



CONVENTION · BA-C-4

Version 1.0 · Stable

Minimum disclosure requirements for AI-visibility claims

Claims about visibility in AI assistants circulate widely — in dashboards, case studies, vendor material, and research summaries — and most cannot be evaluated, because they omit the facts a reader needs to judge them.

DOCUMENT	BA-C-4
VERSION	1.0 (Stable)
EFFECTIVE	2026
SUBJECTS	Measurement & statistics; Conventions & policy
LICENSE	CC-BY-4.0
AUTHOR	Barkhausen AI

ABSTRACT

Claims about visibility in AI assistants circulate widely — in dashboards, case studies, vendor material, and research summaries — and most cannot be evaluated, because they omit the facts a reader needs to judge them. This convention states the minimum any published AI-visibility claim must disclose: the entity and information need; the point estimate with a confidence interval and interval method; the sample size; the engine with interface and version; the measurement channel; the time window; the region; the phrasing representation; refusal handling; and the mention-detection method. A claim that omits an applicable item is not a measurement under this convention. Claims that visibility changed must additionally satisfy the Barkhausen Criterion defined in BA-C-3. The requirements are deliberately minimal: they demand no raw data, no query lists, and no tooling disclosure, so that any measurer — commercial or academic — can comply.

CONTENTS

1	1. Scope	3
2	2. The minimum disclosure set	3
3	3. Change claims	5
4	4. Compliance and citation	5
5	5. What this convention deliberately does not require	6
6	Limitations	6

Claims about AI visibility are now routine commercial and editorial currency: an entity “shows up in ChatGPT,” is “recommended by AI assistants,” “gained 40% AI visibility.” Most such claims cannot be evaluated by the people asked to act on them, because the facts needed to judge a claim — how many observations, over what period, on which system, asked how — are not stated. The public evidence is unambiguous that these facts are load-bearing: rewording a question while preserving its intent replaces much of an answer’s recommendation set [1], and re-issuing an identical question a day later replaces a large share of its cited sources and mentioned brands [2]. Against that instability, an undisclosed claim is not weak evidence; it is no evidence.

This convention states the minimum a published AI-visibility claim must disclose to be evaluable at all. It is the shortest document in this series by design. It defines no metrics (BA-C-2 does), prescribes no sampling design (BA-C-3 does), and takes no position on how a measurement is implemented. It states only what a claim must say.

Normative note. The key words **MUST**, **MUST NOT**, **SHOULD**, and **MAY** carry their established normative meanings. A published claim that violates a **MUST** is not a measurement under this convention. Requirements are stated so that compliance can be checked from the published claim alone.

1 1. Scope

This convention applies to any published claim that a named entity is, is not, became, or ceased to be visible in an AI assistant or answer engine — whether the claim appears in a research report, a dashboard, marketing or sales material, a case study, or press coverage. It applies to point-in-time claims (“mentioned in 62% of answers”) and to change claims (“visibility doubled after the campaign”) alike. It does not apply to statements framed as anecdote and presented as nothing more (“here is one answer we received”), provided no quantitative or comparative conclusion is drawn from them.

2 2. The minimum disclosure set

A published AI-visibility claim **MUST** state every item below that applies to it, either inline or in an unambiguously linked note:

#	DISCLOSURE	WHAT IT FIXES
1	Entity and information need	Who the claim is about, and for which kind of question
2	Point estimate with confidence interval and interval method	The value claimed, and its precision
3	Sample size n , per cell	How many independent observations support it
4	Engine, interface, and version or observation date	Which system, in which deployment, when
5	Measurement channel	The access path the observations came through — consumer interface, official API, or other (BA-C-3)
6	Time window	The period the claim describes
7	Region	The observation locale
8	Phrasing representation	Number of distinct phrasings and how wording variation was covered
9	Refusal handling	Whether refusals were recorded, and the refusal rate where material
10	Detection method	How a “mention” was identified in answer text

A claim that omits an applicable item is not a measurement under this convention; it is an anecdote formatted as a statistic. The single most common omission is item 3: a percentage without its sample size cannot be distinguished from a coin flip reported to two significant figures. The second most common is item 6: engines change without notice, so a claim without a window describes no defined quantity.

A conforming claim fits in one sentence, mapped part by part to the ten disclosures in Figure 1. For example:

Entity X was mentioned in 62% of answers (95% CI 53–70%, Wilson; $n = 120$) for the information need “recommended providers of Y,” sampled across 14 phrasings on Engine Z (consumer interface, model build of 2026-06), 2026-06-01 → 2026-06-28, region DE; refusal rate 3%; mentions detected by lexicon matching (precision/recall 0.97/0.94 on a labeled sample).

A conforming AI-visibility claim

BA-C-4 §2

Entity X¹ was mentioned in 62% of answers (95% CI 53–70%, Wilson²; n = 120³)
 for the information need “recommended providers of Y,”⁴ sampled across
 14 phrasings⁸ on Engine Z⁴ (consumer interface,⁵ model build of 2026-06),⁴
 2026-06-01 → 2026-06-28,⁶ region DE,⁷ refusal rate 3%,⁹ mentions detected
 by lexicon matching (precision/recall 0.97/0.94 on a labeled sample).¹⁰

The ten disclosures, one per numbered part

- | | |
|---------------------------------------|---------------------------|
| 1 Entity and information need | 6 Time window |
| 2 Point estimate, CI, and method | 7 Region |
| 3 Sample size (n) | 8 Phrasing representation |
| 4 Engine, interface, and version/date | 9 Refusal handling |
| 5 Measurement channel | 10 Detection method |

Figure 1. The conforming one-sentence claim above, rendered as a claim card. Superscript ticks map each part of the sentence to the ten minimum-disclosure items: entity and information need (1), point estimate with interval and method (2), sample size (3), engine, interface, and version or date (4), measurement channel (5), time window (6), region (7), phrasing representation (8), refusal handling (9), and detection method (10). A claim that omits an applicable item is not a measurement under this convention.

Barkhausen AI · BA-C-4

The definitions of the quantities named above — Visibility Probability, Share of Voice, Discovery Depth, and their supporting measures — are those of BA-C-2, and metric-specific disclosures required there (competitor set and weighting for SoV, constraint dimensions for DD, classifier accuracy for sentiment) apply in addition to this list.

3. Change claims

A claim that visibility *changed* — improved, declined, jumped after an intervention — asserts more than two point-in-time claims, because engines drift on their own at every time scale [1] [2]. A published change claim **MUST** satisfy the **Barkhausen Criterion** as formalized in BA-C-3: significance established between proper interval estimates; sustainment across at least $K \geq 2$ consecutive windows, with K and the window length disclosed; non-coincidence with a flagged engine change, or re-basing after it; and disclosure of how many cells were monitored, with the multiplicity treatment stated. A before-and-after pair of screenshots, or a comparison of two undated percentages, does not meet this requirement under any circumstances.

4. Compliance and citation

A claim, report, or dashboard **MAY** state compliance as: *disclosed in conformance with BA-C-4 v1.0*. Compliance refers to disclosure only; it is not an endorsement of the underlying design, whose validity is governed by BA-C-3.

Reviewers, journalists, and procurement evaluators **MAY** cite non-compliance neutrally: *the claim does not meet the minimum disclosure requirements of BA-C-4*. This phrasing makes no accusation about the claim’s

truth — it records that the claim, as published, cannot be evaluated. That is the purpose of a disclosure floor: it separates claims that can be checked from claims that must be taken on faith, before any argument about who is right.

5 5. What this convention deliberately does not require

The requirements above are a floor, chosen so that any competent measurer can comply without surrendering legitimate confidential material. This convention does NOT require publication of raw answer data, of the specific query phrasings used, of prompt-construction or variant-engineering methods, or of any tooling. It requires *counts, definitions, and conditions* — the statistical design layer a reader needs to judge whether a conclusion follows — and nothing from the operational layer beneath it. A measurer who cannot state items 1–10 about a claim does not have a disclosure problem; they have a measurement problem.

6 Limitations

Minimum disclosure is necessary for evaluation, not sufficient for validity. A claim can disclose all ten items and still rest on a defective design — too few phrasings, uncontrolled personalization, cherry-picked windows; validity is the province of BA-C-3, and a reader of a fully disclosed claim still has to judge the design the disclosure reveals.

Disclosure can also be false. This convention makes fabricated disclosures checkable in principle — a stated n , window, and engine version are concrete claims a challenger can attempt to reproduce — but it cannot prevent them; no disclosure convention can.

Finally, the disclosure set reflects the engines as they exist in the stated windows of the evidence cited here. These results describe the engines as sampled during their respective windows; engines change without notice, and results should be assumed perishable. If future engine architectures make new conditions load-bearing — new personalization surfaces, new answer formats — the set will grow by a new version of this convention, not by silent reinterpretation.

REFERENCES

1. Jack, Lehman, Maloney, Xu; arXiv:2605.27440. *Paraphrase Brittleness in Production Retrieval-Augmented Commercial Recommendation: Reproducibility Below the Rerun-Stability Baseline* (2026). <https://arxiv.org/abs/2605.27440> Accessed 2026-07-09. [archived]
2. Schulte, Bleeker, Kaufmann; arXiv:2604.07585. *Don't Measure Once: Measuring Visibility in AI Search (GEO)* (2026). <https://arxiv.org/abs/2604.07585> Accessed 2026-07-09. [archived]

HOW TO CITE

Barkhausen AI (2026). Minimum disclosure requirements for AI-visibility claims.
<https://barkhausen.ai/conventions/minimum-disclosure/>

Published under the Creative Commons Attribution 4.0 International (CC-BY-4.0).

