CONTENT CUCUMBER

SEO + AI AUDIT · INTERNAL DEMO SHOWCASE

Cucumber & Co.

A comprehensive audit of the Content Cucumber commerce demo (shopware.contentcucumber.com), a fictional 1826 B2B cucurbit and brining-stock trade catalog. It is fast, well-built, and beautifully written, and it is completely invisible to the AI answer layer. That contrast is the entire point of the demo.

AUDIT DATE **2026-05-17**

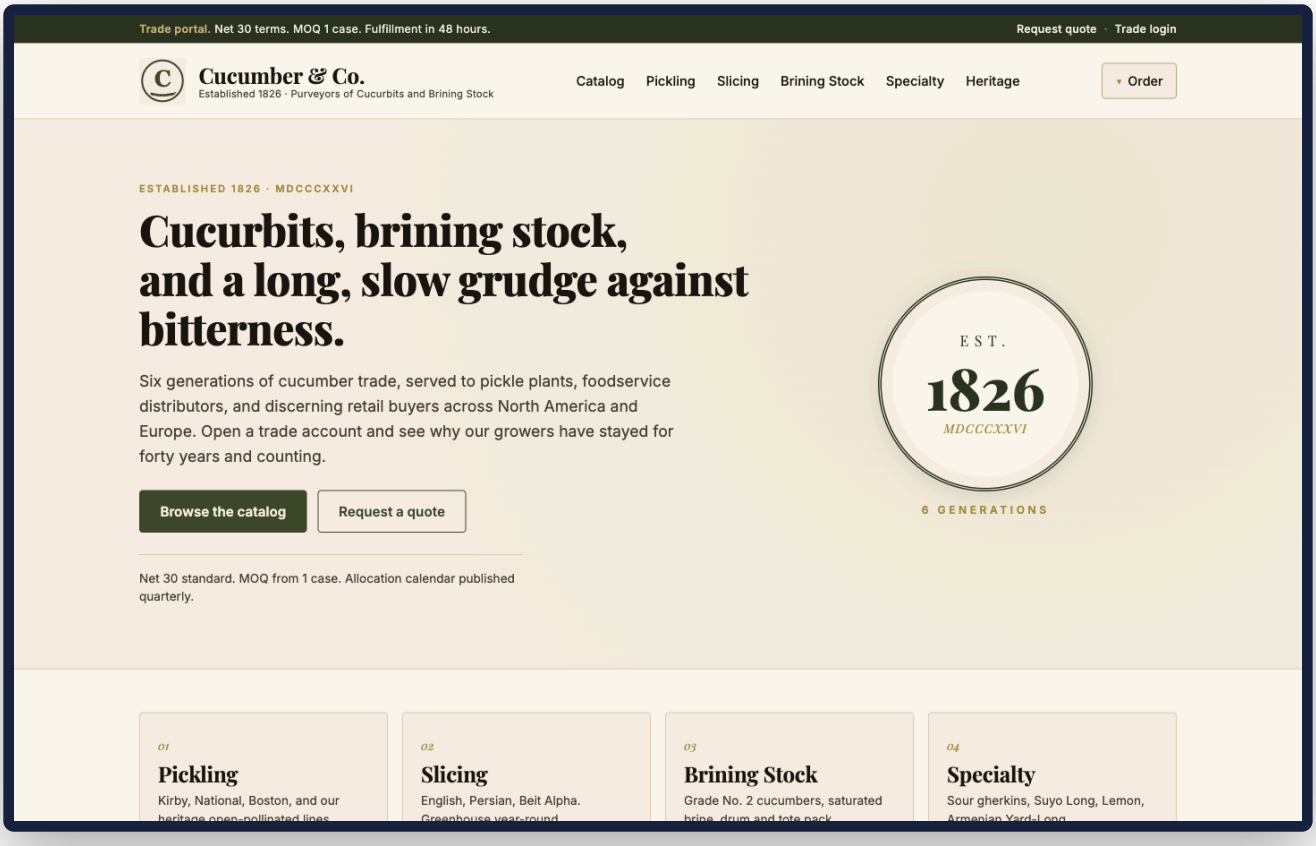

shopware.contentcucumber.com homepage as captured 2026-05-17. Audit findings reference this snapshot.

The short version

AUDITED URL	https://shopware.contentcucumber.com/
AUDIT DATE	2026-05-17
PREPARED FOR	Internal (Content Cucumber commerce demo: "Cucumber & Co.", fictional B2B trade catalog)
PREPARED BY	Content Cucumber
AUDIT TYPE	Comprehensive. External, public-data only. PSI mobile + desktop. Live AI citation baseline across OpenAI, Gemini, Perplexity, and Claude (20 queries, two tiers).
FRAME	Internal demo showcase. No pricing. This is our own demo; the closing is a remediation plan, not an offer.

This is a genuinely good storefront that no AI engine can see. That is the lesson, and it is the cleanest possible illustration of the Content Cucumber thesis on a commerce site. Cucumber & Co. is a server-rendered Astro build behind Caddy, fast (desktop Performance 98, LCP 0.9s; mobile 86, CLS 0), with a real product catalog (five categories, twelve product pages), a clean information architecture (heritage, growers, sustainability, press, EDI, allocations, quote), a strong human-written title and meta description, an open robots.txt with sensible commerce disallows, and a complete sitemap. The craft is real. The copy ("a long, slow grudge against bitterness") is on-brand and excellent.

And the machine-readability layer is essentially absent. **Zero structured data anywhere** (no Organization, no Product, no Offer, no BreadcrumbList on a commerce site), no Open Graph or Twitter cards at all, and llms.txt is a 404. The live four-engine test confirms the consequence: across twenty wholesale-cucurbit category queries, Cucumber & Co. holds **0% share of voice on OpenAI, Gemini, Perplexity, and Claude**, with zero brand mentions. Composites land at OpenAI 21, Gemini 52, Perplexity 41, Claude 21, carried entirely by generic sentiment when the name is supplied. A buyer asking an engine for a wholesale cucumber supplier is handed Sysco, US Foods, and Gordon Food Service. Cucumber & Co. does not appear.

Some of that is by design (a fictional 1826 brand has no real-world authority footprint, and never will). But the technical gaps are real and would apply identically to a live commerce client: a fast, attractive store with no schema, no Open Graph, and no llms.txt is a store the AI answer layer cannot describe, cannot enrich, and will not cite. The demo is the argument.

BRENT'S TAKE

This is the best teaching tool in the set because it is ours and we built it well on purpose. The store is fast, the copy is sharp, the catalog is real, and the AI layer still returns nothing, zero of twenty queries, zero share of voice on all four engines. Part of that is a fictional brand with no authority, which is honest to say. But the technical half is the part that maps straight onto a real merchant. No Product schema on a product site. No Open Graph, so every shared product link is a bare URL. No llms.txt. A beautiful store the engines cannot read is still invisible. That is the whole pitch, and now we can show it instead of explain it.

SEO and AI Health Score

Categorical scores weighted by impact. PSI mobile + desktop captured 2026-05-17 via Google PageSpeed Insights API. AI Search scored on measured live citation outcome.

Category breakdown

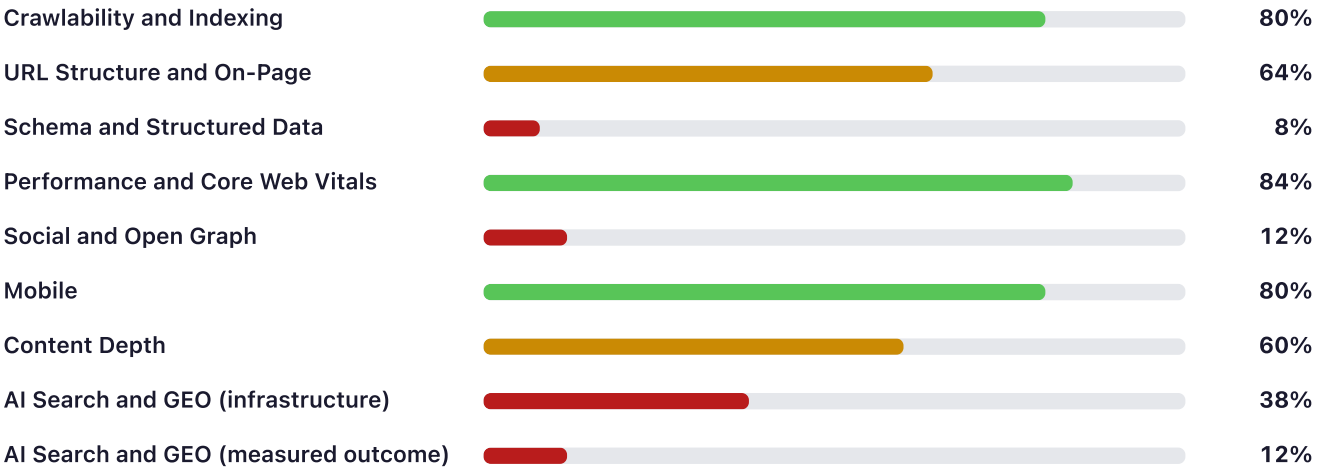

Lighthouse PSI (2026-05-17)

Page / Strategy	Performance	SEO (Lighthouse)	LCP	FCP	CLS	TBT
/ mobile	86	100	3.1s	2.9s	0	0ms
/ desktop	98	100	0.9s	0.9s	0.005	0ms

Performance is strong: zero TBT both strategies, CLS effectively zero, desktop near-perfect. Mobile LCP 3.1s is the one soft metric (the "good" line is 2.5s), driven by FCP 2.9s; a hero-image preload closes it. This is a fast, well-engineered front end. The score is held down by machine-readability and AI visibility, not speed.

What this comprehensive audit covers

The full nine-part deliverable, run here on our own commerce demo.

1. Score and category breakdown	Out-of-100 health score weighted by impact, AI category split into infrastructure and measured outcome.
2. Tech stack	What the site reports: platform, edge, analytics, schema, sitemap, llms.txt, content surface.
3. Top findings	The gaps that hold the score down and what each would add if shipped.
4. Per-check findings detail	Pass / Partial / Fail / Manual per check across the canonical categories.
5. Upgrade code	Copy-paste blocks for the machine-readability gaps (Product / Organization schema, Open Graph, llms.txt).
6. Live AI citation tests	20 category queries, two tiers, run live on OpenAI, Gemini, Perplexity, and Claude on 2026-05-17.
7. 90-day roadmap	What we would sequence on this demo to make it a complete reference build.
8. Internal remediation plan	This is our own demo. No pricing. The fixes we would ship and why.
9. What Content Cucumber produces	The production model that builds the authority surface the demo lacks.

Tech stack

What shopware.contentcucumber.com reports to a direct request (raw HTML, headers, robots, sitemap, llms.txt, route enumeration) on 2026-05-17.

Platform	Astro (static site generator; `/_astro/` assets). Server-rendered HTML, ~17 KB home.
Edge / server	Caddy (`via: 1.1 Caddy`). HTTPS.
Analytics / tools	None detected in static HTML. No GTM, GA4, or pixel.
Structured data	**None.** Zero JSON-LD anywhere: no Organization, no Product, no Offer, no BreadcrumbList. On a commerce site this is the single biggest miss.
Open Graph / Twitter	**Absent.** No `og:*` or `twitter:*` tags. Every shared product or category link renders as a bare URL.
Title / meta description	Strong. Title "Cucumber & Co. · Cucurbits and Brining Stock Since 1826"; a real, human-written meta description.
robots.txt	Present and well-formed. `Allow: /` with sensible commerce disallows (`/cart` , `/checkout` , `/account` , `/admin` , `/api`) and a Sitemap directive.
sitemap.xml	Present and complete. ~28 URLs: home, /products with five category facets, twelve product detail pages, and the heritage/growers/sustainability/press/EDI/allocations/quote/contact set.
llms.txt	**404.** No curated entity/catalog map for language models.
Content surface	A real catalog (five categories, twelve products) plus a genuine brand-story surface (heritage, growers, sustainability, press). No blog or editorial authority layer.

THE INSTRUCTIVE PART

This is a fast, server-rendered, well-structured commerce site with a real catalog and good copy, and it carries no schema, no Open Graph, and no llms.txt. It is the exact profile of a real merchant who invested in design and speed and skipped the machine-readability layer. Section 6 shows what that costs in the AI answer layer: nothing the engines can see, so nothing they cite.

Top findings

1. ZERO STRUCTURED DATA ON A COMMERCE SITE

No JSON-LD anywhere. For a store with a twelve-product catalog and five categories, the absences that matter are Organization, Product + Offer (price, availability, SKU), and BreadcrumbList. Without Product schema there is no rich-result eligibility, no price/availability surface for shopping experiences, and no machine-readable entity for AI engines to ingest. This is the single highest-leverage fix on the site.

What it would add. Rich-result eligibility, a clean entity for AI ingestion, and a price/availability layer commerce surfaces read. Copy-paste Organization + Product blocks in section 5.

2. NO OPEN GRAPH OR TWITTER CARDS ANYWHERE

No `og:*` or `twitter:*` tags. For a B2B catalog whose buyers share product and category links in email, Slack, and LinkedIn, every shared link renders as a bare URL with no title card, image, or description. This depresses click-through on exactly the referral paths a trade catalog runs on.

What it would add. Proper unfurls on every shared product and category link. One head block, sitewide. Section 5.

3. NO LLMS.TXT, AND 0% AI SHARE OF VOICE

llms.txt is a 404, and the live four-engine test returns 0% share of voice with zero brand mentions on all of OpenAI, Gemini, Perplexity, and Claude. Part of this is structural (a fictional 1826 brand has no real authority footprint), but the technical half is real: with no schema and no llms.txt, even a crawlable store gives the engines nothing curated to map. On a live merchant this is the difference between being ingestible and being invisible.

What it would add. A curated catalog map for the crawlers, plus the schema in finding 1, makes the store fully ingestible. Authority (the part a real brand earns) is then the content job. Section 6 and 7.

1. Crawlability, Indexing, and AI Access

Check	Score	If broken	Finding
1.1 robots.txt accessible	PASS		Present, well-formed. `Allow: /` with sensible commerce disallows and a Sitemap directive.
1.2 sitemap.xml present	PASS		Present and complete. ~28 URLs including all twelve product pages and category facets.
1.3 llms.txt	FAIL	Medium	`/llms.txt` 404. No curated catalog/entity map for LLMs. Section 5.
1.4 AI crawler access	PASS		Open robots, no AI-bot blocks. Crawlers can reach everything; there is just little machine-readable structure to ingest.
1.5 Server-rendered content	PASS		Astro static HTML at first contact. No JS-render blind spot.
1.6 HTTPS	PASS		HTTPS via Caddy. Correct.

2. On-Page and Structured Data

Check	Score	If broken	Finding
2.1 Title tag	PASS		"Cucumber & Co. · Cucurbits and Brining Stock Since 1826". Distinctive, category-anchored.
2.2 Meta description	PASS		Real, human-written, on-message. Done right.
2.3 H1	PASS		Single H1, strong voice ("a long, slow grudge against bitterness"). On-brand.
2.4 Open Graph / Twitter	FAIL	High	None at all. Shared product/category links render bare. See finding 2; block in section 5.
2.5 Organization schema	FAIL	High	None. No machine-readable entity for the business. Section 5.
2.6 Product / Offer schema	FAIL	Critical	None on any of the twelve product pages. No price, availability, or SKU surface. The highest-leverage gap on a commerce site.

Check	Score	If broken	Finding
2.7 BreadcrumbList / canonical	MANUAL	Low	Category facets exist (`/products?cat=`); confirm self-referential canonicals and add BreadcrumbList with the schema pass.

3. Performance and Core Web Vitals

Check	Score	If broken	Finding
3.1 Desktop vitals	PASS		Performance 98, LCP 0.9s, FCP 0.9s. Near-perfect.
3.2 Mobile performance	PASS		86 of 100. Strong for a content-rich storefront.
3.3 Mobile LCP	PARTIAL	Low	3.1s, just over the 2.5s "good" line, driven by FCP 2.9s. Hero-image preload closes it.
3.4 CLS	PASS		0 mobile, 0.005 desktop. Excellent layout stability.
3.5 TBT	PASS		0ms both strategies. Minimal JS, as expected from a static Astro build.
3.6 SEO (Lighthouse)	PASS		100 both. On-page basics are correct (the shallow check does not see the schema/llms.txt gaps).

5. Upgrade code

Copy-paste for the three machine-readability gaps. These are the corrective fixes, not marginal polish.

A. Organization JSON-LD (site head)

```
{
  "@context": "https://schema.org",
  "@type": "Organization",
  "name": "Cucumber & Co.",
  "url": "https://shopware.contentcucumber.com/",
  "foundingDate": "1826",
  "description": "Wholesale cucurbits and pickle brining stock for pickle plants, foodservice distributors
}
```

B. Product + Offer JSON-LD (product page template)

```
{
  "@context": "https://schema.org",
  "@type": "Product",
  "name": "{{product_name}}",
  "category": "{{category}}",
  "brand": {"@type":"Brand","name":"Cucumber & Co."},
  "offers": {
    "@type": "Offer",
    "priceCurrency": "USD",
    "price": "{{price}}",
    "availability": "https://schema.org/InStock",
    "url": "{{product_url}}"
  }
}
```

C. Open Graph + Twitter (site head)

```
<meta property="og:type" content="website">
<meta property="og:title" content="Cucumber & Co. Cucurbits and Brining Stock Since 1826">
<meta property="og:description" content="Wholesale cucurbits and pickle brining stock for pickle plants, f
<meta property="og:url" content="https://shopware.contentcucumber.com/">
<meta property="og:image" content="https://shopware.contentcucumber.com/og-card.png">
<meta name="twitter:card" content="summary_large_image">
```

D. /llms.txt (real text file at root)

```
# Cucumber & Co.

> Wholesale cucurbits and pickle brining stock since 1826.
> Trade catalog for pickle plants, foodservice distributors,
> and retail buyers.

## Catalog
- [Products](https://shopware.contentcucumber.com/products)
- [Pickling](https://shopware.contentcucumber.com/products?cat=Pickling)
- [Brining Stock](https://shopware.contentcucumber.com/products?cat=Brining+Stock)

## Company
- [Heritage](https://shopware.contentcucumber.com/heritage)
- [Growers](https://shopware.contentcucumber.com/growers)
- [Contact](https://shopware.contentcucumber.com/contact)
```

A through D are an afternoon on an Astro build. They make the store fully ingestible and eligible. They do not, by themselves, manufacture the real-world authority a live brand earns through content; that is section 7.

6. AI search visibility (live, 4 engines)

20 wholesale-cucurbit category queries (10 strategic, 10 marketplace) run live on 2026-05-17 across OpenAI, Gemini, Perplexity, and Claude. This is what a brand with no authority footprint looks like to the AI answer layer.

Per-engine answer-engine score

OpenAI
gpt-5-search · live web

21 /100

Not in the conversation. Zero share of voice.

Brand Recognition	1/20
Market Score	0/10
Presence Quality	0/20
Brand Sentiment	20/40
Share of Voice	0/10

Gemini
2.5 Pro · live web

52 /100

Generic description only. Zero share of voice.

Brand Recognition	8/20
Market Score	3/10
Presence Quality	16/20
Brand Sentiment	25/40
Share of Voice	0/10

Perplexity
sonar · live web

41 /100

Not surfacing. Zero share of voice.

Brand Recognition	4/20
Market Score	2/10
Presence Quality	15/20
Brand Sentiment	20/40
Share of Voice	0/10

Claude
training data · not live web

21 /100

No training-data footprint. Zero share of voice.

Brand Recognition	1/20
Market Score	0/10
Presence Quality	0/20
Brand Sentiment	20/40
Share of Voice	0/10

THE PATTERN

Four engines, **0% share of voice in every tier, zero brand mentions**. The composites (21, 52, 41, 21) are carried by generic sentiment when the name is in the prompt; brand recognition is 1 to 8 of 20. Asked for a wholesale cucumber or brining-stock supplier, the engines return Sysco, US Foods, Gordon Food Service, and Performance Food Group. Cucumber & Co. never appears. Some of that is the fictional brand having no authority footprint, which is honest. The transferable half is technical: a fast, attractive store with no schema and no llms.txt gives the engines nothing structured to read, so even ingestion adds nothing citable.

What the engines return instead, strategic tier

Cucumber & Co. holds **0% of every pie below**. The slices are the incumbents the engines name for wholesale cucurbit and brining supply.

OpenAI

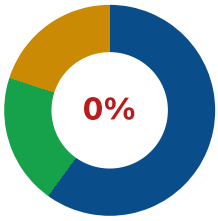

Cucumber & Co. share

- Sysco 60%
- US Foods 20%
- Gordon Food Service 20%

Gemini

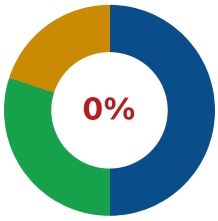

Cucumber & Co. share

- Sysco 50%
- Performance Food Group 30%
- US Foods 20%

Perplexity

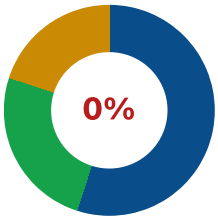

Cucumber & Co. share

- Sysco 55%
- US Foods 25%
- Gordon Food Service 20%

Claude

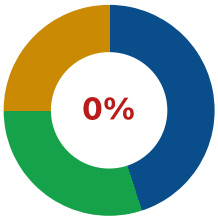

Cucumber & Co. share

- Sysco 45%
- US Foods 30%
- Gordon Food Service 25%

Narrative themes the engines associate with the category

What the engines talk about for wholesale cucurbit and brining supply. Cucumber & Co. is absent from all of it; these are the themes a real brand's content would have to own.

OpenAI

- National distribution scale and reliability
- Foodservice and broadline supply
- Consistent year-round volume
- Established incumbents (Sysco, US Foods)
- Cost and contract terms

Gemini

- Bulk and commercial-grade supply
- Specialty and organic sourcing
- Supply-chain consistency
- Foodservice distributor relationships
- Volume pricing

Perplexity

- Premium and artisan sourcing
- Custom brining blends
- Sustainability and grower relationships
- National vs regional supply
- Vetted, established suppliers

Claude

- No training-data knowledge of Cucumber & Co.
- Defaults to broadline distributors
- Names Sysco, US Foods, GFS instead
- Cannot describe the brand or catalog
- Zero footprint in the model's memory

7. 90-day roadmap (to make this a complete reference build)

Weeks	Workstream	Outcome
1	Machine-readability layer	Organization + per-product Product/Offer schema, BreadcrumbList, full Open Graph + Twitter block, real llms.txt (section 5 A-D). The store becomes fully ingestible and rich-result eligible.
1	Performance polish	Hero-image preload to bring mobile LCP under 2.5s. Everything else is already strong.
2-12	Authority surface (the demo's missing half)	Add the editorial layer a real merchant would need: grower stories, brining technique guides, sourcing and sustainability explainers, all named-author and dated. This is what would move share of voice off zero if the brand were real.
13	Re-baseline	Re-run the 20-query four-engine test + PSI. Track the schema/OG/llms.txt deltas with the new `cc-audit-delta.sh` capability against this 2026-05-17 baseline.

Internal remediation plan

This is our own demo, so there is no pricing and no offer. This is the fix list and the reason the demo is built the way it is.

Ship now (corrective)

Fix	Why	Effort
Product + Organization + Breadcrumb schema	The biggest gap on a commerce site. Rich-result eligibility + a clean entity for AI ingestion.	Afternoon (Astro template)
Open Graph + Twitter block	Every shared product link currently renders bare. One head block, sitewide.	Under an hour
Real /llms.txt	Curated catalog map for the crawlers. Currently a 404.	Under an hour
Mobile hero preload	Closes the one soft performance metric (LCP 3.1s to under 2.5s).	Under an hour

Why the demo is built this way

Choice	Rationale
Fast, beautiful, well-written front end	So the contrast is unmistakable. The store is genuinely good and still scores 49 and 0% AI visibility. Design and speed are not the gap.
Fictional 1826 brand	Honest demonstration that authority is earned in the real world over time. No schema trick manufactures it.
Schema/OG/llms.txt deliberately absent at baseline	So a before/after re-audit with `cc-audit-delta.sh` shows exactly what the machine-readability layer is worth, the same way the aragrow re-audit did.

> **THE POINT OF THE DEMO**
>
> A fast, attractive, well-built commerce site can still be invisible to the AI answer layer. The technical half of that (schema, Open Graph, llms.txt) is a one-afternoon fix and the demo proves the cost of skipping it. The other half, real authority, is the content work, and it cannot be faked, which is exactly why a fictional brand scores zero no matter how good the store looks.

CONTENT CUCUMBER

A beautiful store. Invisible to the AI.

Cucumber & Co. is fast, well-designed, and well-written, and it holds 0% share of voice across OpenAI, Gemini, Perplexity, and Claude with zero brand mentions. The technical half of that gap, no schema, no Open Graph, no llms.txt, is an afternoon of work and the demo exists to prove what skipping it costs. The other half, real authority, is earned through content over time and cannot be faked, which is why a fictional 1826 brand scores zero no matter how good the storefront looks.

The demo, in three lines.

1. **Speed and design are table stakes, not visibility.** This store has both and still scores 49.
2. **Machine-readability is a one-afternoon fix.** Schema, Open Graph, llms.txt. The demo skips it on purpose.
3. **Authority is the content, and it is the only half that cannot be shortcut.** That is what Content Cucumber produces.

This is the reference build. Re-baseline with cc-audit-delta after the fixes ship.

What Content Cucumber produces

The proof. The named-author, schema-wrapped, citable content cadence that builds the authority a storefront alone cannot.

Case studies

Named-client case studies built from interviews, structured, schema-wrapped, written to be cited. The asset class that compounds into category authority.

`contentcucumber.com/our-work`

`contentcucumber.com/services/we-write-case-studies`

Long-form and research

Industry analysis and thought-leadership in a named-author voice. The depth-and-frequency mechanism that moves a brand off zero in the AI answer layer.

`contentcucumber.com/blog`

The demo proves the thesis on commerce: a great store the AI cannot see is still invisible. Schema is the afternoon. Authority is the work.