

OCTOBER 2026

Speaking ChatGPT's Language: How .Best Domains Fit AI Recommendations

How ChatGPT turns a buyer's question into a short list of names, why it speaks the language of "best", and how to build a .Best site it can read, trust and recommend.

24.3%

of ChatGPT advice
searches add "best"

Peec AI, ~5M fan-outs

43.8%

of pages ChatGPT
cites are "best X" lists

Ahrefs, 26,283 URLs

0.664

brand mentions vs AI
visibility, the top
factor

Ahrefs, 75,000 brands

+20%

clicks on identical
.Best ads vs .com

GlobeRunner, small sample

CONTENTS

What this paper covers

- | | |
|----|--|
| 01 | Executive summary
The argument in one page, and the evidence behind each claim |
| 02 | From search results to recommendations
Why being named in the answer is the new first position |
| 03 | How ChatGPT builds a recommendation
The five stages between a buyer's question and a named brand |
| 04 | "Best": the word ChatGPT adds on its own
What 5 million fan-out queries show |
| 05 | What ChatGPT cites
Lists, comparisons and fresh pages, including from new sites |
| 06 | Known is not recommended
Why brand mentions matter more than links |
| 07 | Where a .Best name meets the pipeline
Stage by stage, and what the name does not do |
| 08 | The click after the citation
The only controlled .Best test, read carefully |
| 09 | Building a recommendation-ready .Best site
Readable, verifiable, quotable |
| 10 | Earning the mentions ChatGPT reads
Getting named on the pages that get cited |
| 11 | Use cases
Four business types, four prompt patterns |
| 12 | Test it yourself: the ChatGPT recommendation test
Twenty prompts, a log and a score |
| 13 | Measuring progress over time
Share of recommendation, cadence and common errors |
| 14 | The .Best recommendation study
The open questions, and how we propose to answer them |
| 15 | A 30-day plan, conclusion and sources |

ABOUT THIS PAPER

How to read the evidence

This paper explains how ChatGPT chooses what to recommend, and where a .Best domain fits into that process. It draws on independent research into AI search, citation and naming, and on the one controlled test of .Best published to date.

AI search is young, and much of what is written about it mixes measurement with opinion. We have tried to keep the two apart. Some of what follows is measured on .Best directly; more is measured on AI search in general; and some is our own reasoning about how the two connect. Each claim says which.

THE THREE EVIDENCE TAGS

Every claim carries one of three tags, so you always know how strong the ground is.

MEASURED · .BEST

Measured directly on .Best domains in a controlled test.

MEASURED · GENERAL

Measured by an independent third party on AI search or search behaviour in general, not on .Best specifically.

MECHANISM

Our reasoning about how a .Best name interacts with a measured behaviour. Plausible, not yet measured. Section 14 sets out how to test it.

Disclosure: this paper is published by the registry operator of .Best. Third-party figures are reported as published by their original sources, listed at the end. Readers should consult those sources before quoting any figure.

01 EXECUTIVE SUMMARY

ChatGPT does not rank websites. It recommends a few names.

When a buyer asks ChatGPT which provider, product or service to choose, the answer names a handful of options and links to a few sources. Being one of those names is the new first position. This paper explains how ChatGPT arrives at those names, and where a .Best domain fits into that process.

24.3%

of ChatGPT searches for advice questions add the word "best"

Peec AI · ~5M fan-out queries

43.8%

of pages ChatGPT cites at the buying stage are "best X" lists

Ahrefs · 26,283 URLs

0.664

correlation of brand web mentions with AI Overview visibility, the top factor

Ahrefs · 75,000 brands

~35%

of the "best" lists ChatGPT cites sit on low-authority domains

Ahrefs · same study

Five findings

- 01 ChatGPT searches before it recommends.** For questions that need current information, it rewrites the prompt into its own web searches and builds the answer from what those searches return.
MEASURED · GENERAL
- 02 The word it adds most is "best".** Users rarely type it; ChatGPT inserts it, in about one in four searches for advice questions. **MEASURED · GENERAL**
- 03 What it cites most are "best" lists.** Ranked lists and side-by-side comparisons are the leading cited formats across studies. **MEASURED · GENERAL**
- 04 A .Best name carries that word everywhere it appears:** in the address, the title, every brand mention and every branded link. It cannot make a weak page cited, but it labels a strong one in the vocabulary the search is already using. **MECHANISM**
- 05 When people do see a .Best address, they click it more.** In the only controlled test, identical ads earned 20% more clicks on .Best than on .com, on a small sample. **MEASURED · .BEST**

WHAT WE ARE NOT CLAIMING

There is no public evidence that ChatGPT gives any weight to a top-level domain as such, and no published study yet measures how often ChatGPT recommends .Best sites compared with others. The .Best advantage, where it exists, works through words, mentions and clicks, and only on top of content that deserves to be cited. Section 14 describes the study that would measure it directly.

The evidence at a glance

FINDING	SOURCE	STATUS
ChatGPT has more than 900 million weekly active users	OpenAI, February 2026	MEASURED · GENERAL
ChatGPT adds "best" to 24.3% of searches for advice questions	Peec AI, ~5M fan-out queries, April 2026	MEASURED · GENERAL
"Best X" lists are 43.8% of pages ChatGPT cites at the buying stage	Ahrefs, 750 prompts, 26,283 URLs	MEASURED · GENERAL
Listicles take 63% of AI citations across six engines	Evertune via Search Engine Land, ~25,000 URLs	MEASURED · GENERAL
AI-cited pages are 25.7% fresher than organic results	Ahrefs	MEASURED · GENERAL
Brand web mentions are the top predictor of AI Overview visibility (0.664)	Ahrefs, 75,000 brands	MEASURED · GENERAL
Brands cited in an AI Overview earn 35% more organic clicks	Seer Interactive, 3,119 queries	MEASURED · GENERAL
Statistics, quotations and citations raise visibility in AI answers by up to 40%	Aggarwal et al., KDD 2024	MEASURED · GENERAL
55% of users expect a very clear link between a site and its extension	Nielsen for ICANN, 24 countries	MEASURED · GENERAL
Identical ads earned 20% more clicks and paid 8.1% less per click on .Best	GlobeRunner, 152 clicks	MEASURED · .BEST
A .Best name matches ChatGPT's "best" searches in address, title and mentions	This paper, section 07	MECHANISM
.Best sites are recommended by ChatGPT more often than equivalent sites	Not yet measured	OPEN QUESTION

What to do

NAME

Pair your category with .best: crm.best, ebike.best, movers-lyon.best. The name should state what the site is the best answer to.

BUILD

Make the site readable by AI crawlers, verifiable as one consistent entity, and quotable through answer-shaped comparison pages with sourced figures.

TEST

Run the twenty-prompt test in section 12 before you change anything, then every two weeks with the same prompts. Judge results only against your own baseline.

02 THE SHIFT

From search results to recommendations

For twenty years, visibility meant a position on a page of ten blue links. In an AI answer there is no page of ten. There is a short list of named options, and a few cited sources behind them.

The scale of AI answers

PLATFORM	REPORTED USAGE	SOURCE
ChatGPT	More than 900 million weekly active users	OpenAI, February 2026
Google AI Overviews	More than 2 billion monthly users	Alphabet, Q2 2025 remarks
Google AI Mode	More than 1 billion monthly active users	Alphabet, Q2 2026 remarks
Gemini app	950 million monthly active users	Alphabet, Q2 2026 remarks

Usage figures are self-reported by the platforms and use different definitions (weekly vs monthly users). They describe reach, not how many of those sessions are buying questions.

Fewer clicks, more weight on being named

When an AI summary appears, people click through to ordinary results less often. Pew Research Center tracked the real browsing of 900 US adults in March 2025: searchers clicked a traditional result on 8% of visits with an AI summary, against 15% without one. Ahrefs, comparing 300,000 keywords over two years, found position-one click-through 58% below forecast where an AI Overview appeared.

Share of visits with a click on a traditional result

Google searches by 900 US adults, March 2025

No AI summary shown		15%
AI summary shown		8%

Pew Research Center, 2025. Measured on Google, not ChatGPT; included as the clearest measurement of how AI answers change click behaviour.


The flip side is the citation itself. Seer Interactive, across 3,119 queries from 42 organizations, found that brands cited inside a Google AI Overview earned 35% more organic clicks and 91% more paid clicks than brands left out. **MEASURED · GENERAL**

**In an answer interface, being named is the new first position.
Being unnamed is the new page two.**

03 THE PIPELINE

How ChatGPT builds a recommendation


OpenAI does not publish its ranking logic. What it does document, together with independent studies of ChatGPT's behaviour, describes five stages between a buyer's question and a named brand.



Sources: OpenAI Help Center (ChatGPT search rewrites prompts into searches sent to partner search providers); OpenAI crawler documentation; Peec AI fan-out analysis. Stage 4 is inferred from published citation studies; OpenAI has not described its selection criteria in detail.

Two memories: what the model knows and what it finds

A ChatGPT recommendation draws on two different stores of information, and they behave differently.



For a new or recently renamed site, live retrieval is the realistic path. It rewards pages that exist, can be fetched, and match the searches ChatGPT runs. Trained knowledge follows later, if the site keeps being mentioned. **MECHANISM**

What each stage requires

STAGE	WHAT DECIDES IT	WHAT A SITE OWNER CONTROLS
Prompt	The buyer's own words and intent	Nothing directly. Understand which questions your buyers ask.

STAGE	WHAT DECIDES IT	WHAT A SITE OWNER CONTROLS
Fan-out	The searches ChatGPT writes, which favour words such as "best"	Using the vocabulary those searches use in titles, headings and the name itself
Retrieval	Crawler access, indexing by the search providers, relevance to the query	Allowing OAI-SearchBot, serving full text without JavaScript, keeping pages indexable
Selection	Directness of the answer, specific facts, sources, freshness, agreement with other sources	Answer-shaped pages, published figures, regular updates, consistent facts everywhere
Answer	Retrieved evidence combined with trained knowledge about the brand	Mentions on third-party pages; a name that is easy to recognise and repeat

The crawlers that matter

OpenAI operates separate crawlers for separate purposes. Treating them as one is a common and costly mistake.

USER AGENT	PURPOSE	IF YOU BLOCK IT
OAI-SearchBot	Finds pages for ChatGPT search results	Your pages cannot be retrieved and cited in ChatGPT search
ChatGPT-User	Fetches a page when a user's conversation calls for it	ChatGPT cannot open your page on a user's behalf
GPTBot	Collects content that may be used to train future models	You opt out of future training, not out of search

Based on OpenAI's published crawler documentation at the time of writing. User-agent names and behaviours can change; check OpenAI's current documentation before editing robots.txt.

PRACTICAL CONSEQUENCE

A site that blocks AI agents "to protect its content" mostly removes itself from the live retrieval layer, the one part of the system where a new site can win a recommendation quickly. If you have a reason to opt out of training, block the training agent and allow the search agents. Then check the result at your CDN or firewall, not only in the robots.txt file: bot-protection defaults often refuse these agents without anyone having decided to.

04 THE QUERY

“Best”: the word ChatGPT adds on its own

The fan-out search is the gate every source must pass through. The words ChatGPT puts into those searches decide which pages are even considered. The most common of those words is “best”.

Peec AI analysed about 5 million background (“fan-out”) queries generated by ChatGPT, Perplexity and Grok between 1 and 21 April 2026, and compared them with what users had actually typed.

24.3%

of ChatGPT’s searches for advice questions (“Should I use X?”, “Do you recommend Y?”) added “best”

15.33%

of affected ChatGPT responses contained an added “best”, the word it added most often overall

~5M

fan-out queries analysed across ChatGPT, Perplexity and Grok, 1 to 21 April 2026

Peec AI, via PPC Land

Questions about what to buy, which service to use, or how to choose between options were the most reliable triggers. In other words, ChatGPT translates a buyer’s question into a “best” query before it looks anything up. **MEASURED · GENERAL**

What that translation looks like

WHAT THE USER TYPES	A SEARCH CHATGPT MIGHT RUN	A MATCHING .BEST ADDRESS
Which e-bike should I buy for commuting?	best e-bike for commuting	ebike.best
Do you recommend a CRM for a small agency?	best CRM for small agencies	crm.best
Who should I hire to move house in Lyon?	best movers Lyon	movers-lyon.best
What hotel should I book in Lisbon?	best hotels Lisbon	lisbon-hotels.best

Illustrative examples. Actual fan-out queries vary between runs and model versions, and cannot be observed directly in the ChatGPT interface.

WHY IT MATTERS FOR .BEST

Most domain names are neutral to the question being asked. A .Best name is not: it already contains the word ChatGPT inserts into a quarter of its advice searches. No other namespace carries that word in the address itself. Whether search providers and ChatGPT reward the match is an empirical question; what is certain is that the vocabulary lines up. **MECHANISM**


05 THE SOURCES

What ChatGPT cites

Four independent studies agree on the format AI assistants prefer at the buying stage: ranked lists and comparisons, recently updated, with specific facts. That is exactly the content a .Best domain is named for.

"Best X" lists and comparisons lead the most-cited formats

Share of the most-cited LinkedIn content using each format, six AI assistants, 16 B2B categories. One piece can use several formats.



Meltwater, How LinkedIn Content Wins AI Search, 2026 · 9.5 million AI citations across ChatGPT, Gemini, Claude, Copilot, Google AI Mode and AI Overviews

STUDY	SCOPE	FINDING
Ahrefs	750 buying-stage ChatGPT prompts, 26,283 source URLs	"Best X" blog lists were 43.8% of cited page types, the largest category; the pattern held across the AI platforms examined
Evertune	~25,000 URLs, six AI engines	Listicles took 63% of citations; ranked lists made up 71–86% of those listicles

The door is open to new sites

~35%

of the "best" lists ChatGPT cited sat on low-authority domains. Citation is not reserved for established publishers.

Ahrefs

79.1%

of the 1,100 cited lists Ahrefs could date had been updated in the same year. Freshness is part of the profile.

Ahrefs

A caution on self-promotional lists. A vendor's own "best tools" list that ranks itself first may be cited, but readers and engines can discount it. Lists with published criteria, clear disclosure of ownership and genuine competitors earn more trust and stay on safer ground with advertising rules.

06 THE ENTITY

Known is **not** recommended

Ask ChatGPT about a company by name and it may describe it well. Ask it which company to choose and the same brand can be absent. Recognition and recommendation are different outcomes, driven by different evidence.

THE BRAND QUESTION

"What is Acme Analytics?"

Tests recognition. The model answers from what it knows about this one entity. No buyer types this question when choosing a supplier.

THE CATEGORY QUESTION

"Best attribution software for B2B SaaS?"

Tests recommendation. The model compares the entities it finds and names a few. This is where deals are won or lost.

Mentions predict AI visibility more than links

Ahrefs correlated the web footprint of 75,000 brands with how often Google's AI Overviews mention them. Branded web mentions were the strongest factor (Spearman correlation 0.664), ahead of every classic link-based measure. Branded anchor text was the second-strongest.

Brands mentioned most on the web are mentioned most in AI answers

Average AI Overview mentions, by quarter of brand web mentions

Top quarter for web mentions		169
Next quarter		14

Ahrefs, brand factors vs AI Overview mentions · 75,000 brands. Measured on Google AI Overviews; we assume, but cannot confirm, a similar pattern in ChatGPT.

The correlation does not prove that mentions cause visibility; widely mentioned brands are also bigger and better known. But the direction is consistent with how a model that compares sources would behave: the brands that many independent pages name are the brands it feels safe naming. **MEASURED · GENERAL**

What a .Best name adds to every mention

For a site named `crm.best`, the brand and the topic are the same words. Every forum post, article or social update that names it also says "CRM" and "best". Every branded link reads as a descriptive link. A model that repeatedly sees "`crm.best`" next to CRM comparisons has every reason to associate the name with the category; an invented brand name has to earn that association over years. **MECHANISM**

07 THE FIT

Where a .Best name meets the pipeline

A domain name touches every stage of a ChatGPT recommendation, but its role differs at each one. Here is the honest map.

STAGE	WHAT A .BEST NAME CAN DO	EVIDENCE	STRENGTH
Prompt	Nothing. The buyer's question is the buyer's.	n/a	None
Fan-out	Match the "best" searches ChatGPT writes, in the address and, by extension, in titles built around the name	MEASURED · GENERAL + MECHANISM	Medium
Retrieval	Relevance comes from content. Google states that keywords in a TLD give no ranking advantage; the name starts level with .com	MEASURED · GENERAL	Neutral
Selection	Signal the page's format: a site called ebike.best is expected to compare e-bikes. The page must then deliver	MECHANISM	Low to medium
Answer	Carry category and claim into every third-party mention, the factor most closely tied to AI visibility	MEASURED · GENERAL + MECHANISM	Medium to high
Click	Earn more clicks when the address is shown next to alternatives	MEASURED · .BEST	Directional

WHAT A .BEST NAME DOES NOT DO

- **It does not make ChatGPT prefer a site because of its extension.** There is no public evidence that any AI engine weights TLDs, and Google says it does not.
- **It does not replace content.** A thin page on a perfect name is still a thin page. Section 09 is longer than this one for a reason.
- **It does not guarantee a recommendation.** Recommendations vary between runs, users and model versions. Any single answer is an anecdote.
- **It is not yet measured in ChatGPT.** The click evidence comes from Google Ads. Section 14 proposes the ChatGPT study.

Where the name matters most

The name does the most work where content is otherwise similar: two comparison sites of equal quality, two local providers with equal reviews, two vendors with equal coverage. In those situations the details that remain, including the words a source repeats about you and the address a reader sees, decide who gets named and who gets clicked. MECHANISM

08 THE CLICK


The click **after** the citation

A citation is only worth something if people follow it. The only controlled test of .Best measured exactly that: whether people choose a .Best address over a .com when everything else is identical.

GlobeRunner, a Dallas search marketing agency, ran matched Google Ads campaigns in four local service sectors. Ad copy, keywords, bids and landing pages were identical; only the domain shown differed, such as DallasMovers.Best against DallasMoversBest.com.

Click-through rate was higher on .Best in three of four sectors

Google Ads CTR, identical ads and landing pages



GlobeRunner, .BEST vs .COM Case Study, pp. 11–14 · 11,252 impressions, 152 clicks · scale 0–4.5%

+20%

click-through rate on
.Best, from the same
average position (2.0)

–8.1%

average cost per click
(\$9.14 vs \$9.95)

80 vs 72

clicks, from 5,398 vs
5,854 impressions

–21.1%

cost per click for
movers, the strongest
sector

READ IT CAREFULLY

This test is small: 152 clicks in total. By the standard sample-size calculation for a 20% lift at a 1.23% baseline, each arm would need roughly 35,000 impressions, or about 430 clicks, to confirm the effect at 95% confidence. The result is therefore directional, consistent across three of four sectors, and not yet conclusive. It was measured in Google Ads, where the address appears next to competitors; in ChatGPT, a citation link is shown in a different layout. We report it because it is the only controlled evidence on how people react to a .Best address, and because it points the same way as broader research on keyword domains.

Why an address changes the choice

Independent research on search behaviour explains why the domain shown next to a result matters at all.

STUDY	FINDING	STATUS
leong et al., Microsoft Research, WSDM 2012	Showing the domain flipped users' relevance judgments in about 25% of cases in a 1,000-query blind test	MEASURED · GENERAL
Verisign, 11 million search results	Users were almost twice as likely to click a domain containing their search keyword	MEASURED · GENERAL
Nielsen for ICANN, 5,452 respondents	55% expect a very clear relationship between a site's content and its extension; about eight in ten expect some relationship	MEASURED · GENERAL
Keller, Heckler and Houston, 1998	A brand name that suggests its benefit improves recall of matching advertising claims	MEASURED · GENERAL

From the answer to the visit

In a ChatGPT answer, a reader sees a few names and a few source links. Three things happen after that, and the name plays a part in each.

- 01 The citation click.** A link labelled ebike.best states what the reader will find before they click. MECHANISM
- 02 The follow-up search.** Many readers check a recommended name in a search engine before visiting. A two-word name that says what it is survives being remembered and retyped. MECHANISM
- 03 The return visit.** A remembered address brings a buyer back without another answer, another search or another paid click. MECHANISM

A simple model of AI-driven visits

The value of AI visibility for one site can be written as the sum over the prompts that matter:

$$\text{Visits}_{\text{AI}} = \sum_q N_q \times P(\text{named or cited} \mid q) \times \text{CTR}_{\text{citation}}$$

N_q is how often buyers ask prompt q ; P the probability the site appears in the answer; CTR the share of readers who click through. A .Best name may act on the middle term, through query match and mentions, and on the last term, through the click. Content and third-party coverage remain the largest drivers of P . MECHANISM

09 THE BUILD · LAYER 1

A recommendation-ready .Best site: readable

Three layers make a site recommendable, and the order never changes. A page ChatGPT cannot read cannot be cited. A brand it cannot verify is risky to name. Content nobody else confirms is rarely quoted.

LAYER	WHAT IT ACHIEVES	WITHOUT IT
1 · Readable	ChatGPT's crawlers can fetch and parse your pages	You are invisible, whatever your content says
2 · Verifiable	ChatGPT can confirm you are one real, consistent entity	You are a risk it prefers not to name
3 · Quotable	Your pages answer buyer questions, and other sources name you	You are known but not recommended

Your site in a browser proves nothing

A browser runs JavaScript, accepts cookies and waits for the page to assemble. Many crawlers make one plain request and keep whatever comes back. A site built entirely in client-side JavaScript can look perfect to you and empty to them. Test what a plain request receives:

```
# how many words does a plain fetch return?
curl -sL https://yourname.best/ | sed 's/<[^>]*>/ /g' | wc -w

# does your main headline survive without JavaScript?
curl -sL https://yourname.best/ | grep -i "your headline"

# does your edge answer OpenAI's search agent?
curl -sI -A "OAI-SearchBot" https://yourname.best/ | head -1
```

Sending a user-agent string tests your server and firewall rules, not OpenAI's real infrastructure; some protection services also verify crawler IP ranges. Use it as a first check, then confirm in your server logs that the real agents receive 200 responses.

The readability checklist

- ☐ Every page a buyer could land on is server-rendered or statically generated, and returns its full text to a plain request
- ☐ robots.txt allows OAI-SearchBot and ChatGPT-User, plus the search agents of other engines you want to be cited by
- ☐ No CDN, firewall or bot-protection rule silently refuses those agents
- ☐ Key facts (prices, criteria, rankings) are in text, not only inside images, PDFs, videos or charts
- ☐ Cookie banners and pop-ups do not hide the content from a plain request
- ☐ An XML sitemap is current and referenced in robots.txt; pages are indexable with snippets allowed

robots.txt for a .Best site

```
# Allow AI search and user-triggered agents explicitly
User-agent: OAI-SearchBot
Allow: /

User-agent: ChatGPT-User
Allow: /

User-agent: PerplexityBot
Allow: /

User-agent: Googlebot
Allow: /

# Training agents: allow or disallow according to your own policy
User-agent: GPTBot
Allow: /

Sitemap: https://ebike.best/sitemap.xml
```

Example only. Agent names are as published by OpenAI, Perplexity and Google at the time of writing; verify them against each provider's current documentation.

llms.txt: an optional orientation file

llms.txt is a proposed convention: a short Markdown file at the root of a site that states what the site is and points to its most useful pages. It is young, it is not a standard, and, as far as we know at the time of writing, no major AI engine has committed to reading it. We include it because it costs an hour and forces a discipline that helps everywhere: writing one clear paragraph about what the site is.

```
# ebike.best
> Independent tests and rankings of electric bikes for city commuting.
> Tested on a fixed 18 km route. Updated monthly.

## Core pages
- [Best e-bikes for commuting](/best-commuter-ebikes): ranked list and criteria
- [How we test](/method): route, measurements, scoring
- [Prices](/prices): current prices by model and retailer
- [About](/about): team, ownership and funding disclosure
```

Illustrative example of a fictional site; the test details are invented for the example.

THE .BEST ADVANTAGE AT LAYER 1

None. Readability is the same work on any domain. It is also where most sites fail first, so it comes first. A .Best name only starts to help once the pages behind it can be read. **MECHANISM**

09 THE BUILD · LAYER 2

Verifiable: one entity, one set of facts

A model that cannot confirm who you are will hesitate to recommend you. Verifiability is not a trust badge; it is machine-readable consistency across every place your name appears.

The entity anchor

SIGNAL	WHY IT COUNTS
One name, spelled one way	"ebike.best", "eBike Best" and "Ebike.Best Ltd" look like three entities to a machine. Pick one and use it everywhere.
One canonical domain	Every profile, directory listing and author bio links to the same URL. If you also hold the .com, redirect it.
One category sentence	A single sentence a stranger could repeat correctly: what you are, for whom.
Founding year and location	The cheapest disambiguation signals that exist.
Named people	Founders, editors and authors with real, linked profiles and visible bylines.
Independent confirmation	Several independent sources stating the same facts about you.

The canonical fact sheet

Agree these fields once, then paste them unchanged everywhere the company is described:

Legal name	Founded (year)	Primary domain
Trading name	Headquarters	Services or products
One-line description	Founders	Who it is for
Category	Team size	Pricing shape

The .Best advantage at layer 2

A descriptive .Best name reduces ambiguity by construction. "Acme" may refer to hundreds of companies; `crm-for-agencies.best` refers to one thing, and says what it is. Fewer entities share a name that contains its own category, which should make it easier for a model to keep facts about it together. **MECHANISM**

The flip side: if you run a .Best site alongside an existing brand, decide which is the entity and which is a property of it, and say so explicitly on the about page and in your structured data. Two half-connected identities are worse than one.

Structured data: describe yourself in a machine's language

Schema.org markup in JSON-LD lets you state your facts in a form machines parse without guessing. Start with Organization on the home page, then add the types that match each page: Product or Service on

commercial pages, ItemList on ranked lists, FAQPage on question pages, Article and Person on editorial content.

```
{
  "@context": "https://schema.org",
  "@type": "Organization",
  "name": "ebike.best",
  "url": "https://ebike.best",
  "description": "Independent tests and rankings of electric bikes for city commuting.",
  "foundingDate": "2026",
  "address": { "@type": "PostalAddress",
    "addressLocality": "Lyon", "addressCountry": "FR" },
  "founder": [{ "@type": "Person", "name": "Jane Doe",
    "sameAs": "https://www.linkedin.com/in/janedoe" }],
  "sameAs": [ "https://www.linkedin.com/company/ebike-best",
    "https://www.youtube.com/@ebikebest" ]
}
```

Illustrative example of a fictional site. Validate your own markup with Google's Rich Results Test or the Schema.org validator.

MATCH THE PAGE

Every fact in the markup must match the visible page. A contradiction is worse than no markup.

USE SAMEAS

The most underused field. It ties scattered profiles into one entity.

NAME YOUR AUTHORS

Bylines, bios, Person markup and visible dates are the cheapest trust you can add this week.

A USEFUL DIAGNOSTIC

Ask ChatGPT, in a fresh temporary chat, "Is [your site] legit, and who runs it?" Then read the hedging. Phrases such as "I could not find detailed information" or "there may be several companies with this name" are not politeness. They name the anchor that is missing. Fix what the answer points to, and ask again a few weeks later.

09 THE BUILD · LAYER 3

Quotable: pages shaped like answers

ChatGPT does not read a page the way a person does. It extracts from it. A page is quotable when the answer to a buyer's question can be lifted out of it in one or two sentences, with specifics.

HARD TO QUOTE

"We help riders find the perfect e-bike for their unique lifestyle, with expert advice they can trust."

EASY TO QUOTE

"Our top commuter e-bike for 2026 is Model X: 92 km measured range on our 18 km test loop, 21 kg, €2,499. Updated October 2026."

Illustrative example; the product and figures are invented.

The anatomy of an answer-shaped page

- 01 The headline is the buyer's question**, in the buyer's words, ideally the "best" form ChatGPT searches for: "Best e-bikes for commuting in 2026".
- 02 The first 40 to 60 words answer it completely**, so the answer survives being extracted on its own.
- 03 Specifics throughout**: numbers, dates, prices, names, ranges. A model can only repeat what is stated.
- 04 One idea per section, a table wherever a comparison exists**.
- 05 Published criteria and method**, so the ranking can be checked rather than taken on faith.
- 06 Author, publication date and last-updated date**, visible on the page.

What the research says earns citations

The founding academic study of generative engine optimization (Aggarwal et al., KDD 2024) tested nine content techniques on a 10,000-query benchmark and found they could raise a source's visibility in AI answers by up to 40%. Adding quotations from credible sources, adding statistics and citing sources were among the largest gains; keyword stuffing performed below baseline. **MEASURED · GENERAL**

The lesson for .Best owners is direct: the name puts "best" in the address, so the page does not need to repeat it twenty times. It needs to prove it, with evidence.

The five pages to build first

PAGE	THE BUYER QUESTION IT ANSWERS	WHY CHATGPT NEEDS IT
The ranked list	What is the best [category] for [use]?	The most-cited page type at the buying stage

PAGE	THE BUYER QUESTION IT ANSWERS	WHY CHATGPT NEEDS IT
How to choose	What should I look for in a [category]?	Supplies the criteria the answer uses to compare options
Prices	How much does a [category] cost?	A model can only repeat a price that is published
Head-to-head comparisons	Is A or B better for [use]?	Side-by-side comparisons are the second most-cited format
Method and about	Can I trust this source? Who runs it?	Supports verifiability and disclosure

PUBLISH THE NUMBER YOU ARE NERVOUS ABOUT

Price, timeline, test results, team size. When ChatGPT summarises five options, everything you leave unsaid is a gap it fills with someone else's sentence. And original data, published with its method and sample size, is one of the few assets that keeps earning citations long after it is published.

Keep it fresh

Put the year in titles where it is meaningful, update rankings on a stated schedule, and show the last-updated date. In Ahrefs' data, 79.1% of the cited lists it could date had been updated in the same year, and AI-cited pages were on average 25.7% fresher than organic results. **MEASURED · GENERAL**

The .Best advantage at layer 3

This is where name and content reinforce each other. A ranked list titled "Best e-bikes for commuting in 2026" on ebike.best reads as one consistent answer from address to headline to first paragraph. The same list on a generic domain reads as one page on a site about something else. The expectation the name creates is also a commitment: the site has to rank, compare and update, or the name works against it. **MECHANISM**

10 THE MENTIONS

Earning the mentions ChatGPT reads

Your own site is only part of what ChatGPT reads about you. Most of the sources it cites in a category belong to someone else: review sites, publications, communities, directories, and other people's lists.

Build your citation gap table

Run the prompts in section 12 and collect every domain ChatGPT cites. Then build one table:

CITED DOMAIN	TIMES CITED	YOU APPEAR?	TOP RIVAL APPEARS?	ROUTE IN
example-reviews.com	14	No	Yes	Claim profile, ask customers for reviews
trade-magazine.com	9	No	Yes	Pitch original data
forum.example.org	6	Yes	Yes	Keep answering, named, no links

Illustrative rows. Sort by frequency: in most categories a small number of domains supply a large share of all citations, and those are your targets.

Four routes in, in order

SOURCE TYPE	HOW YOU GET IN	EFFORT
Directories and curated lists	Submit, complete every field with your canonical fact sheet, keep it current	Days
Review platforms	Claim the profile, then ask real customers for reviews over a short, focused period	Weeks
Communities and forums	Answer real questions under your real name, helpfully, over months, without dropping links	Months
Editorial and other lists	Offer journalists and list authors original data, not a product pitch	Needs an asset

WHAT DOES NOT WORK

Paid placements on low-quality sites, link networks, mass guest posts and AI-written roundups. You are not accumulating links; you are trying to appear on a small number of specific pages that ChatGPT already trusts.

THE .BEST ADVANTAGE HERE

Every mention you earn of `crm.best` also says "CRM" and "best", and every link to it reads as a descriptive link. Mentions are the factor most closely tied to AI visibility in the research; a .Best name makes each one carry the category. **MECHANISM**

11 USE CASES

Four businesses, four prompt patterns

Same mechanisms, different prompts and content. The businesses below are illustrative.

COMPARISON PUBLISHER

ebike.best

Buyer asks: "Which e-bike should I buy for commuting?"

Why it fits: the strongest case. The question becomes a "best" search, ranked lists are the most-cited format, and about 35% of cited lists sit on low-authority domains, so a new site can compete on content.

Build: published test method, side-by-side tables, the year in titles, monthly updates, disclosure of affiliate income.

LOCAL SERVICE

movers-lyon.best

Buyer asks: "Who should I hire to move house in Lyon?"

Why it fits: local answers lean on reviews and directories. The name states trade and city in every listing and mention. Movers showed the strongest lift in the GlobeRunner test (+42% CTR).

Build: published price ranges, service areas, review profiles kept current, a consistent name across every directory.

E-COMMERCE BRAND

espresso.best

Buyer asks: "What are good espresso beans for milk drinks?"

Why it fits: a buying-guide site on a .Best name answers the comparison questions ChatGPT turns into "best" searches, and links shoppers to the store.

Build: an honest guide that includes other brands, clear disclosure of ownership, tasting criteria, prices.

B2B VENDOR

crm-for-agencies.best

Buyer asks: "Which CRM should a 20-person agency use?"

Why it fits: B2B buyers ask for shortlists, and comparisons and "best X" lists lead the most-cited B2B content. Every mention of the hub names the niche.

Build: honest comparisons with real competitors, published pricing, migration guides, customer evidence.

Choosing the name

PATTERN	EXAMPLE	BEST FOR
Category + .best	crm.best, ebike.best	Comparison sites and category hubs aiming for AI citations
Category + place + .best	movers-lyon.best	Local services
Use case + .best	crm-for-agencies.best	B2B niches and long-tail buying questions
Brand + .best	yourbrand.best	Campaigns and launches; weaker for AI matching because the category is absent

12 THE TEST

The ChatGPT recommendation test

Most teams test wrong. They type their own name, read a flattering paragraph and conclude they are visible. A buyer never types your name. A buyer types their problem.

The rules

- 01 Use a fresh temporary chat for every run.** Memory and chat history feed your own brand back to you and invalidate the sample.
- 02 Ask category questions, not brand questions.** The brand questions are a separate block, used to find factual errors.
- 03 Run every prompt three times** and record the majority outcome. One run is an anecdote.
- 04 Log the answer as given:** brands named in order, links cited, date, model.
- 05 Test your top rivals with the same set,** so your result has a reference point.
- 06 Freeze the prompt set.** Changing it resets your baseline and hides every trend.

MEASURE THE NOISE BEFORE THE SIGNAL

Answers vary between runs and models change without notice. Run the full set on several separate occasions before changing anything. The natural movement you see is your noise floor; any later improvement smaller than that is weather, not progress.

What to record for each run

FIELD	EXAMPLE
Date and model	2026-10-20 · default ChatGPT model, temporary chat
Prompt number	P04
Named?	Yes / No
Position	3rd of 5 names
All names, in order	Rival A, Rival B, ebike.best, Rival C, Rival D
Cited links	Every URL shown as a source
Factual errors	Anything wrong the answer says about you

Twenty prompts for a .Best site

Replace the brackets with your category, audience, place and rivals. Write them the way your buyers would, not the way your marketing would.

A · Choosing (where you get named)

P01 Which [category] should I pick for [use]?

P02 What's the best [category] for [audience] right now?

P03 Recommend three [category] options for [audience] in [place]

P04 I need a [category] under [budget]. What do you suggest?

P05 Who is the most trusted [category] provider in [place]?

B · Comparing (where you get ranked)

P06 [Rival A] or [Rival B] for [use]?

P07 What are good alternatives to [Rival A]?

P08 Compare the main [category] options on price and quality

P09 Where can I find an independent ranking of [category]?

P10 Which [category] do people like me actually use?

C · Criteria (where you get filtered)

P11 How do I choose a good [category]?

P12 What mistakes do people make when buying [category]?

P13 How much should a [category] cost in [year]?

P14 What features matter most in a [category] for [use]?

P15 What should I ask before choosing a [category]?

D · Identity (what it believes about you)

P16 What is [yourname].best?

P17 Is [yourname].best reliable? Who runs it?

P18 What do people say about [yourname].best?

P19 How does [yourname].best rank [category]?

P20 Is [yourname].best independent?

WHY BLOCK C MATTERS MOST

Prompts P11 to P15 look like content questions, but they produce the criteria ChatGPT then uses to compare options in blocks A and B. A site that publishes the clearest answer to "how do I choose" is, in effect, writing the rubric the answer applies to everyone in the category.

Score yourself, and the name ChatGPT picks first

The test tells you whether you are named. A scorecard tells you why not. Score your site and the competitor ChatGPT names most often, which may not be the rival you worry about most, on five dimensions, out of 20 each.

DIMENSION	WHAT 20 OUT OF 20 LOOKS LIKE	SCORE
Readability	Every key page returns full text to a plain request; AI search agents receive 200 responses	/20
Entity consistency	One name, one description, one fact sheet, identical everywhere; valid structured data	/20
Name–content match	The address, titles and first paragraphs state the same category and the same promise	/20
Third-party footprint	Present on the domains ChatGPT cites most in your category	/20
Answer coverage	A dedicated, answer-shaped page for each of your top buyer questions	/20
Total		/100

Reading the gap

Your score alone is an abstraction. Your score next to the brand named first is a work order. The gap is rarely spread evenly; it usually sits in one or two dimensions, most often third-party footprint and answer coverage.

Reading the pattern across engines

If you also run the prompts in other assistants, the pattern of where you appear is diagnostic. These readings are rules of thumb, not measured laws.

PATTERN	LIKELY CAUSE
Absent everywhere	Readability or entity problem. Start with layer 1.
In engines that search live, not in ChatGPT	Pages are fine; third-party mentions are thin, or ChatGPT is not searching for this prompt
In ChatGPT only for brand questions	Known but not recommended: answer coverage and mentions are missing
Named, but with wrong facts	Inconsistent fact sheet across sources. Fix layer 2.

13 MEASUREMENT

Measuring progress over time

AI visibility has no official dashboard. It has a frozen prompt set, a spreadsheet and a steady rhythm. Every two weeks is a practical cadence: often enough to catch regressions, rarely enough to see past the noise.

Five numbers every two weeks

METRIC	WHAT YOU RECORD
Share of recommendation	Prompts in which you are named ÷ prompts tracked
Average position	Where you sit when you are named
Competitor set	Who else is named, and any new entrant
Citation domains	Which sources the answers draw on, and whether you are on them
Factual errors	Anything wrong ChatGPT says about you, and whether it persists

Alongside the prompt panel

SIGNAL	WHERE TO FIND IT
Referral visits from AI assistants	Analytics referral reports (chatgpt.com, perplexity.ai, gemini.google.com and similar)
Crawler activity	Server logs: requests from OAI-SearchBot and ChatGPT-User, and their response codes
Branded search demand	Search Console impressions for your name: people checking what ChatGPT told them

WHAT TO EXPECT, AND IN WHAT ORDER

Changes that affect live retrieval, such as making pages fetchable or publishing an answer-shaped page, can show up within weeks, once pages are crawled and indexed. Changes that depend on trained knowledge appear only when a new model is trained, on a schedule no one outside OpenAI controls. Expect being named before being ranked first, and expect a crowded category to take longer than a thin one. We do not publish a timeline because we have no measurement that supports one. **MECHANISM**

Common errors

- Testing in a logged-in chat with history, which measures your own conversations, not the market
- Changing the prompt set mid-way, which erases your baseline
- Not keeping a frozen week-zero snapshot: it is the only "before" you will ever have

14 RESEARCH

The .Best recommendation study

This paper rests on measured behaviour of AI search in general and one small controlled test of .Best. The next step is to measure the .Best effect inside ChatGPT directly. Here is how we propose to do it.

Four hypotheses

#	HYPOTHESIS	HOW IT WOULD BE TESTED
H1	Equivalent comparison content is cited more often by ChatGPT on a .Best domain than on a generic domain	Matched sites, same content and launch date, fixed prompt panel tracked monthly
H2	Mentions of a category .Best name are more often associated with that category in AI answers than mentions of an invented brand name	Brand-question prompts (block D) on matched pairs, scored for category accuracy
H3	Readers click a .Best citation link more often than an equivalent .com link	Panel study showing ChatGPT-style answers with matched sources
H4	The paid-search CTR lift holds at adequate sample size and across countries	The GlobeRunner design, at ~35,000 impressions per arm, in several markets

Design principles

- **Isolate the domain.** Same content, same launch date, same promotion; only the name differs. Use self-referencing canonicals on each site.
- **Pre-register the prompt set and the analysis** before the first run, so the result cannot be chosen after the fact.
- **Report nulls.** If .Best shows no effect on a measure, that is published too.
- **Publish the method and the data** so agencies and researchers can check and extend the work.

HOW TO TAKE PART

Businesses already on .Best can contribute by running the twenty-prompt test in section 12 with a frozen set, and by running paid-search split tests on their own accounts. Agencies and researchers interested in co-designing the study can contact the registry through domains.best.

15 ACTION PLAN

A 30-day plan for a .Best site

The order matters more than the speed. Each week depends on the one before.

WEEK 1 · BASELINE

- ☐ Write your twenty prompts and freeze them
- ☐ Run them three times each in temporary chats; log everything
- ☐ Run them for the rival named most often
- ☐ Save the week-zero snapshot

WEEK 2 · READABLE

- ☐ Test six key pages with a plain request
- ☐ Fix rendering on any page that returns little text
- ☐ Allow AI search agents in robots.txt and at the edge
- ☐ Check server logs for 200 responses to them

WEEK 3 · VERIFIABLE

- ☐ Agree the canonical fact sheet
- ☐ Align every profile and directory to it
- ☐ Add Organization and page-level structured data
- ☐ Publish named authors and an about page

WEEK 4 · QUOTABLE

- ☐ Publish or rewrite the ranked list and the "how to choose" page
- ☐ Publish prices and your method
- ☐ Build the citation gap table
- ☐ Start the directory and review routes

IF YOU ONLY DO THREE THINGS

Make your key pages readable without JavaScript. Publish the answer to the question your buyers ask most, with real numbers in it. Get named on the sites ChatGPT already cites in your category. Then measure against your own baseline, every two weeks, with the same prompts.

After day 30

Re-run the prompt set every two weeks. Add one answer-shaped page and one third-party placement each month. Update your ranked content on a stated schedule. Re-score yourself and your top rival each quarter, and move effort to whichever dimension has the largest gap.

CONCLUSION

A name in the language of the answer

ChatGPT turns buying questions into “best” searches, and answers them with a short list of names drawn from “best” lists, comparisons and the brands those pages mention. That process is measured, by several independent studies, across millions of queries and citations.

A .Best domain is built in the same vocabulary. It puts the word ChatGPT adds into the address. It puts the category and the claim into every mention a site earns. And when readers see it beside alternatives, the one controlled test available found they click it more often.

What a .Best name cannot do is replace the work. ChatGPT names sources it can read, entities it can verify and pages that answer the question with evidence. The businesses that will gain most are those that pair the name with content worthy of it: honest rankings, published criteria, real numbers and regular updates.

The direct effect of a .Best name on ChatGPT recommendations has not yet been measured. We have set out the method to measure it, and we intend to publish the results, whichever way they point.

ChatGPT is already searching for “best”. The question is whether your site answers in the same words, with the evidence to back them.

ABOUT .BEST

Type	Generic top-level domain (gTLD), delegated in the internet's root zone
ICANN registry agreement	Signed 19 December 2013
Delegated in the root zone	27 February 2014
Registry operator	Best TLD Pty Ltd
Registry services	CentralNic
Registry website	domains.best

SOURCES

Sources

Third-party figures are reported as published by the sources below. Several were compiled for the registry's October 2026 white paper, .Best Domains: The SEO and GEO Advantage, where full links are given. Verify each figure against its original source before quoting it.

Aggarwal et al. GEO: Generative Engine Optimization. ACM KDD 2024.

Ahrefs. Do self-promotional “best” lists boost ChatGPT visibility? (750 prompts, 26,283 URLs).

Ahrefs. Brand factors and AI Overview mentions (75,000 brands).

Ahrefs. Do AI assistants prefer to cite fresh content?

Ahrefs. AI Overviews reduce clicks (300,000 keywords).

Alphabet. Q2 2025 and Q2 2026 earnings remarks.

GlobeRunner. .BEST vs .COM Case Study, pp. 11–14.

Google Search Central. Google's handling of new top-level domains.

IANA. .best root zone database entry.

ICANN. .best registry agreement.

ICANN / Nielsen. Global Consumer Research, Wave 2 (24 countries, 5,452 respondents).

leong et al. Domain Bias in Web Search. WSDM 2012.

Keller, Heckler and Houston. Journal of Marketing, 1998 (brand name suggestiveness and advertising recall).

Meltwater. How LinkedIn content wins AI search, 2026 (9.5 million citations).

OpenAI. Scaling AI for everyone (weekly active users, February 2026).

OpenAI. Crawlers documentation (OAI-SearchBot, ChatGPT-User, GPTBot).

OpenAI Help Center. ChatGPT search.

Pew Research Center. Google users are less likely to click on links when an AI summary appears, 2025.

PPC Land. Report on Peec AI's analysis of ~5 million ChatGPT fan-out queries, April 2026.

Search Engine Land. Listicles and AI citations (Evertune, ~25,000 URLs).

Seer Interactive. AIO impact on Google CTR, September 2025 update (3,119 queries).

Verisign. Keyword-rich domain names and click-through (11 million search results).

The test method, scorecard and build layers in sections 09 to 13 describe common practice in generative engine optimization and are presented by the registry as guidance, not as measured results. All example sites, products and figures marked “illustrative” are fictional.

.Best

Be the answer ChatGPT is looking for.

Find a .Best name for your category, build the pages behind it, and run the twenty-prompt test to see where you stand today.