

AI Search Visibility in Enterprise Passwordless Authentication

A multi-engine, multi-market benchmark study of how four AI search engines (ChatGPT, Gemini, Google AI Overview and Google AI Mode) respond to buyer-intent questions about enterprise passwordless authentication 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	284,000	4
Gemini	40,000	92,000	4
Google AI Overview	40,000	380,000	4
Google AI Mode	40,000	262,000	4
Total	160,000	1,018,000	4

Published: October 2026

Research period: October 2026

Category: Enterprise passwordless authentication (FIDO2 and WebAuthn passkeys, phishing-resistant MFA, device-bound credentials)

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

Evidence base: 1,018,000 cited sources · 153 distinct product entries · 72 brand entries

Prepared by: GrackerAI · AI Search Visibility Benchmark Series

Abstract

Enterprise buyers increasingly meet their first vendor shortlist inside an AI-generated response rather than on a page of search results, 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 respond to buyer-intent questions about enterprise passwordless authentication. A set of 10,000 buyer-intent queries was issued to ChatGPT, Gemini, Google AI Overview and Google AI Mode in the United States, Canada, India and Germany during October 2026, producing 160,000 analysed AI engine responses and 1,018,000 cited sources. For each response the study recorded length, citations, cited domains, link reachability, hedging, truncation, refusal, and every brand and product mention with its ordinal position. The engines agree on the leading vendor: Microsoft records 148,000 brand mentions, the highest count in every engine, and Microsoft Entra ID is the first-mention entity in three of the four engines. They diverge on evidence. The highest cross-engine source overlap is a Jaccard coefficient of 0.30, and cited-source volume differs by a factor of 4.1 between Google AI Overview and Gemini. Cited links are largely intact, with an aggregate dead-link rate of 1.0%. Visibility below the leading vendor is fragmented across 153 distinct product entries. The results indicate that vendor ordering is stable across engines and markets while the evidence base behind it is not shared, so a single-engine measurement does not describe the category.

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

- 4.2 Brand visibility

- 4.3 Product-level results

- 4.4 Citation volume

- 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

10. Reproducibility

Disclosure · Citation

Appendices A–E

1. Introduction

1.1 Background

An AI search engine does not return a ranked list of documents. It returns a composed response that names a small number of vendors, places them in an order, and attaches reasons to each. For an enterprise security purchase, that single response performs work a buyer once spread across analyst notes, review sites and vendor pages. It sets the category boundary, proposes evaluation criteria and nominates candidates.

The consequence for a vendor is structural. A vendor absent from the returned shortlist is not losing on price or capability. It is absent at the moment the consideration set forms. Conventional search metrics do not detect this condition, since a vendor can hold strong organic rankings and still be omitted from the responses that now precede a click.

Measuring the condition requires observing the engines directly. That means issuing the questions a buyer would ask, recording the responses as returned, and analysing which vendors appear, in what order, and on the strength of which cited evidence. This study applies that method to one category in one collection window.

1.2 Why this category

Enterprise passwordless authentication suits this kind of measurement for four reasons. First, the vendor set is mature and documented at length, so engines have substantial published material to draw on. Second, the evaluation vocabulary is stable. Engine responses in this corpus organise the category around FIDO2 and WebAuthn passkeys, phishing resistance, device-bound versus synced credentials, and integration with existing identity and single sign-on infrastructure. That shared vocabulary makes responses comparable across engines, markets and phrasings.

Third, buying is research-heavy. The core questions behind the query set range from a direct request for the best solutions in 2026 to a practitioner's description of users being phished through push notifications and one-time codes. Fourth, the cost of invisibility is high. An authentication decision touches every workforce login, and the responses recorded here weigh platform breadth against lock-in, hardware assurance against logistics, and overlay products against an additional policy surface. A vendor omitted early in that evaluation is difficult to reinstate later.

1.3 Contribution

This study contributes a directly observed, cross-engine baseline for one enterprise security category. It reports brand visibility, product visibility and ordinal position as separate measures rather than one share-of-voice figure. It quantifies how far the engines draw on shared or disjoint evidence, which determines whether a measurement taken on one engine transfers to another. It also records the construction of engine output itself, including response length, source variety and cited-link reachability, as properties of the category's AI search layer.

2. Research Questions

1. **RQ1.** To what extent do the four AI engines converge on the same vendor shortlist for enterprise passwordless authentication?
2. **RQ2.** How is visibility distributed across brand entries and product entries, and how concentrated is it?
3. **RQ3.** To what extent do the engines draw on a shared evidence base when producing their responses?

4. **RQ4.** Does the first-named entity in a response follow mention volume, or are ordering and volume governed separately?
5. **RQ5.** Does vendor recommendation vary across the four markets tested?

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

3. Methodology

3.1 Study design

The study is observational and cross-sectional. Four commercial AI search engines received an identical set of buyer-intent queries across four markets within a single collection window in October 2026. No engine was primed or given conversational context beyond the query 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. Two of the four, Google AI Overview and Google AI Mode, are surfaces of the same search product, which 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 the engines localise vendor recommendations in response to regional vendor presence, regulatory environment or buyer context. 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 passwordless authentication. The set is built from five core questions together with paraphrase variants of each, issued across repeated collection cycles. The variants test whether vendor recommendations remain stable under rewording or shift with surface phrasing. Every query was issued to all four engines in all four markets, a design of 40,000 responses per engine and 160,000 in total. The core questions are listed in Appendix D.

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 analysed corpus holds 160,000 responses, 40,000 from each of the four engines. Every per-engine rate is computed on that engine's responses or cited sources, and Table 1 states each engine's counts.

3.6 Metric definitions

Every metric used in this report is defined in Table 2. Definitions precede their first use in the results.

Table 2. Metric definitions

Metric	Definition	Unit
Response	One AI engine response to one buyer-intent query in one market; the unit of analysis	count
Mean response length	Mean word count of an engine's responses	words
Source variety	Mean number of distinct cited domains appearing within a single response	mean
Hedging rate	Share of responses containing explicit uncertainty or qualifying language	%
Recency anchoring	Share of responses that anchor the recommendation to an explicit date or recency signal	%
Truncation rate	Share of responses flagged as ending mid-sentence or mid-list. Inferred from response structure by the analysis step, not observed directly	%
Refusal rate	Share of responses declining to address the query	%
Self-citation rate	Share of an engine's citations pointing at domains owned by its own parent ecosystem	%
First-mention entity	The entity named most often in first position across an engine's responses	label
Cited sources	Citations recorded for an engine across the corpus, after de-duplication of each response's source list	count
Distinct domains cited	Number of different domains appearing in one engine's citations. Not a corpus-wide total and not additive across engines	count
Source type	Rule-based class assigned to each cited domain: vendor site, government, academic, news, forum, social, review aggregator, or other	label
Dead-link rate	Share of an engine's cited URLs that returned a not-found, gone or server-error status, or no connection, after a longer retry. Access-control responses are not counted	%
Unreachable citations	The product of cited sources and the dead-link rate, per engine, rounded to a whole citation	count
Median source age	Median age of cited material at collection, from a date in the URL path or a server modification header. Reported only when at least a fifth of an engine's sources carry a known age	days
Brand mention	One occurrence of a brand entry in a response. Brand entries are recorded as extracted and include standards bodies and authentication methods as well as vendors	count
Mean ordinal position	Mean position of a brand's first appearance within a response, where 1.0 is the first entity named	mean
Product mention	One occurrence of a named product entry in a response. Counted separately from brand mentions and not additive with them	count
Mean product rank	Mean ordinal position of a product entry within the responses that name it	mean
Mean sentiment	Arithmetic mean of per-mention polarity scores for a brand, product or market, on a range from -1.0 (unfavourable) to 1.0 (favourable)	-1.0 to 1.0
Jaccard source overlap	Size of the intersection of two engines' cited-domain sets divided by the size of their union	0.00 to 1.00

Metric	Definition	Unit
Leading products by market	The three products ranked highest for a market by the collection pipeline's market-level summary	label

3.7 Data collection and processing

Queries were issued to each engine's public surface for the specified market through a commercial data collection layer. Each response was captured in full, persisted to a columnar analytical database, and converted from its returned markup to plain text before analysis.

Citations were assembled into one canonical set per response. That set combines the engine's listed sources with links in the response text, removes tracking parameters and fragments, and resolves redirect wrappers. Each cited domain was then assigned a source type by fixed rules. Parent-ecosystem domains are matched first, then government and academic domains, then vendor sites, where the registrable domain matches a brand the response names. Forum, social, news, marketplace and review-aggregator domains follow by list and host pattern, and everything else is classed as other.

Each cited URL was requested to test reachability. A URL counted as dead when it returned a not-found, gone or server-error status, or no connection after a longer retry. Access-control responses were not counted as dead. Source age was taken from a date in the URL path or a server modification header. A median age is reported only for an engine where at least a fifth of cited sources carry a known age.

Brand and product mentions, their ordinal positions, hedging, recency anchoring, refusal and suspected truncation were extracted per response by a language-model analysis step working to a fixed rubric. The same step scored sentiment for each mention on a range from -1.0 to 1.0 . Brand and product means are arithmetic means of those per-mention scores. No entity resolution was applied, so variant names are recorded as separate entries. The precise detection rule behind the truncation flag is not documented beyond the rubric's description of a response ending mid-sentence.

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 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. 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, source variety and response construction differ between the four engines, as Table 3 shows.

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

AI engine	Responses	Mean words	Source variety	Hedging	Recency anchoring	Truncation	Refusals	Self-citation
ChatGPT	40,000	552.9	0.9	0%	23%	15%	0%	0%
Gemini	40,000	446.5	1.1	0%	0%	29%	0%	0%
Google AI Overview	40,000	223.3	3.5	0%	22%	0%	0%	0%
Google AI Mode	40,000	339.9	1.9	0%	3%	0%	0%	0%

Mean response length spans a factor of 2.5, from 552.9 words for ChatGPT to 223.3 words for Google AI Overview. Gemini records 446.5 words and Google AI Mode 339.9. Source variety runs in the opposite direction to length. Google AI Overview draws on 3.5 distinct domains per response, against 1.9 for Google AI Mode, 1.1 for Gemini and 0.9 for ChatGPT.

No engine used hedging language in any response, no engine declined a query, and no engine cited its own parent ecosystem. Recency anchoring is concentrated in two engines: ChatGPT anchors 23% of its 40,000 responses to an explicit date and Google AI Overview 22%, against 3% for Google AI Mode and 0% for Gemini. Suspected truncation is recorded for 29% of Gemini responses (n = 40,000) and 15% of ChatGPT responses (n = 40,000). Google AI Overview and Google AI Mode record 0%.

4.2 Brand visibility and share of voice

Brand mentions were recorded per engine with the mean ordinal position at which each brand entry first appeared. Figure 1 shows the ten most-mentioned entries and Table 4 breaks them down by engine.




Figure 1. Brand mentions, ten most-mentioned brand entries. Total brand mentions 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 brand entries

Brand	Total	ChatGPT	Gemini	AI Overview	AI Mode
Microsoft	148,000	32,000 (1.6)	31,000 (1.4)	44,000 (1.6)	41,000 (2.3)
Ping Identity	94,000	21,000 (4.1)	24,000 (3.4)	22,000 (3.4)	27,000 (4.1)
Google Workspace	62,000	11,000	16,000 (1.7)	12,000 (2.0)	23,000 (2.0)
HYPR	60,000	22,000 (3.5)	15,000 (3.9)	6,000 (2.5)	17,000 (3.3)
Cisco	53,000	22,000 (3.0)	9,000 (4.6)	10,000 (3.4)	12,000 (3.3)
Yubico	51,000	—	21,000 (4.2)	8,000 (2.5)	22,000 (3.6)
Okta	35,000	13,000 (2.5)	9,000 (2.5)	6,000 (1.5)	7,000 (2.0)
Beyond Identity	31,000	6,000 (4.3)	4,000 (5.5)	7,000 (3.9)	14,000 (3.6)
Stytch	26,000	8,000 (2.6)	5,000 (2.5)	5,000 (3.5)	8,000 (1.6)
Apple	23,000	—	9,000 (1.5)	7,000 (2.3)	7,000 (2.0)
Ten entries	583,000	135,000	143,000	127,000	178,000

The corpus records 72 brand entries carrying 846,000 brand mentions. Microsoft records 148,000, or 17.5% of the total, followed by Ping Identity at 94,000 and Google Workspace at 62,000. The ten most-mentioned entries account for 583,000 mentions, or 68.9%. Microsoft holds the highest brand-mention count in each of the four engines, from 31,000 in Gemini to 44,000 in Google AI Overview.

Fifteen brand entries are named by all four engines. Fourteen are vendors: Microsoft, Ping Identity, Google Workspace, HYPR, Cisco, Okta, Beyond Identity, Stytch, Descope, ForgeRock, Thales, JumpCloud, FusionAuth and 1Password. The fifteenth is the FIDO2/WebAuthn standard. At the other end, 28 of the 72 entries are named by a single engine only. The brand list also carries standards bodies and agencies as recorded. CISA and NIST each record 6,000 mentions, all of them in ChatGPT.

Coverage is uneven within the leading group. Yubico records 51,000 mentions across Gemini, Google AI Mode and Google AI Overview and none in ChatGPT. Apple likewise appears in three engines and not in ChatGPT. 1Kosmos records 11,000 of its 12,000 mentions in Gemini, and miniOrange 7,000 of its 10,000. Okta records 35,000 brand mentions, seventh among brand entries, while its products are counted separately in Section 4.3.

Table 5. Mean sentiment by brand — range -1.0 to 1.0; twelve brands carried by the brand sentiment summary

Brand	Mean sentiment	Engines naming the brand
Descope	0.71	4 of 4
Stytch	0.69	4 of 4
HYPR	0.68	4 of 4
Beyond Identity	0.63	4 of 4
Cisco	0.61	4 of 4
Yubico	0.57	3 of 4
Okta	0.54	4 of 4

Brand	Mean sentiment	Engines naming the brand
Ping Identity	0.49	4 of 4
Microsoft	0.47	4 of 4
Apple	0.18	3 of 4
ForgeRock	0.17	4 of 4
Google Workspace	0.05	4 of 4

Mean sentiment spans 0.05 for Google Workspace to 0.71 for Descope across the twelve brands in the sentiment summary. Three brands fall below 0.20: Google Workspace at 0.05, ForgeRock at 0.17 and Apple at 0.18. The other nine fall between 0.47 for Microsoft and 0.71 for Descope.

4.3 Product-level results

The corpus records 153 distinct product entries carrying 782,000 product mentions. Product mentions are counted separately from brand mentions and are not additive with them. Table 6 lists the ten most-mentioned entries.

Table 6. Ten most-mentioned product entries — n = 160,000 AI engine responses

#	Product	Brand	Mentions	Mean product rank	Mean sentiment	Engines	Markets
1	Microsoft Entra ID	Microsoft	84,000	1.6	0.65	4	4
2	HYPR	HYPR	53,000	3.5	0.68	4	4
3	YubiKey	Yubico	51,000	3.1	0.59	3	4
4	Ping Identity	Ping Identity	44,000	3.9	0.58	4	4
5	Okta Workforce Identity Cloud	Okta	37,000	2.2	0.66	4	4
6	Okta	Okta	35,000	2.2	0.54	4	4
7	Cisco Duo	Cisco	33,000	3.7	0.64	4	4
8	Beyond Identity	Beyond Identity	30,000	3.9	0.63	4	4
9	Stytch	Stytch	24,000	2.4	0.72	4	4
10	Cisco Duo Passwordless	Cisco	20,000	3.0	0.64	3	4

Microsoft Entra ID records 84,000 mentions at a mean product rank of 1.6, the earliest of the ten most-mentioned entries. HYPR follows at 53,000 and YubiKey at 51,000. Eleven low-volume entries elsewhere in the table record a mean rank of 1.0; none of them exceeds 6,000 mentions. Fourteen entries are named by all four engines and 23 appear in all four markets. 97 entries are named by a single engine, and 87 appear in a single market.




Figure 4. Distinct product entries per brand. Eight brands with the most product entries; n = 153 product entries. Source: [GrackerAI AI Search Visibility Benchmark Series](#), October 2026.

Product naming is fragmented, as Figure 4 shows. Okta appears as 22 separate product entries, Ping Identity as 19, Microsoft as 14 and Yubico as 10. Okta's 22 entries carry 137,000 product mentions against 35,000 brand mentions. Microsoft's 14 entries carry 134,000 and Ping Identity's 19 carry 83,000. Of the 153 entries, 79 record 1,000 mentions, the lowest count in the table.

Auth0, attributed to Okta, records 15,000 mentions at a mean rank of 1.6 and a mean sentiment of 0.75. Product sentiment ranges from -0.10 for Firebase Auth to 0.90. Four entries carry negative mean sentiment: Firebase Auth, Amazon Cognito, Apple iCloud Keychain and Google Password Manager. The last two are described by the engines as synced passkey stores and record -0.05 each. Within Microsoft, Windows Hello for Business records 0.40 and Microsoft Authenticator 0.24, against 0.65 for Microsoft Entra ID.

4.4 Citation volume and source concentration

Across the four engines, 1,018,000 cited sources were recorded. Figure 2 shows the distribution by engine.




Figure 2. Cited sources by AI engine. Total cited sources per engine; n = 1,018,000 cited sources from 160,000 AI engine responses. Source: [GrackerAI AI Search Visibility Benchmark Series](#), October 2026.

Google AI Overview records 380,000 cited sources, ChatGPT 284,000, Google AI Mode 262,000 and Gemini 92,000. The ratio between the highest and lowest volume is 4.1. Google AI Overview records both the highest citation volume and the shortest mean response, at 223.3 words.




Figure 3. Source-type mix by AI engine. Share of each engine's cited sources by source type; segments under 7% are unlabelled and listed in Table B2; n = 1,018,000 cited sources. Source: [GrackerAI AI Search Visibility Benchmark Series](#), October 2026.

Figure 3 shows the source-type mix. Vendor sites account for 76.8% of ChatGPT citations, 57.6% for Gemini, 55.0% for Google AI Mode and 32.4% for Google AI Overview. In Google AI Overview the unclassified category is larger, at 47.9%, and social sources reach 10.0%. ChatGPT draws 8.8% of its citations from government sources, or 25,000. The other three engines record 3,000 government citations between them. Gemini's citations fall into two classes only: vendor site and other.

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

AI engine	Most-cited domains
ChatGPT	learn.microsoft.com (67,000), hypr.com (26,000), duo.com (22,000)
Gemini	1kosmos.com (12,000), miniorange.com (11,000), hypr.com (5,000)
Google AI Overview	linkedin.com (20,000), learn.microsoft.com (17,000), docs.pingidentity.com (13,000)
Google AI Mode	learn.microsoft.com (22,000), startwithidentity.com (14,000), hypr.com (14,000)

Two domains appear in the top three of three engines each. learn.microsoft.com leads ChatGPT with 67,000 citations, 23.6% of that engine's total, and leads Google AI Mode with 22,000. hypr.com appears in the top three of ChatGPT, Gemini and Google AI Mode. Gemini's two most-cited domains are 1kosmos.com at 12,000 and miniorange.com at 11,000, together 25.0% of its citations. Google AI Overview's most-cited domain is linkedin.com, at 20,000.

4.5 Source quality and link decay

Table 7 reports cited-link reachability, domain breadth and source age by engine.

Table 7. Citation behaviour and source quality by AI engine — n = 1,018,000 cited sources

AI engine	Cited sources	Distinct domains cited	Dead-link rate	Unreachable citations	Median source age	Self-citation
ChatGPT	284,000	34	0.7%	1,988	24 days	0%
Gemini	92,000	39	0.0%	0	n.r.	0%

AI engine	Cited sources	Distinct domains cited	Dead-link rate	Unreachable citations	Median source age	Self-citation
Google AI Overview	380,000	122	1.8%	6,840	n.r.	0%
Google AI Mode	262,000	72	0.4%	1,048	n.r.	0%
Total	1,018,000	—	1.0%	9,876	—	0%

Of the 1,018,000 cited URLs tested, 9,876 did not resolve, an aggregate dead-link rate of 1.0%. Google AI Overview records the highest rate at 1.8%, or 6,840 unreachable citations. ChatGPT records 0.7% and Google AI Mode 0.4%. Gemini records 0.0%.

Distinct domain counts range from 34 for ChatGPT to 122 for Google AI Overview, a factor of 3.6. Gemini records 39 and Google AI Mode 72. These counts are per engine and are not additive. ChatGPT is the only engine with a reported median source age, at 24 days. For the other three engines fewer than a fifth of cited sources carried a known age, and no median is reported (n.r.).

4.6 Cross-engine source overlap

Jaccard source overlap measures the proportion of cited domains two engines hold in common. Table 9 reports every pair.

Table 9. Jaccard source overlap between AI engines — 0.00 denotes no shared cited domains, 1.00 identical sets

	ChatGPT	Gemini	AI Overview	AI Mode
ChatGPT	1.00	0.18	0.12	0.18
Gemini	0.18	1.00	0.15	0.23
Google AI Overview	0.12	0.15	1.00	0.30
Google AI Mode	0.18	0.23	0.30	1.00

Shading denotes overlap magnitude: 0.10–0.14 0.15–0.19 0.20–0.24 0.25–0.30

The highest overlap between two different engines is 0.30, between Google AI Mode and Google AI Overview. The next highest is 0.23, between Gemini and Google AI Mode. The lowest is 0.12, between ChatGPT and Google AI Overview. ChatGPT records 0.18 with both Gemini and Google AI Mode. All six pairs fall between 0.12 and 0.30.

4.7 First-mention position

The entity named most often in first position was recorded for each engine. Table 10 sets it against Microsoft's mentions and mean position in the same engine.

Table 10. First-mention entity and Microsoft position by AI engine

AI engine	First-mention entity	Microsoft mentions	Microsoft mean position	Earliest mean position, ten most-mentioned brand entries
ChatGPT	Auth0	32,000	1.6	Microsoft (1.6)
Gemini	Microsoft Entra ID	31,000	1.4	Microsoft (1.4)

AI engine	First-mention entity	Microsoft mentions	Microsoft mean position	Earliest mean position, ten most-mentioned brand entries
Google AI Overview	Microsoft Entra ID	44,000	1.6	Okta (1.5)
Google AI Mode	Microsoft Entra ID	41,000	2.3	Stytch (1.6)

Microsoft Entra ID is the first-mention entity in Gemini, Google AI Overview and Google AI Mode. ChatGPT opens most often with Auth0. Auth0 does not appear as a brand entry; it is recorded as a product of Okta, with 15,000 product mentions at a mean rank of 1.6. Microsoft's mean ordinal position ranges from 1.4 in Gemini to 2.3 in Google AI Mode.

The earliest mean position among the ten most-mentioned brand entries belongs to Microsoft in ChatGPT and Gemini. In Google AI Overview it belongs to Okta at 1.5, on 6,000 mentions. In Google AI Mode it belongs to Stytch at 1.6, on 8,000 mentions, while Microsoft records 41,000 mentions at 2.3. Ping Identity holds the second-highest brand total at 94,000 with mean positions between 3.4 and 4.1.

4.8 Geographic variance

Table 11 lists the three leading products recorded for each market, with mean sentiment across that market's responses.

Table 11. Leading products and mean sentiment by market — three leading products per market as recorded

Market	First	Second	Third	Mean sentiment
United States	Microsoft Entra ID (Microsoft)	Okta Workforce Identity Cloud (Okta)	HYPR (HYPR)	0.48
Canada	Microsoft Entra ID (Microsoft)	Okta (Okta)	Stytch (Stytch)	0.46
India	Okta Workforce Identity Cloud (Okta)	Microsoft Entra ID (Microsoft)	HYPR (HYPR)	0.42
Germany	Microsoft Entra ID (Microsoft)	Okta / Auth0 (Okta)	Yubico (Yubico)	0.50

Microsoft Entra ID leads in the United States, Canada and Germany and is second in India. In India, Okta Workforce Identity Cloud holds first position. An Okta product appears in the leading three of every market. The third position varies: HYPR in the United States and India, Stytch in Canada and Yubico in Germany. Mean sentiment ranges from 0.42 in India to 0.50 in Germany.

4.9 Inter-engine disagreement on factual claims

The engines agree on the shape of the category. All four position FIDO2 and WebAuthn passkeys and device-bound credentials as the primary replacement for passwords and for phishable methods such as push notifications and one-time codes. All four present hardware keys and platform authenticators as complements for high-assurance use.

One substantive disagreement was recorded. On the strength of device-bound enforcement, ChatGPT repeatedly stressed explicit policy controls to block synced passkeys. Gemini treated platform authenticators as sufficient without the same level of enforcement detail.

ChatGPT alone stated three dated claims that no other engine reproduced. It stated that WebAuthn Level 3 became a W3C Recommendation in August 2026. It stated that Microsoft would make passkeys the default Entra ID experience in September 2026, with SMS retirement in February 2027. This report records these claims as observed without assessing their accuracy.

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 vendor shortlist?

At the top of the list, they do. Microsoft holds the highest brand-mention count in all four engines and Microsoft Entra ID is the first-mention entity in three of them (§4.2, §4.7). Fourteen vendors are named by all four engines. No engine nominates a leading vendor that the others omit.

Below the leader, convergence weakens. Yubico and Apple are absent from ChatGPT, and 28 of 72 brand entries appear in one engine only (§4.2). ChatGPT opens with Auth0 rather than Microsoft Entra ID. The data is consistent with a shared head and an engine-specific tail. The head is the part a buyer sees first, which is why the agreement matters more than its narrow width suggests.

5.2 RQ2 — How concentrated is brand visibility?

It is concentrated at the brand level and fragmented at the product level. The ten most-mentioned brand entries hold 68.9% of brand mentions, and Microsoft alone 17.5% (§4.2). Against that, 153 product entries are recorded, and 79 of them carry 1,000 mentions, the lowest count recorded (§4.3).

Okta shows how far the two views can diverge. It records 35,000 brand mentions, seventh among brand entries, yet its 22 product entries carry 137,000 product mentions, more than any other brand's products (§4.3). Auth0 is ChatGPT's first-mention entity while never registering as a brand entry. A plausible explanation is that engines name Okta's offerings by product rather than by company, so a brand-only count understates the vendor. The study measures this fragmentation; it does not establish its cause.

Sentiment separates some brands in this category. Google Workspace, ForgeRock and Apple all fall below 0.20, and the four products with negative scores include both synced passkey stores (§4.2, §4.3). The data is consistent with engines treating synced, consumer-oriented credentials as weaker for enterprise use. The sentiment model is a judgement of framing, though, not a measurement of the products.

5.3 RQ3 — Do the engines share an evidence base?

Largely they do not. The highest overlap between two different engines is 0.30, between the two Google search surfaces, and the lowest is 0.12 (§4.6). Even the closest pair shares less than a third of its combined domain set.

The engines also differ in how they build evidence. Google AI Overview cites 4.1 times as many sources as Gemini, across 3.6 times as many distinct domains as ChatGPT, and leans on unclassified and social sources (§4.4, §4.5). ChatGPT concentrates on vendor sites and official documentation, with 23.6% of its citations from one Microsoft documentation domain. Gemini's two most-cited domains belong to 1Kosmos and miniOrange, two brands it names far more often than the other engines do (§4.2, §4.4). The data is consistent with citation sources shaping which secondary vendors an engine names. It does not establish the direction of that relationship.

5.4 RQ4 — Does ordering follow mention volume?

Partly. Microsoft leads on both measures in Gemini and ChatGPT (§4.7). In Google AI Overview, Okta records an earlier mean position than Microsoft on far fewer mentions. In Google AI Mode, Stytc records the earliest mean position among the ten most-mentioned entries on 8,000 mentions. Ping Identity, second by volume, sits between positions 3.4 and 4.1 in every engine.

A plausible interpretation is that volume reflects how widely a vendor is documented across the sources an engine reads. Position, on this reading, reflects which framing an engine adopts for the query, for example a developer-oriented framing that opens with Stytc or Auth0. Volume and position are therefore best treated as separate objectives. The data establishes that they diverge; it does not establish the mechanism.

5.5 RQ5 — Does recommendation vary by market?

Membership of the leading group is stable; ordering varies at the margin. Microsoft Entra ID and an Okta product occupy the first two positions in every market, and only India reverses their order (§4.8). The third position changes between HYPR, Stytc and Yubico. Mean sentiment varies within a narrow band, from 0.42 to 0.50.

Many product entries are market-specific: 87 of 153 appear in one market only (§4.3). These are low-volume entries. The leading recommendations do not localise, while the long tail varies between markets in a way that this corpus cannot separate from ordinary response-to-response variation.

5.6 What these results do not resolve

The study observes engine behaviour but not its causes. It cannot say why Gemini cites 1Kosmos and miniOrange domains so heavily, why ChatGPT opens with Auth0, or why Okta leads in India. It cannot separate a vendor absent because it is thinly documented from one absent because an engine's retrieval did not reach its documentation. It does not test narrower queries, such as strict device-bound requirements, budget constraints, or customer identity as distinct from workforce identity. Because the engines are not version-pinned, it 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 stands in for prominence. It assumes that an earlier mention carries more weight with a reader, which is plausible but not measured here.

Sentiment is a language-model judgement of how favourably a response describes an entity, scored per mention and averaged. The project's own guidance treats two values within about 0.2 of each other as indistinguishable. Differences between most vendors in Table 5 fall inside that margin. The low values for Google Workspace,

ForgeRock and Apple fall outside it.

Truncation is inferred from response structure rather than observed. A response ending mid-list is flagged, but the study cannot distinguish an output-length limit from an engine legitimately concluding. Brand entries are recorded as extracted, without entity resolution. Okta and Auth0, Microsoft and Azure AD, and Beyond Identity and BeyondIdentity therefore appear as separate entries, which understates consolidated visibility for those vendors.

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. 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 markets and one collection window. The query set is purposive rather than randomly sampled from a defined population of buyer questions. For this reason 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 passwordless authentication 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 structure is a single leading vendor across engines, low cross-engine source overlap, and low dead-link rates. It would not be expected to reproduce exact mention counts, exact positions or the specific most-cited domains, which depend on retrieval conditions at collection. Categorical findings such as the first-mention entity per engine are likely to be more stable than continuous ones, but have not been tested across windows.

7. Limitations

- One category, four markets and one collection window.
- Engine responses were not fact-checked. Where engines disagree on a factual claim, the disagreement is recorded but not adjudicated.
- No entity resolution was applied, so variant brand and product names are counted as separate entries.
- A median source age could be reported for one engine only, so source freshness is not compared across engines.

8. Practical Implications

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

- **Measure more than one engine.** Source overlap between engines runs from 0.12 to 0.30 (§4.6). A programme watching one engine observes an evidence base largely unrelated to the others. ChatGPT, Gemini and one Google search surface cover three substantially different source populations.
- **Consolidate product naming.** Okta's visibility is spread over 22 product entries and Ping Identity's over 19 (§4.3). Use one consistent product name in published material, and map legacy names such as Azure AD and ForgeRock explicitly to the current product.
- **Track position separately from volume.** Stytl and Okta hold the earliest mean positions among leading entries in the two Google surfaces on low mention counts (§4.7). A single share-of-voice number would hide both results.
- **Publish to the documentation domains engines already cite.** learn.microsoft.com supplies 23.6% of ChatGPT's citations, and vendor sites supply most citations in three of four engines (§4.4). Technical documentation that states device-bound versus synced behaviour, phishing resistance and identity integration plainly is the material these engines draw on.
- **Do not neglect secondary engines.** Gemini names 1Kosmos and miniOrange far more often than the other engines, alongside heavy citation of their domains (§4.2, §4.4). A smaller vendor can hold visibility in one engine that it lacks elsewhere.
- **Address the synced-versus-device-bound question directly.** It is the one substantive point on which engines disagree (§4.9), and synced passkey stores carry negative sentiment (§4.3). Vendors whose products support device-bound enforcement should document how it is configured.

9. Conclusion

Across 160,000 AI engine responses and 1,018,000 cited sources, four AI search engines converge on Microsoft as the leading vendor for enterprise passwordless authentication. Microsoft holds the highest brand-mention count in every engine, and Microsoft Entra ID opens the response in three of them.

Below the leader the picture fragments. Brand visibility is concentrated in ten entries while product visibility spreads across 153 entries, and Okta's presence is recorded mostly under product names. The engines also read different evidence. Source overlap peaks at 0.30, and citation volume and domain breadth vary several-fold between engines. Cited links, by contrast, are largely intact.

For a vendor in this category, the practical consequence is that visibility is engine-specific below the first position. Consistent product naming, documentation on the domains engines already cite, and measurement across several engines are the levers the results point to. Market localisation is not among them, since the leading recommendations barely change between the four markets.

10. Reproducibility and Data Availability

Queries were issued to each engine's public surface through a commercial data collection layer. Each response, its cited sources and the per-response analysis were persisted to a columnar analytical database. Retained records include the response text, the extracted citations with resolved domains, source types and reachability status, and the brand and product mentions with ordinal position and sentiment. The per-engine aggregate metrics reported here are also retained.

Reproducing the study requires the 10,000-query set described in Appendix D, 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 target is the structure of the findings rather than their values. Three follow-up designs arise from the data: queries restricted to device-bound credentials, queries adding budget constraints, and queries separating customer identity from workforce identity.

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

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.

How to cite this report

GrackerAI. (October 2026). *AI Search Visibility in Enterprise Passwordless Authentication: 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-passwordless-authentication-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 entry was not recorded in that engine. Entries are listed as extracted, including standards bodies, agencies and authentication methods.

Table A1. Brand mentions and mean ordinal position, all brand entries — 72 entries, ordered by total mentions

Brand	Total	ChatGPT	Gemini	AI Overview	AI Mode
Microsoft	148,000	32,000 (1.6)	31,000 (1.4)	44,000 (1.6)	41,000 (2.3)
Ping Identity	94,000	21,000 (4.1)	24,000 (3.4)	22,000 (3.4)	27,000 (4.1)
Google Workspace	62,000	11,000	16,000 (1.7)	12,000 (2.0)	23,000 (2.0)
HYPR	60,000	22,000 (3.5)	15,000 (3.9)	6,000 (2.5)	17,000 (3.3)
Cisco	53,000	22,000 (3.0)	9,000 (4.6)	10,000 (3.4)	12,000 (3.3)
Yubico	51,000	—	21,000 (4.2)	8,000 (2.5)	22,000 (3.6)
Okta	35,000	13,000 (2.5)	9,000 (2.5)	6,000 (1.5)	7,000 (2.0)
Beyond Identity	31,000	6,000 (4.3)	4,000 (5.5)	7,000 (3.9)	14,000 (3.6)
Stytch	26,000	8,000 (2.6)	5,000 (2.5)	5,000 (3.5)	8,000 (1.6)
Apple	23,000	—	9,000 (1.5)	7,000 (2.3)	7,000 (2.0)
Descope	20,000	8,000 (2.9)	3,000 (1.3)	2,000	7,000 (2.4)
ForgeRock	16,000	5,000	3,000 (4.0)	1,000	7,000 (5.0)
Thales	13,000	6,000 (5.8)	3,000 (5.5)	1,000 (3.0)	3,000 (5.0)
1Kosmos	12,000	—	11,000 (4.8)	1,000 (4.0)	—

Brand	Total	ChatGPT	Gemini	AI Overview	AI Mode
Active Directory	12,000	2,000	—	5,000	5,000
Clerk	10,000	6,000 (4.5)	4,000 (3.0)	—	—
FIDO Alliance	10,000	6,000 (1.3)	1,000 (1.0)	—	3,000
JumpCloud	10,000	2,000 (6.5)	3,000 (6.0)	2,000 (2.5)	3,000 (5.5)
miniOrange	10,000	—	7,000 (6.0)	1,000 (4.0)	2,000 (4.5)
FIDO2/WebAuthn	9,000	4,000	3,000	1,000	1,000
OneLogin	9,000	7,000	—	—	2,000
Bitwarden	8,000	—	2,000	6,000 (4.5)	—
FusionAuth	7,000	3,000	1,000 (3.0)	1,000 (4.0)	2,000 (4.0)
1Password	6,000	1,000	3,000	1,000 (6.0)	1,000
CISA	6,000	6,000	—	—	—
NIST	6,000	6,000	—	—	—
WorkOS	6,000	4,000 (4.0)	2,000 (3.0)	—	—
OneSpan	5,000	—	3,000 (6.0)	—	2,000 (5.0)
Secret Double Octopus	5,000	—	—	2,000 (2.0)	3,000 (3.0)
CyberArk	4,000	1,000	1,000 (4.0)	—	2,000 (5.0)
Face ID	4,000	3,000	1,000	—	—
HID Global	4,000	—	1,000	2,000 (4.0)	1,000 (4.0)
LoginRadius	4,000	—	1,000	3,000 (2.0)	—
ManageEngine	4,000	—	1,000	1,000 (5.0)	2,000 (5.5)
MojoAuth	4,000	—	—	3,000 (4.3)	1,000 (4.0)
Azure AD	3,000	—	1,000	1,000	1,000
Dashlane	3,000	—	1,000	—	2,000
Double Octopus	3,000	—	—	2,000	1,000
Keycloak	3,000	1,000	1,000 (4.0)	—	1,000 (7.0)
SecureAuth	3,000	—	1,000 (2.0)	—	2,000 (3.5)
Amazon	2,000	2,000	—	—	—
Authgear	2,000	—	2,000 (2.0)	—	—
AWS	2,000	—	1,000	—	1,000
BeyondTrust	2,000	—	—	—	2,000 (5.0)
IBM	2,000	—	—	2,000 (5.0)	—
OLOID	2,000	—	—	1,000 (1.0)	1,000 (5.0)
RSA	2,000	1,000	—	—	1,000 (7.0)
Salesforce	2,000	—	1,000	1,000	—
TraitWare	2,000	—	1,000 (3.0)	1,000 (6.0)	—
Trusona	2,000	—	—	2,000 (4.5)	—
W3C	2,000	2,000 (3.0)	—	—	—
Zitadel	2,000	—	1,000 (4.0)	—	1,000 (7.0)
AD FS	1,000	1,000	—	—	—
ADFS	1,000	1,000	—	—	—
BeyondIdentity	1,000	—	1,000 (5.0)	—	—
Cloudflare	1,000	—	1,000	—	—
Entrust	1,000	—	1,000	—	—
Hardware FIDO2 security keys	1,000	1,000	—	—	—
HID DigitalPersona	1,000	1,000	—	—	—

Brand	Total	ChatGPT	Gemini	AI Overview	AI Mode
KuppingerCole	1,000	—	—	—	1,000
LoginID	1,000	—	—	1,000 (4.0)	—
NCSC (National Cyber Security Centre)	1,000	1,000	—	—	—
PCMag	1,000	—	—	—	1,000
PingOne for Workforce	1,000	1,000	—	—	—
Push-to-approve MFA	1,000	1,000	—	—	—
Recovery codes	1,000	1,000	—	—	—
SAP	1,000	—	1,000	—	—
Servnet UK	1,000	—	1,000	—	—
SMS/email OTP	1,000	1,000	—	—	—
Token2	1,000	—	1,000	—	—
TOTP/authenticator-app codes	1,000	1,000	—	—	—
Transmit Security	1,000	—	—	—	1,000
Total	846,000	221,000	213,000	171,000	241,000

Appendix B. Source mix and most-cited domains

Table B1. Source-type distribution by AI engine — cited sources by classified source type

AI engine	Vendor site	Other	Social	Forum	Government	News	Review aggregator	Academic	Total
ChatGPT	218,000	34,000	—	7,000	25,000	—	—	—	284,000
Gemini	53,000	39,000	—	—	—	—	—	—	92,000
Google AI Overview	123,000	182,000	38,000	18,000	1,000	6,000	9,000	3,000	380,000
Google AI Mode	144,000	95,000	5,000	4,000	2,000	5,000	5,000	2,000	262,000
Total	538,000	350,000	43,000	29,000	28,000	11,000	14,000	5,000	1,018,000

Table B2. Source-type shares by AI engine — share of each engine's cited sources

AI engine	Vendor site	Other	Social	Forum	Government	News	Review aggregator	Academic
ChatGPT	76.8%	12.0%	—	2.5%	8.8%	—	—	—
Gemini	57.6%	42.4%	—	—	—	—	—	—
Google AI Overview	32.4%	47.9%	10.0%	4.7%	0.3%	1.6%	2.4%	0.8%
Google AI Mode	55.0%	36.3%	1.9%	1.5%	0.8%	1.9%	1.9%	0.8%

The three most-cited domains for each engine are listed in Table 8 in Section 4.4.

Appendix C. Product entries

All 153 product entries as recorded, with mentions, mean product rank, mean sentiment, and the number of engines and markets in which each was named. An em dash indicates no rank was recorded.

Table C1. All product entries — 153 entries, in the order recorded by the collection pipeline

#	Product	Brand	Mentions	Mean rank	Sentiment	Engines	Markets
1	Microsoft Entra ID	Microsoft	84,000	1.6	0.65	4	4
2	Okta Workforce Identity Cloud	Okta	37,000	2.2	0.66	4	4
3	YubiKey	Yubico	51,000	3.1	0.59	3	4
4	HYPR	HYPR	53,000	3.5	0.68	4	4
5	Okta	Okta	35,000	2.2	0.54	4	4
6	Ping Identity	Ping Identity	44,000	3.9	0.58	4	4
7	Stytch	Stytch	24,000	2.4	0.72	4	4
8	Auth0	Okta	15,000	1.6	0.75	4	4
9	Cisco Duo	Cisco	33,000	3.7	0.64	4	4
10	Microsoft Authenticator	Microsoft	11,000	—	0.24	4	4
11	Descope	Descope	19,000	2.4	0.70	4	4
12	Beyond Identity	Beyond Identity	30,000	3.9	0.63	4	4
13	Windows Hello for Business	Microsoft	18,000	2.2	0.40	3	4
14	Cisco Duo Passwordless	Cisco	20,000	3.0	0.64	3	4
15	Okta Workforce Identity Cloud (FastPass)	Okta	8,000	1.8	0.69	3	4
16	Microsoft Entra ID (with Windows Hello & Passkeys)	Microsoft	4,000	1.0	0.70	1	3
17	Okta FastPass	Okta	14,000	3.4	0.61	4	4
18	Okta Identity Engine	Okta	8,000	2.3	0.68	2	4
19	Apple iCloud Keychain	Apple	6,000	1.0	-0.05	3	3
20	Google Password Manager	Google	6,000	1.0	-0.05	3	3
21	Google Workspace	Google	5,000	—	0.00	3	3
22	Apple Touch ID/Face ID	Apple	7,000	2.0	0.39	3	4
23	FIDO2/WebAuthn Passkeys	FIDO Alliance	2,000	1.0	0.85	2	1
24	Clerk	Clerk	9,000	3.9	0.51	2	3
25	Microsoft Entra ID (formerly Azure AD)	Microsoft	2,000	1.0	0.70	1	1
26	1Kosmos BlockID	1Kosmos	9,000	4.6	0.70	1	4
27	YubiKey / YubiEnterprise	Yubico	5,000	3.0	0.76	3	2
28	Thales SafeNet	Thales	2,000	—	0.35	2	2
29	Okta Identity Engine (FastPass)	Okta	3,000	2.0	0.70	2	3
30	HYPR Workforce Passwordless Platform	HYPR	3,000	2.0	0.67	2	2
31	Secret Double Octopus	Secret Double Octopus	5,000	3.0	0.50	2	3
32	Ping Identity / PingOne	Ping Identity	6,000	3.8	0.48	2	4
33	Ping Identity (PingOne / PingID)	Ping Identity	5,000	3.5	0.62	3	3
34	FusionAuth	FusionAuth	6,000	3.7	0.34	4	3
35	Ping Identity (PingOne / PingFederate)	Ping Identity	4,000	3.0	0.63	1	2
36	PingOne	Ping Identity	4,000	3.3	0.68	3	3
37	Okta Customer Identity Cloud & Workforce Identity Cloud	Okta	3,000	2.3	0.57	1	2
38	miniOrange	miniOrange	7,000	5.5	0.56	3	4
39	Microsoft Entra ID Passwordless	Microsoft	4,000	3.3	0.63	3	2
40	Synced passkeys (FIDO2/WebAuthn)	FIDO Alliance	1,000	1.0	0.90	1	1
41	Amazon Cognito	Amazon	2,000	—	-0.05	1	2
42	FIDO2 / WebAuthn Passkeys and Security Keys	FIDO Alliance	1,000	1.0	0.90	1	1

#	Product	Brand	Mentions	Mean rank	Sentiment	Engines	Markets
43	FIDO2 Passkeys	FIDO Alliance	1,000	1.0	0.90	1	1
44	JumpCloud	JumpCloud	6,000	5.2	0.57	4	4
45	Okta Identity Platform	Okta	1,000	1.0	0.80	1	1
46	LoginRadius Identity Platform	LoginRadius	1,000	1.0	0.80	1	1
47	Microsoft Entra ID (formerly Azure AD) & Windows Hello for Business	Microsoft	1,000	1.0	0.80	1	1
48	WorkOS	WorkOS	5,000	3.7	0.28	2	3
49	Google Titan Security Key	Google	2,000	2.0	0.70	2	2
50	HYPR Passwordless SSO	HYPR	1,000	—	0.70	1	1
51	PingOne Authenticator	Ping Identity	1,000	—	0.70	1	1
52	1Kosmos 1Key	1Kosmos	1,000	—	0.70	1	1
53	Google Titan	Google	2,000	2.0	0.60	1	1
54	Transmit Security Mosaic	Transmit Security	1,000	—	0.60	1	1
55	Microsoft Entra ID Passkeys (FIDO2)	Microsoft	1,000	—	0.60	1	1
56	PingFederate	Ping Identity	1,000	—	0.60	1	1
57	Microsoft Entra ID Conditional Access	Microsoft	1,000	—	0.60	1	1
58	Entrust Identity as a Service	Entrust	1,000	—	0.60	1	1
59	Ping Identity Passwordless	Ping Identity	1,000	—	0.60	1	1
60	Okta Customer Identity Cloud (formerly Auth0)	Okta	1,000	—	0.50	1	1
61	Token2	Token2	1,000	—	0.50	1	1
62	MojoAuth	MojoAuth	4,000	4.3	0.55	2	3
63	LoginRadius	LoginRadius	2,000	2.5	0.78	1	2
64	Google Workspace / Cloud Passkeys	Google	1,000	—	0.40	1	1
65	RSA ID Plus	RSA	1,000	—	0.40	1	1
66	Okta FastPass / Workforce Identity	Okta	1,000	—	0.40	1	1
67	PingID	Ping Identity	1,000	—	0.40	1	1
68	HID Authentication Service	HID Global	1,000	—	0.40	1	1
69	Ping Identity (ForgeRock)	Ping Identity	3,000	3.5	0.63	2	3
70	Ping Identity (PingFederate / PingOne)	Ping Identity	3,000	3.7	0.63	2	3
71	OneLogin Workforce Identity	OneLogin	1,000	—	0.30	1	1
72	1Password Extended Access / Unlock with SSO	1Password	1,000	—	0.30	1	1
73	YubiKey 5 Series	Yubico	1,000	—	0.30	1	1
74	1Password Extended Access Management	1Password	1,000	—	0.30	1	1
75	JumpCloud Go	JumpCloud	4,000	5.0	0.60	3	3
76	Duo Verified Push	Duo	1,000	—	0.20	1	1
77	OLOID	OLOID	2,000	3.0	0.75	2	2
78	Thales SafeNet Trusted Access	Thales	5,000	5.8	0.32	1	4
79	Microsoft Entra External ID + Entra ID	Microsoft	4,000	5.0	0.42	2	3
80	PingOne / PingFederate	Ping Identity	2,000	3.0	0.60	2	2
81	ForgeRock	ForgeRock	3,000	4.7	0.57	2	3
82	Bitwarden	Bitwarden	4,000	5.5	0.38	2	3
83	Thales Access Management (SafeNet)	Thales	2,000	3.5	0.75	2	1
84	Okta Verify	Okta	1,000	—	0.00	1	1
85	Okta Workforce	Okta	1,000	—	0.00	1	1
86	OneLogin	OneLogin	1,000	—	0.00	1	1

#	Product	Brand	Mentions	Mean rank	Sentiment	Engines	Markets
87	Bitwarden Passwordless.dev	Bitwarden	2,000	3.5	0.70	1	2
88	Hardware-backed device-bound passkeys / security keys	FIDO Alliance	1,000	2.0	0.85	1	1
89	SecureAuth	SecureAuth	2,000	3.5	0.60	1	2
90	Okta (Customer Identity Cloud / Auth0 & Workforce Identity)	Okta	1,000	2.0	0.80	1	1
91	Firebase Auth	Google	1,000	—	-0.10	1	1
92	ManageEngine ADSelfService Plus	ManageEngine	3,000	5.3	0.57	2	2
93	Okta Workforce Identity Cloud (with Okta FastPass)	Okta	1,000	2.0	0.70	1	1
94	Secret Double Octopus (Octopus Authentication Platform)	Secret Double Octopus	1,000	2.0	0.70	1	1
95	Okta FastPass / Advanced Server Access	Okta	1,000	2.0	0.70	1	1
96	HYPR Identity Assurance Platform	HYPR	1,000	2.0	0.70	1	1
97	Okta Workforce Identity Cloud (Auth0)	Okta	1,000	2.0	0.70	1	1
98	Authgear	Authgear	1,000	2.0	0.65	1	1
99	HID Global	HID Global	2,000	4.0	0.65	2	2
100	Apple Platform SSO	Apple	1,000	2.0	0.60	1	1
101	Okta Adaptive MFA / Okta Workforce Identity	Okta	1,000	2.0	0.60	1	1
102	SecureAuth Workforce Authority	SecureAuth	1,000	2.0	0.60	1	1
103	Google Cloud / Firebase / Workspace	Google	1,000	2.0	0.50	1	1
104	Trusona	Trusona	2,000	4.5	0.60	1	1
105	CyberArk Identity	CyberArk	2,000	4.5	0.55	2	2
106	TraitWare	TraitWare	2,000	4.5	0.55	2	2
107	Ping Identity PingID	Ping Identity	2,000	4.5	0.50	1	2
108	1Kosmos	1Kosmos	2,000	5.0	0.65	2	2
109	Yubico YubiEnterprise (YubiKey)	Yubico	2,000	5.5	0.80	1	2
110	IBM Verify	IBM	2,000	5.0	0.60	1	2
111	BeyondTrust	BeyondTrust	2,000	5.0	0.55	1	2
112	OneSpan Digipass	OneSpan	2,000	5.0	0.55	2	2
113	Yubico (YubiKeys)	Yubico	2,000	5.5	0.70	1	2
114	WebAuthn/FIDO2	W3C / FIDO Alliance	1,000	3.0	0.80	1	1
115	Thales Access Management (SafeNet Trusted Access)	Thales	2,000	5.5	0.65	1	1
116	Okta Identity Engine / Auth0	Okta	1,000	3.0	0.70	1	1
117	PingOne / PingID	Ping Identity	1,000	3.0	0.70	1	1
118	Microsoft Entra passwordless	Microsoft	1,000	3.0	0.70	1	1
119	WorkOS AuthKit	WorkOS	1,000	3.0	0.70	1	1
120	Okta Customer Identity	Okta	1,000	3.0	0.70	1	1
121	Keycloak	Keycloak	2,000	5.5	0.55	2	1
122	Zitadel	Zitadel	2,000	5.5	0.55	2	1
123	Ping Identity (incorporating ForgeRock)	Ping Identity	1,000	3.0	0.60	1	1
124	Okta FastPass / Passkeys	Okta	1,000	3.0	0.50	1	1
125	Cisco Secure Access by Duo	Cisco	1,000	3.0	0.50	1	1
126	Apple Touch ID	Apple	1,000	3.0	0.50	1	1
127	1Password	1Password	2,000	6.0	0.35	2	2
128	Okta Workforce & Customer Identity Cloud	Okta	1,000	4.0	0.70	1	1

#	Product	Brand	Mentions	Mean rank	Sentiment	Engines	Markets
129	Ping Identity (including ForgeRock)	Ping Identity	1,000	4.0	0.70	1	1
130	PingOne / Ping Identity	Ping Identity	1,000	4.0	0.60	1	1
131	LoginID Passkey & MFA	LoginID	1,000	4.0	0.60	1	1
132	HYPR Workforce Access	HYPR	1,000	4.0	0.60	1	1
133	Ping Identity Enterprise Connect Passwordless	Ping Identity	1,000	4.0	0.60	1	1
134	miniOrange Passwordless Authentication	miniOrange	1,000	4.0	0.60	1	1
135	Yubico YubiKey (Hardware-Backed Authentication)	Yubico	1,000	4.0	0.60	1	1
136	Ping Identity (PingOne / DaVinci)	Ping Identity	1,000	4.0	0.60	1	1
137	Microsoft Entra ID Authenticator	Microsoft	1,000	3.0	0.20	1	1
138	Microsoft Entra ID (External ID & Workforce ID)	Microsoft	1,000	4.0	0.50	1	1
139	Yubico YubiKey (Yubico Enrollment Suite)	Yubico	1,000	5.0	0.80	1	1
140	Microsoft Authenticator / Entra ID	Microsoft	1,000	5.0	0.60	1	1
141	Yubico (YubiKey ecosystem)	Yubico	1,000	5.0	0.60	1	1
142	HYPR Enterprise Passkeys	HYPR	1,000	5.0	0.60	1	1
143	Beyond Identity Secure Workforce	Beyond Identity	1,000	5.0	0.60	1	1
144	BeyondIdentity	BeyondIdentity	1,000	5.0	0.60	1	1
145	OneSpan Digipass FIDO2 Series	OneSpan	1,000	5.0	0.60	1	1
146	OneSpan Digipass FIDO2	OneSpan	1,000	5.0	0.50	1	1
147	Yubico YubiKey Hardware Keys	Yubico	1,000	6.0	0.80	1	1
148	Thales Passwordless 360° / SafeNet	Thales	1,000	6.0	0.70	1	1
149	Yubico (YubiKey + Yubico Enterprise Service)	Yubico	1,000	6.0	0.70	1	1
150	Auth0 / Okta Customer Identity	Auth0	1,000	6.0	0.70	1	1
151	HID DigitalPersona	HID Global	1,000	6.0	0.60	1	1
152	RSA Unified Passwordless	RSA	1,000	7.0	0.60	1	1
153	OneSpan (Digipass Authentication)	OneSpan	1,000	7.0	0.60	1	1

Appendix D. Query set and markets

Table D1. 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 (five core questions with paraphrase variants, issued across repeated collection cycles)
Responses per engine	40,000 (10,000 queries × 4 markets)
Total responses	160,000 (10,000 queries × 4 engines × 4 markets)
Total cited sources	1,018,000
Brand entries recorded	72

Parameter	Value
Product entries recorded	153
Collection window	October 2026
Dataset generated	October 2026

Table D2. Core questions anchoring the query set

#	Core question
Q1	Best passwordless authentication solutions for enterprises in 2026
Q2	Which passwordless authentication platforms support FIDO2 passkeys, phishing-resistant MFA and device-bound credentials for employees?
Q3	Which workforce authentication tools can replace passwords and one-time codes while integrating with an existing IdP and SSO?
Q4	Our users keep getting phished through MFA push notifications and one-time codes. What should we use instead of passwords and codes?
Q5	Which passwordless login solutions work for customer-facing apps as well as employees, with SDKs for web and mobile?

The query set combines these core questions with paraphrase variants. Every query was issued to every engine in every market.

Appendix E. Metric computation

Table E1. Formulas for derived metrics

Metric	Computation
Responses per engine	queries × markets
Total responses	queries × engines × markets
Total cited sources	Σ cited sources across the four engines
Unreachable citations	cited sources × dead-link rate, per engine, rounded to the nearest whole citation
Aggregate dead-link rate	Σ unreachable citations ÷ Σ cited sources
Source-type share	cited sources of one type ÷ the engine's cited sources
Brand mention total	Σ brand mentions across the four engines, per brand entry
Share of brand mentions	brand mentions of one or more entries ÷ brand mentions across all entries
Product entries per brand	count of distinct product entries attributed to a brand
Jaccard source overlap	$ A \cap B \div A \cup B $, where A and B are two engines' cited-domain sets
Mean ordinal position	mean index of a brand's first appearance within the responses that name it

Metric	Computation
Mean sentiment	arithmetic mean of per-mention polarity scores, range -1.0 to 1.0