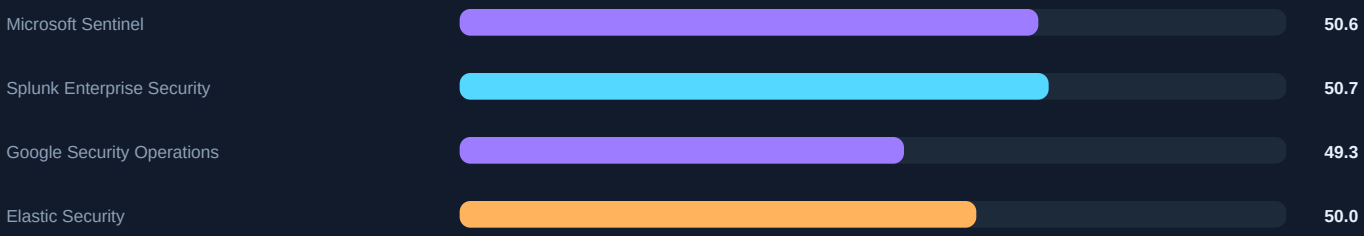

CURRENT RANK STATUS: WITHHELD PENDING CONTROLLED PROOF.

Evidence confidence is separate from capability

E5 — Independent repeatable validation
confidence modifier 1.00

E4 — Production evidence in adopting environment
confidence modifier 0.91

E3 — Controlled Mythos laboratory result
confidence modifier 0.78

E2 — Official documentation / release / source
confidence modifier 0.64

E1 — Vendor assertion or marketing
confidence modifier 0.38

E0 — Unsupported or unverified
confidence modifier 0.00

CURRENT EVIDENCE STATE

Microsoft Sentinel	E2 REFRESHED
Historical adjusted confidence: 58.2%	
E3 LAB: NOT REPRODUCED	
Splunk Enterprise Security	E2 REFRESHED
Historical adjusted confidence: 58.2%	
E3 LAB: NOT REPRODUCED	
Google Security Operations	E2 REFRESHED
Historical adjusted confidence: 58.2%	
E3 LAB: NOT REPRODUCED	
Elastic Security	E2 REFRESHED
Historical adjusted confidence: 58.2%	
E3 LAB: NOT REPRODUCED	

ASSURANCE RULE

A current release note can change capability context, but it cannot raise assurance beyond E2. Current recommendation strength is therefore bounded until the test plan produces reproducible artifacts.

Lifecycle risk is evaluated independently of features

The benchmark models migration, upgrade, compatibility, support, staffing, data/state export, downtime and exit. Current release changes below are explicit proof triggers.

Microsoft Sentinel

MODERATE REVALIDATION

Data lake, graph, MCP-assisted investigation, detections-as-code, and expanded UEBA materially change the platform boundary.

RETEST: C01 / C05 / C06 / C07 / C08 / C12 / C14

Splunk Enterprise Security

MODERATE REVALIDATION

8.7.0 adds AI SOC Analyst and synchronized CIM updates; known issues and upgrade compatibility remain decision inputs.

RETEST: C01 / C07 / C10 / C11 / C14

Google Security Operations

MODERATE REVALIDATION

Parser updates plus preview reaction triggers and case playbooks expand response automation and lifecycle considerations.

RETEST: C01 / C05 / C06 / C07 / C11 / C14

Elastic Security

MODERATE REVALIDATION

Current 9.x line continues rapid detection, endpoint, AI-assistant, and migration changes; exact fit depends on deployment model.

RETEST: C01 / C05 / C07 / C12 / C13 / C14

IRREVERSIBLE COMMITMENT GATE

Before migration or procurement lock-in, prove rollback, state portability, operational reconstruction, and supplier-exit assumptions.

Controlled proof is the next decision-producing step

REPRODUCTION STATE / NOT EXECUTED IN THIS EVIDENCE-REFRESH RELEASE

01

CONTROLLED SCENARIO

Replay a representative 30-day telemetry corpus with known attacks and operational noise.

02

CONTROLLED SCENARIO

Measure ingest latency, normalization quality, query latency, detection delay, and analyst workload.

03

CONTROLLED SCENARIO

Implement equivalent detections, risk scoring, cases, and response playbooks.

04

CONTROLLED SCENARIO

Model hot, warm, archive, replay, and legal-hold economics.

05

CONTROLLED SCENARIO

Exercise regional loss, collector outage, schema change, and source backpressure.

06

CONTROLLED SCENARIO

Export detections, cases, evidence, and normalized data to prove portability.

EVIDENCE PACKAGE

Preserve exact versions, architecture, configuration, scenario, expected result, instrumentation, raw output, timing, failures, operator actions, deviations, acceptance state, owner, and artifact hashes.

Current decision posture

The July 0.1-point Sentinel/Splunk difference is not decision certainty. Current data-lake, graph, AI, SOAR, and endpoint changes are material. Reproduce ingestion, detection, investigation, response, retention, migration, and cost behavior before profile selection.

CURRENT RANKING: NOT REISSUED / CONTROLLED LAB REQUIRED

CURRENT OFFICIAL-SOURCE REGISTER / CHECKED 2026-09-17

Microsoft Sentinel	E2 / OFFICIAL	
Microsoft Sentinel - What is new		learn.microsoft.com
Splunk Enterprise Security	E2 / OFFICIAL	
Splunk Enterprise Security 8.7 release notes		help.splunk.com
Google Security Operations	E2 / OFFICIAL	
Google SecOps SIEM/SOAR release notes		docs.cloud.google.com
Elastic Security	E2 / OFFICIAL	
Elastic Security release notes		www.elastic.co

REVISION LINEAGE

1.2.0-candidate / 2026-09-17 - current-source evidence refresh; no lab rescore issued.

1.1.0-candidate / 2026-07-24 - baseline benchmark using the final comparative evaluation system.

Trigger earlier review after major release, acquisition, licensing change, critical vulnerability, material outage, failed PoC, or workload change.

BASELINE ANALYTIC RECORD CONTINUES ON PAGE 11 FOR TRACEABILITY.

ATOMIC CRITERIA MODEL

WEIGHTED CAPABILITY WITH EXPLICIT MINIMUM EVIDENCE

ID	CRITERION	WEIGHT	GATE	MINIMUM EVIDENCE
C01	SIEM, detection, investigation, hunting, and response coverage	5	YES	Feature documentation, APIs, configuration exports, and scenario demonstration.
C02	Security and trust architecture	5	YES	Threat model, identity flows, RBAC tests, audit records, and security configuration.
C03	Governance and approval controls	5	YES	Approval workflow, policy controls, separation-of-duties tests, and decision records.
C04	Telemetry, schema, detection, and case-state integrity	5	YES	Backup, restore, rollback, drift, and corruption-recovery evidence.
C05	Automation and API quality	4	NO	API specifications, authentication tests, rate limits, event samples, and automation runs.
C06	Integration ecosystem	4	NO	Supported integrations, connector tests, failure behavior, and lifecycle ownership.
C07	Observability and evidence	4	YES	Logs, traces, metrics, reports, export tests, and evidence-retention controls.
C08	Ingestion, retention, query, detection, and case scale	4	NO	Controlled load tests, topology scale, saturation behavior, and capacity model.
C09	Resilience and continuity	5	YES	Failure injection, HA tests, recovery timing, degraded-mode operation, and runbooks.
C10	Usability and operator cognition	3	NO	Task observation, error-rate measures, navigation tests, and operator interviews.
C11	Change safety and rollback	5	YES	Plan or preview output, canary tests, rollback execution, and post-change verification.
C12	Extensibility and programmability	3	NO	Plugin, module, provider, workflow, or code-extension tests and upgrade impact analysis.
C13	Deployment and sovereignty options	3	NO	Deployment architecture, data-flow mapping, residency evidence, and offline tests.
C14	Lifecycle and upgrade discipline	4	NO	Release notes, upgrade test, compatibility matrix, support policy, and migration plan.
C15	Ingestion, retention, compute, staffing, and response economics	3	NO	Bill-of-materials, licensing model, staffing estimates, metering, and sensitivity analysis.
C16	Portability and exit	4	YES	Export tests, format analysis, replacement workflow, and decommission evidence.
C17	Vendor and ecosystem risk	3	NO	License analysis, roadmap evidence, bus-factor indicators, and dependency inventory.
C18	Assurance confidence	5	YES	Source provenance, lab artifacts, production evidence, independent replication, and review date.

SCORE SCALE

0 absent; 1 materially deficient; 2 partial; 3 adequate with limits; 4 strong; 5 leading for the defined profile. Scores remain provisional until the required evidence grade is achieved.

RAW CAPABILITY SCORECARD

DOCUMENTED CAPABILITY BEFORE EVIDENCE DISCOUNT

CRITERION	SENTINEL	SPLUNK ES	GOOGLE SECOPS	ELASTIC
C01 SIEM, detection, investigation, hunting, and response coverage	4.6	4.7	4.5	4.4
C02 Security and trust architecture	4.5	4.4	4.4	4.3
C03 Governance and approval controls	4.3	4.2	4.2	4.0
C04 Telemetry, schema, detection, and case-state integrity	4.4	4.4	4.3	4.2
C05 Automation and API quality	4.5	4.6	4.4	4.6
C06 Integration ecosystem	4.7	4.7	4.4	4.3
C07 Observability and evidence	4.6	4.7	4.6	4.4
C08 Ingestion, retention, query, detection, and case scale	4.7	4.6	4.8	4.5
C09 Resilience and continuity	4.5	4.4	4.5	4.2
C10 Usability and operator cognition	4.4	4.5	4.3	4.2
C11 Change safety and rollback	4.4	4.2	4.3	4.1
C12 Extensibility and programmability	4.2	4.6	4.1	4.6
C13 Deployment and sovereignty options	3.8	4.2	3.7	4.6
C14 Lifecycle and upgrade discipline	4.3	4.0	4.0	3.9
C15 Ingestion, retention, compute, staffing, and response economics	3.6	3.0	3.5	4.0
C16 Portability and exit	3.9	4.0	3.7	4.4
C17 Vendor and ecosystem risk	4.0	3.9	3.8	4.1
C18 Assurance confidence	3.6	3.6	3.5	3.5

WEIGHTED BASELINE / 100

Sentinel	Splunk ES	Google SecOp	Elastic
86.0	85.5	83.9	84.4

EVIDENCE-ADJUSTED SCORECARD

CAPABILITY DISCOUNTED BY CURRENT EVIDENCE CONFIDENCE

CANDIDATE	RAW	EVIDENCE CONFIDENCE	ADJUSTED	CURRENT STATE
Sentinel	86.0	58.2%	50.6	CONTROLLED LAB PENDING
Splunk ES	85.5	58.2%	50.5	CONTROLLED LAB PENDING
Google SecOps	83.9	58.2%	49.2	CONTROLLED LAB PENDING
Elastic	84.4	58.2%	49.6	CONTROLLED LAB PENDING

EVIDENCE SCALE

E0 / 0.00
Unsupported or unverified

E1 / 0.38
Vendor assertion or marketing statement

E2 / 0.64
Official documentation, release notes, or source repository

E3 / 0.78
Controlled Mythos laboratory result

E4 / 0.91
Production evidence from the adopting environment

E5 / 1.00
Independent repeatable validation

CURRENT CONFIDENCE BOUNDARY

Most candidate criteria are supported at E2: official documentation or source evidence. Environment-specific laboratory, production, and independent evidence remain future gates.

PROFILE-SPECIFIC RANKINGS

EVIDENCE-ADJUSTED SENSITIVITY ANALYSIS

MICROSOFT-CENTRIC SOC

Prioritizes Microsoft identity, endpoint, cloud, collaboration, data-lake, and automation integration.

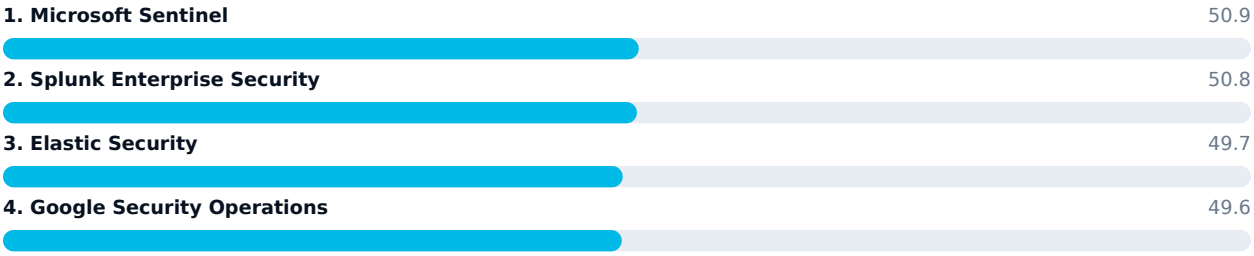

HIGH-SCALE SECURITY DATA PLATFORM

Prioritizes ingestion, retention, search, detection throughput, data architecture, and reliability.

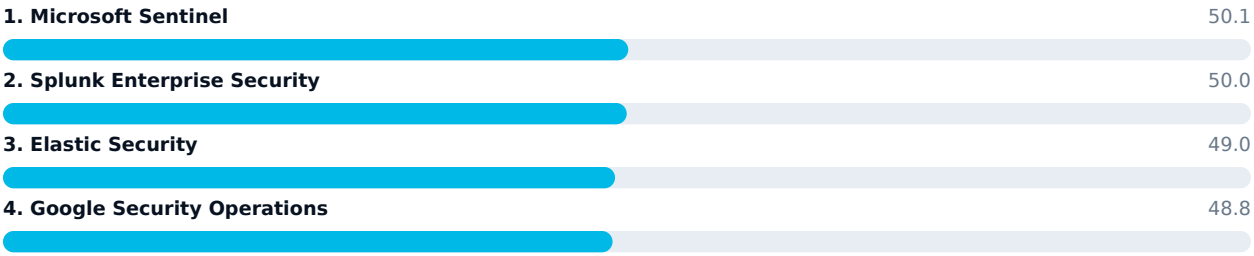

SOVEREIGN AND EXTENSIBLE SOC

Prioritizes deployment control, data portability, open integration, customization, and exit.

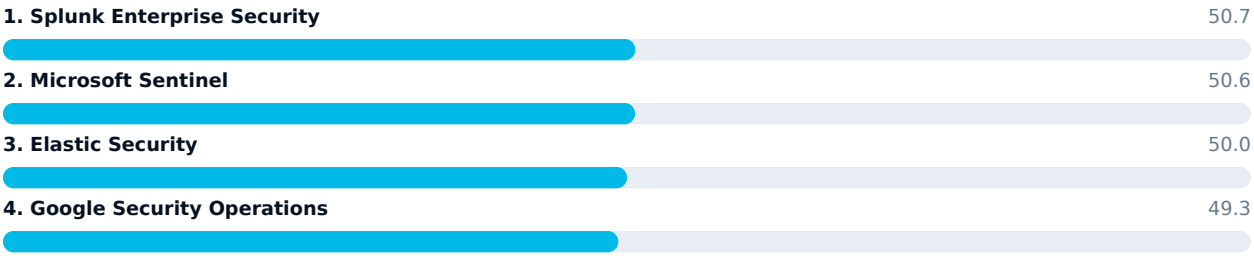

RANK SENSITIVITY

Small score differences are not decision certainty. Treat near-ties as a requirement for deeper PoC, commercial, migration, and operating-model evidence.

CANDIDATE DEEP DIVE / 01

MICROSOFT SENTINEL

OPERATING MODEL

Cloud-native Microsoft security operations platform spanning SIEM, data lake, analytics, investigation, automation, and Defender integration.

DOCUMENTED STRENGTHS

- Strong Microsoft ecosystem integration.
- Cloud-native analytics and automation.
- Broad connector and content ecosystem.
- Unified trajectory through the Defender portal and security data lake.

CAUTIONS AND PROOF OBLIGATIONS

- Cost depends heavily on ingestion, retention, query, and architecture.
- Portal and platform transitions require operational planning.
- Non-Microsoft data and response depth must be validated per source.

BASELINE SCORE PROFILE

Raw documented capability 86.0

Evidence-adjusted capability 50.6

GATE DISPOSITION

PASS - Microsoft-centric and cloud-native profiles | CONDITIONAL - ingestion economics and portal transition

CANDIDATE DEEP DIVE / 02

SPLUNK ENTERPRISE SECURITY

OPERATING MODEL

Search- and analytics-centric enterprise SIEM operating on Splunk platform capabilities with risk-based alerting and broad ecosystem support.

DOCUMENTED STRENGTHS

- Powerful search, data exploration, and extensibility.
- Mature enterprise SIEM content and operations.
- Risk-based alerting supports entity-centered detection.
- Flexible deployment and rich ecosystem.

CAUTIONS AND PROOF OBLIGATIONS

- Cost and platform engineering can be substantial.
- Data-model and content discipline are prerequisites.
- Scale and search performance require experienced architecture and operations.

BASELINE SCORE PROFILE

Raw documented capability 85.5

Evidence-adjusted capability 50.5

GATE DISPOSITION

PASS - search-intensive mature SOC profile | CONDITIONAL - cost and platform engineering

CANDIDATE DEEP DIVE / 03

GOOGLE SECURITY OPERATIONS

OPERATING MODEL

Cloud-native SIEM and SOAR platform with high-scale normalization, threat intelligence, detection, investigation, and playbooks.

DOCUMENTED STRENGTHS

- Strong cloud-scale data processing and search.
- Integrated threat intelligence and security analytics.
- SIEM and SOAR convergence.
- High-volume security-data orientation.

CAUTIONS AND PROOF OBLIGATIONS

- Commercial, regional, integration, and migration details require validation.
- Analyst and content migration can be significant.
- Operating model may depend on Google Cloud adjacency and partner ecosystem.

BASELINE SCORE PROFILE

Raw documented capability 83.9

Evidence-adjusted capability 49.2

GATE DISPOSITION

PASS - high-scale cloud security-data profile | CONDITIONAL - integration and operating-model fit

CANDIDATE DEEP DIVE / 04

ELASTIC SECURITY

OPERATING MODEL

Unified SIEM, XDR, endpoint, and cloud-security solution built on Elasticsearch and Kibana with flexible deployment options.

DOCUMENTED STRENGTHS

- Strong search and analytics foundation.
- Flexible self-managed and cloud deployment.
- Integrated endpoint and cloud-security capabilities.
- Open ecosystem and extensibility.

CAUTIONS AND PROOF OBLIGATIONS

- Enterprise operations require Elastic engineering capability.
- Feature and subscription boundaries need validation.
- Content, tuning, and response maturity vary by operating model.

BASELINE SCORE PROFILE

Raw documented capability 84.4

Evidence-adjusted capability 49.6

GATE DISPOSITION

PASS - sovereign and extensible profile | CONDITIONAL - engineering and content maturity

ARCHITECTURE AND INTEGRATION FIT

THE PRODUCT IS ONLY ONE PART OF THE OPERATING SYSTEM

REFERENCE OPERATING LAYERS

01 Telemetry acquisition and routing

02 Normalization and security data platform

03 Detection, analytics, hunting, and intelligence

04 Case, investigation, automation, and response

05 Evidence retention, reporting, and assurance

INTEGRATION BOUNDARIES

- Identity, endpoint, cloud, network, application, and threat-intelligence sources
- SOAR and response tools
- ITSM and collaboration
- Data lake, archive, and eDiscovery
- Detection-as-code and CI/CD
- Asset, vulnerability, and exposure systems

ARCHITECTURE GATE

Every external dependency requires an owner, identity boundary, failure mode, evidence path, version policy, recovery procedure, and exit strategy.

SECURITY, OPERATIONS, LIFECYCLE, ECONOMICS, AND EXIT

CROSS-DOMAIN DECISION RECORD

SECURITY AND AUTHORITY

Identity, credentials, secrets, roles, privileged actions, signing, approvals, isolation, audit, and incident response must be tested as an integrated system.

RELIABILITY AND RECOVERY

Control-plane, data-plane, dependency, region, database, queue, network, and operator failures require defined degraded modes and recovery objectives.

CHANGE AND LIFECYCLE

Version, compatibility, migration, canary, rollback, content, plugin, provider, firmware, and decommission procedures are part of product fitness.

ECONOMICS AND CAPACITY

Licenses, metering, data, compute, hardware, storage, support, staffing, training, integration, migration, and opportunity cost require sensitivity analysis.

SOVEREIGNTY AND PRIVACY

Data location, support access, telemetry, subprocessors, encryption, key control, disconnected operation, and legal obligations must align with policy.

PORTABILITY AND EXIT

Configuration, policy, data, state, evidence, workflows, identities, and operational knowledge must be exportable and reconstructable.

DECISION OWNERSHIP

Architecture, security, operations, finance, procurement, legal, data, and service owners jointly accept the final profile, evidence, residual risk, and exit obligations.

RISK AND FAILURE REGISTER

SCENARIOS THAT CAN INVALIDATE A FEATURE-BASED SELECTION

ID	SEVERITY	FAILURE SCENARIO	REQUIRED TREATMENT
R-01	CRITICAL	The cheapest ingest design removes telemetry required for detection or investigation.	PREVENT / DETECT / RECOVER / EVIDENCE
R-02	HIGH	A high query benchmark conceals weak schema, detection, and case workflows.	PREVENT / DETECT / RECOVER / EVIDENCE
R-03	HIGH	Automation executes destructive response without adequate approval and rollback.	PREVENT / DETECT / RECOVER / EVIDENCE
R-04	HIGH	Data residency or retention assumptions violate legal or contractual obligations.	PREVENT / DETECT / RECOVER / EVIDENCE
R-05	MEDIUM	Detection content is imported without validation against local log semantics.	PREVENT / DETECT / RECOVER / EVIDENCE
R-06	MEDIUM	Migration loses searches, cases, data models, risk history, and analyst knowledge.	PREVENT / DETECT / RECOVER / EVIDENCE

RISK RULE

A candidate with an unresolved critical failure path cannot advance because optional capability, market position, or commercial discount cannot compensate for missing safety or recovery evidence.

CONTROLLED PROOF-OF-CAPABILITY PLAN

REPEATABLE SCENARIOS, INSTRUMENTATION, AND ACCEPTANCE

TEST 01

Replay a representative 30-day telemetry corpus with known attacks and operational noise.

TEST 02

Measure ingest latency, normalization quality, query latency, detection delay, and analyst workload.

TEST 03

Implement equivalent detections, risk scoring, cases, and response playbooks.

TEST 04

Model hot, warm, archive, replay, and legal-hold economics.

TEST 05

Exercise regional loss, collector outage, schema change, and source backpressure.

TEST 06

Export detections, cases, evidence, and normalized data to prove portability.

EVIDENCE PACKAGE

Preserve versions, architecture, configuration, data set, scenario, expected result, instrumentation, raw output, timing, failures, operator actions, deviations, acceptance state, owner, and artifact hashes.

CONDITIONAL RECOMMENDATION AND REVALIDATION

DECISION RECORD, ACCEPTANCE, AND NEXT EVIDENCE

RECOMMENDATION POSITION

- Select by workload profile, data architecture, and analyst operating model.
- Use a controlled attack-and-noise replay rather than vendor demos.
- Model full data, compute, content, staffing, and migration economics.
- Revalidate detection quality and cost monthly during pilot and quarterly after adoption.

CANDIDATE	ADJ. SCORE	GATE POSITION	LAB STATE
Sentinel	50.6	PASS - Microsoft-centric and cloud-native profiles; CONDITIONAL - ingestion economics and portal transition	PENDING
Splunk ES	50.5	PASS - search-intensive mature SOC profile; CONDITIONAL - cost and platform engineering	PENDING
Google SecOps	49.2	PASS - high-scale cloud security-data profile; CONDITIONAL - integration and operating-model fit	PENDING
Elastic	49.6	PASS - sovereign and extensible profile; CONDITIONAL - engineering and content maturity	PENDING

REVALIDATION SCHEDULE

Review no later than 2026-10-24. Review earlier after a major release, commercial change, critical vulnerability, material outage, architecture transition, failed acceptance test, or change to the target workload profile.

SOURCE AUTHORITY AND REVISION

CURRENT EVIDENCE SNAPSHOT AND PUBLICATION HISTORY

SOURCE ID	PUBLISHER	TITLE	CHECKED
NIST-CSF-20	NIST	Cybersecurity Framework 2.0	2026-07-24
NIST-SP800-53	NIST	Security and Privacy Controls for Information Systems and Organizations	2026-07-24
CIS-CONTROLS	Center for Internet Security	CIS Critical Security Controls v8.1	2026-07-24
SENTINEL	Microsoft	Microsoft Sentinel Overview	2026-07-24
SENTINEL-AUTO	Microsoft	Microsoft Sentinel Automation Rules	2026-07-24
SPLUNK-ES	Splunk	Splunk Enterprise Security Documentation	2026-07-24
SPLUNK-RBA	Splunk	Splunk Risk-Based Alerting	2026-07-24
GOOGLE-SECOPS	Google Cloud	Google Security Operations Overview	2026-07-24
GOOGLE-SOAR	Google Cloud	Google Security Operations SOAR Overview	2026-07-24
ELASTIC-SEC	Elastic	Elastic Security Documentation	2026-07-24

SOURCE POSITION

Official documentation and source repositories support the current capability model. Source records preserve publisher, title, URL, purpose, and review date in the machine-readable authority register. Commercial terms, performance, and environment-specific behavior remain unverified until controlled proof.

REVISION HISTORY

1.1.0-candidate / 2026-07-24 - permanent-publication rebuild using the final benchmark system, profile-specific rankings, evidence adjustment, mandatory gates, and controlled PoC requirements.