

The Content *Authority Audit*

A working diagnostic rubric for evaluating how much authority any piece of content actually establishes – not whether it's technically correct, adequately long, or keyword-optimized, but whether it earns the trust of a reader and the citation of an AI system. Score a single piece, a sample of ten, or your entire library. The score tells you where the real problem is.

8

DIMENSIONS

32

MAX POINTS

4

SCORE BANDS

0–4

POINTS PER DIMENSION

How to use this audit

FOR A SINGLE PIECE

- Score each of the 8 dimensions 0–4
- Total the scores (max 32)
- Read the score band to understand overall authority level
- The lowest individual dimension score tells you where to start
- Use the action items in each dimension to prioritize fixes

FOR A CONTENT LIBRARY

- Sample 10–15 pieces across topic clusters and content types
- Score each piece on all 8 dimensions
- Average dimension scores across the sample
- Dimensions with average below 2 are systemic problems – not piece-level fixes
- Systemic problems require a strategy change, not a revision pass

SCORE BANDS — OVERALL CONTENT AUTHORITY

0–12

CRITICAL

Content is producing active harm: eroding brand credibility, occupying URLs that should be doing work, and signaling to search systems that the site is not a reliable source. Immediate triage required.

13–19

WEAK

Content is technically acceptable but editorially invisible. It ranks for nothing competitive, earns no citations, and gives readers no reason to trust the source over anyone else's median output.

20–25

SOLID

Content is earning its place. Specific gaps exist but the foundation is sound. Targeted improvements in the weakest dimensions will produce measurable gains in authority and visibility.

26–32

STRONG

Content is operating at citation-worthy level. It gives readers something they can't get from AI-generated summaries and gives search systems a specific, credible source to attribute. Protect and extend this.

DIMENSION

1 Information Gain

4
MAX POINTS

Does this piece add something that doesn't already exist in the information landscape for this query? The question isn't whether the content is accurate – it's whether it's worth reading given that accurate content on this topic already exists everywhere. AI systems synthesizing answers default to the most information-dense, specific, credible source. Content with no information gain is what the AI generates on demand; there's no reason to click through for it.

SCORE

CRITERIA

0

Zero gain – fully substitutable

Everything in this piece exists in nearly identical form elsewhere. No unique data, no first-person account of doing the thing, no specific claim that couldn't be generated by AI from the top five ranking pages. Reading this piece adds nothing to the reader's understanding beyond what they'd get from a Wikipedia summary or an AI Overview.

Signals: entirely hedged claims, generic examples, no named sources, no specific data, no position the author couldn't reverse without consequence.

1

Marginal gain – one original element

There is one moment in the piece that adds something – a specific statistic, a named example, a personal observation – but it is surrounded by content that adds nothing. The gain is real but would not make someone cite this piece specifically.

Signals: one section with original detail, rest of piece is competent but generic.

2

Meaningful gain – distinctive angle or data

The piece has a clear angle that differs from how competitors treat the topic, or includes specific data, examples, or case material that doesn't appear elsewhere. A reader who has already read several pieces on this topic would find something new here.

Signals: cited original research, counter-intuitive framing, specific named client/case material, data not repeated from primary sources.

SCORE	CRITERIA
3	Strong gain – piece fills a real gap The piece addresses a question or angle that is not adequately covered by existing content in this space. A reader who has read everything else on this topic would find this piece specifically valuable. The information gain is the reason to cite it. <i>Signals: original research or primary data, first-hand experience narrated specifically, contrarian position with specific evidence, synthesis connecting things not previously connected.</i>
4	Definitive gain – this is the piece to cite The piece is the best available treatment of this topic from this author's position. It is cited by others, quoted in industry conversation, and cannot be substituted by an AI-generated summary without losing something material. The information gain is the reason it exists. <i>Signals: piece receives inbound links, is referenced in industry discussions, contains specific proprietary insight, has been shared as a resource by credible practitioners.</i>
IF SCORE IS 0-1, FIX BY ↑ Identify what this piece knows that other pieces don't – if the answer is nothing, consider whether it should exist ↑ Add one first-person observation that is specific and verifiable ↑ Find the consensus position on this topic and take a specific position on whether it's right ↑ Include one example your competitors have never used – something specific to a real situation IF SCORE IS 3-4, PROTECT BY ✓ Keep the original data current – outdated statistics erode the gain ✓ Add FAQPage schema so the information gain is machine-extractable ✓ Mention the piece in outreach and distribution as the citation source ✓ Cross-link from related posts so it accumulates authority	

DIMENSION

2 Author Authority & Attribution

4
MAX POINTS

Is there a specific, credentialed, verifiable person behind this content – and is that person's expertise legible to a reader who doesn't already know them? Google's E-E-A-T framework rewards content from identifiable authors with documented experience. AI systems citing sources prefer content with clear, verifiable authorship over anonymous or corporate-attributed content. Author authority is not about fame – it's about specificity. "Jacob Clifton, who spent fourteen years as a staff writer at Television Without Pity and has written 40+ books" is more citable than "The Clifton Creative Team."

SCORE

CRITERIA

0

No attribution

Content is attributed to a brand or organization, or has no attribution at all. There is no person behind it. No credentials, no biography, no external verification possible.

Signals: "Published by [Brand]," no author field, generic corporate voice, no author bio.

1

Named but uncredentialed

A person's name appears on the content, but there is no information about who that person is, why they are qualified to write this, or where to find more from them. The name is present; the authority is not.

Signals: author name only, no bio, no links, no verifiable credentials.

2

Named with basic credentials

The author is named and a brief bio exists that establishes some relevant context – role, organization, general area of expertise. A reader could verify that this person exists and has some relevant background.

Signals: author bio with role and organization, link to author page or LinkedIn, general credential statement.

3

Named with specific, verifiable credentials

The author is named with credentials specific to this topic – not just "content strategist" but "content strategist who built three verticals to 1M

SCORE	CRITERIA
	readers." The credentials are specific enough to be checked and are directly relevant to the content. <i>Signals: specific named accomplishments in bio, external mentions or citations, dedicated author page on site with publication history.</i>
4	Author as named entity with cross-platform authority The author is a recognized entity across platforms – their name appears in their bylines, their LinkedIn, their About page, their author schema, their external mentions, and their bio consistently matches. A search for their name returns verifiable results relevant to this topic. Author schema (Person type with @id, name, url, sameAs) is implemented on the page. <i>Signals: Person schema with sameAs links, author page with consistent credentials, external citations by name, LinkedIn profile corroborates bio claims.</i>
IF SCORE IS 0-1, FIX BY ↑ Add a named author to every piece – no more "The Team" ↑ Write a specific bio that names credentials relevant to the topic, not general ones ↑ Create a dedicated author page and link to it from each byline ↑ Add Person schema to author pages with sameAs links to LinkedIn IF SCORE IS 3-4, EXTEND BY ✓ Add Article schema with author @id pointing to the Person schema ✓ Build external mentions – guest posts, interviews, citations – that reinforce the author entity ✓ Ensure bio language is consistent across site, LinkedIn, and social profiles	

DIMENSION

3 Specificity & Evidence Quality

4
MAX POINTS

Are the claims in this piece specific, attributable, and verifiable – or are they the kind of thing anyone could say? Specificity is the primary signal that distinguishes first-hand expertise from synthesized summary. "Studies show that content with images gets 94% more views" is a floating statistic. "According to the 2024 HubSpot State of Marketing Report, blog posts with at least one image receive 94% more total views than text-only posts" is a specific, attributable claim. Only the second one earns trust. AI can generate the first type endlessly; it cannot generate the second type accurately.

SCORE

CRITERIA

0

No specific evidence – assertion only

Claims are made without support. No data, no examples, no citations, no named cases. The content asserts things confidently and provides no basis for the reader to evaluate them. Everything could have been written without knowing anything specific about the subject.

1

Vague evidence – floating statistics or unnamed examples

Some evidence is present but it is unattributed ("studies show," "research indicates," "many companies find") or the examples are hypothetical and generic ("imagine a company that..."). The evidence is present in form but provides no actual information gain.

2

Attributed evidence – cited sources, named examples

Claims are supported by named sources with enough information to verify. Examples use real companies or real situations. Statistics are attributed to specific reports or studies. A skeptical reader could check the claims.

3

First-person evidence – specific experience narrated

The evidence includes first-hand accounts of doing the thing: a real client situation described specifically (not hypothetically), a specific outcome from a specific decision, a specific

SCORE	CRITERIA
	failure and what it cost. The author is present in the evidence, not just citing others.
4	Primary evidence – original data or research The piece contains data that exists only here – a survey the author conducted, a test they ran, a case study with real numbers. Readers cite this piece because it's the primary source, not because it synthesizes other sources well.
IF SCORE IS 0-1, FIX BY ↑ Find the actual source for every statistic in the piece – if you can't, cut it ↑ Replace one hypothetical example with a real, named one ↑ Add one sentence about what you have personally observed on this topic ↑ Audit for "studies show" – every instance must be replaced or cut IF SCORE IS 3-4, PROTECT BY ✓ Update statistics annually – outdated primary data is worse than no data ✓ Build more first-person examples from client work into the content calendar ✓ Mark primary data clearly in the piece – "from our analysis of X" – so it's citable	

DIMENSION

4 Structural Clarity & Argument

4
MAX POINTS

Does the piece make an argument – a sequence of moves that changes what the reader thinks or can do – or does it list topics? Most content fails this dimension not because the writing is bad but because the structure is a topic list rather than an argument. "Five benefits of X, three mistakes to avoid, how to get started" is a structure. "Why the standard approach to X fails at the moment that matters most, and what to do instead" is an argument. Only arguments earn authority. Topic lists earn pageviews, maybe.

SCORE

CRITERIA

0

No structure – wall of text or random assembly

Content has no discernible organizational logic. Sections don't build on each other. The reader could read in any order without losing meaning. There is no argument because there is no sequence.

1

List structure – topics enumerated, not argued

The piece is organized as a list of topics or tips. Each section is standalone. Nothing builds toward a conclusion. A reader who reads only the headers understands the content as well as one who reads the full piece. The structure conveys no argument.

2

Logical structure – sections build, conclusion follows

Sections lead from each other in a sequence that makes sense. The conclusion follows from the earlier content rather than merely restating it. A reader who reads in order understands more than one who reads selectively.

3

Argument structure – piece takes a position and defends it

The piece has a specific thesis – something that could be argued against – and builds toward it. The opening establishes a tension or problem, the body makes the case, the close resolves to a position. A reader who finishes the piece has been moved to believe something or do something they weren't doing before.

SCORE	CRITERIA
4	Definitive argument – the piece is the best case for this position The piece takes a defensible position and makes it comprehensively. It anticipates and addresses the strongest counter-arguments. It earns its conclusion rather than asserting it. Readers who disagree with the position still find it valuable because the argument is so clearly made.
IF SCORE IS 0-1, FIX BY ↑ Identify the single sentence this piece is arguing – if you can't state it, the piece doesn't have one ↑ Rewrite the opening to establish a tension, not introduce a topic ↑ Ensure the close resolves to a position, not a summary ↑ Check that section sequence is driven by argument, not by "what comes next"	IF SCORE IS 3-4, EXTEND BY ✓ Add FAQPage schema – well-argued pieces answer clear questions ✓ Extract the thesis as the meta description and title tag ✓ Create a pull-quote social card with the core argument for distribution

DIMENSION

5 Brand Voice Distinctiveness

4
MAX POINTS

Would a reader who knows your brand recognize this piece without a byline? Voice is not tone. Tone adjusts by context; voice is consistent across all contexts. A brand with a strong voice makes the same argument the same way whether it's in a case study, a blog post, or an email subject line. A brand without a voice sounds like whatever was convenient to write. AI-generated content defaults to the median internet voice – competent, even, neutral, unrecognizable. If your content could have been written by anyone, it was effectively written by no one.

SCORE

CRITERIA

0

No voice – interchangeable with any competitor

Remove the logo and the byline. There is no signal left that identifies the source. The vocabulary, the sentence rhythms, the default examples, the way opinions are framed – all are generic. This piece could appear on any brand's blog without editing.

1

Trace voice – one or two moments of distinctiveness

There are one or two sentences in the piece that sound like a specific person or brand. The rest is generic. The voice is present as an exception, not as the default.

2

Emerging voice – consistent in sections, inconsistent overall

The voice is recognizable in some sections – particularly the opener or the close – but lapses into generic in the body. A reader who knows the brand would recognize it intermittently. The voice exists; it hasn't been applied with discipline.

3

Consistent voice – brand is recognizable throughout

The piece sounds like this brand, distinctly, in nearly every paragraph. The vocabulary is specific, the sentence rhythms are characteristic, the way the brand positions its argument is recognizable. A reader who knows the brand would identify it without the byline.

4

Definitive voice – the piece could only have been written by this brand

SCORE	CRITERIA
	The voice is so specifically the brand's that it would be immediately identifiable not just to existing readers but to a newcomer who had just read two or three other pieces. The content and the voice are inseparable – the argument is made in a way that only this brand could make it.
IF SCORE IS 0-1, FIX BY ↑ Read the piece aloud – every sentence that sounds "written" is a candidate for rewrite ↑ Identify the two or three things this brand always says and find where they belong in this piece ↑ Replace one generic example with one from your actual experience ↑ Apply a brand voice document if you have one – if you don't, that's a separate problem	IF SCORE IS 3-4, PROTECT BY ✓ Use this piece as a calibration example for your brand voice document ✓ Build an annotated example library from high-voice pieces for writer onboarding ✓ Require new writers to calibrate to this piece before publishing under the brand

DIMENSION

6 AEO / GEO Readiness

4
MAX POINTS

Is this piece structured to be cited and quoted by AI systems – ChatGPT, Perplexity, Claude, Google's AI Overviews? The AI answer layer now sits above most organic results for informational queries. Content that earns positions in that layer needs to be: answer-first (the answer appears before the explanation), explicitly structured as Q&A where applicable, attributed to a verifiable author entity, and marked up with schema that declares what the content is. Content that doesn't meet these criteria will be paraphrased (anonymously) or ignored.

SCORE

CRITERIA

0

Not AEO-ready – buries the answer, no schema, no Q&A

The main answer to the query this page targets appears in paragraph three or later. There is no FAQ structure. There is no schema of any kind. An AI system cannot efficiently extract the answer from this page without reading the whole thing and guessing.

1

Partially ready – answer accessible but not structured

The answer appears relatively early but is embedded in prose rather than structured explicitly. No FAQ blocks. No schema. A diligent AI can extract the answer; it won't be cited specifically.

2

Structured – answer-first, FAQ blocks present

The page opens with the direct answer to the query. FAQ questions and explicit answers appear in the content body. The page is optimized for human readers to scan and for AI to extract. Schema is absent or partial.

SCORE	CRITERIA
3	Schema-complete – FAQPage and Article schema implemented FAQPage schema is valid and matches the visible FAQ content. Article schema is present with author @id, datePublished, and dateModified. The page makes explicit machine-readable declarations that allow AI systems to cite it accurately.
4	Citation-ready – full schema, verified author entity, answer-first throughout Full schema stack: Article, WebPage, BreadcrumbList, FAQPage. Author @id linked to a Person entity with sameAs. Every major claim in the piece can be extracted as a quotable, attributable statement. The page has been validated in Google's Rich Results Test. AI systems citing this page will attribute it accurately to a specific, credentialed author.
IF SCORE IS 0-1, FIX BY ↑ Move the direct answer to the query into the first paragraph – before any context or framing ↑ Identify three questions the piece already answers and write explicit Q&A blocks ↑ Add FAQPage schema – this is the single highest-ROI schema implementation ↑ Add Article schema with author @id IF SCORE IS 3-4, PROTECT BY ✓ Validate schema monthly in Rich Results Test – updates can break it ✓ Monitor AI Overview appearances for this query in GSC ✓ Update FAQ answers when the underlying facts change	

DIMENSION 7 Internal Architecture & Linking

4
MAX POINTS

Is this piece connected to the rest of the content library, or is it an island? A page with no internal links receives no link equity from the rest of the site and sends no signals about its relationship to related content. A page that is part of a deliberate topic cluster – with links to and from its pillar, and contextual links to related spokes – tells search engines that the site has coherent, organized expertise on this topic. Topical authority is built at the cluster level, not the page level. This dimension evaluates how well this piece is wired into that structure.

SCORE

CRITERIA

0

Orphaned – no internal links in or out

This page is not linked from any other page on the site, and does not link to any other page on the site. It is invisible to crawlers traversing the site and receives no link equity. It exists as a standalone document, not as part of a content library.

1

Minimally linked – navigational links only

The page appears in site navigation or a category archive but has no contextual internal links – no links within the prose to related pages, and no other pages linking to it contextually. It is accessible but not integrated.

2

Partially linked – some contextual links but not cluster-integrated

The piece links contextually to some related pages and receives some contextual links from others. But it is not clearly associated with a topic cluster – there is no pillar page it points to and no cluster structure it participates in.

3

Cluster-integrated – links to pillar and receives spoke links

The piece is a spoke in a deliberate topic cluster. It links back to its pillar page with a cluster attribution line. The pillar page links to this piece. Contextual internal links to related spokes

SCORE	CRITERIA
	are present where relevant. The piece is part of an architecture.
4	Fully integrated – bidirectional cluster links, contextual linking throughout The piece is fully wired into the content architecture: cluster attribution to pillar, contextual links to related spokes, inbound links from the pillar and from related pieces. Internal anchor text is specific and descriptive rather than generic ("click here"). The page is a node in an intentional network, not a leaf on a branch.
IF SCORE IS 0-1, FIX BY ↑ Identify which topic cluster this piece belongs to – if none, that's the root problem ↑ Add a cluster attribution line: "This post is part of the [Pillar Topic] guide" with a link to the pillar ↑ Add the pillar link to this piece from the pillar page itself ↑ Find two related spoke pages and add contextual links in each direction IF SCORE IS 3-4, EXTEND BY ✓ Audit anchor text – replace "read more" and "click here" with descriptive text ✓ Add new spoke links as cluster grows ✓ Check for broken internal links quarterly	

DIMENSION

8 Reader Utility & Actionability

4
MAX POINTS

Does a reader who finishes this piece know something they didn't know, or can do something they couldn't do? "Helpful content" is a phrase Google uses to describe content that serves people – but it describes something specific: content that leaves readers better equipped than when they arrived. Thought leadership that makes no practical claims, explainers that cover the topic without enabling the reader to do anything with the information, and inspiration pieces that produce feeling without producing capability all score low on this dimension. Utility is not about adding a bullet list of "takeaways" – it is about whether the piece was designed to change what the reader can do.

SCORE

CRITERIA

0

No utility – reader finishes having gained nothing actionable

The piece may be accurate and well-written but leaves the reader exactly where they were. No decision has been made easier, no skill has been acquired, no misconception has been corrected. The content exists to be read, not to be used.

1

Nominal utility – concepts explained, no application

The piece explains what something is without enabling the reader to do anything with that understanding. Awareness is raised; capability is not. The reader now knows the term; they still don't know how to apply it.

2

Functional utility – reader understands how to approach the problem

The piece gives a reader enough to begin. It explains the approach, identifies the variables, and describes what good looks like. A motivated reader could start with this, though they'd need more.

3

High utility – reader can do something specific after reading

The piece is designed around an outcome: the reader who finishes it can make a specific decision, execute a specific step, or evaluate a specific

SCORE	CRITERIA
	situation differently. The utility is the purpose of the piece, not a by-product.
4	Reference utility – piece becomes a resource readers return to The piece is so useful – as a diagnostic tool, a framework, a checklist, a reference – that readers bookmark it, share it as a resource, and return to it. It is not consumed once and forgotten. Its value increases with use.
IF SCORE IS 0-1, FIX BY ↑ State explicitly what the reader can do, decide, or evaluate after reading this piece ↑ Add one concrete step-by-step instruction to what is currently an explanation ↑ Add a "what to do if you scored X" or "what to do next" section ↑ Build a companion tool – checklist, rubric, template – that extends the utility	IF SCORE IS 3-4, EXTEND BY ✓ Create a downloadable version of any framework, checklist, or rubric in the piece ✓ Add a HowTo schema block for any step-by-step section ✓ Track return visits – high-utility pages show repeat users in GA4

SCORE SHEET – FILL IN YOUR TOTALS

Dimension 1 – Information Gain / 4

Dimension 2 – Author Authority & Attribution / 4

Dimension 3 – Specificity & Evidence Quality / 4

Dimension 4 – Structural Clarity & Argument / 4

Dimension 5 – Brand Voice Distinctiveness / 4

Dimension 6 – AEO / GEO Readiness / 4

Dimension 7 – Internal Architecture & Linking / 4

Dimension 8 – Reader Utility & Actionability / 4

Total Content Authority Score / 32

Interpretation Guide — What the Score Means

The single lowest dimension is your first priority — always

Overall scores are useful for tracking direction over time. The dimension-level scores tell you where to work. A piece that scores 28/32 overall but 0/4 on Author Authority has a fundamental problem that no amount of improvement in the other dimensions can fix. AI systems cannot cite an anonymous piece, regardless of how well-structured or evidence-rich it is.

When scoring a library sample, average each dimension across all pieces scored. Any dimension averaging below 2.0 is a systemic problem. Those dimensions require policy changes, not piece-by-piece revisions.

What systemic dimension failures mean

- Dimension 1 avg < 2.0 – The content strategy has no information gain requirement. The library is producing substitutable content by design.
- Dimension 2 avg < 2.0 – Authors are not being credited specifically. This is a brand policy issue, not a writer issue.
- Dimension 3 avg < 2.0 – The brief process does not require evidence. Writers are producing assertion without support because no one has told them not to.
- Dimension 4 avg < 2.0 – Briefs specify topics, not arguments. The editorial direction is missing.
- Dimension 5 avg < 2.0 – There is no brand voice document, or the one that exists doesn't constrain editorial decisions. Voice training has not been provided to writers.
- Dimension 6 avg < 2.0 – Schema markup is not part of the publishing workflow. AEO is not part of the content strategy.
- Dimension 7 avg < 2.0 – The content library has no cluster architecture. Posts are being produced in isolation rather than as part of a coherent topic map.
- Dimension 8 avg < 2.0 – Content is being produced for its own sake rather than to change what readers can do. The strategy doesn't specify audience outcomes.

The 20/32 threshold is the editorial floor

Content below 20/32 has a net-negative effect on the site's authority. It occupies URLs, consumes crawl budget, sends weak topical signals, and trains both readers and search systems to expect low-quality output from this domain. Content that consistently scores below this threshold should be audited for consolidation, improvement, or removal before new content is produced.

Content above 26/32 is citation-worthy and authority-building. A library that consistently scores above this threshold is a compounding asset – it earns links, earns citations, and earns rankings that make the next piece easier to place.

This rubric does not evaluate

- Word count or length – these are implementation variables, not quality signals
- Keyword density – keyword targeting is separate from content authority
- Grammar and mechanics – proofreading is downstream of editorial quality
- Visual design or image quality – this rubric evaluates content, not presentation
- Publishing frequency – volume is not a substitute for authority, and this rubric doesn't pretend otherwise

Clifton Creative Agency · The Content Authority Audit · cliftoncreative.agency

[Full audit service](#) · [E-E-A-T guide](#) · [Book a consultation](#)