

# Demand-Side Content Audit

Eight dimensions across ownership, citability, and distribution. The audit that separates content calendars from compounding engines.

## WHEN TO USE

Pick your five highest-trafficked or most-cited pieces and score the portfolio across all 8 dimensions. Use the shared 1-5 legend. Total the score. Read the band recommendation. The goal is not to fix the worst piece — it is to identify which dimension is the systemic leak. Fix one dimension this quarter; do not thrash across all eight.

**Preamble.** The Ehrenberg-Bass / LinkedIn B2B Institute 95/5 rule: 95% of a category is out-of-market in any given quarter. Rented attention (paid) does nothing for the 95%. Owned and earned attention accumulate in the 95% so that when they enter the 5% window, you surface first. This audit scores whether your content portfolio is building the 95% asset or running in place.

## Source lineage:

- LinkedIn B2B Institute / Ehrenberg-Bass — 95/5 (dimensions 3, 5, 7)
- Ahrefs 2026 AI Overview study (75,000 brands, 25M responses) — branded mentions across the citation pool predict AI visibility more than backlinks (dimensions 2, 4)
- Christensen/JTBD — content serves pain, not personas (dimension 5)
- Schwartz awareness spectrum — query classification maps to awareness level (dimension 3)

**Scoring legend (1-5).** Every dimension uses the same 1-5 scale:

- 1 — Absent:** You do not do this. Either you have not built it, you built it and retired it, or the discipline does not exist in the org.
- 2 — Ad-hoc:** Someone does this, sometimes, when they remember. No cadence, no owner, no artifact. Cannot be audited.
- 3 — Documented:** An owner exists. A cadence exists. An artifact exists. Not yet measured; not yet feeding other layers.
- 4 — Measured:** Outputs are tracked against a baseline. Reviewed at least quarterly. Decisions are made from the numbers, not vibes.
- 5 — Compounding:** Outputs feed back into upstream layers. Every cycle makes the next cycle sharper. The discipline survives the person who built it.

*Total possible: 40. Your band determines where to look first.*

Score range	Band	Where to look first
≤ 16	Foundation missing	Your content is indistinguishable from category filler. No named entities, no PSP mapping, no citation pool. The fix is not more content — it is inventing 2-3 named entities and rebuilding the next 10 pieces around them. Pause the content calendar until you have entities.
17 – 26	Leaking between layers	You have entities and PSPs but the distribution and refresh layers are missing. Pick the distribution layer (dimension 6) for the next quarter. One piece per week with a 6-surface distribution plan outperforms 4 pieces per week that get published and forgotten.
27 – 34	Working, not compounding	Content is earning attention but not yet compounding. The gap is usually the citation pool (dimension 4) and backlink anchoring (dimension 8). 90 days of intentional external placement — 2-3 podcast appearances, 5-8 guest posts — measurably shifts AI Overview citation share.

35 – 40	Compounding	Your portfolio is a demand asset. The risk now is refresh cadence (dimension 7) breaking as the portfolio grows. Put the rolling-refresh ritual on the calendar — content portfolios collapse when they stop getting tended, and the collapse shows up 6-12 months later when everyone has forgotten.
------------	-------------	---

---

 DECISION CRITERIA
 

---

### 1. Topic entity mapping

*Does each flagship piece own a named entity — a concept, framework, or term that only your work defines with precision? Entity ownership is what AI engines retrieve on; commoditized content gets paraphrased, proprietary entities get cited by name.*

**Score 1 (low):** Every piece covers broad topics already owned by HubSpot, Gartner, or 50 other domains. No named entities unique to your work.

**Score 5 (high):** Every flagship piece owns 1-3 named entities (framework, term, diagnostic) that the piece defines. Entities are used consistently across the portfolio.

**Recommendation:** List the named entities in your top 5 pieces. If fewer than 3 total, you are writing category content, not owned content. Invent and define 2-3 entities this quarter; use them consistently across the next 10 pieces.

### 2. Passage citability

*Can an LLM quote a 30-word passage cleanly, with attribution that still makes sense outside the article? Ahrefs' 2026 study found branded mentions across the citation pool correlate with AI visibility more strongly than backlinks or DR. Citability is a craft discipline at the paragraph level.*

**Score 1 (low):** Passages are vague, hedged, or too long to quote. Claims are not attributable to a source. LLMs paraphrase instead of citing.

**Score 5 (high):** Every flagship piece has 3-5 citable passages (25-40 words each) with specific claims, named sources, and enough context to stand alone. Pieces are structured so LLMs have to cite the passage to preserve meaning.

**Recommendation:** Audit your top 5 pieces for citable passages. If any piece has zero, rewrite the opening 500 words with one specific claim per paragraph (numbers, names, dates). Short sentences. Named sources. This is the single-highest-lift content edit.

### 3. Query classification coverage

*Schwartz awareness spectrum mapped to content: informational (unaware/problem-aware), commercial-investigation (solution-aware), transactional (product-aware/most-aware). A portfolio skewed to one end cannot serve the 95/5 rule.*

**Score 1 (low):** Portfolio is 80%+ one query type. Usually: all informational (ranking play) OR all transactional (product pages).

**Score 5 (high):** Balanced coverage: ~50% informational, ~30% commercial, ~20% transactional. Intentional ratios per PSP. Awareness-level progression mapped across pieces.

**Recommendation:** Tag every flagship piece with its query type. If the portfolio is skewed, you are either missing the 95% (too transactional) or never closing (too informational). Rebalance the next 10 pieces.

### 4. Citation pool composition

*Earned mentions outside your own domain — Reddit, Quora, vertical newsletters, YouTube transcripts, podcast appearances, specific forums. Ahrefs' 2026 study isolated this as the top correlate of AI Overview visibility. Backlinks matter less than scattered, specific mentions.*

**Score 1 (low):** Zero or minimal branded mentions outside owned properties. No tracking of where the brand is referenced.

**Score 5 (high):** 30+ earned mentions across 8+ surface types (Reddit, Quora, newsletters, podcasts, YouTube, forums). Monthly citation-pool monitoring. Named entities referenced in external discussion.

**Recommendation:** Install branded-mention monitoring this week (Mention, Google Alerts, Reddit/HN RSS). Target 2-3 earned mentions per month in the next 90 days. Podcast appearances and vertical-newsletter guest posts are the two highest-leverage surfaces.

### 5. Content-to-PSP mapping

**Score 1 (low):** Content planned by keyword volume or topical relevance. PSPs not referenced in briefs.

**Score 5 (high):** Every flagship brief references one PSP and one awareness level. Pieces retired when their PSP is retired. Content calendar and PSP review share the same quarterly ritual.

**Recommendation:** Tag every flagship piece with its PSP. Any piece that cannot be tagged is a candidate for retirement or redirect. Future content briefs start with PSP + awareness level as required fields.

## 6. Distribution layer

*How intentionally is each piece distributed across LinkedIn / podcast / newsletter / guest surfaces? HubSpot-era research put the distribution-to-creation ratio at ~89% distribution; most teams invert this and over-invest in production.*

**Score 1 (low):** Publish and move on. Distribution = sharing the link once on LinkedIn the day of publish.

**Score 5 (high):** Every flagship piece has a 6-8 surface distribution plan: atomized LinkedIn posts, newsletter send, podcast angle, Reddit/Quora comment, guest post pitch. Distribution runs for 6+ weeks post-publish.

**Recommendation:** Write a distribution template with 6 mandatory surfaces. Require every new piece to ship with the distribution plan attached before publish. Retroactively distribute 1-2 top pieces per quarter.

## 7. Refresh cadence

*How old is each flagship piece? Content decay starts at 12 months for most topics; beyond 18 months the piece is a liability (stale stats, retired frameworks, dead links). The 95/5 rule demands the long-lived pieces are kept evergreen.*

**Score 1 (low):** Most flagship pieces are 2+ years old with no refresh. Stale data, dead links, retired framework references.

**Score 5 (high):** Every flagship piece is reviewed every 12 months minimum. Quarterly refresh cycle runs. Each refresh re-validates stats, refreshes examples, updates frameworks. Published date is visible.

**Recommendation:** Audit every flagship piece for age. Any piece >18 months with no refresh is a candidate for either refresh or retirement. Schedule a rolling 90-day refresh ritual — one piece per week, never batched.

## 8. Backlink profile quality

*Anchored to named frameworks or generic? "Click here" backlinks are low-signal; anchor text referencing your named entities compounds. Ahrefs' 2026 data found anchor relevance was a stronger citation driver than raw backlink count.*

**Score 1 (low):** Backlinks are "click here," company name only, or random anchor text. No named entity anchoring.

**Score 5 (high):** Backlinks increasingly anchor to named entities from your portfolio. Outreach targets specific anchor strategies. Brand-term backlinks tracked separately from entity-term backlinks.

**Recommendation:** Segment backlink reports by anchor type. If <20% anchor to named entities, your content is citable but your entities are not spreading. Prioritize guest posts and interview appearances where you can name-check your frameworks.