

Who is AI-friendly?

281 trades businesses, measured.

We fetched robots.txt, /llms.txt and the home page on every site in two sets of tradesmen: 157 businesses across 24 trades in Iowa, and 124 of the same trades in fourteen other states. Then we read every body that came back and scored it against nine checks. Four of the 281 publish a file an assistant can quote. None of the 281 tells a crawler the file is there.

281

businesses fetched, on one day, in 24 trades

109

answer at /llms.txt, which is 39% of the set

4

publish a file with a phone number, a place and a price or the hours

0

point at that file from their own robots.txt

Every URL fetched 17 September 2026

Iowa: 157 businesses, 24 trades

Nationwide: 124 businesses, 14 states

pinelandtech.com

Why this was worth counting

A customer asking an assistant for a plumber in Ankeny gets one answer with one or two names in it. The assistant does not browse the way a person does: it goes looking for the shortest trustworthy account of a business it can find, and the first address it tries is **yoursite.com/llms.txt**, a plain text page stating what the business does, what it costs, where it is and how to reach it.

Six months ago almost nobody had one. That has changed, and not in the way it looks. The file is now generated for you by website builders and SEO plugins, so **39% of these businesses answer at that address**. Read the bodies and the number that matters collapses: **17 of 281** publish something that is about the business rather than a machine-written index of the site's own page links, and **4 of 281** put a phone number, a place and a price in it.

That gap is the opportunity, and it has a shelf life. The businesses that get quoted first are the ones the assistants keep quoting.

Nine checks, and what each market scored

Being readable by a machine is not a matter of taste, so we wrote it down as nine checks anybody can run on any address in about a minute. Every one is a fetch and a read. Nothing here depends on liking the site, and nothing needs access to it.

01	robots.txt answers, in plain text The address every crawler tries first returns a text file, not an error and not a styled page. Twenty-eight sites have no usable one: eleven answer 404 with a full HTML error page, seven answer 200 with an empty body, and the rest refuse or never connect.	IOWA 141 / 157	US 112 / 124
02	The assistants are not blocked Nothing in robots.txt shuts out the crawlers reading for ChatGPT, Claude, Perplexity or Gemini. Three sites block one by name. One of them shuts out the two agents that fetch a page because a person has just asked a question about it, which is a customer standing at the door.	IOWA 156 / 157	US 122 / 124
03	They are admitted by name GPTBot, ClaudeBot, PerplexityBot, Google-Extended and the rest, each named in its own group. A wildcard admits them anyway, so this changes nothing about what is fetched: it states a position that can be read, and its absence reads as not having thought about it.	IOWA 9 / 157	US 8 / 124
04	The sitemap is declared A Sitemap line in robots.txt, so a crawler on a budget is handed the page list instead of guessing. The oldest rule here and still the one most often missed.	IOWA 119 / 157	US 99 / 124
05	robots.txt points at the llms.txt There is no directive for it, so it goes in as a comment on its own line after the sitemap. Not one site in either market carries it, including every site that went to the trouble of publishing a file.	IOWA 0 / 157	US 0 / 124
06	/llms.txt answers with a real body A 200, plain text, and something in it. Two Iowa sites answer 200 with zero bytes, which passes every check that stops at the status line, and twelve more hand back a web page with a stylesheet on it.	IOWA 64 / 157	US 45 / 124
07	The file is the business, not the sitemap It reads as an account of the business rather than a list of its own page links: no generator's credit on the first line, no builder placeholders, no staging domain, and under 50 KB, past which a model stops reading. This is where nearly every file that exists falls over.	IOWA 8 / 157	US 9 / 124
08	It carries the facts a buyer asks for A phone number, a place, and either a price or the opening hours: what somebody is actually asking for when your name comes up. Four of the 281 have all three in the file.	IOWA 1 / 157	US 3 / 124

09 The page says the same thing in structured data

JSON-LD on the home page describing the business with the telephone in it, so the older reader and the newer one get the same record. A quarter of the set publishes no structured data of any kind.

IOWA

62 / 157



US

67 / 124



Nobody in either market passes all nine. The best score in Iowa is six of nine, held by three businesses. Two sites nationwide reach eight: an HVAC contractor in Denver and a roofing company in Tampa. Both miss the same check, the one nobody passes: telling a crawler in robots.txt that the file exists.

WHAT WE FOUND

The file is now a checkbox somebody else ticks for you

The single most useful finding is not that businesses have no file. It is who wrote the ones that exist. Sorting the 281 sites by the platform they are built on separates cleanly:

PLATFORM	SITES	WITH A FILE	FILE IS ABOUT THE BUSINESS
WordPress (with an SEO plugin)	137	36	4
Duda	41	41	1
Wix	10	10	0
GoDaddy Website Builder	9	6	3
Squarespace	11	1	1
Custom or unknown	66	12	7
Shopify, Webflow, other	7	3	1

Platform read from the generator meta tag and the asset hosts on the home page, 17 September 2026. Four home pages could not be reached.

Every Duda site in the set has a file. Every Wix site has one. Not one of the businesses on them asked for it, and only one of the 51 files reads as an account of the business. What the builder generates is an index of the site's own pages, opening with the domain name where the business name should be:

```
# www.examplecompany.com

## Pages

- [Location One](https://www.examplecompany.com/location-one)
- [Thank You](https://www.examplecompany.com/thank-you)
- [Blog](https://www.examplecompany.com/blog)
- [Home](https://www.examplecompany.com/): We provide bath and kitchen remodels, water heater and sump pump repairs.
```

The shape of 77 of the 109 files in the set, paraphrased. No phone number, no address, no price, and a page called Location One the builder left behind.

Beside that, a second family: the WordPress SEO plugins that credit themselves on line one and then dump the whole sitemap. Those files run to **107 KB, 141 KB, 234 KB and 305 KB** in this set. A model reading an answer does not get through them.

What was actually in the 109 files

OF THE 109 FILES THAT ANSWERED	COUNT	SHARE
Carry a phone number	43	39%
Carry a street address	7	6%
Carry any price	15	14%
Carry opening hours	9	8%
Are mostly a list of the site's own page links	77	71%

OF THE 109 FILES THAT ANSWERED	COUNT	SHARE
Open with a plugin's credit rather than the business	38	35%
Still contain the builder's placeholder pages	13	12%
Point mostly at a domain that is not the business's own	5	5%
Run past 50 KB	11	10%

Five things a machine read on a live site this week

- **Two files that are 200 OK and zero bytes.** A gutter company in Cedar Rapids and a detailer in the Des Moines metro. Every automated check that stops at the status line calls these a pass.
- **A towing company's file that describes a different company.** Different name, different brand, in the file served at their own address.
- **An HVAC company's file listing a staging domain** of the form *blueviolet-gazelle-459437.hostingersite.com*, which is what an assistant would be asked to cite.
- **A tire and service shop's file with the template markers still in it:** "[ND:ClientName] is not only the best store, it's THE STORE", and the service area left as a placeholder.
- **Three sites blocking assistant crawlers by name:** a welding shop in Austin, a pest control company in Nashville, and a tire and service shop in Coralville that blocks the two agents which fetch a page because a customer has just asked about it.

It can be done, and four businesses are doing it

These are the only four of the 281 whose file carries a phone number, a place and a price or the hours. Every one is public: open the address and read it. Three of the four are outside Iowa.

BUSINESS	TRADE AND PLACE	SCORE	WHAT THE FILE DOES
Denver's Best Heating and AC Repair denversbestheatingandacrepair.com	HVAC, Denver CO	8 / 9	Names one canonical phone number and address, states the diagnostic fee and an installation range, dates the file, and names the assistants' crawlers in robots.txt.
RoofX roofxincfl.com	Roofing, Tampa FL	8 / 9	Legal name, two offices, licence number, a text-only line marked as such, county-by-county service area, and price ranges per roof type.
Performance Electrical Contractors performanceelectricalllc.com	Electrical, Atlanta GA	7 / 9	Written as labelled fields: business name, legal name, site, address, hours. Reads like a record rather than a page.
DM Repair dmrepair.net	Appliance repair, Des Moines IA	5 / 9	The only one in Iowa. A builder's page index, but with a real contact block and the opening hours under it, which is what carries it.

The Iowa result in one line. Across 157 businesses in 24 trades, one has a file a machine can quote, and it is an appliance repair shop that got there by accident. Every trade in the state is open: plumbing, electrical, HVAC, roofing, concrete, tree work, detailing, welding, towing, chimney, flooring, fencing, painting, drain and sewer. Not one business in any of them has done this on purpose.

Iowa, trade by trade

TRADE	SITES	FILE	ABOUT THE BUSINESS	CARRIES THE FACTS	BUSINESS SCHEMA
Electrical	18	4	0	0	10
Plumbing	14	6	1	0	7
HVAC	12	4	0	0	5
Auto detailing	10	3	1	0	4
Landscaping and lawn care	8	5	0	0	3
Appliance repair	8	5	2	1	3
Welding and fabrication	8	1	0	0	2
Remodeling	8	4	0	0	2
Tree service	7	5	0	0	3
Concrete	6	4	1	0	1
Gutters	6	3	0	0	3
Small engine repair	6	6	2	0	1

TRADE	SITES	FILE	ABOUT THE BUSINESS	CARRIES THE FACTS	BUSINESS SCHEMA
Roofing	5	3	0	0	1
Auto repair	5	2	0	0	2
Fencing	5	2	0	0	4
Chimney and masonry	5	0	0	0	2
Garage doors	4	1	0	0	0
Painting	4	0	0	0	2
Towing	4	2	0	0	1
Drain and sewer	4	0	0	0	2
Septic and excavation	3	2	1	0	1
Flooring	3	0	0	0	0
Pest control	2	1	0	0	1
Windows and siding	2	1	0	0	2

The position is open, and it will not stay open

Search took twenty years to teach everyone that the business at the top of the page gets the call. The assistants are forming the same habits now, on a much shorter clock, and the habit forms around whoever is easiest to quote. Being easy to quote is a day of work that almost nobody has done.

What a business gets from doing it

- **Named in the answer** when somebody asks who does this work here, because there is a citable account of the business at the address the assistant checks.
- **Quoted accurately.** The price, the phone number and the service area come from the business rather than from a directory page that ranked that day.
- **Quoted consistently.** Generated from the site's own content, so the file cannot disagree with the pages, and a price change is one edit.
- **Held.** A source that has been cited before gets cited again. This is the cheapest incumbency a small business will ever be offered.

What it takes, honestly

- **A single content source.** The phone number, hours and prices written once and read by every page, the structured data and the file.
- **A route, not a file.** Generated on request from that source. A hand-typed file in a public folder goes stale in silence and is then repeated back to customers by a machine.
- **A robots.txt that lets them in, names them, and says where the file is.** Three lines of work that nobody in 281 businesses has done.
- **Verification on the live address.** Read the body after deploying, not the status code. Two sites in this study serve zero bytes and pass every other check.

Every site Pineland builds passes all nine on the day it launches, and the build fails if it would not. After each deploy the live address is fetched and the body is read, because a file that answers 200 with nothing in it is the failure this study found twice.

Files you can open right now: kjsineworks.com/llms.txt, jeromeprintshop.com/llms.txt, deafinitelydogs.org/llms.txt, the-poo-patrol.com/llms.txt, pinelandtech.com/llms.txt.

Method, and what this study does not claim

Both sets were built the way a customer finds a tradesman: a search for the trade and the town, with directories, franchise location pages, manufacturer finders and dealership departments removed, leaving independent businesses with their own domain. Iowa covers 24 trades in more than 30 towns, from Sioux City to Dubuque. The nationwide set covers the same trades in Phoenix, Atlanta, Dallas, Denver, Tampa, Charlotte, Portland, Nashville, Columbus, Seattle, Minneapolis, Los Angeles, Boston, Houston, Austin, Toledo and rural Pennsylvania.

Every URL was fetched once, on 17 September 2026, with a desktop browser user agent, following redirects. Eight Iowa home pages and four nationwide did not answer (timeouts, a 403, a 500, a 521); those sites are still counted in the totals, and a site that will not answer a browser will not answer a crawler either. Checks 7 and 8 are measured on the body of the file: check 7 looks for a generator's credit, a builder placeholder, an off-domain link set, a bare link index or a size past 50 KB, and check 8 looks for a phone number, a postal marker and a price or an hours block. A file could in principle be an excellent account of a business and be scored down by a formatting choice; reading the 109 bodies, we did not find one.

This is a snapshot of one day. The platform-generated files in particular are new and are spreading fast, which is the reason the counts are dated on every page of this report.