

AI Search Visibility in Enterprise Access Management (Single Sign-On)

A multi-engine, multi-market benchmark study of which vendors four AI search engines (ChatGPT, Gemini, Google AI Overview and Google AI Mode) name, rank and cite when asked buyer-intent questions about enterprise single sign-on, across the United States, Canada, India and Germany.

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

AI engine	Responses	Cited sources	Markets
ChatGPT	40,000	306,000	4
Gemini	40,000	90,000	4
Google AI Overview	40,000	495,000	4
Google AI Mode	40,000	203,000	4
Total	160,000	1,094,000	4

Published October 2026

Research period October 2026

Scope Access management (enterprise single sign-on and workforce identity), four AI engines, four markets

Evidence base 160,000 AI engine responses, 1,094,000 cited sources, 934,000 brand mentions, 871,000 product mentions

Prepared by GrackerAI, GrackerAI AI Search Visibility Benchmark Series

Abstract

Enterprise buyers increasingly meet their first vendor shortlist inside an AI engine response rather than on a search results page or an analyst grid. This study measures which access management vendors four AI engines name, in what order, and on what evidence, when asked buyer-intent questions about enterprise single sign-on. 10,000 buyer-intent queries were issued to ChatGPT, Gemini, Google AI Overview and Google AI Mode in the United States, Canada, India and Germany in October 2026. The analysed corpus contains 160,000 responses, 1,094,000 cited sources and 934,000 brand mentions. Microsoft and Okta account for 37.8% of all brand mentions, and both appear in every engine and every market. Twelve of 67 brand entities appear in all four engines, while 38 appear in only one. Evidence bases diverge sharply: vendor sites supply 82.4% of ChatGPT's cited sources but 14.1% of Google AI Overview's, and pairwise Jaccard overlap of cited domains ranges from 0.05 to 0.24. The leading two brands are identical in all four markets. The results describe a stable two-vendor head on a long, engine-specific tail, built from evidence bases that the four engines barely share.

Abstract

1. Introduction

2. Research Questions

3. Methodology

4. Results

5. Discussion

6. Threats to Validity

7. Limitations

8. Practical Implications

9. Conclusion

Disclosure

10. Reproducibility and Data Availability

How to cite this report

References

Appendices A-E

1. Introduction

1.1 Background

Software evaluation used to begin with a list of links. A buyer searched, opened several vendor pages and analyst summaries, and assembled a shortlist by hand. AI engines compress that step. A single response now names a handful of vendors, orders them, attaches a sentence of rationale to each, and cites a small set of pages as support.

The structural consequence for a vendor is binary at the first step. A vendor named in the response enters the consideration set before any sales contact. A vendor not named has to be discovered by some other route, after a shortlist already exists. Position inside the response matters as well, since a reader encounters the first-named vendor before the others.

These responses are not uniform. Each engine draws on its own retrieval index, applies its own synthesis, and answers differently by market. Measuring one engine therefore describes one slice of the buyer's exposure. This report measures four engines in four markets under one query set.

1.2 Why this category

Enterprise single sign-on is a suitable instrument for three reasons drawn from the corpus itself. First, the vendor set is mature: the engines converge on a recognisable group of established vendors, led by Microsoft Entra ID and Okta Workforce Identity Cloud, with JumpCloud, Ping Identity and OneLogin named as secondary options. Second, the evaluation criteria are stable and shared. All four engines frame the decision around centralised SSO, automated SCIM provisioning, SAML and OIDC federation, and adaptive or phishing-resistant MFA. Third, the buying question is research-heavy. The query set includes multi-requirement questions (protocol support, provisioning, MFA) and a problem statement written from the IT team's perspective, the kind of question a buyer brings to an AI engine before contacting vendors.

Stable criteria make differences between engines easier to read. When every engine agrees on what matters, divergence in which vendors are named reflects retrieval and synthesis rather than disagreement about the category.

1.3 Contribution

This study provides a side-by-side measurement of four AI engines answering the same enterprise SSO questions in the same four markets within one collection window. It reports brand visibility, product-level naming, first-mention position, citation volume, source mix, link decay and cross-engine source overlap on a common set of definitions. It quantifies how much of the vendor shortlist is shared across engines and how much is specific to one engine. A single-engine measurement cannot show either the shared head or the engine-specific tail, and cannot show that the engines rest their responses on largely separate evidence.

2. Research Questions

1. **RQ1.** To what extent do the four AI engines agree on which vendors form the enterprise SSO shortlist?
2. **RQ2.** How concentrated is brand visibility across the vendors named?
3. **RQ3.** How far do the engines diverge in the volume, type and overlap of the sources they cite?
4. **RQ4.** Which vendors are named first, and does first position differ by engine?
5. **RQ5.** Does the vendor shortlist vary across the four markets?

3. Methodology

3.1 Study design

The study is observational and cross-sectional. One query set was issued to each engine in each market within a single collection window in October 2026. No intervention was applied to the engines, and no engine output was edited.

3.2 Engines under test

Four AI engines were studied: ChatGPT, Gemini, Google AI Overview and Google AI Mode. All four are commercial services that change without notice. None was version-pinned, and the results describe the engines as they behaved during the collection window.

3.3 Markets

Queries were issued in four markets: the United States, Canada, India and Germany. Market variance was tested because AI engines can localise retrieval, and a vendor's visibility in one market need not carry to another. The four markets span North America, South Asia and Western Europe.

3.4 Query construction

The query set comprises 10,000 buyer-intent queries. It combines five core category questions with paraphrase variants of each. The first core question is the headline buyer query, "Best enterprise single sign-on (SSO) solutions in 2026". The remaining four restate the same need from different angles: protocol and provisioning requirements, ease of rollout without a dedicated identity team, a problem statement about per-app logins and password resets, and passwordless with phishing-resistant MFA. Together they test whether the shortlist holds under rewording. The five core questions are listed in Appendix D.

Every query was issued to every engine in every market. The design therefore comprises 40,000 query runs per engine (10,000 queries × 4 markets) and 160,000 in total (10,000 queries × 4 engines × 4 markets).

3.5 Corpus and unit of analysis

The unit of analysis is one engine response to one query in one market. The analysed corpus contains 160,000 responses, 40,000 from each of ChatGPT, Gemini, Google AI Overview and Google AI Mode (Table 1). Those responses carry 1,094,000 cited sources, 934,000 brand mentions and 871,000 product mentions.

3.6 Metric definitions

Table 2 defines every metric used in the results, the discussion and the appendices. Rates are shares of the stated base; means and medians are reported as recorded.

Table 2. Metric definitions — every metric used in this report

Metric	Definition	Unit
Response	One AI engine response to one query in one market; the unit of analysis.	response
Mean words	Arithmetic mean word count of an engine's responses, excluding code blocks and markdown syntax.	words
Brand mention	One occurrence of a brand entity in a response, as extracted by the analysis model.	count
Brand entity	A distinct name recorded in the brand field; includes vendors, standards bodies, analyst firms and two generic category terms.	count
Share of mentions	A brand's mentions divided by all 934,000 brand mentions.	%
Brand mentions per response	An engine's brand mentions divided by its responses.	mentions
Mean ordinal position	Mean position of a brand among the brands named in a response, where 1 is the first brand named.	position
Modal first mention	The entity most often named first in an engine's responses.	entity
Mean sentiment	Mean of per-mention sentiment scores assigned by the analysis model on a scale from -1 (unfavourable) to +1 (favourable).	-1 to +1

Metric	Definition	Unit
Product entry	A distinct product name recorded in the product field; spelling and naming variants are recorded as separate entries.	count
Product mention	One occurrence of a product entry in a response.	count
Mean rank	Mean position of a product within the ordered list in which it appears, where 1 is first.	position
Hedging rate	Share of responses containing at least one hedging phrase from a fixed pattern list.	%
Recency - anchored rate	Share of responses containing an explicit time anchor, such as a year or "as of" phrase.	%
Refusal rate	Share of responses containing a refusal phrase from a fixed pattern list.	%
Suspected truncation rate	Share of responses of 200 or more characters whose final character is not a sentence-closing mark.	%
Self-citation (responses)	Share of an engine's responses citing at least one domain belonging to that engine's parent ecosystem.	%
Cited source	One distinct cited page in a response, after normalising tracking parameters and resolving redirect wrappers.	count
Distinct domains	Number of distinct registrable domains cited by that engine across its responses.	count
Source type	Rule-based class of a cited domain: vendor site, government, news, forum, social, review aggregator, parent ecosystem, or other.	category
Self-citation (sources)	Share of an engine's cited sources classified as parent ecosystem.	%
Dead-link rate	Share of cited sources returning HTTP 404, 410 or 5xx, or not connecting after a longer retry; 401, 403 and 429 are not counted as dead.	%
Dead links	Cited sources × dead-link rate, per engine.	count
Median source age	Median age of cited sources with a known date, reported only when at least one fifth of an engine's sources carry a known date.	days
Jaccard overlap	Size of the intersection of two engines' cited-domain sets divided by the size of their union.	0 to 1

3.7 Data collection and processing

Responses were captured from each engine through a commercial data-collection service and stored with their engine, market and query. Response HTML was converted to markdown text. Word counts exclude code blocks and markdown syntax.

Citations were assembled as one canonical set per response: the sources the engine attached plus pages linked in the response text, with junk links removed, tracking parameters and fragments normalised, and search-engine redirect wrappers resolved. Each cited domain was classified by deterministic rules: parent-ecosystem domains first, then government and education domains, then vendor site where the domain matches a brand named in the response, then forum, social, news, marketplace and review-aggregator patterns, and otherwise "other". Every cited URL was requested to test reachability. Source age was taken from a date in the URL path or, where absent, from a sufficiently old Last-Modified header.

Hedging, refusal, recency anchoring and suspected truncation were detected by fixed text patterns, without a model. Brand and product mentions, ordinal positions and sentiment were extracted per response by a language model under a fixed rubric, then aggregated by deterministic code. Sentiment is a model judgement of how favourably the response describes the entity. The exact model version used for extraction during this run is not documented in the dataset.

3.8 Scope exclusions

Engine responses were not paraphrased or edited. Claims made by the engines were not fact-checked against external sources, and responses were not scored for accuracy. No product was evaluated or tested. No vendor paid for, requested or was offered placement in the study.

4. Results

4.1 Response characteristics by engine

Table 3 reports response length and response construction behaviour for each engine.

Table 3. Response characteristics by AI engine — n = 160,000 responses; rates are shares of each engine's responses

AI engine	Responses	Mean words	Modal first mention	Hedging	Recency anchored	Refusal	Suspected truncation	Self-citation (responses)
ChatGPT	40,000	519.5	Microsoft Entra ID	5%	21%	0%	31%	0%
Gemini	40,000	363.8	Okta	7%	0%	0%	25%	0%
Google AI Overview	40,000	276.5	Okta	3%	21%	0%	0%	3%
Google AI Mode	40,000	281.8	Okta	0%	3%	0%	0%	10%

ChatGPT produces the longest responses at a mean of 519.5 words, the highest of the four engines. Google AI Overview produces the shortest at 276.5 words, with Google AI Mode at 281.8 and Gemini at 363.8. Suspected truncation is recorded in 31% of ChatGPT responses and 25% of Gemini responses, and in none from the two Google engines.

No engine records a refusal. Hedging is low throughout, from 0% for Google AI Mode to 7% for Gemini. Recency anchoring reaches 21% for both ChatGPT and Google AI Overview, against 3% for Google AI Mode and 0% for Gemini. Self-citation of the parent ecosystem appears in 10% of Google AI Mode responses and 3% of Google AI Overview responses.

4.2 Brand visibility and share of voice

Figure 1 and Table 4 report brand mentions for the ten most-mentioned brands; Appendix A gives all 67 brand entities.




Figure 1. Brand mentions for the ten most-mentioned brands. Mentions summed across four AI engines; n = 160,000 responses, 934,000 brand mentions in total. Source: GrackerAI AI Search Visibility Benchmark Series, October 2026.

Table 4. Brand mentions by AI engine, ten most-mentioned brands — mentions, with mean ordinal position in parentheses; n = 160,000 responses, 934,000 brand mentions

Brand	Total	Share of mentions	ChatGPT	Gemini	AI Overview	AI Mode	Mean sentiment
Microsoft	204,000	21.8%	48,000 (1.5)	54,000 (1.6)	52,000 (1.7)	50,000 (1.8)	0.46
Okta	149,000	16.0%	35,000 (1.7)	39,000 (2.3)	36,000 (1.5)	39,000 (2.2)	0.63
JumpCloud	108,000	11.6%	32,000 (4.4)	26,000 (3.5)	21,000 (4.1)	29,000 (3.2)	0.59
Google	97,000	10.4%	20,000 (5.6)	26,000 (3.0)	21,000 (2.7)	30,000 (3.0)	0.26
Ping Identity	90,000	9.6%	18,000 (3.1)	24,000 (3.1)	23,000 (3.0)	25,000 (3.2)	0.61
CyberArk	36,000	3.9%	12,000 (4.0)	11,000 (4.2)	2,000 (4.0)	11,000 (4.2)	0.56
Cisco	26,000	2.8%	17,000 (4.5)	5,000 (3.5)	1,000	3,000 (3.0)	0.54
Rippling	20,000	2.1%	6,000 (3.7)	6,000 (3.8)	3,000 (3.0)	5,000 (3.3)	0.58
Slack	19,000	2.0%	6,000	6,000	2,000	5,000	0.00
Salesforce	17,000	1.8%	6,000	6,000	2,000	3,000	0.00

The 160,000 responses contain 934,000 brand mentions across 67 brand entities. Microsoft records 204,000 mentions (21.8%) and Okta 149,000 (16.0%), together 37.8% of all mentions. JumpCloud (108,000), Google (97,000) and Ping Identity (90,000) follow. The ten most-mentioned brands account for 766,000 mentions, 82.0% of the total. Microsoft is the most-mentioned brand in each of the four engines, and Okta is second in each.

Mean sentiment for the ten most-mentioned brands ranges from 0.00 (Slack, Salesforce) to 0.63 (Okta). Microsoft records 0.46 and Google 0.26.




Figure 2. Brand entities by the number of AI engines that mention them. n = 67 brand entities recorded across 160,000 responses from four AI engines. Source: GrackerAI AI Search Visibility Benchmark Series, October 2026.

Figure 2 shows the distribution of brand entities by the number of engines that name them. Twelve of the 67 entities appear in all four engines: Microsoft, Okta, JumpCloud, Google, Ping Identity, CyberArk, Cisco, Rippling, Slack, Salesforce, GitHub and Workday. Eight appear in three engines, nine in two, and 38 in only one. Fifty entities record 5,000 mentions or fewer, and 27 record the minimum of 1,000.

Table 5. Brand breadth by AI engine — brand entities named, entities named by that engine alone, and brand mentions; n = 160,000 responses

AI engine	Responses	Brand entities named	Named by this engine only	Brand mentions	Brand mentions per response
ChatGPT	40,000	24	9	233,000	5.83
Gemini	40,000	37	14	258,000	6.45
Google AI Overview	40,000	33	7	203,000	5.08
Google AI Mode	40,000	34	8	240,000	6.00
Total	160,000	67	38	934,000	5.84

Table 5 shows that Gemini names the most brand entities (37) and the most that no other engine names (14). ChatGPT names the fewest (24), of which 9 are named by ChatGPT alone. Brand mentions per response range from 5.08 for Google AI Overview to 6.45 for Gemini.

4.3 Product-level results

Table 6 lists the 18 product entries with 10,000 or more mentions; Appendix C gives all 86 entries.

Table 6. Product entries with 10,000 or more mentions — n = 86 product entries, 871,000 product mentions; 18 entries shown

Product	Brand	Mentions	Mean rank	Mean sentiment	AI engines	Markets
Microsoft Entra ID	Microsoft	147,000	1.6	0.65	4	4
OneLogin	OneLogin	115,000	4.0	0.57	4	4
JumpCloud	JumpCloud	103,000	3.9	0.58	4	4
Okta Workforce Identity Cloud	Okta	102,000	1.6	0.75	4	4
Okta	Okta	39,000	2.4	0.46	4	4
Google Workspace	Google	26,000	3.1	0.30	4	4
Cisco Duo	Cisco	26,000	4.3	0.58	4	4
PingOne for Workforce	Ping Identity	25,000	3.1	0.70	3	4
Microsoft 365	Microsoft	20,000	—	0.03	4	4
Google Cloud Identity	Google	19,000	3.8	0.48	4	4
CyberArk Identity	CyberArk	18,000	4.1	0.61	4	4
Rippling	Rippling	16,000	3.7	0.59	4	4
CyberArk Workforce Identity	CyberArk	16,000	4.1	0.53	3	4
Ping Identity (PingFederate / PingOne)	Ping Identity	14,000	3.1	0.63	3	4
WorkOS	WorkOS	13,000	5.2	0.68	3	4
Auth0	Okta	12,000	4.7	0.51	3	4
Keycloak	Keycloak	11,000	6.1	0.29	3	4
miniOrange	miniOrange	10,000	5.2	0.62	3	3

The 86 product entries carry 871,000 product mentions. Microsoft Entra ID leads with 147,000 mentions, followed by OneLogin (115,000), JumpCloud (103,000) and Okta Workforce Identity Cloud (102,000). Microsoft Entra ID and Okta Workforce Identity Cloud share the earliest mean rank among the 18 entries in Table 6, at 1.6. Okta Workforce Identity Cloud records the highest mean sentiment among those 18 entries, at 0.75; Microsoft 365 records the lowest, at 0.03.

Eleven of the 86 entries appear in all four engines, and 47 appear in only one engine. Twenty-two entries appear in all four markets. Thirty-seven entries record the minimum of 1,000 mentions. Naming is fragmented: Okta and Ping Identity each hold nine separate entries, Microsoft and Google six each, and JumpCloud four. Six entries carry a negative mean sentiment, all with 1,000 or 2,000 mentions.

4.4 Citation volume and source concentration

Figure 3 and Table 7 report citation volume and the most-cited domain per engine; Table 8 reports the source mix.




Figure 3. Cited sources by AI engine. Distinct cited pages per response, summed per engine; n = 1,094,000 cited sources across 160,000 responses. Source: GrackerAI AI Search Visibility Benchmark Series, October 2026.

Table 7. Citation volume and concentration by AI engine — n = 1,094,000 cited sources

AI engine	Cited sources	Distinct domains	Most-cited domain	Hits	Share of engine's sources
ChatGPT	306,000	29	learn.microsoft.com	81,000	26.5%
Gemini	90,000	30	startwithidentity.com	13,000	14.4%
Google AI Overview	495,000	167	linkedin.com	37,000	7.5%
Google AI Mode	203,000	53	startwithidentity.com	27,000	13.3%
Total	1,094,000	—	—	—	—

The four engines cite 1,094,000 sources. Google AI Overview cites 495,000, the most of the four, from 167 distinct domains. ChatGPT cites 306,000 from 29 distinct domains, Google AI Mode 203,000 from 53, and Gemini 90,000 from 30. ChatGPT's most-cited domain, learn.microsoft.com, supplies 81,000 of its 306,000 sources (26.5%), the highest single-domain share among the four engines. The domain startwithidentity.com is among the three most-cited domains for Gemini, Google AI Overview and Google AI Mode.




Figure 4. Vendor-site share of cited sources by AI engine. Percentage of each engine's cited sources classified as vendor site; n = 1,094,000 cited sources (ChatGPT 306,000, Gemini 90,000, Google AI Overview 495,000, Google AI Mode 203,000). Source: GrackerAI AI Search Visibility Benchmark Series, October 2026.

Table 8. Source mix by AI engine — cited sources by source type; n = 1,094,000 cited sources

Source type	ChatGPT	Gemini	AI Overview	AI Mode	Total	Share
Vendor site	252,000	31,000	70,000	90,000	443,000	40.5%
Other	29,000	53,000	317,000	99,000	498,000	45.5%
Social	0	0	52,000	3,000	55,000	5.0%
Forum	1,000	0	34,000	4,000	39,000	3.6%
Review aggregator	6,000	6,000	13,000	3,000	28,000	2.6%
Government	18,000	0	3,000	0	21,000	1.9%
News	0	0	5,000	0	5,000	0.5%
Parent ecosystem	0	0	1,000	4,000	5,000	0.5%
Total	306,000	90,000	495,000	203,000	1,094,000	100.0%

Vendor sites supply 443,000 of the 1,094,000 cited sources (40.5%), and unclassified “other” sources 498,000 (45.5%). The vendor-site share differs by engine: 82.4% for ChatGPT, 44.3% for Google AI Mode, 34.4% for Gemini and 14.1% for Google AI Overview (Figure 4). Google AI Overview accounts for 52,000 of the 55,000 social citations and 34,000 of the 39,000 forum citations. Government sources account for 21,000 citations, 18,000 of them from ChatGPT.

4.5 Source quality and link decay

Table 9 reports dead-link rates, source age and self-citation at the source level.

Table 9. Source quality and link decay by AI engine — n = 1,094,000 cited sources

AI engine	Cited sources	Dead-link rate	Dead links	Median source age (days)	Self-citation (sources)
ChatGPT	306,000	0.0%	0	42	0%
Gemini	90,000	1.1%	990	—	0%
Google AI Overview	495,000	2.0%	9,900	—	0%
Google AI Mode	203,000	0.0%	0	—	2%
Total	1,094,000	1.0%	10,890	—	—

10,890 of the 1,094,000 cited sources are dead links, an aggregate dead-link rate of 1.0%. Google AI Overview records a 2.0% dead-link rate (9,900 dead links), the highest of the four engines, and Gemini 1.1% (990 dead links). ChatGPT and Google AI Mode record none. ChatGPT is the only engine with a reported median source age, at 42 days; for the other three, fewer than one fifth of sources carry a known date. Parent-ecosystem sources make up 2% of Google AI Mode’s cited sources and less than 0.5% for each other engine.

4.6 Cross-engine source overlap

Table 10 reports pairwise Jaccard overlap of the engines' cited-domain sets.

Table 10. Jaccard overlap of cited-domain sets between AI engines — pairwise; 1.00 on the diagonal; n = 1,094,000 cited sources

	ChatGPT	Gemini	Google AI Mode	Google AI Overview
ChatGPT	1.00	0.05	0.19	0.07
Gemini	0.05	1.00	0.20	0.11
Google AI Mode	0.19	0.20	1.00	0.24
Google AI Overview	0.07	0.11	0.24	1.00

Shading: below 0.10 0.10–0.14 0.15–0.19 0.20 and above diagonal

All six engine pairs record overlap at or below 0.24. The highest value is between Google AI Mode and Google AI Overview (0.24); the lowest is between ChatGPT and Gemini (0.05). Google AI Mode is the closest neighbour of each of the other three engines: 0.19 with ChatGPT, 0.20 with Gemini and 0.24 with Google AI Overview. The mean of the six pairwise values is 0.14.

4.7 First-mention position

Table 11 reports which entity each engine most often names first, and mean ordinal positions for the two leading brands.

Table 11. First-mention position by AI engine — modal first-mentioned entity and mean ordinal positions; n = 160,000 responses

AI engine	Modal first mention	Microsoft mean position	Okta mean position	Earliest mean position in engine
ChatGPT	Microsoft Entra ID	1.5	1.7	Microsoft (1.5)
Gemini	Okta	1.6	2.3	Microsoft (1.6)
Google AI Overview	Okta	1.7	1.5	Okta (1.5)
Google AI Mode	Okta	1.8	2.2	Microsoft (1.8)

ChatGPT most often names Microsoft Entra ID first; Gemini, Google AI Overview and Google AI Mode most often name Okta first. Microsoft's mean ordinal position ranges from 1.5 (ChatGPT) to 1.8 (Google AI Mode), and Okta's from 1.5 (Google AI Overview) to 2.3 (Gemini). Microsoft holds the earliest mean position in three engines, and Okta in Google AI Overview. In Gemini and Google AI Mode, Okta is the modal first mention while Microsoft holds the earlier mean position.

Below the two leaders, mean positions spread out. Among the ten most-mentioned brands, Google ranges from 2.7 (Google AI Overview) to 5.6 (ChatGPT), and JumpCloud from 3.2 (Google AI Mode) to 4.4 (ChatGPT). Ping Identity (3.0 to 3.2) and CyberArk (4.0 to 4.2) record the narrowest ranges across the four engines.

4.8 Geographic variance

Table 12 reports the five leading brands in each market, mean sentiment, and the number of product entries observed per market.

Table 12. Leading brands by market — five leading brands by mentions, mean sentiment and product entries observed per market

Market	First	Second	Third	Fourth	Fifth	Mean sentiment	Product entries
United States	Microsoft	Okta	JumpCloud	Ping Identity	Google	0.47	49
Canada	Microsoft	Okta	Google	JumpCloud	Ping Identity	0.45	49
India	Microsoft	Okta	JumpCloud	Ping Identity	Google	0.49	41
Germany	Microsoft	Okta	Google	JumpCloud	Ping Identity	0.45	44

Microsoft is first and Okta second in all four markets. The same five brands occupy the first five places in every market, in two orders: JumpCloud, Ping Identity, Google in the United States and India, and Google, JumpCloud, Ping Identity in Canada and Germany. Mean sentiment ranges from 0.45 (Canada, Germany) to 0.49 (India). The number of product entries observed per market ranges from 41 (India) to 49 (United States, Canada). Thirty-nine of the 86 product entries appear in a single market.

4.9 Inter-engine disagreement on factual claims

Claims were not adjudicated for accuracy (§3.8), so this subsection reports framings present in one engine and absent from the other three (Table 13).

Table 13. Engine-specific framings observed in the corpus — qualitative; n = 160,000 responses

AI engine	Framing observed in that engine and not in the other three
ChatGPT	References NIST and CISA guidance on phishing-resistant MFA; names NIST (4,000 mentions) and CISA (3,000 mentions).
Gemini	Places YubiKey and miniOrange among primary options in several US and India responses; several responses carry no citations while making product claims.
Google AI Overview	Includes Bitwarden and 1Password as password-manager complements to SSO.
Google AI Mode	Includes links to its parent ecosystem in 4,000 of 40,000 responses.

All four engines share the core framing of the category: Microsoft Entra ID and Okta Workforce Identity Cloud as the leading options, with SCIM provisioning, SAML and OIDC federation and phishing-resistant MFA as the evaluation criteria. The divergences in Table 13 sit outside that shared core.

5. Discussion

5.1 RQ1: Do the engines agree on the shortlist?

The engines agree on the head of the shortlist and disagree on the tail. §4.2 shows Microsoft first and Okta second by mentions in every engine, and twelve brand entities present in all four engines. The same section shows 38 of 67 entities named by one engine only, and §4.3 shows 47 of 86 product entries confined to one engine. The data is

consistent with a shared canon of five to six vendors, around which each engine adds its own selection of smaller vendors. Gemini contributes the largest engine-specific set (14 entities), which a plausible explanation links to its distinct evidence base (§4.6).

5.2 RQ2: How concentrated is brand visibility?

Visibility is concentrated. Two brands hold 37.8% of 934,000 mentions and ten hold 82.0% (§4.2). Fifty of 67 entities record 5,000 mentions or fewer. For a vendor outside the leading group, the typical outcome in this corpus is occasional mention by one engine. Product-level naming adds fragmentation (§4.3): Okta and Ping Identity are each spread across nine entries. A plausible explanation is that engines restate product names in varying forms, which disperses a vendor's product-level count even where its brand-level count is high.

5.3 RQ3: How far do evidence bases diverge?

Evidence bases diverge more than shortlists do. §4.6 shows pairwise domain overlap between 0.05 and 0.24, and §4.4 shows a vendor-site share ranging from 82.4% (ChatGPT) to 14.1% (Google AI Overview). ChatGPT builds its responses mainly from vendor documentation, with learn.microsoft.com alone supplying 26.5% of its sources. Google AI Overview draws on 167 distinct domains, including social and forum sources. The data is consistent with engines reaching similar vendor conclusions from largely different evidence. It does not show which evidence base produces more accurate responses, since accuracy was not scored. Link decay is low in every engine (§4.5), at an aggregate 1.0%.

5.4 RQ4: Who is named first?

First position is split between Microsoft and Okta, and the split depends on the engine and the measure (§4.7). ChatGPT most often opens with Microsoft Entra ID; the other three most often open with Okta. Mean position favours Microsoft in three engines. The two measures differ because the modal first mention counts only the opening entity, whereas mean position averages across every response naming the brand. Either way, no vendor other than Microsoft or Okta holds the earliest mean position in any engine. A plausible explanation for ChatGPT's Microsoft lead is its reliance on learn.microsoft.com (§4.4), though the data does not distinguish this from other causes.

5.5 RQ5: Does the shortlist vary by market?

Market variance is small at the head. §4.8 shows the same two leaders and the same five-brand set in every market, with only the order of positions three to five changing. Mean sentiment varies by 0.04 across markets. Variance appears in the tail: 39 of 86 product entries were observed in a single market. The data does not show whether those single-market entries reflect localisation or the small number of responses per market.

5.6 What the results do not resolve

The study does not establish why engines select the sources they do, whether vendor-heavy or community-heavy evidence produces better guidance, or whether the shortlist would hold across a longer window. It does not measure how buyers act on a response. Ordinal position is a proxy for prominence, not a measure of persuasion.

6. Threats to Validity

6.1 Construct validity

Ordinal position measures the order in which brands appear, not the emphasis a response places on them; a vendor named second with a detailed paragraph may be more prominent than one named first in a list. Sentiment is a model judgement on a -1 to +1 scale, and differences of about 0.2 or less between entities are within the resolution of the method. Suspected truncation is inferred from the final character of a response and not observed directly; a response ending in a table or list item can be flagged without being cut off. Brand-level and product-level counts come from separate extraction fields and are not subtotals of one another. OneLogin illustrates the gap: it records 3,000 brand mentions but 124,000 product mentions across three product entries, so its brand-level count understates its presence in the corpus. The brand-entity list also includes non-vendor names such as standards bodies, analyst firms and generic category terms.

6.2 Internal validity

All responses were collected within one window, so any event during that window that shifted engine behaviour affects all results equally and cannot be separated from them. AI engines are non-deterministic, and the same query can return different responses minutes apart. Personalisation and session state were not controlled beyond issuing queries through a collection service. No engine was version-pinned.

6.3 External validity

The study covers one category, four markets and one collection window. The query set is purposive: it was written to represent buyer intent, not sampled at random from real buyer queries. Because neither the queries nor the responses are a random sample from a defined population, the results support description of this corpus and do not support statistical inference to all buyer queries or all time periods. No inferential statistics are reported for this reason.

6.4 Reliability

A repeat run would be expected to reproduce the head of the results: Microsoft and Okta leading in every engine and market, and low cross-engine domain overlap. Individual counts would be expected to move, and the composition of the long tail of single-mention brands and products would be expected to change substantially between runs.

7. Limitations

- One category, enterprise single sign-on, studied in one collection window in October 2026.
- Four markets, all queried in the same language framing; other markets and languages were not covered.
- Four AI engines; other AI search products were not part of the analysed corpus.
- Engine claims were not adjudicated for accuracy, so factual disagreement is reported only as divergent framing.
- The model version used for mention and sentiment extraction is not documented in the dataset.

8. Practical Implications

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

- **Measure every engine, not one.** 38 of 67 brand entities appear in only one engine (§4.2). A vendor tracking a single engine can misjudge its visibility in either direction.
- **Treat evidence as engine-specific.** Vendor documentation supplies 82.4% of ChatGPT's cited sources but 14.1% of Google AI Overview's (§4.4). Keep product documentation current for engines that read vendor sites, and seek coverage on third-party and community sources for engines that read beyond them.
- **Use the shared criteria.** All four engines evaluate the category on SCIM provisioning, SAML and OIDC federation and phishing-resistant MFA (§4.9). State capability against those criteria explicitly and in the same terms.
- **Consolidate product naming.** Okta and Ping Identity are each spread across nine product entries (§4.3). Use one consistent product name in owned content so engines are less likely to split it.
- **Do not expect market-specific shortlists at the head.** The five leading brands are identical in all four markets (§4.8). Market-specific effort is more likely to matter for vendors in the tail.

9. Conclusion

This study examined 160,000 responses from four AI engines (ChatGPT, Gemini, Google AI Overview and Google AI Mode) to 10,000 buyer-intent queries about enterprise single sign-on, across the United States, Canada, India and Germany, in October 2026.

The engines agree on the head of the category. Microsoft and Okta lead in every engine and every market, together holding 37.8% of brand mentions, and twelve brand entities appear in all four engines. Below that head the shortlist is engine-specific: 38 of 67 brand entities appear in only one engine.

The engines reach similar shortlists from different evidence. Vendor sites supply 82.4% of ChatGPT's cited sources and 14.1% of Google AI Overview's, and no pair of engines shares more than 0.24 of its cited domains. Link decay is low, at 1.0% overall.

For a vendor in this category, the practical change is that visibility has to be managed per engine. A position in one engine's response says little about the others, and the evidence each engine reads differs in kind as well as in volume.

Disclosure

GrackerAI produces this research and also sells AI search visibility measurement to vendors, including vendors in the category covered here. The study measured how AI engines describe vendors; it did not test or evaluate any vendor's products. No vendor paid for, requested or was offered placement in the study or in this report.

10. Reproducibility and Data Availability

Responses were collected through a commercial data-collection service, converted to markdown, and stored with their engine, market, query and collection cycle. Behaviour metrics were computed by deterministic pattern-matching code. Brand, product, position and sentiment fields were extracted by a language model under a fixed rubric and aggregated by deterministic code. Source types were assigned by rule, and link reachability by HTTP request.

The retained dataset holds the raw responses, the extracted citations with their classifications and reachability results, and the per-response extraction output. Reproducing the study requires the query set in Appendix D, the four engines, the four market settings and the metric definitions in Table 2 and Appendix E. Because the engines are not version-pinned, a reproduction would produce a comparable rather than identical corpus. The dataset was generated on October 2026. Requests for the underlying data can be made to GrackerAI through gracker.ai.

How to cite this report

GrackerAI. (October 2026). *AI Search Visibility in Enterprise Access Management (Single Sign-On): A Multi-Engine, Multi-Market Benchmark Study*. GrackerAI AI Search Visibility Benchmark Series. <https://gracker.ai/data-and-research-reports/ai-search-visibility-access-management-sso-2026/>

References

- OASIS. (2005). *Assertions and Protocols for the OASIS Security Assertion Markup Language (SAML) V2.0*. OASIS Standard.
- OpenID Foundation. (2014). *OpenID Connect Core 1.0*.
- Hunt, P., Grizzle, K., Wahlstroem, E., & Mortimore, C. (2015). *System for Cross-domain Identity Management: Core Schema*. RFC 7643. IETF.
- Hunt, P., Grizzle, K., Ansari, M., Wahlstroem, E., & Mortimore, C. (2015). *System for Cross-domain Identity Management: Protocol*. RFC 7644. IETF.

Appendix A. Brand visibility by engine

Table A1 lists every brand entity recorded, with mentions per engine and mean ordinal position in parentheses where one was recorded.

Table A1. Brand mentions by AI engine, all brand entities — mentions, with mean ordinal position in parentheses where recorded; n = 67 entities, 934,000 mentions

Brand entity	Total	ChatGPT	Gemini	Google AI Overview	Google AI Mode
Microsoft	204,000	48,000 (1.5)	54,000 (1.6)	52,000 (1.7)	50,000 (1.8)
Okta	149,000	35,000 (1.7)	39,000 (2.3)	36,000 (1.5)	39,000 (2.2)
JumpCloud	108,000	32,000 (4.4)	26,000 (3.5)	21,000 (4.1)	29,000 (3.2)
Google	97,000	20,000 (5.6)	26,000 (3.0)	21,000 (2.7)	30,000 (3.0)
Ping Identity	90,000	18,000 (3.1)	24,000 (3.1)	23,000 (3.0)	25,000 (3.2)

Brand entity	Total	ChatGPT	Gemini	Google AI Overview	Google AI Mode
CyberArk	36,000	12,000 (4.0)	11,000 (4.2)	2,000 (4.0)	11,000 (4.2)
Cisco	26,000	17,000 (4.5)	5,000 (3.5)	1,000	3,000 (3.0)
Rippling	20,000	6,000 (3.7)	6,000 (3.8)	3,000 (3.0)	5,000 (3.3)
Slack	19,000	6,000	6,000	2,000	5,000
Salesforce	17,000	6,000	6,000	2,000	3,000
WorkOS	13,000	—	5,000 (4.2)	3,000 (5.7)	5,000 (6.0)
Keycloak	12,000	—	5,000 (6.0)	5,000 (5.7)	2,000 (6.5)
GitHub	11,000	4,000	4,000	1,000	2,000
miniOrange	10,000	—	6,000 (4.7)	2,000 (6.0)	2,000 (6.5)
Workday	9,000	3,000	2,000	1,000	3,000
AWS	8,000	6,000	—	1,000	1,000
Zoom	7,000	—	2,000	2,000	3,000
1Password	5,000	—	1,000	3,000	1,000
BambooHR	5,000	2,000	—	2,000	1,000
Beyond Identity	5,000	—	1,000 (5.0)	4,000 (4.0)	—
Bitwarden	4,000	—	—	3,000	1,000
NIST	4,000	4,000	—	—	—
OpenIAM	4,000	—	2,000 (6.0)	—	2,000 (6.0)
PingOne for Workforce	4,000	4,000	—	—	—
SecureAuth	4,000	—	4,000 (5.3)	—	—
CISA	3,000	3,000	—	—	—
Frontegg	3,000	—	2,000 (4.5)	—	1,000 (6.0)
OneLogin	3,000	—	—	—	3,000 (4.2)
RSA	3,000	—	—	—	3,000 (6.0)
Secret Double Octopus	3,000	—	3,000 (6.3)	—	—
Yubico	3,000	—	1,000	1,000	1,000
Authentik	2,000	—	1,000 (8.0)	1,000	—
Clerk	2,000	—	—	1,000	1,000 (6.0)
Dashlane	2,000	—	—	1,000	1,000
Entrust	2,000	—	2,000 (6.0)	—	—
IBM	2,000	1,000	—	—	1,000 (7.0)
Keeper	2,000	—	1,000	1,000	—
OLOID	2,000	—	—	2,000 (3.0)	—
PropelAuth	2,000	—	2,000 (4.5)	—	—
Vercel	2,000	—	2,000	—	—
1Kosmos	1,000	—	1,000 (4.0)	—	—
Apple	1,000	1,000	—	—	—
Authgear	1,000	—	1,000	—	—
Azure	1,000	—	—	—	1,000
Descope	1,000	—	—	—	1,000 (5.0)
Forrester	1,000	1,000	—	—	—
FreeIPA	1,000	—	—	1,000	—
FusionAuth	1,000	—	1,000 (6.0)	—	—
Gartner	1,000	1,000	—	—	—
Guideflow	1,000	—	1,000	—	—
HID Authentication Platform	1,000	1,000	—	—	—
HYPR	1,000	—	—	1,000	—

Brand entity	Total	ChatGPT	Gemini	Google AI Overview	Google AI Mode
Identity and Access Management (IAM)	1,000	—	—	1,000	—
IdentityServer	1,000	—	—	1,000	—
LoginRadius	1,000	—	1,000	—	—
Office 365	1,000	—	—	—	1,000
OneSpan	1,000	—	1,000 (8.0)	—	—
Palo Alto Networks	1,000	—	1,000	—	—
PingOne (Ping Identity)	1,000	1,000	—	—	—
Red Hat	1,000	—	—	—	1,000
SailPoint	1,000	1,000 (8.0)	—	—	—
Single Sign-On (SSO)	1,000	—	—	1,000	—
Stytch	1,000	—	—	—	1,000 (6.0)
SuccessFactors	1,000	—	1,000	—	—
Thales	1,000	—	—	1,000	—
Windows Hello for Business	1,000	—	—	—	1,000
WSO2	1,000	—	1,000 (8.0)	—	—
Total	934,000	233,000	258,000	203,000	240,000

Appendix B. Source mix and top cited domains by engine

Table B1. Source mix by AI engine, percentage — share of each engine's cited sources; n = 1,094,000 cited sources

Source type	ChatGPT	Gemini	AI Overview	AI Mode
Vendor site	82.4%	34.4%	14.1%	44.3%
Other	9.5%	58.9%	64.0%	48.8%
Social	0.0%	0.0%	10.5%	1.5%
Forum	0.3%	0.0%	6.9%	2.0%
Review aggregator	2.0%	6.7%	2.6%	1.5%
Government	5.9%	0.0%	0.6%	0.0%
News	0.0%	0.0%	1.0%	0.0%
Parent ecosystem	0.0%	0.0%	0.2%	2.0%

Table B2. Three most-cited domains by AI engine — hits and share of the engine's cited sources

AI engine	Rank	Domain	Hits	Share
ChatGPT	1	learn.microsoft.com	81,000	26.5%
	2	okta.com	33,000	10.8%
	3	jumpcloud.com	31,000	10.1%
Gemini	1	startwithidentity.com	13,000	14.4%
	2	oloid.com	11,000	12.2%
	3	learn.g2.com	6,000	6.7%
Google AI Overview	1	linkedin.com	37,000	7.5%
	2	startwithidentity.com	19,000	3.8%
	3	reddit.com	17,000	3.4%
Google AI Mode	1	startwithidentity.com	27,000	13.3%
	2	okta.com	27,000	13.3%
	3	learn.microsoft.com	15,000	7.4%

Appendix C. Full product ranking

Table C1. Full product ranking — all product entries ordered by mentions, then mean rank; n = 86 entries, 871,000 mentions

#	Product	Brand	Mentions	Mean rank	Mean sentiment	AI engines	Markets
1	Microsoft Entra ID	Microsoft	147,000	1.6	0.65	4	4
2	OneLogin	OneLogin	115,000	4.0	0.57	4	4
3	JumpCloud	JumpCloud	103,000	3.9	0.58	4	4
4	Okta Workforce Identity Cloud	Okta	102,000	1.6	0.75	4	4
5	Okta	Okta	39,000	2.4	0.46	4	4
6	Google Workspace	Google	26,000	3.1	0.30	4	4
7	Cisco Duo	Cisco	26,000	4.3	0.58	4	4
8	PingOne for Workforce	Ping Identity	25,000	3.1	0.70	3	4
9	Microsoft 365	Microsoft	20,000	—	0.03	4	4
10	Google Cloud Identity	Google	19,000	3.8	0.48	4	4
11	CyberArk Identity	CyberArk	18,000	4.1	0.61	4	4
12	Rippling	Rippling	16,000	3.7	0.59	4	4
13	CyberArk Workforce Identity	CyberArk	16,000	4.1	0.53	3	4
14	Ping Identity (PingFederate / PingOne)	Ping Identity	14,000	3.1	0.63	3	4
15	WorkOS	WorkOS	13,000	5.2	0.68	3	4
16	Auth0	Okta	12,000	4.7	0.51	3	4
17	Keycloak	Keycloak	11,000	6.1	0.29	3	4
18	miniOrange	miniOrange	10,000	5.2	0.62	3	3
19	Google Workspace Identity	Google	9,000	2.5	0.46	3	4
20	Ping Identity (PingOne / PingFederate)	Ping Identity	7,000	3.0	0.66	2	3
21	OneLogin Workforce Identity	OneLogin	7,000	4.2	0.46	1	4
22	Slack	Slack	7,000	—	0.00	3	4
23	Salesforce	Salesforce	6,000	—	0.00	3	4
24	1Password Business	1Password	5,000	—	0.42	3	4
25	Bitwarden	Bitwarden	4,000	—	0.33	2	3
26	GitHub	GitHub	4,000	—	0.00	2	3
27	YubiKey	Yubico	4,000	—	0.00	3	3
28	JumpCloud Directory Platform	JumpCloud	3,000	3.3	0.73	2	3
29	Google Cloud Identity / Workspace	Google	3,000	4.5	0.57	2	2
30	Frontegg	Frontegg	3,000	5.0	0.63	2	2
31	OpenIAM	OpenIAM	3,000	6.0	0.57	2	2
32	Secret Double Octopus	Secret Double Octopus	3,000	6.3	0.57	1	2
33	Microsoft Entra ID (formerly Azure AD)	Microsoft	2,000	2.0	0.70	2	2
34	PingFederate	Ping Identity	2,000	3.0	0.65	1	2
35	OLOID	OLOID	2,000	3.0	0.60	1	2
36	Rippling IT	Rippling	2,000	3.5	0.70	2	2
37	CyberArk	CyberArk	2,000	4.0	0.50	1	2
38	PropelAuth	PropelAuth	2,000	4.5	0.60	1	1
39	OneLogin Workforce IAM	OneLogin	2,000	5.0	0.60	1	2
40	Entrust Workforce Identity	Entrust	2,000	6.0	0.65	1	2
41	RSA Unified Identity Platform	RSA	2,000	6.0	0.50	1	2

#	Product	Brand	Mentions	Mean rank	Mean sentiment	AI engines	Markets
42	Clerk	Clerk	2,000	6.0	-0.05	2	2
43	IBM Verify	IBM	2,000	7.0	0.40	2	2
44	Authentik	Authentik	2,000	8.0	-0.15	2	2
45	Okta FastPass	Okta	2,000	—	0.65	2	1
46	Keeper	Keeper	2,000	—	0.40	2	2
47	Dashlane	Dashlane	2,000	—	0.35	2	2
48	Okta Identity Cloud	Okta	2,000	—	0.30	2	2
49	Zoom	Zoom	2,000	—	0.00	1	2
50	Okta Workforce Identity Cloud (WIC)	Okta	1,000	1.0	0.80	1	1
51	Rippling Identity	Rippling	1,000	2.0	0.80	1	1
52	Google Workspace (Built-in SAML/SSO)	Google	1,000	2.0	0.60	1	1
53	Ping Identity (including ForgeRock)	Ping Identity	1,000	3.0	0.60	1	1
54	Ping Identity / ForgeRock	Ping Identity	1,000	3.0	0.60	1	1
55	PingOne Workforce	Ping Identity	1,000	3.0	0.60	1	1
56	1Kosmos BlockID	1Kosmos	1,000	4.0	0.70	1	1
57	Descope	Descope	1,000	5.0	0.70	1	1
58	SecureAuth	SecureAuth	1,000	5.0	0.70	1	1
59	JumpCloud (Open Directory Platform)	JumpCloud	1,000	5.0	0.60	1	1
60	SecureAuth Workforce Authority	SecureAuth	1,000	5.0	0.60	1	1
61	ForgeRock / Ping Identity (Ping Identity Platform)	Ping Identity	1,000	5.0	0.60	1	1
62	Managed Keycloak	Keycloak	1,000	5.0	0.50	1	1
63	ForgeRock Autonomous Identity	ForgeRock	1,000	5.0	0.50	1	1
64	SecureAuth Workforce	SecureAuth	1,000	6.0	0.60	1	1
65	FusionAuth	FusionAuth	1,000	6.0	0.50	1	1
66	Stytch	Stytch	1,000	6.0	0.50	1	1
67	OneSpan	OneSpan	1,000	8.0	0.60	1	1
68	WSO2 Identity Server	WSO2	1,000	8.0	0.40	1	1
69	SailPoint	SailPoint	1,000	8.0	0.30	1	1
70	Okta Identity Engine (Okta Verify)	Okta	1,000	—	0.70	1	1
71	JumpCloud (JumpCloud Go)	JumpCloud	1,000	—	0.60	1	1
72	Thales SafeNet Trusted Access	Thales	1,000	—	0.60	1	1
73	RSA	RSA	1,000	—	0.60	1	1
74	PingID	Ping Identity	1,000	—	0.50	1	1
75	HYPR	HYPR	1,000	—	0.40	1	1
76	Google Workspace identity services	Google	1,000	—	0.40	1	1
77	HID Authentication Platform	HID Global	1,000	—	0.30	1	1
78	Okta SWA	Okta	1,000	—	0.20	1	1
79	BambooHR	BambooHR	1,000	—	0.00	1	1
80	Auth0 / Okta Customer Identity	Okta	1,000	—	0.00	1	1
81	Jira	Atlassian	1,000	—	0.00	1	1
82	Microsoft Entra ID (Azure AD)	Microsoft	1,000	—	0.00	1	1
83	Microsoft Active Directory	Microsoft	1,000	—	-0.30	1	1
84	Microsoft Intune	Microsoft	1,000	—	-0.30	1	1
85	IdentityServer	IdentityServer	1,000	—	-0.70	1	1
86	FreeIPA	FreeIPA	1,000	—	-0.70	1	1
Total			871,000				

Appendix D. Query set and market list

Table D1. Core questions — the five core questions from which the query set is built

Query text

1	Best enterprise single sign-on (SSO) solutions in 2026
2	Which access management platforms support SAML, OIDC, SCIM provisioning and adaptive MFA for workforce users?
3	Which SSO solutions are easiest to roll out across hundreds of SaaS apps without a dedicated identity team?
4	Our employees have a separate login for every app and IT spends hours each week on password resets and offboarding. What should we put in place?
5	Which workforce identity and access management platforms offer passwordless login and phishing-resistant MFA together with SSO?

Table D2. Study parameters

Parameter	Value
AI engines	ChatGPT, Gemini, Google AI Overview, Google AI Mode
Markets	United States, Canada, India, Germany
Buyer-intent queries	10,000 (5 core questions with paraphrase variants)
Query runs per engine	40,000
Analysed responses	160,000
Cited sources	1,094,000
Collection window	October 2026
Dataset generated	October 2026

Appendix E. Metric computation

Table E1. Formulas for derived metrics

Metric	Computation
Rate metrics (hedging, recency, refusal, truncation, self-citation by response)	responses flagged ÷ engine responses
Dead links	engine cited sources × engine dead-link rate, rounded to whole links
Aggregate dead-link rate	Σ dead links ÷ Σ cited sources = 10,890 ÷ 1,094,000
Share of mentions	brand mentions ÷ 934,000
Brand mentions per response	engine brand mentions ÷ engine responses
Source-type share	cited sources of that type ÷ engine cited sources

Metric	Computation
Self-citation (sources)	parent-ecosystem sources ÷ engine cited sources
Jaccard overlap	$ A \cap B \div A \cup B $, where A and B are two engines' cited-domain sets
Mean ordinal position	mean position of a brand among the brands named in each response that names it
Mean sentiment	arithmetic mean of per-mention sentiment scores