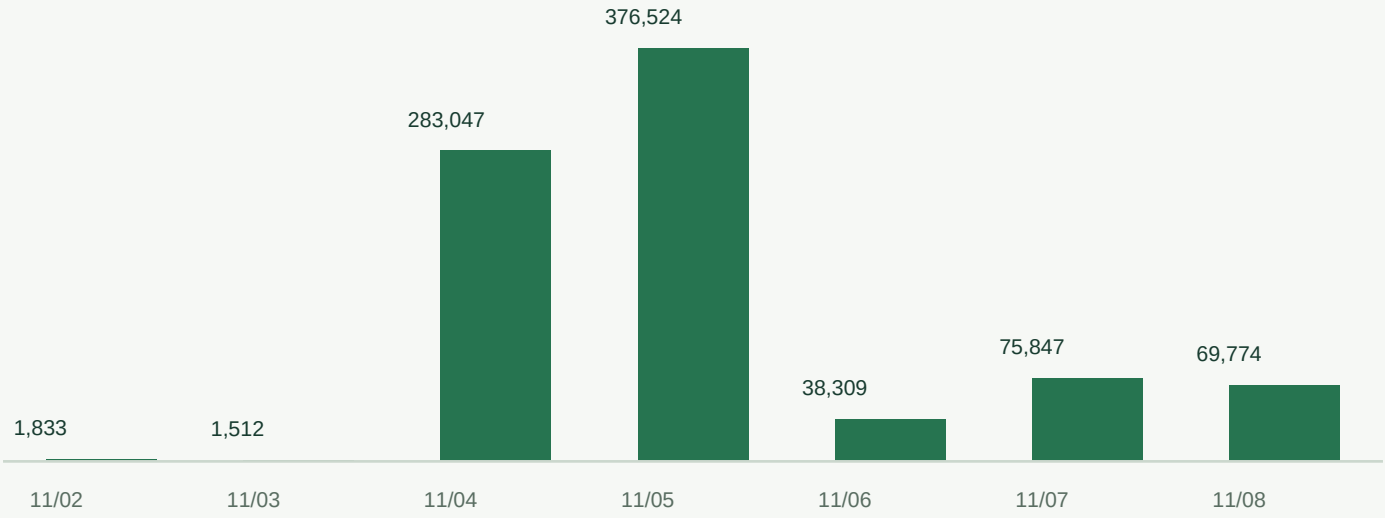
OVERVIEW / FIRST-PARTY SITE OBSERVATIONS

AI and search crawlers:
Which brands are they visiting?

From 846,846 site requests, we identified 190,043 visited brand, company and organization profiles. Linking sources, profiles and content types reveals how crawlers access brand information.

Weekly requests	Matched profiles	Source companies	Error rate
846,846	190,043	15	0.122%
GET, HEAD and failures	Brands and organizations / 2xx	25 crawler names	HTTP status >=400

Daily crawler requests



77.89% of requests occurred on November 4-5. These two days contributed 659,571 requests. Monday and Tuesday each recorded fewer than 2,000; activity surged on Wednesday, peaked on Thursday, then fell.

The busiest hour: November 5, 14:00-15:00, Beijing time (UTC+8), with 39,248 requests. The peak does not represent sustained activity throughout the day.

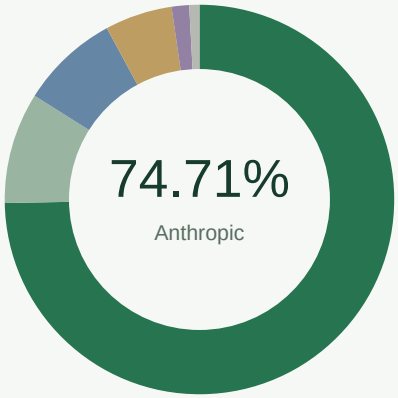
No week-over-week comparison. Logging began October 28 at 23:17, Beijing time, leaving the preceding week incomplete. This report establishes a complete-calendar-week baseline.

SOURCES / VOLUME AND BREADTH

Anthropic accounts for 74.71%.
Coverage varies across companies.

Request counts measure frequency; distinct profile counts measure breadth. Different crawler names also represent different types of activity.

Sources of 846,846 weekly requests



Anthropic	632,680	74.71%
Amazon	78,186	9.23%
OpenAI	69,190	8.17%
Google	47,354	5.59%
Perplexity	12,301	1.45%
10 others	7,135	0.84%

OpenAI: 69,190 requests. GPTBot contributed 50,764, OAI-SearchBot 18,365 and ChatGPT-User 61. Anthropic recorded 632,668 ClaudeBot requests and 12 Claude-User requests.

Source	Distinct matched profiles
Anthropic	189,178
Amazon	33,757
Google	8,533
Perplexity	3,734
OpenAI	1,562

Profile counts include only matched visits with 2xx responses. The same profile may be visited by multiple companies, so source-level coverage counts must not be added together.

Crawler names or inferred purposes do not prove that content was used for training, cited or fetched at a user's request.

PROFILE COVERAGE / BREADTH AND OVERLAP

190,043 profiles received visits.

One in five had multiple sources.

At least two companies visited 40,692 profiles, representing 21.41% of the covered population. Profiles include brands, companies and organizations.



One company: 149,351 | Exactly two: 35,706 | Three or more: 4,986

Top profiles x source companies / dots indicate recorded 2xx visits

	Anthropic	Amazon	OpenAI	Google	Perplexity	Meta	Microsoft	ByteDance	Exa	Shape	Requests
Amazon	●	●	●	●	●	●	●	●	●	●	498
Google	●	●	●	●	●	●	●	●	●	●	380
Tesla	●	●	●	●	●	●	●	●	·	●	309
Samsung	●	●	●	·	●	·	●	·	●	●	293
Apple	●	●	●	●	●	●	●	●	·	●	260
Samsung Electronics	●	●	●	·	●	·	●	·	·	●	232
Amazon Web Services	●	●	●	●	●	·	●	·	·	●	204
Universal Music Group	●	●	●	·	●	●	●	●	·	●	197
Microsoft	●	●	●	●	●	●	●	●	·	●	193
Facebook	●	●	●	●	●	·	●	●	●	●	187

Amazon and Google each recorded visits from ten companies. Dots indicate presence, not intensity, citations or recommendations. Samsung and Samsung Electronics remain distinct database entities.

188,317 profiles first appeared in the available logs this week. The short logging history strongly affects this metric. First observed does not mean first discovered by AI.

CONTENT AND SECTORS / WHAT GETS FETCHED

Controversy and recent-event paths received 104,985 requests.

Controversy paths recorded 59,566 requests and recent-event paths 45,419. Path categories include failures; profile analysis uses matched 2xx responses only.

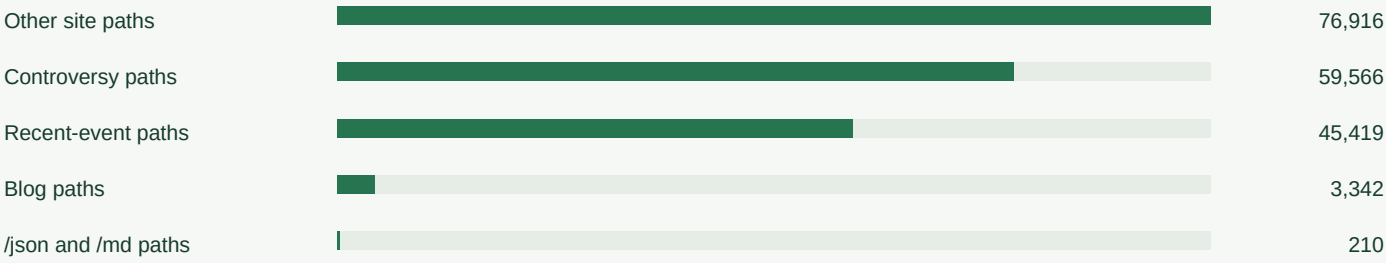
Other Brand Wiki paths: 661,393 requests, or 78.10%. This includes main profiles and some site entry points. The detail chart below compares the remaining 21.90% on its own linear scale.

Other Brand Wiki

78.10% / 661,393 requests

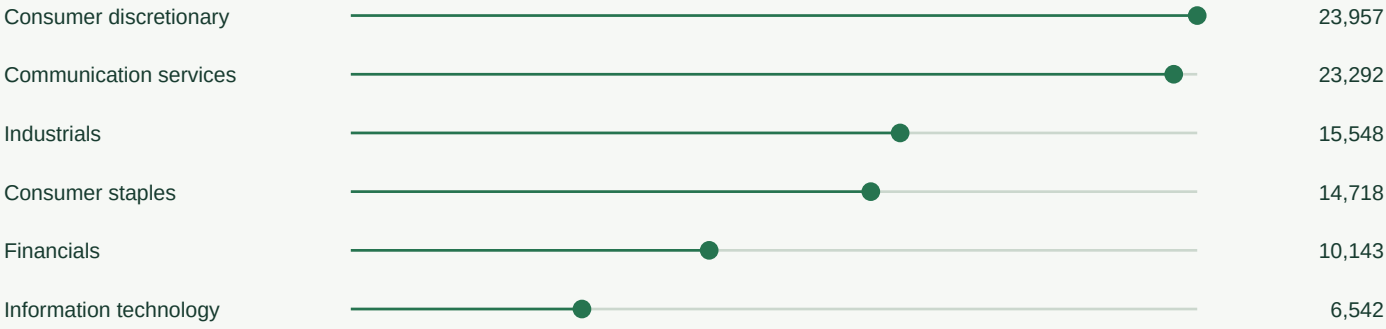


Other content paths / requests



Blog requests involved 28 distinct paths; /json and /md requests involved 97. Path counts are not article or profile counts. Few explicit-format requests do not imply that crawlers cannot read ordinary pages.

Six major classified sectors / distinct profiles



45.09% of covered profiles are classified as Other or lack a classification. Sector populations differ. These six selected sectors describe distribution, not preference. Classifications reflect the database at extraction time.

QUALITY AND METHODS / READING THE EVIDENCE

An error rate of 0.122%.

Almost all errors were missing pages.

Among 1,034 error responses, approximately 99.42% were 404s. Follow-up work should examine old links, path spellings and requests for nonexistent addresses.

Error composition / among 1,034 error responses only



404: 1,028 / 99.42%

403: 6 / 0.58%

HTTP status	Meaning	Requests
200	OK	845,705
308	Redirect	106
304	Not modified	1
404	Not found	1,028
403	Forbidden	6

Watch next week: whether concentrated crawls continue, how many profiles receive return visits, and whether multiple-source coverage becomes more common.

Methodology and limitations

Window and sources: November 2, 00:00 to November 9, 00:00, Beijing time (UTC+8), end excluded. Site logs are classified by User-Agent and deduplicated. SEO tools are excluded; GET, HEAD and failures are included.

Profile matching: Only matched 2xx responses are included. Paths are merged by database identity at extraction time, yielding 747,883 profile requests. Another 1,650 unmatched path keys, representing 4,101 requests, are excluded from profile and sector analysis.

History: Logging began October 28 at 23:17, Beijing time. The preceding week is incomplete, so no week-over-week comparison is provided. Without an end-of-week inventory snapshot, whole-database coverage is not calculated. Sectors use current classifications.

Interpretation: Identity verification results were not stored. Requests are not unique users, citations, training adoption or business conversions. First recorded does not mean first discovered.

Snapshot: Extracted around November 10, 2026, at 12:21, Beijing time (UTC+8). Later deduplication or profile revisions may slightly change historical totals.