

RESEARCH REPORT 003 / VERSION 1.0

Beyond Citation Count.

What 35,310 reported AI citations reveal about time, volume, concentration, relative presence, and the pages supplying the evidence.

JESSE KILLIAN

580 DIGITAL INFRASTRUCTURE RESEARCH

AUGUST 12, 2026

35.3K

**35,310 reported
citations across 181
days.**

Exploratory
observational case
study. First-party
Microsoft reporting.
Not peer reviewed.

THE RECORD AT A GLANCE

The headline number was only the beginning.

The daily overview, page export, and grounding-query export are separate reporting views. They are analyzed independently and never combined into one denominator.

35,310

aggregate citations reported

181-day overview export

1,354

cited URLs

Page-level export

55.1%

held by the top 10 query groups

Query-export concentration

14.7%

held by the top 10 pages

Page-export concentration

94.4%

from structured page families

/fragrance, /clone, /clone-of


OBSERVED PATTERN

Concentrated information
needs. Distributed
structured evidence.

FOUR VIEWS OF THE RESULT

The distribution is the finding.

Time, citation volume, relative presence, query concentration, and page participation described different parts of Scentonym's visibility.




The aggregate total came from a changing daily pattern, not a steady level of citation activity.

FIGURE 2 / RELATIVE PRESENCE

Citation volume and Citation Share diverged.

Each point is one exported grounding-query group. N = 72.




Citation volume and Citation Share were only weakly aligned within the sampled grounding-query export.

FIGURE 3 / CONCENTRATION

One property. Two different distribution shapes.

Percent of citations represented in each separate sampled export.




Query activity concentrated faster than page-level activity.

FIGURE 4 / PAGE ARCHITECTURE

Structured pages carried the citation activity.

Researcher-defined URL families within the 22,745-citation page export.



94.4% of page-export citations came from /fragrance, /clone, and /clone-of pages.

Structured entity and relationship pages supplied most represented citations.

Abstract

AI search visibility is often reduced to citation totals, share-of-voice scores, or screenshots of favorable answers. Those measurements can be useful, but they do not describe the same outcome. A citation total says how often a source was displayed. It does not show how broadly that activity was distributed, where it was concentrated, how much relative citation presence the source held, or whether the associated brand was recommended.

This report examines six months of first-party Bing Webmaster Tools AI Performance data from Scentonym, an independent fragrance-information website that was approximately eight months old at the end of the observation period.

The complete daily overview export contains **181 consecutive dates from February 11 through August 10, 2026**, with no missing days. It reports **35,310 citations**, a daily mean of **195.1 citations**, and a mean of **59.75 cited pages per day**. An earlier dashboard capture recorded August 9 displayed a rounded **34.5K citations** and **59 average cited pages**. Microsoft notes that reported results may be refined as additional data are processed, so the later export supplies the exact aggregate used in this report while the screenshot preserves the earlier interface state.

Two additional exports provide narrower analytical views. The page export contains **1,354 cited URLs accounting for 22,745 sampled page-level citations**. The grounding-query export contains **72 grouped query representations accounting for 1,901 sampled query-level citations**, together with Microsoft-generated intent, topic, and Citation Share classifications.

The data reveal temporal variation and two different distribution patterns.

Daily citation activity was not steady. The seven-day average began at **69.1 citations per day**, reached **401.7** for the period ending July 26, and ended at **266.4**. The final 30 days averaged **254.5 daily citations**, compared with **132.2** during the first 30 days. This describes variation within one observed window, not permanent growth.

Query-level activity was concentrated. The ten highest-volume grounding-query groups accounted for **55.1%** of citations represented in the query export. Commercial and Comparison classifications together represented **72.9%** of that sampled activity. Citation count and Citation Share were only weakly aligned, with a Pearson correlation of approximately **0.028** and a Spearman rank correlation of approximately **0.181**.

Page-level activity was considerably broader. The ten most-cited URLs accounted for only **14.7%** of citations represented in the page export, while the top 100 accounted for **47.9%**. Citation activity overwhelmingly involved structured fragrance and relationship pages: URLs under `/fragrance`, `/clone`, and `/clone-of` accounted for **94.4%** of page-export citations. Blog pages accounted for **2.9%**.

The central observation is therefore not simply that Scentonym received a large number of citations. It is that the activity changed substantially over time while **a concentrated set of recurring information needs drew from a broad base of structured pages**.

This report does not establish that page structure caused citation activity. It does not prove recommendation, traffic, sales, or equivalent visibility in another AI system. Its contribution is narrower:

an original, reproducible case showing why citation visibility should be measured as a profile rather than collapsed into one number.

What this report contributes

This report makes six contributions.

It publishes first-party documentation of substantial citation activity from a relatively young independent website.

It publishes a complete 181-day overview export and measures change within the six-month window.

It separates the aggregate daily view from sampled page-level and query-level exports.

It measures concentration at both the page and grounding-query levels.

It documents the difference between absolute citation volume and relative Citation Share.

It shows that citation activity can be concentrated by information need while remaining distributed across a broad page base.

The report extends the argument developed in [Measuring AEO](#), which separates technical eligibility, retrieval, citation, information use, brand presence, recommendation, preference, and business outcome.¹

The central claim is straightforward:

A large citation number establishes source use. It does not describe the complete shape or consequence of that visibility.

1. Why citation count is not enough

Traditional search reporting commonly revolves around rankings, impressions, clicks, and conversions. Generative systems can retrieve information from multiple pages and synthesize those sources into an answer.

The foundational GEO research presented at KDD 2024 treated visibility within generated answers as an optimization problem distinct from conventional ranking.² More recent work has separated citation selection from citation absorption, recognizing that a system can cite a page without relying heavily on it for the final answer.³ A 2026 critical survey describes generative visibility as a multistage, partially observable process rather than one ranking event.⁴

Microsoft's AI Performance reporting makes a related distinction in practice. The public preview introduced total citations, cited pages, grounding queries, page-level activity, and trends.⁵ The expanded report also includes intent classifications, topics, Citation Share, and comparisons over time. Microsoft explicitly states that the data are aggregated and sampled, that totals can differ between views, and that citation reporting does not measure rankings, authority, traffic, or page importance.⁶

Citation count answers one legitimate question:

How often was content from this property displayed as a source?

It does not answer several others:

How many pages participated?

Did a few pages produce most of the activity?

Did a few information needs produce most of the activity?

How much of the available citation space did the property occupy?

Was the brand named or recommended?

Did the visibility produce traffic or business?

This study examines what becomes visible when the headline number is separated into its underlying reporting views.

2. Study context and disclosure

The subject property is **Scentonym**, an independently operated fragrance-information website.

Scentonym organizes information about fragrances, comparable products, clone relationships, brands, notes, performance, and value. At the end of the observed six-month period, the website was approximately eight months old.

The property was selected because it generated enough first-party citation activity to support descriptive analysis and because its operator had access to the underlying Bing Webmaster Tools reports.

The author of this report also operates Scentonym. That is a direct conflict of interest.

The conflict is disclosed rather than minimized. The original dashboard capture, source exports, analysis materials, and checksums are included so the descriptive results can be inspected without relying entirely on the author's interpretation.

This is an observed case. It is not a controlled test of a service or optimization technique. No untreated control property, randomized intervention, or counterfactual design is available.

3. Three reporting views, not one denominator

The study uses three related artifacts from Bing Webmaster Tools.

Reporting view	Observed value	Correct interpretation
Daily overview export	35,310 citations across 181 days	Aggregate daily citation activity from February 11 through August 10
Page export	22,745 citations across 1,354 URLs	Sampled page-level activity represented in the export
Grounding-query export	1,901 citations across 72 groups	Sampled query-level activity represented in the export

These totals are not treated as parts of one reconciled dataset.

Microsoft states that AI Performance data are aggregated and summarized, that the information shown is not a complete log, and that totals may differ between page, grounding-query, and time-series views. Microsoft also notes that different views and filters can use samples from slightly different portions of the selected period.⁶

Accordingly:

22,745 is not reported as 64.4% of 35,310.

1,901 is not reported as 5.4% of 35,310.

Page-level and grounding-query counts are not joined as though every exported row has a complete one-to-one mapping.

The three views answer different descriptive questions.

3.1 Daily overview export and dashboard documentation

The daily overview export contains **181 consecutive dates** from February 11 through August 10, 2026. It has no missing dates and no duplicate rows.

Daily overview measure	Observed value
Total Citations	35,310
Daily citation mean	195.1
Daily citation median	178
Mean Cited Pages	59.75 per day
Selected reporting view	6 months

The earlier six-month dashboard capture, recorded August 9, 2026, reported:

Dashboard measure	Observed value
Total Citations	34.5K
Average Cited Pages	59
Selected reporting view	6 months

Microsoft defines Total Citations as the number of times content was visibly referenced or shown as a source in supported AI-generated answers during the selected range. Average Cited Pages is the average number of unique pages cited per day.⁶

The dashboard screenshot and later overview export are dated source records, not interchangeable measurements. The screenshot documents an earlier rounded interface total. The export extends through August 10 and provides the exact daily values used for aggregate and temporal analysis. Microsoft states in the interface that results may be refined as additional data are processed.

3.2 Page-level export

The page export contains two fields:

Page
Citations

It includes **1,354 unique URL rows** and **22,745 represented citations**.

This view is used to measure page participation, page-level concentration, and citation distribution across researcher-defined URL families.

3.3 Grounding-query export

The grounding-query export contains five fields:

Grounding Query
Intent
Topic
Citations
Citation Share

It includes **72 grouped query rows** and **1,901 represented citations**.

Microsoft describes grounding queries as grouped phrases associated with retrieval activity. They are not complete user prompts and should not be interpreted as direct estimates of consumer search volume.⁶

Citation Share is the percentage of citations attributed to the measured site out of all citations shown for the same grounding-query group. It measures relative citation presence, not rank, traffic, quality, or recommendation.

4. Methods

All three source CSV files were preserved without alteration.

For the daily overview analysis:

- dates were parsed as calendar dates and checked for gaps and duplicates;
- daily citation and cited-page values were parsed as integers;
- total, mean, median, minimum, maximum, and standard deviation were calculated;
- seven-day rolling means were calculated for citations and cited pages;
- the first and final 30-day daily means were compared descriptively;
- Pearson correlation was calculated between daily citations and daily cited pages.

For the grounding-query analysis:

- citation counts were parsed as integers;
- Citation Share percentages were converted to numeric percentages;
- Microsoft's Intent and Topic classifications were retained as supplied;
- missing classifications remained unclassified;
- no rows were removed;
- Pearson and Spearman correlations were calculated descriptively between Citations and Citation Share;
- repeated named product concepts were grouped through case-insensitive phrase matching for an illustrative semantic-family analysis.

For the page analysis:

- citation counts were parsed as integers;
- no URL rows were removed;
- pages were ordered from highest to lowest citation count;
- top-N concentration was calculated against the 22,745 citations represented in that export;
- each URL was assigned to a path family using its first path segment.

The path-family classification is researcher-defined:

- `/fragrance` pages are treated as fragrance entity pages;
- `/clone` pages are treated as pages identifying alternatives to a fragrance;
- `/clone-of` pages are treated as pages organizing products associated with a target fragrance;
- `/blog` pages are treated as conventional editorial content.

The combined `/fragrance`, `/clone`, and `/clone-of` category is referred to as **structured entity and relationship pages**. That label describes the site's URL and content architecture. It is not a Microsoft

classification.

All calculations are descriptive. The export's sampling process is not sufficiently documented to support population-level inference from the rows.

5. Results

5.1 Daily activity changed substantially within the six-month window

The complete overview export contains **35,310 citations across 181 consecutive days**. Daily citations averaged **195.1**, with a median of **178** and a standard deviation of **84.4**.

The lowest daily count was **45** on February 14. The highest was **533** on May 31.

The seven-day rolling average changed substantially:

Seven-day period ending	Average daily citations
February 17	69.1
July 26	401.7
August 10	266.4

The first 30 days averaged **132.2 citations per day**. The final 30 days averaged **254.5**, an increase of **92.5%** within the observed window. Average daily cited pages increased from **61.23** in the first 30 days to **71.03** in the final 30 days.

Daily citation count and cited-page count were positively associated, with a descriptive Pearson correlation of **0.685**. That does not mean one caused the other. Both measures can rise when more of the property is being used across supported answers.

The time series also shows sharp changes. The seven-day citation average fell from its July peak before rising again at the end of the export. A six-month total therefore does not describe a steady level of daily activity.

The aggregate establishes scale. The daily export shows that the scale was produced by a changing process.

5.2 The aggregate record establishes substantial source use

For an approximately eight-month-old property, the **35,310-citation** record establishes a narrow but important fact:

Scentonym's content was repeatedly displayed as source material across supported Microsoft AI experiences within its first year.

It does not establish why the property was selected. It does not show that young sites generally receive comparable visibility.

5.3 Grounding-query activity was concentrated

The 72 grounding-query groups represented 1,901 citations.

Highest-volume query groups	Citations	Share of query-export citations
Top 1	173	9.1%
Top 3	503	26.5%
Top 5	744	39.1%
Top 10	1,048	55.1%
Top 20	1,386	72.9%

More than half of the represented query-level activity came from ten of the 72 groups.

An aggregate citation total can therefore grow while depending heavily on a relatively small set of recurring information needs.

5.4 Commercial and Comparison classifications dominated the query sample

Microsoft's Intent field classified the represented citations as follows:

Microsoft Intent classification	Citations	Share
Commercial	1,188	62.5%
Informational	250	13.2%
Comparison	198	10.4%
Research	142	7.5%
Other or unclassified	123	6.5%
Total	1,901	100%

Commercial and Comparison classifications together represented 72.9% of the query-export citations.

That result shows that the sampled citation activity was not confined to general informational contexts. It does not show that a purchase occurred. Microsoft's intent labels are machine-generated classifications and can be imperfect, particularly in specialized categories.⁶

5.5 Citation volume and Citation Share were weakly aligned

Across the 72 query groups:

- median Citation Share was **26.2%**;
- the first quartile was **18.8%**;
- the third quartile was **35.0%**;
- the minimum was **7.86%**;
- the maximum was **66.67%**.

The descriptive association between citation count and Citation Share was weak:

CITATION COUNT AND CITATION SHARE

Pearson $r = 0.028$ | Spearman $\rho = 0.181$

Several source rows make the distinction clear:

Grounding-query group	Citations	Citation Share
imagination dupe	173	14.08%
zara tobacco rich warm addictive	166	37.30%
hawas ice dupe of	106	43.62%
varakh silver dupe	10	66.67%
Hawas Ice clone	10	66.67%

The highest-volume group produced 173 citations but held only 14.08% Citation Share. Two ten-citation groups each reached 66.67%.

Eight groups had Citation Share of at least 50%. Together, they represented only **123 citations**, or **6.5%** of query-export activity.

Citation Count measures absolute source-use activity. Citation Share measures relative source presence for one grounding-query group. Neither substitutes for the other.

5.6 Page-level activity was broadly distributed

The page export shows a different distribution from the query export.

Highest-volume pages	Citations	Share of page-export citations
Top 1	825	3.6%
Top 3	1,797	7.9%
Top 5	2,306	10.1%
Top 10	3,336	14.7%
Top 20	4,983	21.9%
Top 50	7,883	34.7%
Top 100	10,902	47.9%

The median cited page received **6 citations** in the export. The most-cited page received 825, or 3.6% of represented page-level activity.

The top ten grounding-query groups held more than half of query-export citations. The top ten pages held less than one-sixth of page-export citations.

These percentages should not be compared as if the exports were matched samples. They do, however, describe a clear difference in the shape of each reporting view: query activity was concentrated, while page participation extended across a large collection of URLs.

5.7 Structured entity and relationship pages accounted for most page citations

Citation activity by URL family was:

URL family	Cited pages	Citations	Share of page-export citations
/fragrance	692	10,266	45.1%
/clone	420	6,934	30.5%
/clone-of	154	4,265	18.8%
/blog	53	668	2.9%
/brand	27	385	1.7%
Homepage	1	193	0.8%
Other paths	7	34	0.1%

Combined, `/fragrance`, `/clone`, and `/clone-of` contained:

- 1,266 cited pages;
- 21,465 represented citations;
- 94.4% of all page-export citation activity.

This is one of the study's strongest observations.

Scentonym's citation activity was not primarily concentrated in conventional articles about broad topics. It was distributed across pages that describe specific fragrance entities and explicit relationships among products.

The result does not establish that the URL structure or page format caused the citations. It establishes where the citation activity appeared.

5.8 Repeated information needs appeared under multiple formulations

Closely related product concepts appeared across multiple grounding-query strings.

Seven Hawas Ice-related groups represented **274 citations**. Six Imagination-related groups represented **207**.

Illustrative concept family	Query groups	Represented citations
Hawas Ice	7	274
Imagination	6	207
God of Fire	3	153
Afnan 9PM	5	86
Rayhaan Obsidian	3	58
Khamrah	3	45

These are researcher-created groupings based on repeated entity phrases. They are not Microsoft topic clusters.

The observation reinforces an important warning:

A grounding-query phrase is a retrieval representation, not an independent measure of human demand.

6. Concentrated demand, distributed evidence

Taken together, the two exports reveal the study's most useful pattern.

At the query level, a relatively small collection of recurring information needs accounted for much of the represented activity.

At the page level, citations were spread across hundreds of specific fragrance and relationship pages. No single page dominated the export. Even the top 100 URLs accounted for less than half of represented page-level citations.

The pattern can be described as:

OBSERVED CITATION PATTERN

Concentrated information needs + distributed structured evidence

This is not a causal model. The exports cannot show that one query group retrieved every page within a related family, and their sampled totals should not be joined directly.

It is a descriptive account of the property visible through two different first-party reporting surfaces.

The distinction matters for strategy. A site can depend on a relatively narrow group of recurring questions without depending on only a handful of pages. Conversely, a large content library does not necessarily imply diversified demand.

Both dimensions should be measured.

7. Citation visibility is a profile

This case supports reporting citation visibility through six separate dimensions.

Volume

How often was content from the property displayed as a source?

Breadth

How many pages, entities, or subject areas participated?

Concentration

How much activity depended on the highest-volume pages or grounding-query groups?

Relative share

How much of the citation space did the property occupy within each reported grounding-query group?

Intent

In what types of platform-classified contexts did the citation activity occur?

Persistence

Did the pattern continue, grow, decline, or shift across repeated observation periods?

This dimension requires repeated measurement. One recorded period cannot establish that visibility is durable, particularly when generated answers and source selection can vary over time.⁷

These dimensions form a **Citation Visibility Profile**.

Dimension	Scentonym observation in this study
Volume	35,310 citations represented in the 181-day overview export
Breadth	1,354 URLs represented in the page export
Concentration	Top 10 query groups held 55.1%; top 10 pages held 14.7% within their respective exports
Relative share	Median query-level Citation Share of 26.2%
Intent	72.9% of query-export citations classified as Commercial or Comparison
Persistence	181 consecutive daily observations; 7-day average ranged from 69.1 to 401.7 citations per day

No universal composite score is necessary. Collapsing the profile into one number would discard the distinctions the data reveal.

8. What the citations establish

The dashboard and exports establish that Scentonym's pages were repeatedly used as visible sources across supported Microsoft AI experiences.

They do not establish:

- recommendation rate;
- brand preference;
- click-through rate;
- customer acquisition;
- sales;
- citation absorption;
- equivalent visibility in Google, ChatGPT, Perplexity, or another product;
- the cause of the observed citation activity.

Research distinguishing citation selection from citation absorption is especially relevant here. A system can cite a page while using little of its language, evidence, or factual content in the generated response.³

The 35,310 figure therefore belongs in the **citation layer** of AEO measurement. It should not be promoted into a claim about recommendation or business impact.

9. Practical implications

Report more than the headline total

Citation volume provides scale. It does not provide breadth, concentration, intent, or relative context.

Preserve the time series

The daily export showed a large difference between the beginning, July peak, and end of the observed period. A dashboard total cannot show whether citation activity is stable, rising, falling, or volatile.

Analyze pages and queries separately

The two views can produce different distribution patterns. A site may have concentrated query activity and broad page participation at the same time.

Preserve absolute and relative measurements

Citation Count and Citation Share behaved differently in this dataset. Reporting only one would conceal part of the result.

Treat platform classifications as analytical aids

Intent and Topic add context, but they remain machine-generated labels. They should not be treated as verified descriptions of individual people.

Measure content architecture without claiming causation

The page export shows that structured entity and relationship pages accounted for most represented citations. That is a useful architectural observation. It is not proof that copying Scentonym's page structure will reproduce its performance.

Microsoft describes grounding systems as selecting supportable information for generated answers rather than simply presenting a ranked page.⁸ Controlled research also suggests that topical relevance and context position can influence citation selection, but those experiments do not establish that any one website architecture will cause durable organic visibility.⁹

Do not equate citation with recommendation

A source can help an AI system answer a question without the associated brand or business becoming the recommended choice.

10. Limitations

This report has substantial limitations.

Single-property case study

The observations come from one fragrance-information website. They should not be assumed to generalize to another property, industry, or level of brand recognition.

Single reporting ecosystem

The source data come from Bing Webmaster Tools and supported Microsoft AI experiences. They do not measure other AI products as equivalent systems.

Sampled and summarized data

Microsoft states that AI Performance is a representative, aggregated view rather than a complete event log. Dashboard, daily overview, page, and grounding-query totals can differ. The interface also states that results may be refined as additional data are processed.

Dated records can change

The August 9 screenshot displayed a rounded 34.5K citations. The later overview export contains 35,310 daily citations through August 10. The report preserves both artifacts and uses the later complete export for exact aggregate calculations. Neither artifact should be treated as an immutable event log.

Unknown inclusion mechanism

The exports do not disclose enough about row inclusion, sampling, suppression, or grouping to support population-level statistical inference.

No direct page-to-query join

The two source exports used here do not provide a complete row-level mapping between every grounding query and every cited page. Their distributions are analyzed separately.

Platform-generated classifications

Intent and Topic are produced by evolving Microsoft classifiers. Some labels in this dataset appear broad or counterintuitive for a fragrance-specific property.

No answer-level content

The data do not include complete generated answers. Citation absorption, brand presence, recommendation, preference, and factual influence cannot be determined from these exports.

No demand data

Grounding-query activity cannot be interpreted as query volume or a count of people asking a question.

No causal design

The study cannot isolate why Scentonym was cited. It cannot establish that structured pages, internal architecture, website age, content depth, SEO, or another factor caused the observed result.

Author conflict of interest

The author operates both the subject property and the organization publishing the report.

11. Data and reproducibility

The Version 1.0 publication package contains:

Supplementary Material S1

Original daily AI Performance overview CSV

The unmodified 181-row export used for aggregate and temporal analysis.

SHA-256:

```
4597ce5cc018ca3273d1b96124b02899c1de28ce2f2203ccba80685387d394d1
```

Supplementary Material S2

Original Bing Webmaster Tools dashboard capture

The six-month AI Performance view reporting approximately 34.5K Total Citations and 59 Average Cited Pages.

SHA-256:

```
1d9d1fb1d84c62bfd3f455b0830f4975a064e1f989df4614c8fdc1c5a8b83963
```

Supplementary Material S3

Original grounding-query CSV

The unmodified 72-row source export used for query-level analysis.

SHA-256:

```
fc78c7cfd86711f3c0a0ac3fa4a5bae4ec2b92ce5642af0d78731835ffeb200f
```

Supplementary Material S4

Original page-level CSV

The unmodified 1,354-row source export used for page-level analysis.

SHA-256:

```
7a0400b9d192cb791d8f94a34d420625c5d84fbfcb1d33937203822b67e7c1fa
```

Supplementary Material S5

Reproducibility script and derived tables

The analysis materials reproduce:

- export totals;
- date continuity and daily descriptive statistics;
- seven-day rolling averages;
- first-versus-final 30-day comparisons;
- daily citation and cited-page correlation;
- top-N concentration;
- intent distribution;
- Citation Share summary statistics;
- Pearson correlation;
- Spearman correlation;
- URL-family totals;
- illustrative concept-family counts.

Original source files remain separate from all derived materials.

12. Future research

This case creates several testable questions.

- Do other information sites show concentrated query activity and distributed page participation?
- Does the balance between structured entity pages and editorial content vary by industry?
- Does page-level breadth predict future citation persistence?
- Does grounding-query concentration reveal fragile dependence on a narrow set of information needs?

Does Citation Share become more or less concentrated as absolute citation volume grows?

Do citation-rich page families also demonstrate measurable citation absorption?

Does citation breadth precede brand presence or recommendation?

Which citation-profile dimensions predict traffic, branded search, leads, or other business outcomes?

A multi-property, repeated-window study would be necessary before making broad claims about content architecture, website age, or citation acquisition.

13. Conclusion

The obvious result in this case is that the daily Bing Webmaster Tools overview export reported **35,310 citations** for Scentonym across 181 consecutive days.

The more useful result is what appears underneath that number.

The ten highest-volume grounding-query groups accounted for **55.1%** of query-export citations.

The seven-day citation average ranged from **69.1** to **401.7 citations per day** within the observed period.

Commercial and Comparison classifications accounted for **72.9%** of represented query activity.

Citation Count and Citation Share were only weakly aligned.

The ten most-cited pages accounted for just **14.7%** of page-export citations.

Structured fragrance and relationship pages accounted for **94.4%** of page-export citation activity.

The property therefore showed a combination of concentrated information needs and distributed structured evidence.

That pattern would disappear inside one visibility score.

AEO requires more than counting visibility. It requires identifying the kind, distribution, and limits of the visibility being measured.

Author disclosure

Jesse Killian is the owner of 580 Digital Infrastructure and the operator of Scentonym, the property analyzed in this report.

This is not independent third-party research. No claim is made that the observed performance was caused by services provided by 580 Digital Infrastructure.

The source documentation and analysis materials are published so readers can inspect the descriptive findings directly.

Suggested citation

Killian, Jesse. “Beyond Citation Count: What 35,310 Reported AI Citations Reveal About Generative Search Visibility.” 580 Digital Infrastructure Research, Research Report 003, Version 1.0, August 2026.

<https://580di.com/research/beyond-citation-count>

License

The report text and original figures are licensed under [Creative Commons Attribution 4.0 International](#). The analysis and figure-generation scripts are available under the MIT License. Microsoft Bing Webmaster Tools source records remain attributed to Microsoft and are governed by the rights statement included with the research package.

References

Official platform documentation

-
- 6 Microsoft Bing. “[AI Performance in Bing Webmaster Tools](#).” Accessed August 2026. Official definitions and limitations for Total Citations, Average Cited Pages, page-level citation activity, grounding queries, Intent, Topic, Citation Share, sampling, and differences among reporting views.
 - 5 Madhavan, Krishna; Meenaz Merchant; Fabrice Canel; and Saral Nigam. “[Introducing AI Performance in Bing Webmaster Tools Public Preview](#).” Microsoft Bing Blogs, February 10, 2026.
 - 8 Madhavan, Krishna; Knut Risvik; and Meenaz Merchant. “[Evolving Role of the Index: From Ranking Pages to Supporting Answers](#).” Microsoft Bing Blogs, May 6, 2026.
-

Published and preprint research

-
- 2 Aggarwal, Pranjal; Vishvak Murahari; Tanmay Rajpurohit; Ashwin Kalyan; Karthik Narasimhan; and Ameet Deshpande. “[GEO: Generative Engine Optimization](#).” Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 2024.
 - 3 Zhang, Kai; He Xinyue; and Yao Jingang. “[From Citation Selection to Citation Absorption: A Measurement Framework for Generative Engine Optimization Across AI Search Platforms](#).” arXiv:2604.25707, April 2026. Preprint.
 - 7 Schulte, Julius; Malte Bleeker; and Philipp Kaufmann. “[Don't Measure Once: Measuring Visibility in AI Search \(GEO\)](#).” arXiv:2604.07585, April 2026. Preprint.
 - 4 Martinez, Olivier. “[Optimizing Visibility in Generative Engines: A Critical Survey of Generative Engine Optimization \(2023-2026\)](#).” arXiv:2607.14035, July 2026. Preprint.
 - 9 Vishwakarma, Rahul; Shushant Kumar; and Ratnesh Jamidar. “[What Gets Cited: Competitive GEO in AI Answer Engines](#).” arXiv:2605.25517, May 2026. Preprint.
-

Prior 580 Digital research

1 Killian, Jesse. "Measuring AEO: A Practical Framework for AI Search Visibility, Citation, and Recommendation." 580 Digital Infrastructure Research, Version 1.0, August 2026.