

AI Search Visibility in Enterprise Security Information and Event Management (SIEM): A Multi-Engine, Multi-Market Benchmark Study

How four AI search engines, ChatGPT, Gemini, Google AI Overview and Google AI Mode, answer buyer-intent questions about enterprise SIEM and security operations platforms across the United States, Canada, India and Germany.

Table 1. Study size by AI engine — responses, cited sources, distinct domains cited and markets

AI engine	Responses	Cited sources	Distinct domains cited	Markets
ChatGPT	40,000	216,000	33	4
Gemini	40,000	124,000	47	4
Google AI Overview	40,000	393,000	139	4
Google AI Mode	40,000	221,000	70	4
Total	160,000	954,000	—	4

Published: October 2026

Research period: October 2026

Category: Enterprise security information and event management (SIEM) and security operations platforms

Scope: 10,000 buyer-intent queries · 160,000 AI engine responses · 4 AI engines · 4 markets

Evidence base: 954,000 cited sources · 921,000 brand mentions · 121 distinct product entries

Prepared by: GrackerAI · GrackerAI AI Search Visibility Benchmark Series

Abstract

Enterprise security buyers increasingly meet their first vendor shortlist inside an AI-generated response, so the vendors an AI engine names now shape the consideration set before any vendor is contacted. This study measures how four commercial AI search engines answer buyer-intent questions about enterprise security information and event management (SIEM). A set of 10,000 buyer-intent queries, built from five core buyer questions and reworded variants of each, was issued to ChatGPT, Gemini, Google AI Overview and Google AI Mode in the United States, Canada, India and Germany during October 2026. The 160,000 AI engine responses carried 954,000 cited sources. For every response the study recorded length, citations, cited domains, source reachability, hedging, recency anchoring, truncation, and each brand and product mention with its ordinal position. Microsoft is the most-mentioned brand in every engine and every market, with 161,000 of 921,000 brand mentions, and Microsoft Sentinel is the most-mentioned product entry at 118,000 mentions. The engines agree far less on evidence than on vendors: no pair of engines shares more than 0.29 of its cited domains, measured by the Jaccard coefficient. Citation behaviour diverges sharply, from 0.6 citations per 100 words in Gemini to 4.6 in Google AI Overview, and from a 94.9% vendor-site share of ChatGPT citations to 30.5% for Google AI Overview. The results show a stable shortlist produced from largely separate evidence bases, which means a measurement taken on one engine does not describe the others.

Contents

1. Introduction

- 1.1 Background
- 1.2 Why this category
- 1.3 Contribution

2. Research Questions

3. Methodology

- 3.1 Study design
- 3.2 Engines under test
- 3.3 Markets
- 3.4 Query construction
- 3.5 Corpus and unit of analysis
- 3.6 Metric definitions
- 3.7 Data collection and processing
- 3.8 Scope exclusions

4. Results

- 4.1 Response characteristics by engine
- 4.2 Brand visibility and share of voice

4.3 Product-level results

4.4 Citation volume and source concentration

4.5 Source quality and link decay

4.6 Cross-engine source overlap

4.7 First-mention position

4.8 Geographic variance

4.9 Inter-engine disagreement

5. Discussion

6. Threats to Validity

7. Limitations

8. Practical Implications

9. Conclusion

Disclosure

10. Reproducibility and Data Availability

How to cite this report

Appendices A-E

1. Introduction

1.1 Background

An AI search engine does not return a list of documents for the reader to rank. It returns a composed response that names a small number of vendors, places them in an order and attaches a reason to each. For a high-consideration enterprise security purchase, that single response performs work a buyer once spread across analyst notes, peer conversations and many search results. It sets the category boundary, states the evaluation criteria and nominates the candidates.

The consequence for a vendor is structural rather than competitive. A vendor absent from the returned shortlist is not losing on price, capability or references. It is simply not present when the consideration set forms. Conventional search metrics do not detect this condition: a vendor can rank well for its category terms and still be missing from the responses that increasingly come before a click.

1.2 Why this category

Security information and event management is a suitable instrument for four reasons. First, the vendor set is mature and heavily documented: the engines in this study named 100 distinct brand entries, and a group of 20 brands appears in all four engines. Second, the evaluation criteria are stable and publicly articulated. The engine responses recorded here repeatedly organise the category around the same capabilities: detection mapped to MITRE ATT&CK, user and entity behaviour analytics (UEBA), threat intelligence integration, built-in security orchestration, automation and response (SOAR), and the cost of log ingestion and retention. That shared vocabulary makes responses comparable across engines, markets and phrasings.

Third, the buying process is research-heavy: the core questions used here describe replacing a legacy SIEM, equipping a small security operations centre (SOC), and consolidating alerts from many tools. Fourth, a SIEM sits at the centre of security operations and is costly to replace, so omission from an early shortlist is hard to recover from later in the cycle.

1.3 Contribution

This study contributes a directly measured, cross-engine baseline for a single enterprise security category. It reports brand visibility and ordinal position separately rather than merging them into one share-of-voice figure. It quantifies how far the engines draw on shared or separate evidence, which determines whether a measurement on one engine generalises to another. It records length, citation density, source type and cited-source reachability as properties of the category's AI search layer. Finally, it tests whether the leading vendors hold their position when the buyer question changes.

2. Research Questions

1. RQ1. To what extent do the four AI engines converge on the same SIEM vendor shortlist?
2. RQ2. How is brand and product visibility distributed across the vendor population, and how concentrated is it?
3. RQ3. To what extent do the engines draw on a shared evidence base, and how does the composition of that evidence differ?
4. RQ4. Does the vendor named first, and the top pick, hold when the buyer question changes?

5. RQ5. Does vendor visibility vary across the four markets tested?

Each question is answered explicitly in Section 5, with reference to the results that support the answer.

3. Methodology

3.1 Study design

The study is observational and cross-sectional. Four commercial AI search engines were issued an identical set of buyer-intent questions across four markets within a single collection window in October 2026. No engine was primed or given conversational context beyond the question itself. Each response was captured as returned and analysed without modification.

3.2 Engines under test

Four engines were tested: ChatGPT, Gemini, Google AI Overview and Google AI Mode. All four are commercial services. None exposes a pinned model version to callers, and none guarantees deterministic output. Findings therefore describe engine behaviour during the collection window and are not attributable to a specific model release. Three of the four engines are Google surfaces; this composition is relevant to the overlap results in Section 4.6.

3.3 Markets

Four markets were tested: the United States, Canada, India and Germany. Market was varied to test whether engines localise vendor visibility in response to differing regulatory environments, data-residency expectations or regional vendor presence. Query text was held constant across markets; only the market context of the request varied.

3.4 Query construction

The study used a set of 10,000 enterprise buyer-intent queries about SIEM and security operations platforms, built from five core buyer questions, listed in Appendix D, and reworded variants of each. Every query was issued to all four engines in all four markets, giving 40,000 responses per engine and 160,000 in total. The core questions range from a direct request for the best SIEM platforms to a problem statement that never names the category. The combination tests whether vendor visibility is stable under rewording or sensitive to surface phrasing.

3.5 Corpus and unit of analysis

The unit of analysis is a single engine response to a single query in a single market. The corpus comprises 160,000 responses, distributed evenly at 40,000 per engine. All per-engine rates and means in this report are computed over each engine's responses.

3.6 Metric definitions

Every metric used in this report is defined below. Definitions precede their first use in the results.

Table 2. Metric definitions

Metric	Definition	Unit
Responses	Engine responses recorded and analysed, one per query per market per engine	count

Metric	Definition	Unit
Mean response length	Mean word count of an engine's responses	words
Citation density	Mean number of cited sources per 100 words of response text	per 100 words
Source variety	Index of the breadth of distinct domains an engine cites relative to the number of responses it gives. Higher values denote a wider spread of domains	index
Hedging rate	Share of responses containing at least one qualifying phrase from a fixed list (for example "it depends on" or "there is no single best")	%
Recency anchoring	Share of responses that tie their content to an explicit date, such as "as of" followed by a month and year	%
Truncation rate	Share of responses of 200 characters or more whose final character is not closing punctuation. Inferred from response structure, not observed directly	%
Refusal rate	Share of responses declining to answer the question	%
Total cited sources	Count of distinct cited pages recorded for an engine across the corpus	count
Distinct domains cited	Number of different domains appearing in an engine's citations. Not a corpus-wide total and not additive across engines	count
Source type	Classification of each cited domain into one of: vendor site, social, forum, review aggregator, news, government, marketplace, parent ecosystem, or other	category
Dead-link rate	Share of an engine's cited URLs that did not resolve when tested (not found, gone, server error, or unreachable host)	%
Unreachable citations	Total cited sources times the dead-link rate, per engine	count
Median source age	Median age, in days, of the dated material an engine cited at time of collection	days
Self-citation rate	Share of an engine's cited sources pointing at domains owned by its own parent ecosystem	%
Brand mention	One occurrence of a brand entry in a response. Counted separately from product mentions and not additive with them	count
Mean ordinal position	Mean position of a brand among the brands named in a response, where 1.0 is the first brand named	mean
Modal first-mention entity	The brand or product most often named first in an engine's responses	label
Product mention	One occurrence of a named product entry in a response. A brand may hold many product entries	count
Mean product rank	Mean ordinal position of a product entry within a response	mean
Mean sentiment	Mean polarity of the language surrounding a brand or product mention, on a range from -1.0 (negative) to 1.0 (positive)	-1.0 to 1.0
Top pick	The product or approach an engine most often placed first for a given core question	label
Top-pick stability	Share of an engine's question phrasings on which it returned its most frequent top pick	%
Pairwise agreement	Share of pairs of question phrasings on which an engine returned the same top pick	%
Jaccard source overlap	Size of the intersection of two engines' cited-domain sets divided by the size of their union. 1.00 denotes identical sets, 0.00 no shared domains	0.00 to 1.00

3.7 Data collection and processing

Queries were issued to each engine's public surface for the specified market and the returned response was captured in full. Captured responses were normalised to plain text before analysis. Citations were taken from the sources an engine attached to its response and from pages it linked within the text, and were de-duplicated by normalised URL. Each cited domain was assigned a source type against curated domain lists, with unlisted domains assigned to "other". Each cited URL was then requested to test whether it still resolved. Pages returning not-found, gone or server-error status, and hosts that could not be reached after a retry, were recorded as dead links. Pages behind an access wall were recorded as reachable, so access restrictions do not raise the dead-link rate.

Hedging, recency anchoring, refusal and truncation were detected with fixed text patterns applied to each response. Brand and product mentions were extracted by a model-based analysis pass over each response, which also recorded each mention's ordinal position and scored the sentiment of its surrounding language on the -1.0 to 1.0 range. Extraction recorded entities as named in the response. The brand table therefore includes some entries that are not SIEM vendors, such as framework names, adjacent tools and generic category terms; they are reported as recorded. Product entries were not consolidated: variant names for one product are kept as separate entries. The model and prompt used for the analysis pass, and the validation applied to its output, are not documented in the source dataset and are therefore not described here.

3.8 Scope exclusions

The study did not paraphrase or edit engine output. It did not fact-check engine claims against external sources, and it did not score responses for technical accuracy. It is not a product evaluation: vendor and product names appear only as observed in engine output, their ordering reflects engine behaviour rather than any assessment by the authors, and no placement was paid for or solicited. Pricing and specification statements made by engines were not verified and are not reported. Where engines disagree on a factual claim, this report records the disagreement without adjudicating it.

4. Results

This section reports what was observed. Interpretation is deferred to Section 5, and implications for practice to Section 8.

4.1 Response characteristics by engine

Response length, citation density and response construction differ between the four engines.

Table 3. Response characteristics by AI engine — n = 160,000 AI engine responses

AI engine	Responses	Mean words	Citation density	Source variety	Hedging	Recency anchoring	Truncation	Refusals
ChatGPT	40,000	444.5	1.3	1.0	6%	15%	42%	0%
Gemini	40,000	481.6	0.6	1.2	7%	3%	17%	0%
Google AI Overview	40,000	224.1	4.6	3.8	0%	24%	0%	0%
Google AI Mode	40,000	341.1	1.4	1.7	3%	5%	3%	0%

Mean response length spans a factor of 2.1, from 481.6 words for Gemini to 224.1 words for Google AI Overview. ChatGPT averages 444.5 words and Google AI Mode 341.1. Citation density runs in a different order. Google AI Overview records 4.6 cited sources per 100 words, against 1.4 for Google AI Mode, 1.3 for ChatGPT and 0.6 for Gemini. Google AI Overview therefore gives the shortest responses and the most citation-dense ones, while Gemini gives the longest responses and the least citation-dense ones. The single most citation-dense Google AI Overview response recorded 35.85 cited sources per 100 words.

Source variety is also highest for Google AI Overview, at 3.8, against 1.7 for Google AI Mode, 1.2 for Gemini and 1.0 for ChatGPT. Recency anchoring is highest in Google AI Overview at 24% of responses, followed by ChatGPT at 15%. Hedging language appears in between 0% and 7% of responses: 7% for Gemini, 6% for ChatGPT, 3% for Google AI Mode and 0% for Google AI Overview. No engine refused any query.

Response endings also differ. ChatGPT, at 42%, and Gemini, at 17%, record the highest truncation rates. These are the two engines with the longest mean responses. Google AI Mode records 3%, and Google AI Overview records no truncated responses.

4.2 Brand visibility and share of voice

The four engines produced 921,000 brand mentions across 100 distinct brand entries. Brand mentions were recorded per engine together with the mean ordinal position at which each brand appeared.




Figure 1. Brand mentions, ten most-mentioned brands. Brand mentions summed across four AI engines; n = 160,000 AI engine responses. Source: [GrackerAI AI Search Visibility Benchmark Series](#), October 2026.

Table 4. Brand mentions by AI engine — mentions, with mean ordinal position in parentheses; ten most-mentioned brands; n = 160,000 AI engine responses

Brand	Total	ChatGPT	Gemini	AI Overview	AI Mode
Microsoft	161,000	33,000 (1.2)	46,000 (1.7)	35,000 (1.5)	47,000 (1.6)
Splunk	87,000	21,000 (2.6)	27,000 (2.1)	20,000 (1.7)	19,000 (3.0)
Google	86,000	21,000 (2.7)	25,000 (3.2)	18,000 (2.8)	22,000 (2.8)

Brand	Total	ChatGPT	Gemini	AI Overview	AI Mode
CrowdStrike	66,000	10,000 (3.7)	21,000 (2.4)	14,000 (2.9)	21,000 (2.3)
Elastic	50,000	12,000 (4.4)	12,000 (3.7)	10,000 (3.9)	16,000 (4.4)
Securonix	45,000	3,000 (6.0)	15,000 (4.3)	12,000 (4.5)	15,000 (4.1)
Exabeam	40,000	4,000 (4.0)	13,000 (3.0)	11,000 (3.4)	12,000 (3.6)
Palo Alto Networks	36,000	6,000 (4.0)	15,000 (3.0)	6,000 (3.7)	9,000 (4.7)
IBM	30,000	15,000 (4.6)	7,000 (5.0)	4,000 (6.0)	4,000 (4.5)
Snowflake	25,000	1,000	11,000	7,000 (3.5)	6,000

Figure 1 and Table 4 show the ten most-mentioned brands. Together they account for 626,000 brand mentions, 68.0% of the 921,000 recorded. Microsoft records 161,000 mentions, 17.5% of all brand mentions and 74,000 more than the second-placed brand. Splunk follows at 87,000 and Google at 86,000, then CrowdStrike at 66,000 and Elastic at 50,000. Microsoft is the most-mentioned brand in each of the four engines, between 33,000 mentions in ChatGPT and 47,000 in Google AI Mode.

Twenty brands appear in all four engines, including all ten in Table 4. At the other end, 60 of the 100 brand entries appear in only one engine, and 47 record the minimum recorded count of 1,000 mentions. Gemini names the most distinct brand entries at 57, against 34 for ChatGPT, and 22 brand entries appear only in Gemini, among them Fortinet and Secure.com.

Coverage within the leading group varies by engine. Securonix records 15,000 mentions in Gemini and 3,000 in ChatGPT. Snowflake records 11,000 in Gemini and 1,000 in ChatGPT. IBM runs the other way, with 15,000 mentions in ChatGPT and 4,000 in each of the two Google AI surfaces. Gemini records the highest count of the four engines for five of the ten brands in Table 4 outright, and shares the highest count for two more.

Microsoft holds a mean ordinal position between 1.2 and 1.7 in every engine. In ChatGPT, Google AI Overview and Google AI Mode this is the earliest mean position of any brand in the full table. In Gemini the earliest mean position belongs to Rapid7, at 1.0 on 5,000 mentions; Microsoft's 1.7 is the earliest among the ten brands in Table 4.

Table 5. Mean sentiment by brand — twelve brands with recorded sentiment; possible range -1.0 to 1.0

Brand	Brand mentions	Mean sentiment	Engines naming the brand
Google	86,000	0.62	4 of 4
Securonix	45,000	0.61	4 of 4
Exabeam	40,000	0.59	4 of 4
Elastic	50,000	0.56	4 of 4
Microsoft	161,000	0.54	4 of 4
CrowdStrike	66,000	0.53	4 of 4
Rapid7	18,000	0.52	4 of 4
Splunk	87,000	0.43	4 of 4
Palo Alto Networks	36,000	0.41	4 of 4

Brand	Brand mentions	Mean sentiment	Engines naming the brand
Mandiant	24,000	0.30	4 of 4
Snowflake	25,000	0.21	4 of 4
IBM	30,000	0.09	4 of 4

Mean sentiment varies more widely than mention volume would imply. Among the twelve brands with recorded sentiment, Google holds the highest value at 0.62 and IBM the lowest at 0.09, a spread of 0.53 on a possible range of 2.0. Microsoft records 0.54 and Splunk 0.43. No brand in Table 5 carries negative mean sentiment, and all twelve appear in all four engines.

4.3 Product-level results

The corpus records 121 distinct product entries. Product mentions are counted separately from brand mentions and are not additive with them.

Table 6. Ten most-mentioned product entries — of 121 product entries recorded

#	Product	Brand	Mentions	Mean product rank	Mean sentiment	Engines	Markets
1	Microsoft Sentinel	Microsoft	118,000	1.6	0.71	4	4
2	Splunk Enterprise Security	Splunk	61,000	2.1	0.56	4	4
3	Elastic Security	Elastic	46,000	4.2	0.59	4	4
4	CrowdStrike Falcon Next-Gen SIEM	CrowdStrike	34,000	3.0	0.67	4	4
5	Securonix Unified Defense SIEM	Securonix	30,000	4.4	0.60	4	4
6	Google Security Operations (Chronicle)	Google	23,000	3.4	0.70	4	4
7	Palo Alto Cortex XSIAM	Palo Alto Networks	23,000	3.7	0.59	4	4
8	Google Security Operations	Google	21,000	3.0	0.66	4	4
9	IBM QRadar SIEM	IBM	21,000	4.8	0.30	4	4
10	Panther	Panther	17,000	2.7	0.71	4	4

Microsoft Sentinel records 118,000 product mentions at a mean product rank of 1.6, the earliest mean rank of the ten entries in Table 6. Splunk Enterprise Security follows at 61,000 mentions and a mean rank of 2.1, then Elastic Security at 46,000 and a mean rank of 4.2. Elastic Security ranks third by mentions with a mean rank of 4.2; among the ten entries in Table 6 only Securonix Unified Defense SIEM, at 4.4, and IBM QRadar SIEM, at 4.8, are later. IBM QRadar SIEM also records a mean sentiment of 0.30, the lowest in Table 6.

Of the 121 entries, 59 record the minimum recorded count of 1,000 mentions and 75 record 2,000 or fewer. Sixteen entries appear in all four engines and 22 in all four markets. Seven entries hold a mean product rank of 1.0; Microsoft Defender XDR is the only one of them with more than 2,000 mentions, at 10,000. Three entries carry negative mean sentiment: ArcSight at -0.40, and Datadog and ELK Stack at -0.20 each.

Table 7. Product-entry fragmentation by brand — brands recorded under five or more distinct product entries

Brand	Product entries	Combined mentions	Largest single entry	Its mentions	Its share
Exabeam	13	37,000	Exabeam Security Operations Platform	8,000	21.6%
Google	12	85,000	Google Security Operations (Chronicle)	23,000	27.1%
Splunk	10	84,000	Splunk Enterprise Security	61,000	72.6%
Microsoft	9	138,000	Microsoft Sentinel	118,000	85.5%
CrowdStrike	8	58,000	CrowdStrike Falcon Next-Gen SIEM	34,000	58.6%
Cribl	7	12,000	Cribl	6,000	50.0%
Rapid7	5	17,000	Rapid7 InsightIDR	12,000	70.6%
Securonix	5	46,000	Securonix Unified Defense SIEM	30,000	65.2%
SentinelOne	5	7,000	SentinelOne Singularity Data Lake	2,000	28.6%

Product naming is fragmented for several brands. Exabeam is recorded under 13 product entries totalling 37,000 mentions, the largest of which holds 21.6% of that total. Google is recorded under 12 entries totalling 85,000 mentions, all of them variant names for one product family: Google Security Operations, Google SecOps and Chronicle. The largest single Google entry holds 23,000 mentions, 27.1% of the brand's product total. Microsoft shows the opposite pattern: Microsoft Sentinel holds 85.5% of the 138,000 product mentions spread across nine Microsoft entries.

4.4 Citation volume and source concentration

The 160,000 responses carried 954,000 cited sources.




Figure 2. Citation volume by AI engine. Total cited sources per engine; n = 954,000 cited sources across 160,000 responses. Source: [GrackerAI AI Search Visibility Benchmark Series](#), October 2026.

Google AI Overview supplies 393,000 cited sources, 41.2% of the evidence base (Figure 2). Google AI Mode supplies 221,000 (23.2%), ChatGPT 216,000 (22.6%) and Gemini 124,000 (13.0%). The highest citation volume is 3.2 times the lowest, from the same 40,000 responses per engine. Distinct domains cited follow the same order at the extremes: 139 for Google AI Overview, 70 for Google AI Mode, 47 for Gemini and 33 for ChatGPT.




Figure 4. Source-type composition of citations by AI engine. Share of each engine's cited sources by classified source type; segments under 6% carry no label and are listed in Table B1; n = 954,000 cited sources. Source: [GrackerAI AI Search Visibility Benchmark Series](#), October 2026.

The composition of that evidence differs more than its volume (Figure 4). Vendor sites account for 500,000 of the 954,000 cited sources, or 52.4%. ChatGPT draws 94.9% of its citations from vendor sites. That share falls to 54.3% for Google AI Mode, 44.4% for Gemini and 30.5% for Google AI Overview. Social and forum sources account for 16.5% of Google AI Overview citations, against 3.6% for Google AI Mode, 0.8% for Gemini and none for ChatGPT. All 41,000 citations classified as social come from Google AI Overview. The residual “other” class is largest for Gemini at 54.0% of its citations.

Table 8. Three most-cited domains by AI engine — citation counts in parentheses

AI engine	First	Second	Third
ChatGPT	learn.microsoft.com (52,000)	splunk.com (27,000)	cloud.google.com (24,000)
Gemini	paloaltonetworks.com (12,000)	rapid7.com (8,000)	sentinelone.com (8,000)
Google AI Overview	linkedin.com (23,000)	splunk.com (18,000)	seceon.com (17,000)
Google AI Mode	exabeam.com (13,000)	learn.microsoft.com (12,000)	elastic.co (11,000)

The most-cited domains differ by engine. ChatGPT cites learn.microsoft.com 52,000 times, the highest single-domain count in Table 8, followed by splunk.com and cloud.google.com. Its three most-cited domains belong to the three most-mentioned brands. Gemini's three most-cited domains are vendor sites for Palo Alto Networks, Rapid7 and SentinelOne. Google AI Overview's most-cited domain is linkedin.com at 23,000. Google AI Mode's three most-cited domains are vendor sites for Exabeam, Microsoft and Elastic. Only learn.microsoft.com and splunk.com appear in the top three of more than one engine.

4.5 Source quality and link decay

Every cited URL was tested for reachability. Table 9 reports the result alongside the other source-quality measures.

Table 9. Citation volume and source quality by AI engine — n = 954,000 cited sources

AI engine	Total cited sources	Distinct domains cited	Dead-link rate	Unreachable citations	Median source age	Self-citation rate
ChatGPT	216,000	33	0.5%	1,080	42 days	0%
Gemini	124,000	47	3.2%	3,968	5 days	0%

AI engine	Total cited sources	Distinct domains cited	Dead-link rate	Unreachable citations	Median source age	Self-citation rate
Google AI Overview	393,000	139	5.1%	20,043	—	1%
Google AI Mode	221,000	70	6.3%	13,923	—	2%
Total	954,000	—	4.1%	39,014	—	—




Figure 3. Dead-link rate by AI engine. Share of each engine's cited URLs that did not resolve when tested; n = 954,000 cited sources. Source: [GrackerAI AI Search Visibility Benchmark Series](#), October 2026.

Of the 954,000 cited URLs tested, 39,014 did not resolve, an aggregate dead-link rate of 4.1%. The rate varies by a factor of 12.6 between engines, from 0.5% for ChatGPT to 6.3% for Google AI Mode (Figure 3). Google AI Overview records 5.1% and Gemini 3.2%. Google AI Overview contributes 20,043 of the 39,014 unreachable citations, the largest share of any engine, combining the highest citation volume with the second-highest dead-link rate.

Median source age is 42 days for ChatGPT and 5 days for Gemini. No median source age was recorded for Google AI Overview or Google AI Mode. Self-citation is low throughout: 2% of Google AI Mode citations and 1% of Google AI Overview citations point to Google-owned domains, and ChatGPT and Gemini record 0%. The source dataset also records an instance of a Google AI Mode response that explicitly promoted Google products.

4.6 Cross-engine source overlap

Jaccard source overlap measures the proportion of cited domains two engines hold in common.

Table 10. Jaccard source overlap between AI engines — overlap of cited-domain sets; 1.00 denotes identical sets

	ChatGPT	Gemini	Google AI Overview	Google AI Mode
ChatGPT	1.00	0.16	0.11	0.21
Gemini	0.16	1.00	0.18	0.29
Google AI Overview	0.11	0.18	1.00	0.24
Google AI Mode	0.21	0.29	0.24	1.00

Shading denotes overlap magnitude: 0.10–0.14 0.15–0.19 0.20–0.24 0.25–0.29

The highest overlap between two different engines is 0.29, between Gemini and Google AI Mode. The lowest is 0.11, between ChatGPT and Google AI Overview. All six engine pairs fall below 0.30, and the mean across the six pairs is 0.20. The two Google AI surfaces share 0.24 of their cited domains. ChatGPT records 0.16 with Gemini and 0.21 with Google AI Mode.

4.7 First-mention position

The entity most often named first in each engine's responses was recorded, together with each engine's top pick for each of the five core questions.

Table 11. Modal first-mention entity by AI engine — with Microsoft brand mentions and mean ordinal position in that engine

AI engine	Modal first-mention entity	Microsoft mentions in this engine	Microsoft mean ordinal position
ChatGPT	Microsoft	33,000	1.2
Gemini	Microsoft	46,000	1.7
Google AI Overview	Microsoft Sentinel	35,000	1.5
Google AI Mode	Microsoft Sentinel	47,000	1.6

Microsoft is the modal first-mention entity in all four engines. ChatGPT and Gemini most often name the Microsoft brand first; Google AI Overview and Google AI Mode most often name the Microsoft Sentinel product first. Microsoft's mean ordinal position ranges from 1.2 in ChatGPT to 1.7 in Gemini.

Table 12. Top pick by question phrasing and AI engine — Q1–Q5 are the five core questions listed in Appendix D

AI engine	Top-pick stability	Pairwise agreement	Q1	Q2	Q3	Q4	Q5
ChatGPT	80%	75%	Microsoft Sentinel	Microsoft Sentinel	Microsoft Sentinel	Microsoft Sentinel + XDR	Microsoft Sentinel
Gemini	60%	50%	Splunk	CrowdStrike LogScale	Microsoft Sentinel	XDR/SOAR	Microsoft Sentinel
Google AI Overview	75%	70%	Splunk	Microsoft Sentinel	Microsoft Sentinel	XDR/SOAR	Microsoft Sentinel
Google AI Mode	70%	65%	Microsoft Sentinel	Google SecOps	Microsoft Sentinel	XDR/SOAR	Microsoft Sentinel

Microsoft Sentinel is the top pick in 12 of the 20 engine-question combinations in Table 12, and in a 13th combination ChatGPT pairs it with XDR. All four engines place Microsoft Sentinel first for Q3, on SOAR and case management for a small SOC, and for Q5, on MITRE ATT&CK coverage, UEBA and threat intelligence. For Q4, a problem statement about alert volume that does not name the category, three engines place a product category, XDR or SOAR, first rather than a named product. Q1 and Q2 produce the widest spread: Splunk, CrowdStrike LogScale and Google SecOps each take the top position at least once.

Top-pick stability is 80% for ChatGPT, 75% for Google AI Overview, 70% for Google AI Mode and 60% for Gemini, a mean of 71% across the four engines. Pairwise agreement follows the same order, from 75% for ChatGPT to 50% for Gemini.

4.8 Geographic variance

The same three brands lead in all four markets.

Table 13. Leading brands by market — brands ordered by recorded mentions within each market, with market mean sentiment

Market	First	Second	Third	Fourth	Fifth	Mean sentiment
United States	Microsoft	Splunk	Google	CrowdStrike	Securonix	0.39
Canada	Microsoft	Google	Splunk	CrowdStrike	Elastic	0.44
India	Microsoft	Splunk	Google	CrowdStrike	Elastic	0.46
Germany	Microsoft	Google	Splunk	CrowdStrike	Elastic	0.46

Table 14. Leading product entries by market — three leading product entries per market, as recorded in the source dataset

Market	First	Second	Third
United States	Microsoft Sentinel (Microsoft)	Splunk Enterprise Security (Splunk)	Google Security Operations (Google)
Canada	Microsoft Sentinel (Microsoft)	CrowdStrike Falcon Next-Gen SIEM (CrowdStrike)	Elastic Security (Elastic)
India	Microsoft Sentinel (Microsoft)	Splunk Enterprise Security (Splunk)	CrowdStrike Falcon Next-Gen SIEM (CrowdStrike)
Germany	Microsoft Sentinel (Microsoft)	Splunk Enterprise Security (Splunk)	Google Security Operations (Google)

Microsoft leads in all four markets, and Microsoft, Splunk and Google form the leading three brands in all four. Only the order of second and third place changes: Google is second in Canada and Germany, Splunk in the United States and India. CrowdStrike is fourth everywhere. The fifth place is held by Elastic in three markets and by Securonix in the United States. Market mean sentiment ranges from 0.39 in the United States to 0.46 in India and Germany.

At product level, Microsoft Sentinel leads in all four markets and Splunk Enterprise Security is second in three (Table 14). Canada differs most: CrowdStrike Falcon Next-Gen SIEM and Elastic Security take second and third place. No engine produced market-specific vendor lists, and no market introduced a leading brand absent from the others.

4.9 Inter-engine disagreement

Two substantive disagreements between engines were recorded in the corpus.

On the best SIEM for a small SOC that needs built-in SOAR, Gemini frequently listed Fortinet FortiSIEM alongside Microsoft Sentinel, while Google AI Overview more often named Elastic Security and Google SecOps. On the primary choice for replacing a legacy SIEM, Gemini highlighted CrowdStrike Falcon LogScale and Panther, while

Google AI Mode more often led with Google SecOps or Microsoft Sentinel.

This report records both disagreements as observed without assessing which position is correct.

5. Discussion

This section answers each research question from the results in Section 4. Statements about mechanism are marked as interpretation.

5.1 RQ1: Do the engines converge on the same SIEM vendor shortlist?

They do, more strongly at the top than in the tail. Microsoft is the most-mentioned brand in every engine, and the ten most-mentioned brands all appear in all four engines (§4.2). Microsoft is also the modal first-mention entity in every engine (§4.7). Splunk and Google occupy the next positions in three of the four engines, and CrowdStrike and Elastic follow. A buyer consulting any of the four engines meets substantially the same leading group.

Convergence weakens below that group: 60 of the 100 brand entries appear in only one engine (§4.2). The data is consistent with a shared head of widely documented vendors and an engine-specific tail.

5.2 RQ2: How concentrated is visibility?

Visibility is concentrated at brand level. Ten brands hold 68.0% of brand mentions, and Microsoft alone holds 17.5%, almost twice the next brand (§4.2). At product level the concentration is stronger still for Microsoft, whose Sentinel entry holds 118,000 mentions against 61,000 for the next entry (§4.3).

Product-level counts are shaped by naming as well as by visibility. Google's 85,000 product mentions are spread across 12 entries, most of them variants of one product family, while Microsoft's are concentrated in one entry (§4.3). A plausible interpretation is that consistent product naming in a vendor's published material consolidates recorded visibility, while a recent rebrand, such as Chronicle to Google Security Operations, spreads it across several names. The study measures the fragmentation; it does not establish its cause.

5.3 RQ3: Do the engines share an evidence base?

Largely they do not. No pair of engines shares more than 0.29 of its cited domains, and the mean across all six pairs is 0.20 (§4.6). Even the two Google AI surfaces share only 0.24. The composition of evidence also diverges: ChatGPT relies almost entirely on vendor sites, while Google AI Overview draws close to half its citations from the residual “other” class and a sixth from social and forum sources (§4.4).

Read together with RQ1, this is the central result. The engines reach a similar shortlist from substantially different evidence. A plausible explanation is that the leading vendors are documented widely enough to surface from almost any retrieval path, so agreement on the head does not require shared sources. The practical consequence is methodological: a visibility programme instrumented on one engine observes a source population largely different from the one another engine reads.

5.4 RQ4: Do the first mention and the top pick hold when the question changes?

For the category leader, largely yes. Microsoft Sentinel is the top pick in 12 of 20 engine-question combinations, and in all four engines for two of the five core questions (§4.7). Mean top-pick stability is 71%. Where the top pick changes, it changes with the buyer's framing. The question about alert volume, which does not name the category,

draws a product category rather than a named product from three engines. The question about legacy replacement draws CrowdStrike LogScale or Google SecOps.

The data is consistent with engines matching the question's framing to a vendor's published positioning, with a problem-led question retrieving category-level explanations rather than vendor comparisons. The study cannot separate phrasing effects from run-to-run variation.

5.5 RQ5: Does visibility vary by market?

Membership of the leading group does not vary; ordering below first place varies slightly. Microsoft, Splunk and Google form the leading three in all four markets, and Microsoft Sentinel leads at product level everywhere (§4.8). Canada shows the largest departure, with CrowdStrike and Elastic products in the leading three product entries. Market mean sentiment varies within a narrow band from 0.39 to 0.46.

No engine produced market-specific vendor lists; the resource consequences are discussed in Section 8.

5.6 What these results do not resolve

The study observes engine behaviour but not its causes. It cannot say why an engine cites one domain rather than another, why ChatGPT draws almost exclusively on vendor sites, or why Google AI Overview alone cites social sources. It cannot distinguish a vendor absent because it is poorly documented from one absent because an engine's retrieval did not reach its documentation. It does not establish whether ordinal position affects buyer behaviour. Because the engines are not version-pinned, the study also cannot separate stable engine characteristics from conditions specific to October 2026.

6. Threats to Validity

6.1 Construct validity

Several metrics are proxies for the concepts they represent.

Mean ordinal position is a proxy for prominence; that earlier mention confers more influence on a reader is plausible but not measured here. The modal first-mention entity mixes brand and product names (Table 11).

Sentiment is scored by a model-based analysis pass whose model, prompt and validation are not documented in the source dataset (§3.7). The metric may partly reflect the descriptive register of vendor summaries rather than a judgement about the vendor.

Truncation is inferred from how a response ends. A response ending in a table or list item without closing punctuation can be recorded as truncated, so the rate is an upper indication rather than a count of incomplete responses.

Brand extraction records entities as named, including some non-vendor terms, and product variants are kept separate (§3.7). Tail brand counts and product-entry counts per brand carry this caveat.

6.2 Internal validity

All data comes from a single collection window in October 2026. Engine output is non-deterministic, so repeating the identical query set would not reproduce identical responses. The engines are commercial services without pinned model versions, and any of them may have changed behaviour during or after the window without notice.

Personalisation and session state were not controlled beyond issuing each query without prior conversational context; residual market-level personalisation cannot be excluded.

6.3 External validity

The study covers one category, four engines, four markets and one collection window. Three of the four engines are Google surfaces, which weights the corpus toward one vendor family. The query set is purposive rather than randomly sampled from a defined population of buyer questions, and is built from five core questions. For these reasons no statistical inference is performed anywhere in this report. There are no significance tests, no confidence intervals and no claims of generalisation beyond the observed corpus. All figures are descriptive statistics of what was observed.

Findings for SIEM should not be transferred to other security categories without measurement. Categories differ in vendor-set maturity, volume of published material and stability of terminology, all of which plausibly affect engine behaviour.

6.4 Reliability

A repeat study would be expected to reproduce the broad structure of these results. That includes Microsoft at the head of the shortlist, a leading group of Microsoft, Splunk and Google, low cross-engine source overlap, and the contrast between vendor-site-led and mixed-source citation profiles. It would not be expected to reproduce exact mention counts, exact ordinal positions or the specific most-cited domains, all of which depend on retrieval conditions at the moment of collection. Categorical findings, such as the leading brand by engine and by market, are more stable than continuous ones, but have not been tested across collection windows.

7. Limitations

- One category, four engines, four markets and one collection window in October 2026; the query set is built from five core buyer questions and their reworded variants.
- Engine responses were not fact-checked; where engines disagree on a factual claim, the disagreement is recorded but not adjudicated.
- Brand and product entities were not consolidated, so variant product names are reported as separate entries.
- The model, prompt and validation used for mention extraction and sentiment scoring are not documented in the source dataset.
- Median source age was recorded for two of the four engines only.

8. Practical Implications

This section is derived from the results rather than measured by them.

- Measure more than one engine, and include at least one non-Google engine. No pair of engines shares more than 0.29 of its cited domains (§4.6). A programme measuring one engine observes a source population largely different from the others.
- Match your evidence to each engine's source mix. ChatGPT draws 94.9% of its citations from vendor sites, while Google AI Overview draws 30.5% from vendor sites and 16.5% from social and forum sources (§4.4). Product documentation serves the first; third-party discussion and comparison content serve the second.
- Write for the problem-led question, not only the category question. For the alert-volume question, which never names the category, three engines put a product category first rather than any named product (§4.7). A vendor whose content does not connect the buyer's problem to its product is absent from that response.
- Use one product name consistently. Google's 85,000 product mentions are spread across 12 entries, while Microsoft Sentinel concentrates 85.5% of Microsoft's product mentions in one (§4.3). Consistent naming across your own and third-party material is the precondition for consolidated recorded visibility.
- Do not fund market-specific variants on the assumption that engines localise. The leading three brands are identical in all four markets (§4.8). Resources aimed at locale variants are likely to return more if directed at documentation depth and problem-led content.
- Keep cited pages reachable. Google AI Mode records a 6.3% dead-link rate and Google AI Overview 5.1% (§4.5). A cited page that no longer resolves cannot carry a vendor's claims.

9. Conclusion

Across 160,000 AI engine responses and 954,000 cited sources, four AI search engines converge on one leader for enterprise SIEM. Microsoft is the most-mentioned brand and the modal first-mention entity in every engine, and Microsoft Sentinel is the most-mentioned product and the most common top pick across the five core questions. Splunk and Google complete a leading group that is identical in all four markets.

Agreement on vendors sits alongside disagreement on evidence. No two engines share more than 0.29 of their cited domains, and the engines differ in what kind of evidence they read: ChatGPT reads vendor documentation, Google AI Overview reads a broad mix that includes social and forum content, and Gemini cites sparingly. A similar shortlist is being produced from substantially different sources.

For a vendor inside the leading group, the position appears durable across engines, markets and most question phrasings. For a vendor outside it, the route to visibility differs by engine, and runs through the source types each engine reads, consistent product naming, and content that answers the buyer's problem as well as the category question.

Disclosure

GrackerAI publishes this research and sells AI search visibility measurement for the category it covers. The study measured how AI engines describe vendors in their responses. It did not test, evaluate or rank any vendor's product, and no vendor paid for or requested placement. Vendor and product names appear only as observed in engine output.

10. Reproducibility and Data Availability

Queries were issued to each engine's public surface through a data collection layer, and each response, its cited sources and its per-response analysis were stored in an analytical database. Retained records include response text, citations with domain, source type and reachability, mentions with position and sentiment, and the per-engine aggregates reported here.

Reproducing the study requires the five core questions in Appendix D and their reworded forms, the four market contexts, access to the four engine surfaces during a comparable window, and the metric definitions in Table 2. Because the engines are not version-pinned and their output is non-deterministic, an exact reproduction of these figures is not achievable. The appropriate reproduction target is the structure of the findings rather than their values.

The underlying dataset for this report was generated on October 2026. Requests for the aggregate dataset supporting the tables in this report may be directed to GrackerAI.

How to cite this report

GrackerAI. (October 2026). *AI Search Visibility in Enterprise Security Information and Event Management (SIEM): A Multi-Engine, Multi-Market Benchmark Study*. GrackerAI AI Search Visibility Benchmark Series. Available at: <https://gracker.ai/data-and-research-reports/ai-search-visibility-siem-2026/>

Appendix A. Brand visibility by AI engine

Figures are brand mentions, with mean ordinal position in parentheses where recorded. An em dash indicates the brand was not recorded in that engine. Mean ordinal position is stated on a range where 1.0 denotes the first brand named in a response. Brand entries are reported as extracted and include some entities that are not SIEM vendors.

Table A1. Brand mentions and mean ordinal position by AI engine — the 27 brand entries recording 5,000 mentions or more

Brand	Total	ChatGPT	Gemini	AI Overview	AI Mode
Microsoft	161,000	33,000 (1.2)	46,000 (1.7)	35,000 (1.5)	47,000 (1.6)
Splunk	87,000	21,000 (2.6)	27,000 (2.1)	20,000 (1.7)	19,000 (3.0)
Google	86,000	21,000 (2.7)	25,000 (3.2)	18,000 (2.8)	22,000 (2.8)
CrowdStrike	66,000	10,000 (3.7)	21,000 (2.4)	14,000 (2.9)	21,000 (2.3)
Elastic	50,000	12,000 (4.4)	12,000 (3.7)	10,000 (3.9)	16,000 (4.4)
Securonix	45,000	3,000 (6.0)	15,000 (4.3)	12,000 (4.5)	15,000 (4.1)
Exabeam	40,000	4,000 (4.0)	13,000 (3.0)	11,000 (3.4)	12,000 (3.6)
Palo Alto Networks	36,000	6,000 (4.0)	15,000 (3.0)	6,000 (3.7)	9,000 (4.7)
IBM	30,000	15,000 (4.6)	7,000 (5.0)	4,000 (6.0)	4,000 (4.5)
Snowflake	25,000	1,000	11,000	7,000 (3.5)	6,000
Mandiant	24,000	7,000	10,000	1,000	6,000
Panther	18,000	3,000 (3.3)	7,000 (3.1)	4,000 (1.8)	4,000 (2.7)
Rapid7	18,000	3,000 (4.3)	5,000 (1.0)	4,000 (4.8)	6,000 (3.8)
SentinelOne	15,000	2,000 (4.5)	7,000 (3.0)	3,000 (3.5)	3,000
AWS	14,000	1,000	6,000	2,000	5,000
Cribl	12,000	4,000 (1.8)	4,000	3,000 (5.0)	1,000
Cisco Talos	11,000	1,000	5,000	3,000	2,000
Sumo Logic	11,000	2,000 (7.0)	1,000 (5.0)	3,000 (4.3)	5,000 (4.3)
Amazon S3	10,000	—	4,000	2,000 (4.0)	4,000

Brand	Total	ChatGPT	Gemini	AI Overview	AI Mode
VirusTotal	9,000	5,000	3,000	—	1,000
Databricks	6,000	1,000	2,000	2,000 (3.5)	1,000
Fortinet	6,000	—	6,000 (2.7)	—	—
Tines	6,000	1,000	4,000	1,000	—
Wazuh	6,000	—	3,000	1,000 (6.0)	2,000 (4.0)
Siemplify	5,000	—	1,000	—	4,000
Stellar Cyber	5,000	—	1,000 (2.0)	—	4,000 (2.3)
Torq	5,000	1,000	2,000	1,000	1,000

The matrix makes coverage gaps explicit: Fortinet records 6,000 mentions, all of them in Gemini.

Appendix B. Source mix and most-cited domains

Table B1 gives the full source-type distribution behind Figure 4. The three most-cited domains for each engine are given in Table 8 (§4.4).

Table B1. Source-type distribution by AI engine — cited sources by classified source type; n = 954,000

AI engine	Vendor site	Other	Social	Forum	Review aggregator	News	Parent ecosystem	Government	Marketplace	Total
ChatGPT	205,000	7,000	—	—	1,000	2,000	—	1,000	—	216,000
Gemini	55,000	67,000	—	1,000	1,000	—	—	—	—	124,000
Google AI Overview	120,000	191,000	41,000	24,000	13,000	2,000	2,000	—	—	393,000
Google AI Mode	120,000	84,000	—	8,000	3,000	—	5,000	—	1,000	221,000
Total	500,000	349,000	41,000	33,000	18,000	4,000	7,000	1,000	1,000	954,000

Parent-ecosystem citations total 7,000, all from the two Google AI surfaces.

Appendix C. Product entries

The corpus records 121 distinct product entries, all listed below in descending order of mentions. Mean product rank is the mean ordinal position of the entry within a response, where 1.0 denotes first named; an em dash indicates no rank was recorded. Entries are reported as extracted, without consolidating variant names.

Table C1. All product entries, ordered by mentions — 121 entries; engines and markets give the number of the four of each in which the entry was recorded

#	Product	Brand	Mentions	Mean product rank	Mean sentiment	Engines	Markets
1	Microsoft Sentinel	Microsoft	118,000	1.6	0.71	4	4
2	Splunk Enterprise Security	Splunk	61,000	2.1	0.56	4	4
3	Elastic Security	Elastic	46,000	4.2	0.59	4	4
4	CrowdStrike Falcon Next-Gen SIEM	CrowdStrike	34,000	3.0	0.67	4	4
5	Securonix Unified Defense SIEM	Securonix	30,000	4.4	0.60	4	4
6	Google Security Operations (Chronicle)	Google	23,000	3.4	0.70	4	4
7	Palo Alto Cortex XSIAM	Palo Alto Networks	23,000	3.7	0.59	4	4
8	Google Security Operations	Google	21,000	3.0	0.66	4	4
9	IBM QRadar SIEM	IBM	21,000	4.8	0.30	4	4
10	Panther	Panther	17,000	2.7	0.71	4	4
11	Rapid7 InsightIDR	Rapid7	12,000	4.0	0.63	4	4
12	Sumo Logic Cloud SIEM	Sumo Logic	11,000	4.9	0.59	4	4
13	Microsoft Defender XDR	Microsoft	10,000	1.0	0.40	4	4
14	Google SecOps (Chronicle)	Google	9,000	2.4	0.72	2	4

#	Product	Brand	Mentions	Mean product rank	Mean sentiment	Engines	Markets
15	Splunk SOAR	Splunk	9,000	2.7	0.38	3	4
16	Google Security Operations (Google SecOps)	Google	8,000	2.5	0.68	3	4
17	Exabeam Security Operations Platform	Exabeam	8,000	2.6	0.74	3	4
18	FortiSIEM	Fortinet	8,000	3.0	0.68	2	4
19	CrowdStrike Falcon	CrowdStrike	8,000	3.5	0.36	3	4
20	Securonix	Securonix	8,000	5.1	0.56	4	3
21	CrowdStrike Falcon LogScale	CrowdStrike	7,000	1.7	0.71	4	4
22	Google Security Operations (SecOps)	Google	7,000	2.6	0.66	3	3
23	Exabeam Fusion	Exabeam	7,000	2.7	0.69	3	3
24	Cribl	Cribl	6,000	2.0	0.60	3	4
25	Securonix Next- Gen SIEM	Securonix	6,000	3.3	0.74	3	3
26	Wazuh	Wazuh	6,000	4.7	0.05	3	3
27	Tines	Tines	6,000	—	0.22	3	4
28	CrowdStrike Falcon Next- Gen SIEM (LogScale)	CrowdStrike	5,000	1.3	0.78	2	3
29	Google Chronicle	Google	4,000	2.0	0.30	3	3
30	Splunk Enterprise Security (ES)	Splunk	4,000	2.8	0.50	1	3
31	Google SecOps	Google	4,000	3.0	0.40	2	2
32	Palo Alto Networks Cortex XDR	Palo Alto Networks	4,000	3.0	0.40	2	2
33	Exabeam New - Scale Fusion	Exabeam	4,000	3.3	0.65	2	3
34	Snowflake	Snowflake	4,000	3.5	0.41	2	2
35	Exabeam New - Scale	Exabeam	4,000	4.5	0.63	4	3
36	Palo Alto Cortex XSOAR	Palo Alto	4,000	—	0.28	2	3
37	Torq	Torq	4,000	—	0.28	3	3
38	Google Security Operations (formerly Chronicle)	Google	3,000	1.3	0.75	3	2
39	Stellar Cyber Open XDR	Stellar Cyber	3,000	2.0	0.78	2	2
40	Exabeam Security Operations Platform (Fusion SIEM)	Exabeam	3,000	3.0	0.67	1	3
41	LogRhythm	LogRhythm	3,000	3.3	0.73	1	2
42	Exabeam	Exabeam	3,000	4.7	0.58	2	3
43	IBM QRadar + QRadar SOAR	IBM	3,000	5.3	0.27	1	3
44	Azure Logic Apps	Microsoft	3,000	—	0.23	2	2
45	Shuffle	Shuffle	3,000	—	0.30	1	2
46	TheHive	TheHive	3,000	—	0.30	1	2
47	Microsoft Sentinel + Defender	Microsoft	2,000	1.0	0.75	1	2
48	Google Security Operations (Google SecOps / Chronicle)	Google	2,000	2.0	0.80	1	2
49	Google Chronicle Security Operations	Google	2,000	3.0	0.70	2	1
50	SentinelOne Singularity Data Lake	SentinelOne	2,000	3.0	0.70	2	2
51	Anomali Unified Security Data Lake	Anomali	2,000	3.5	0.65	1	2
52	Databricks	Databricks	2,000	3.5	0.63	1	1
53	Rapid7 Incident Command (InsightIDR)	Rapid7	2,000	3.5	0.75	2	2
54	Splunk Enterprise Security Premier	Splunk	2,000	3.5	0.45	2	1
55	Exabeam Security Operations Platform (New - Scale Fusion)	Exabeam	2,000	4.0	0.65	1	2
56	Gurukul SIEM	Gurukul	2,000	4.0	0.72	2	1
57	LimaCharlie	LimaCharlie	2,000	4.0	0.65	1	2
58	Splunk Enterprise Security + Splunk SOAR	Splunk	2,000	4.0	0.50	2	2
59	Datadog Security Monitoring	Datadog	2,000	5.0	0.60	2	2
60	SentinelOne Singularity AI SIEM	SentinelOne	2,000	5.0	0.55	2	2
61	Splunk User Behavior Analytics (UBA)	Splunk	2,000	—	0.50	2	2
62	Swimlane	Swimlane	2,000	—	0.25	2	1
63	Cribl Detect + Cribl Lake	Cribl	1,000	1.0	0.85	1	1
64	CrowdStrike Falcon Next- Gen SIEM (LogScale / Humio)	CrowdStrike	1,000	1.0	0.80	1	1
65	Gurukul AI - Powered Next- Gen SIEM	Gurukul	1,000	1.0	0.85	1	1
66	Microsoft Sentinel / Sentinel Data Lake	Microsoft	1,000	1.0	0.70	1	1
67	Stellar Cyber Open XDR / Next- Gen SIEM	Stellar Cyber	1,000	1.0	0.90	1	1

#	Product	Brand	Mentions	Mean product rank	Mean sentiment	Engines	Markets
68	Cribl / Cribl Detect	Cribl	1,000	2.0	0.80	1	1
69	Elastic Security (Elastic Workflows)	Elastic	1,000	2.0	0.80	1	1
70	Exabeam New - Scale / LogRhythm SIEM	Exabeam	1,000	2.0	0.70	1	1
71	Google Security Operations (Google Chronicle)	Google	1,000	2.0	0.60	1	1
72	Graylog	Graylog	1,000	2.0	0.70	1	1
73	Matano	Matano	1,000	2.0	0.90	1	1
74	Rapid7 Incident Command	Rapid7	1,000	2.0	0.70	1	1
75	SentinelOne Singularity SIEM	SentinelOne	1,000	2.0	0.70	1	1
76	Anomali Security Analytics / Modernized SIEM	Anomali	1,000	3.0	0.80	1	1
77	LogRhythm NextGen SIEM	LogRhythm	1,000	3.0	0.60	1	1
78	AWS Security Lake	Amazon	1,000	4.0	0.60	1	1
79	Coralogix	Coralogix	1,000	4.0	0.80	1	1
80	Elastic Security (Elasticsearch Platform)	Elastic	1,000	4.0	0.70	1	1
81	Exabeam New - Gen SIEM	Exabeam	1,000	4.0	0.60	1	1
82	Exabeam New - Gen SIEM (Fusion)	Exabeam	1,000	4.0	0.60	1	1
83	Exabeam New - Gen SIEM Portfolio	Exabeam	1,000	4.0	0.50	1	1
84	Google SecOps (formerly Chronicle SIEM)	Google	1,000	4.0	0.70	1	1
85	Gurukul Security Analytics	Gurukul	1,000	4.0	0.60	1	1
86	Panther Labs (Security Data Platform)	Panther Labs	1,000	4.0	0.70	1	1
87	Securonix EON (with BYO Snowflake)	Securonix	1,000	4.0	0.60	1	1
88	SentinelOne Singularity AI SIEM / Data Lake	SentinelOne	1,000	4.0	0.55	1	1
89	Splunk Enterprise Security (with Splunk UEBA)	Splunk	1,000	4.0	0.50	1	1
90	Splunk Security Operations Suite	Splunk	1,000	4.0	0.40	1	1
91	Stellar Cyber	Stellar Cyber	1,000	4.0	0.70	1	1
92	Cribl (Telemetry Pipeline)	Cribl	1,000	5.0	0.60	1	1
93	Exabeam New - Scale Security Operations Platform	Exabeam	1,000	5.0	0.70	1	1
94	Exabeam Security Log Management & SIEM	Exabeam	1,000	5.0	0.60	1	1
95	ManageEngine Log360	ManageEngine	1,000	5.0	0.80	1	1
96	Splunk UEBA	Splunk	1,000	5.0	0.50	1	1
97	AWS Security Lake + analytics	AWS	1,000	—	0.30	1	1
98	AlienVault USM	AlienVault / AT&T Cybersecurity	1,000	—	0.40	1	1
99	Amazon S3	Amazon	1,000	—	0.40	1	1
100	ArcSight	OpenText	1,000	—	-0.40	1	1
101	Axoflow	Axoflow	1,000	—	0.60	1	1
102	Charlotte AI	CrowdStrike	1,000	—	0.50	1	1
103	Cribl Pipeline	Cribl	1,000	—	0.50	1	1
104	Cribl Stream	Cribl	1,000	—	0.70	1	1
105	Cribl Stream / Lake	Cribl	1,000	—	0.70	1	1
106	CrowdStrike	CrowdStrike	1,000	—	0.00	1	1
107	CrowdStrike Next - Gen SIEM	CrowdStrike	1,000	—	0.00	1	1
108	D3 Security	D3 Security	1,000	—	0.50	1	1
109	Datadog	Datadog	1,000	—	-0.20	1	1
110	Dropzone AI	Dropzone AI	1,000	—	0.60	1	1
111	ELK Stack	Elastic	1,000	—	-0.20	1	1
112	FortiSOAR	Fortinet	1,000	—	0.40	1	1
113	Microsoft 365 Defender	Microsoft	1,000	—	0.00	1	1
114	Microsoft Azure	Microsoft	1,000	—	0.20	1	1
115	Microsoft Defender Threat Intelligence	Microsoft	1,000	—	0.30	1	1
116	Microsoft Unified SecOps Platform	Microsoft	1,000	—	0.30	1	1
117	Rapid7 InsightConnect	Rapid7	1,000	—	0.00	1	1
118	Rapid7 Threat Command	Rapid7	1,000	—	0.30	1	1
119	Securonix UEBA	Securonix	1,000	—	0.70	1	1
120	SentinelOne	SentinelOne	1,000	—	0.50	1	1
121	Splunk User Behavior Analytics	Splunk	1,000	—	0.10	1	1

Appendix D. Query set and markets

Table D1. Core buyer questions

Ref.	Core question
Q1	Best SIEM platforms for enterprise security operations in 2026
Q2	Which security analytics platforms can replace a legacy SIEM with cloud-native log ingestion and lower storage costs?
Q3	Which SIEM tools include built-in SOAR automation and case management for a small SOC team?
Q4	Our SOC gets thousands of alerts a day from different security tools and analysts cannot keep up. What kind of platform should we buy to centralize detection and response?
Q5	Which SIEM software supports MITRE ATT&CK detection coverage, UEBA and threat intelligence integration out of the box?

Table D2. Collection parameters

Parameter	Value
AI engines	ChatGPT, Gemini, Google AI Overview, Google AI Mode
Markets	United States, Canada, India, Germany
Buyer-intent queries	10,000
Core questions	5
Responses per engine	40,000
Total responses	160,000
Total cited sources	954,000
Collection window	October 2026
Dataset generated	October 2026

Appendix E. Metric computation

Table E1. Formulas for derived metrics

Metric	Computation
Total responses	queries × engines × markets (10,000 × 4 × 4 = 160,000)
Responses per engine	queries × markets (10,000 × 4 = 40,000)
Citation density	mean over responses of (cited sources ÷ words × 100)
Unreachable citations	total cited sources × dead-link rate, per engine
Aggregate dead-link rate	Σ unreachable citations ÷ Σ total cited sources (39,014 ÷ 954,000)

Metric	Computation
Share of cited sources	engine total cited sources $\div$ Σ total cited sources
Source-type share	cited sources of a source type $\div$ engine total cited sources
Brand mention total	Σ brand mentions across the four engines, per brand
Share of brand mentions	brand mentions $\div$ Σ brand mentions across all 100 brand entries (921,000)
Mean ordinal position	mean position of a brand among the brands named in each response in which it appears
Top-pick stability	phrasings returning the engine's most frequent top pick $\div$ phrasings tested
Jaccard source overlap	$ A \cap B \div A \cup B $, where A and B are two engines' cited-domain sets