AI Visibility (GEO) for WordPress & WooCommerce

Documentation, version 1.7.11

What this plugin does

When someone asks ChatGPT, Claude, Gemini or Perplexity about a product, or asks for a comparison, the assistant answers from whatever it can find and understand about you. This plugin makes sure what it finds is complete, structured and consistent.

It does five things:

1. Describes every product, page and listing in a form assistants can quote.
2. Publishes clean, machine-readable versions of your content.
3. Tells AI crawlers they are allowed in, and shows you whether they actually get through.
4. Fills the whole catalogue in one pass, so you are not editing hundreds of items by hand.
5. Describes your images in words, because most catalogues leave them silent.

What it does not do: promise rankings. Nobody can. What it does is remove the reasons an assistant would skip you.

Step 1: Install and activate your licence

1. In WordPress, go to **Plugins → Add New → Upload Plugin**.
2. Choose the `.zip` file you received and click **Install Now**, then **Activate**.
3. Go to **AI Visibility → Licence** in the left menu.
4. Paste the licence key from your purchase email and click **Activate**.

The status turns to **Active** and shows how many sites your licence covers.

Moving to another site? Click **Remove from this site** first. That frees the seat immediately.

If the licence server is unreachable, the plugin keeps working normally for 45 days on its last verified check. The check itself runs once a month, because a one time licence has nothing that can lapse between one day and the next. A network problem will never break your website.

Without a licence the plugin still writes your structured data, still shows the editor panel and still manages robots.txt. Only the extras stop: the markdown mirror, the AI sitemap, llms.txt, bulk fill and image descriptions.

Agency and vendor installs. If you were given a master key, add it to `wp-config.php` instead of activating in the admin:

```
define( 'AVW_MASTER_KEY', 'your-master-key' );
```

Step 2: Two minutes of setup

Business tab

Fill in what applies. Blank fields are simply left out.

Field	Why it matters
How do you sell	Online store, local business, or content site. This decides how prices and availability are described. Pick the one that matches reality.
Business name and one-sentence description	Used to tell AI who you are.
Official profiles	One URL per line: Facebook, Instagram, LinkedIn, YouTube, TikTok, Google Business, Wikipedia. Published as sameAs so an assistant can confirm your identity and tell you apart from a business with a similar name. Follower counts are irrelevant; the link itself is the signal.
Address, phone, opening hours	For a local business this is what answers "where can I see this in person".
Currency	For prices that do not come from WooCommerce.
Price in the AI summary	Leave on. It makes answers read "Model X, 3000 RON, available to try on in person" instead of "contact them for pricing".
Delivery time	Write a range: "2-4 working days" is read as a range, "fast" is read as nothing. Tick **delivery varies by product** if bulky and small items ship differently, and no single figure is published.
Delivered to	A country code, or a short list.
Delivery cost	Tick **this is a starting price** to publish it as "From 25 RON".
Return window	In days. Leave at 0 if you would rather not state one.
Warranty	Write the unit too. "24 months", not "24". The plugin publishes exactly what you type and never guesses.
Payment methods	Plain words: "Card, bank transfer, cash on delivery".

These last six answer the questions shoppers ask an assistant *before* they ask about the product itself. Anything you leave empty is left out entirely.

Modules tab

- Tick the content types that hold your catalogue.
- Set the **"Describe it to AI as"** column. If your products live in a custom content type rather than WooCommerce products, set it to **Product**. This one setting is what gives them price, availability and comparable attributes.
- Review the **attribute dictionary**. See Step 4.

Which products are published

Drafts, private, scheduled, trashed and password protected products are never published to AI, whatever these switches say. Beyond that, WooCommerce lets a product be technically published while still being hidden or restricted, and three switches on the Modules tab decide what happens to those:

- **Products hidden from the catalog.** Off by default: a product you have hidden stays hidden from AI too.
- **Products set to search results only.** On by default: still published, since you chose to keep it reachable.
- **Out of stock products.** On by default: still published, so an assistant can tell a shopper it will be back.

AI Crawlers tab

Three presets in plain language. **Maximum visibility** is the default and the right answer for most shops.

Worth knowing: blocking training crawlers does not affect whether you appear in AI answers. Blocking *search* crawlers removes you entirely.

Step 3: Fill the catalogue

AI Visibility → Fill everything

1. Click **Check what is missing** first. You get a count of items, how many prices can be detected in your text, and the state of your images.
2. Click **Fill everything now**. It walks the catalogue in batches of 20 with a progress bar. A few hundred items take a couple of minutes.

It writes: the AI summary, the price, key facts, starter questions, and alt text for images that have none.

With the overwrite box off, nothing you wrote by hand is touched. Ticking the box is different: it replaces every field on every item, including summaries, questions and attributes edited by hand, and alt text a person wrote in the media library. There is no undo, so leave it off unless you genuinely want to start from zero.

Why fields say Custom after a bulk pass. Once the bulk fill has written a field, the value is stored, so the panel shows Custom even though the plugin wrote it. The next bulk pass skips stored fields. That is why running an update can appear to change nothing until the overwrite box is ticked.

Check your shared images afterwards. Alt text belongs to the image, not to the product. If the same image is attached to several products, the description written for the first one stays with it, and every other product using that image inherits it. On most catalogues this is invisible, because each product has its own images. On catalogues that reuse a stock shot across a range, it means one product can end up described as another.

After the first bulk pass, open **Media → Library**, switch to list view and sort by date. Scan the alt text on any image you know is used by more than one product. Anything wrong, correct it there once: the plugin never touches alt text a human has written, so your correction is permanent.

Everything stays editable afterwards. Open any item and the **AI / GEO** panel sits right under the editor, next to Yoast or Rank Math.

Step 4: The attribute dictionary

Assistants build comparison tables from attributes. An item with one attribute never makes it into the table, however good its description is.

Most catalogues already state their attributes inside the description, in prose. The dictionary finds them.

Format: one attribute per line, a label, a colon, then the terms to look for.

```
Material: lace, tulle, satin, silk, crepe, beads, sequins
Cut: A-line, mermaid, princess, straight, fitted
Length: mini, midi, short, long, maxi
```

Wildcards. End a term with `*` to capture whatever follows it:

```
Neckline: neckline *
```

This picks up "neckline straight", "neckline V", "neckline sweetheart" without listing every variation. Use it for anything phrased differently on every item.

A wildcard only captures a phrase that reads as an attribute. If the word straight after the term is a verb or a connector, the term was being used in a sentence rather than as a label, and nothing is written. "The table has an oak top" produces no attribute, which is the correct outcome: a fragment in structured data is worse than a gap.

Measurements. A measurement cannot be captured by following a word, because the meaning sits in the numbers and the unit. There are three instructions, one per family, so that the label a value appears under is the label it belongs to.

`#size` reads lengths:

```
Dimensions: #size
```

This picks up "l 80, L 130, h 79 cm", "80x200 cm", "4 mm" and "30 inch" as they are written. Lengths only: centimetres, millimetres, metres and inches.

`#weight` reads weights:

```
Weight: #weight
```

This picks up "500g", "500 g", "1.5 kg" and "250 grame", written as "500 g" and "1.5 kg". The unit you wrote is kept, so a weight in grams is never republished in kilograms.

`#volume` reads volumes:

```
Volume: #volume
```

This picks up "250 ml", "1 l" and "2 litri", written as "250 ml" and "2 l".

If your dictionary says `Dimensions: #size` and your descriptions state weights, add a `Weight: #weight` line: `#size` reads lengths and will not answer with a weight under a heading that says Dimensions.

A number that states a ratio is not a measurement of the product and none of the three read it. "70g protein per 100g" and "12 g/100 ml" describe what is inside the product, not how much of it you get, so they are left alone.

Counts. A quantity is written *before* the noun, so put the star first:

```
Capacity: * seats, * people, * drawers
```

This picks up "6 seats" and "4 people".

Defaults. Add `| default: value` to use a value when the text says nothing:

```
Colour: white, ivory, nude | default: ivory
```

Only do this where it is true for **every** item. A wrong attribute in a comparison is worse than a missing one.

Diacritics are ignored on both sides, so `croiala` matches `croială`.

Starting points. The buttons above the box load a ready-made list for clothing, furniture, electronics, services or food. The built-in default is the clothing one, if you sell something else, load the closest list and edit it. Five minutes here is worth more than anything else in this plugin.

Language. The lists ship in English because the plugin does. The terms are matched against your product descriptions, so write them in the language those descriptions are actually written in. Diacritics are handled either way, a dictionary entry of `croiala` still matches `croială` in the text, and the reverse.

Click **Test the dictionary** to sample your catalogue and see what it finds. It reads whatever is in the box at that moment, so you can try a list before saving it.

Step 5: What else gets published

Three things sit in WooCommerce already and never reach an assistant on their own. Each is a module you can switch off.

Ratings

Your star rating and review count. Nothing is published unless real reviews exist, because an aggregate rating with nothing behind it is markup a validator accepts and a shopper cannot trust.

Options and variations

Sizes, colours and anything else a product varies by, each with its own price and stock. A variable product then publishes a price span rather than one figure.

This is what lets an assistant answer "do you have it in grey". Without it, the answer is silence.

Related items

The upsells and cross-sells you already set in WooCommerce. Your own associations, not our guesses.

Factual comparison, off by default

Places an item beside related ones using only attributes both actually state, printed without judgement. No "better than", no "more durable", nothing you would have to defend if a customer challenged it.

Switch it on only if your related products are genuinely comparable. Two items that share one attribute do not make a comparison, and the plugin skips those.

Buyer questions

The questions panel used to repeat the price and the stock status, both of which the structured data already carried. It now answers what a shopper actually asks: dimensions, materials, capacity, options, delivery, returns, warranty and how the item is rated.

Three rules govern it:

- A question is generated only when there is a real answer. No answer, no question.
- Commercial questions are only asked about things you sell. A blog post never gets a delivery time or a return window.
- Anything you type by hand stays. The panel switches from **Auto** to **Custom** the moment you edit it.

The Google "Preferred Source" button, off by default

In August 2026 Google added preferred sources: a visitor can tell Google that your site is one of their preferred sources, and from then on your content is more likely to appear for that person with a "preferred" badge - in AI Mode, in AI Overviews, and in Top Stories. The choice belongs to the visitor and affects only the visitor who made it.

This module renders a small button that takes the visitor to Google's confirmation page with your domain already filled in. It is the deeplink form Google documents for sites that do not want Google's JavaScript on their pages: a plain link, nothing loaded from Google, no effect on page speed. The one trade-off is that the visitor lands on Google's tool rather than being returned to your page automatically.

Three honest things to know before enabling it:

- Only the people who tap and confirm are affected. A store with no returning visitors gains nothing, and nothing will happen in the first weeks either way.
- Top Stories rarely applies to a shop. The value for a store is the badge in AI Mode and AI Overviews.
- Google publishes no numbers for this, so the plugin cannot show how many people added you, and does not invent a substitute.

Enable it in the Features tab, choose light or dark, then either place the [ai_visibility_preferred_source] shortcode where you want the button, or let the plugin render it centered in the footer of every page. The label defaults to Google's own wording, "Add as Preferred Source"; you can reword it, but the confirmation screen the visitor lands on will still use Google's words.

Only domain-level and subdomain-level sites are eligible. The Diagnostics tab has the check: a link that opens Google's source preferences tool with your domain, and a button to record what you saw there, dated. This is also the first thing the plugin ever renders for

your human visitors - everything else it publishes is for machines. If that matters to how you present your site, it is one more reason the module ships off.

Step 6: The /md/ route

Every item gets a clean text version at:

```
https://yoursite.com/product-name/md/
```

Open one in a browser. You should see plain text: a YAML block with title, price, currency and availability, then the summary, a key-facts table, who it suits, the images described in words, the description, and the questions.

This is what a retrieval system reads best, no menus, no scripts, no layout. The address is announced from every page through a `<link rel="alternate">` tag and an HTTP `Link` header, so crawlers find it on their own.

The same page also answers at `https://yoursite.com/product-name/index.md`, which is the address form the llms.txt specification uses. Both work. The `/md/` form is the one the plugin advertises, because some Nginx and LiteSpeed setups intercept anything ending in `.md` before WordPress ever sees the request.

Your `llms.txt` links to these text versions, not to the normal pages. That is what the specification asks for: the index file stays small enough to fit in an assistant's context, and the detail sits behind links it can read directly.

If it returns 404: go to **Settings → Permalinks** and click **Save Changes** without changing anything. That refreshes WordPress rewrite rules. See also the cache rules below.

Step 7: Check that AI crawlers can actually reach you

AI Visibility → Status → Run the check

This requests one of your own pages five times, identifying as GPTBot, OAI-SearchBot, ChatGPT-User, Claude-SearchBot and PerplexityBot, and reports what came back.

A correct `robots.txt` proves nothing on its own. Hosting firewalls, security plugins and CDN bot rules block AI crawlers silently, with no visible drop in traffic. **Cloudflare now blocks several AI crawlers by default.** This test is how you find out.

What the results mean:

Result	What to do
Reached the page, HTTP 200	Nothing. This is correct.
Blocked, HTTP 403	Refused, not slow. Check your CDN or firewall bot rules. On Cloudflare: Security → Bots → allow AI crawlers.
Could not complete	The request timed out rather than being refused. On small hosting plans a site testing itself runs out of PHP workers. Run it again; if it persists after two or three runs, then look at your firewall.
Response time over 600 ms	Crawlers throttle slow sites and time out after a few seconds. Worth fixing.
Very little text without JavaScript	Your content is being added by JavaScript. No AI crawler runs JavaScript, so that content does not exist for them.
A nosnippet directive was found	Restricting snippets removes the page from Google's AI features. Check your SEO plugin.

You can paste any URL of your own site into the field to test a specific page.

The visit counter

Below the check, the Status tab counts the AI crawler requests your site actually receives: how many in the last 7 days, split by crawler and by surface, and which pages they read most over 30 days. A visit to a `/md/` mirror is the strongest signal there: it means the crawler is reading the clean version this plugin publishes for it.

Verified, and why the column exists

Since 1.7.7 there is a **Verified** column. A user agent is a claim, not proof: anyone can send a request that says "GPTBot" with a single line of curl, and a counter that takes that at face value will happily show you traffic from OpenAI that never happened.

So each request is checked against the crawler owner's own network. The plugin asks the address what hostname it has, then asks that hostname what address it has, and only counts the request as genuine when the answer comes back to the same place. The first half can be faked by whoever controls the address block; the second half cannot, because it is served by the operator's own DNS. This is the method Google, Apple, OpenAI and Amazon each document for their crawlers.

What you see	What it means
A number	That many requests came from the crawler's real network.
3 of 7	Three were genuine; the rest claimed a name they could not prove.
Not checked	Counted before 1.7.7. Verification did not exist yet, so nothing is assumed either way.

Nothing is ever blocked on the strength of this check. A slow or unavailable DNS resolver would otherwise turn into a refused crawl, and losing a real crawler costs far more than counting a fake one. One lookup per address per day is cached, so a crawl of a thousand pages costs a single check.

Two names can never be verified, on purpose. **Google-Extended and Applebot-Extended are not crawlers**; they are robots.txt settings that tell Google and Apple what may be done with content their real crawlers already fetched. No request legitimately arrives under those names, so anything that does is something else wearing the label.

Two honest limits. The count comes from the user agent each request declares, which can be faked, though nobody spoofs GPTBot at scale against a shop. And a full-page cache that answers before WordPress runs is not counted, so on heavily cached sites the real number is higher than the one shown, never lower.

The AI Traffic tab

Everything above answers one half of the question: whether AI can read you. This tab answers the other half, on one screen, in the order that makes sense to somebody who has never heard of structured data.

Is AI reading your shop. The requests AI crawlers actually made, for whatever period you choose, split by which crawler and by which surface: your normal pages, the plain-text copies, llms.txt, the AI sitemap. Under it, the pages they read most.

Who each crawler belongs to. A table naming the company behind every crawler and what it is for, because "Bytespider" means nothing on its own. Two entries carry a warning: Google-Extended and Applebot-Extended are not crawlers at all. They are settings in robots.txt that tell Google and Apple what they may do with pages their real crawlers already took, and nothing ever arrives at a server under those names. So a request claiming to be one of them is something else borrowing the label, usually a scanner. The plugin counts it and says so, because hiding it would hide something worth seeing.

One more thing worth knowing, since it confuses almost everybody: Google AI Overviews and AI Mode have no crawler of their own. They read Google's ordinary index. What matters for them is that Googlebot can reach you.

Is AI sending you visitors. When somebody arrives on your site from ChatGPT, Perplexity, Claude, Copilot or Gemini, that visit is counted here. This happens entirely on your own server: no cookie is set, no IP address is read, no record of an individual visit is stored. Only a daily total per source. That is why there is nothing to add to a consent banner and nothing to explain in a privacy policy, though a line saying you count which sites send you visitors never hurts.

Is your catalogue ready. Three rings showing how much of your catalogue has an AI summary, has attributes, and has buyer questions. These fill in as you run Fill everything.

Other places AI activity shows up. Bing publishes its own report on how Copilot used your pages, and Google shows AI Mode impressions inside Search Console. The Search Console report is real and every site now has it, but Google exposes it to no API at all, so no plugin can pull those numbers in. The button takes you straight there.

Choosing a period, and taking the numbers with you

Above everything sits a date range: the last 7, 14, 30, 90 or 180 days, or any window you pick. It drives every figure on the screen, so the labels never disagree with the charts.

Download spreadsheet saves the period you are looking at, one row per day, with a column for each crawler and each assistant that appeared, and a totals row. A second sheet lists your most read pages and the pages AI visitors land on. It opens in Excel, Numbers and Google Sheets.

History is kept for 180 days, so you can report on a quarter and compare it with the one before. If you have just updated from an earlier version, that history starts building now: days the old 30 day limit already removed cannot be recovered.

The two tables of most read pages are the exception to the date range. They hold a running total and a last seen date per page rather than a day by day history, so they always show everything kept, whatever period you selected. The screen says so under each one.

The Report tab

The second tab, right after Status. Everything on it is read from data the plugin already stored; nothing on this screen re-runs the attribute dictionary to produce a number.

Top row. A half-circle dial with the readable score out of 100 and a tier word (poor, fair, good, strong), and underneath it how many products are complete, partial or thin. Complete means at least four canonical facts were found; partial means one to three; thin means

none. Next to it a 30-day sparkline of that score, with a hollow ring marking the day of the last bulk pass and a filled dot at today. On the right, a donut of the last 7 days of AI crawler requests split into verified, unverified and unknown, with a line stating how many of those requests actually reached a plain-text surface (the mirror or llms.txt) rather than a normal page.

Crawler matrix. One row per crawler: whether robots.txt allows it, whether the last live check reached it, how many requests it made in 7 days, how many of those were verified, and who operates it. The two robots.txt control tokens, Google-Extended and Applebot-Extended, are marked "not a crawler" in the owner column, because no request ever arrives under those names. Under the table, the robots.txt block the plugin would write (or already writes), with a Copy button.

Findings. An ordered list, worst first: blocking (red), needs you (amber), one click (blue), worth knowing (grey). Up to eight rules fire on live data - crawlers being blocked, a live check that has never run, products with zero facts, delivery, returns or warranty left unset, crawlers visiting but never reading a text surface, products never touched by a bulk pass, requests that claimed a crawler name they could not prove, and a full-page cache plugin being active. Each has a "How to fix" link that opens a modal: what it is, what was measured, what to do, and for the two crawler-reachability findings, a Check again button that reruns the same live check as the Status tab.

Dictionary families and exact values. Two panels: how many products carry a fact under each dictionary label, and how many carry each distinct value string, both sorted by count, top 8 shown with a "See all" to expand the rest.

What an assistant can now answer. The product with the most stored facts and the product with the fewest, each showing the summary sentence actually published (quoted) and how many facts it was assembled from. A "Pick another product" list lets you check any of your 50 most recently edited products the same way.

Unanswered questions, categories, since last pass. Which buyer questions still get no answer, with a link to set the three global ones (delivery, returns, warranty) once, or to fix the rest per product. A readable score per category. A three-box strip: products edited since the last bulk pass, the score then and now, and how many crawlers are still blocked.

Modules and weakest products. One card per module, on or off, with one figure that shows it did something (entries in llms.txt, mirror pages, FAQ blocks, alt texts written, comparison tables, and so on). Below it, a table of the weakest products, sorted by sales when WooCommerce reports them, each row linking straight to the editor.

Export and print. "Process N products" jumps to the bulk screen when any are thin or partial. "Print or save as PDF" hides the admin chrome and prints the panels in one column. "Catalogue, XLSX or CSV" downloads one row per product: id, name, category, state, facts found, facts missing, summary present, FAQ count, alt texts written, last edited, edit link. A checkbox schedules a plain-text weekly email to the site's admin address with the score,

the findings titles and the since-last-pass strip; unticking it stops the email. A demo account sees the whole tab but none of the buttons that change or export anything.

Step 8: Validate your structured data

1. Open Google Rich Results Test or the Schema Markup Validator.
2. Paste a product URL.
3. You should see a **Product** with name, description, price, availability and image, plus your attributes under `additionalProperty`.

Target: **0 errors**. Warnings about optional fields are normal and harmless.

If you see **two Product blocks**, another plugin is also outputting product markup. Check your SEO plugin's schema settings for that content type and set it to "Off" or "Article", this plugin handles the product side.

Cache compatibility

Dynamic routes and caching plugins do not always get along. Add these exclusions once and forget about them.

WP Rocket

Settings → WP Rocket → Advanced Rules → Never Cache URL(s)

```
/(.*)/md/?$
/ai-sitemap.xml
/llms.txt
```

LiteSpeed Cache

LiteSpeed Cache → Cache → Excludes → Do Not Cache URIs

```
/md/
/ai-sitemap.xml
/llms.txt
```

LiteSpeed matches a path fragment, so `/md/` covers every item.

Also, after any change to modules: **LiteSpeed Cache → Toolbox → Purge All**.

WP Fastest Cache

WP Fastest Cache → Exclude → Add New Rule, one rule each:

- Prefix: `/ai-sitemap.xml`
- Prefix: `/llms.txt`
- Contains: `/md/`

W3 Total Cache

Performance → Page Cache → Advanced → Never cache the following pages

```
.*/md/?$
ai-sitemap.xml
llms.txt
```

Cloudflare

Two things, both in the dashboard.

1. Cache rules, create a rule with the expression:

```
(http.request.uri.path contains "/md/") or (http.request.uri.path eq "/ai-sitemap.xml") or (http.r
```

Set the action to **Bypass cache**.

2. Bot protection, this is the important one. **Security → Bots → AI Scrapers and Crawlers**: make sure it is not blocking the crawlers you want. Cloudflare's defaults block AI training and agent crawlers on many plans. Run the crawler check in the plugin after any change.

Nginx

Some Nginx setups intercept anything that looks like a static file. This is why the markdown route is `/md/` rather than `/index.md`. If you still get 404s, make sure requests fall through to WordPress:

```
location / {
    try_files $uri $uri/ /index.php?$args;
}
```

Frequently asked questions

Do I still need Yoast or Rank Math? Yes. They handle titles, meta descriptions, canonicals and sitemaps. This plugin handles what AI assistants read. It detects them and merges into their structured data rather than competing with it.

Will this get me into ChatGPT? It removes the reasons you would be skipped: missing structured data, no price, no attributes, unreadable pages, blocked crawlers. Whether you get cited also depends on things off your site, reviews, mentions, third-party pages. Around 85% of brand mentions come from sites you do not own. This plugin fixes your half.

Is llms.txt important? Less than you have been told. Google has publicly stated it does not use llms.txt for Search. Other tools do read it, and it costs nothing, so the plugin publishes one, as one optional module among several, not the headline. What the plugin does do is follow the specification properly: the file links to the plain text version of each page rather than the normal page, and every page points back at it with a `describedby` link relation.

Will it slow my site down? No. Everything heavy runs in the admin or from cache. The public pages get some extra structured data and, if you enable it, a link tag.

What happens if my licence expires? Structured data, the editor panel and robots.txt keep working. The markdown mirror, AI sitemap, llms.txt, bulk fill and image descriptions stop. Your site does not get worse than it was before you bought the plugin.

I have a multilingual site. WPML and Polylang are detected. Fields are per language.

My products are not WooCommerce products. That is supported and common. Modules tab → set your content type to **Product**. Everything else works the same.

The AI summary starts with the price. Can I change it? Business tab → turn off "Price in the AI summary". You can also edit any summary by hand; the plugin stops managing a field the moment you type in it.

Where do I report a problem? Include: WordPress version, PHP version, which SEO plugin you use, and a screenshot of **AI Visibility → Status**. That is almost always enough to find the cause.

What leaves your site

Three outbound connections, all of them necessary, none of them tracking you.

To	What	When
Lemon Squeezy	Your licence key and your domain name	On activation, and once every 30 days to re-check
api.indexnow.org	The address of a page you just published or changed	Only if the IndexNow module is on
Your own site	A request to one of your own pages, as each AI crawler	Only when you click "Run the check"

No analytics, no usage tracking, no content ever sent anywhere.

One thing worth stating plainly: the bulk pass **writes alt text into your media library**. That is a real change to your data, it is the point of the feature, and it is not reversed when you uninstall. Existing alt text is never touched.

A short honest note

Two things are worth stating plainly, because plenty of tools in this space will not.

First, AI answers are unstable. Repeated identical prompts return different brands more often than not. There is no "position 1" to win, only the probability of being included, which you raise by being the clearest, most complete, most machine-readable source in your category.

Second, most of what decides whether an assistant recommends you happens off your website. Reviews, forums, comparison articles, your presence on the platforms these models trust. This plugin makes your own site as good as it can be. That is necessary, and it is not sufficient.

Release notes

1.7.9 The Report tab. One screen that puts the findings first: a readiness score with its history, a crawler matrix, the rules that flag what is blocking, what needs a setting, what is one click away and what is only worth knowing, each with the fix beside it, dictionary and value coverage, the product an assistant can answer best and the one it cannot answer at all, unanswered questions, scores by category, module status, the weakest products, export to XLSX or CSV, print to PDF and an optional weekly digest. Every figure comes from data the plugin already stored; nothing is extracted again to draw the screen.

Only products you have actually published reach an AI surface. Drafts, private, scheduled, trashed and password protected products are never published, in any surface. The

WooCommerce edge cases are yours to decide, on the Modules tab: products hidden from the catalog (off by default), products set to search results only (on by default) and out of stock products (on by default, published with their real availability). Changing a product's visibility, or deleting it for good, clears the cached files at once rather than leaving a stale answer for hours.

1.7.8 Crawler verification learned a second method. Reverse DNS works for Google, Apple, Amazon and Bing, but OpenAI's crawlers carry no reverse DNS record at all, so GPTBot, OAI-SearchBot and ChatGPT-User came back "unknown" on every real shop, which reads as a feature that does not work. OpenAI and Perplexity publish the address ranges their crawlers use instead, as five public files; the plugin reads them at most once a day, caches them for 24 hours, keeps the previous list when a read fails, and matches the visiting address against them when there is no reverse DNS to check. The order is reverse DNS first when a record exists, published ranges when it does not, and "unknown" when the operator publishes neither.

Visits from assistants are now also recognised by the utm_source tag on the address a visitor lands on, not only by the Referer header. ChatGPT adds utm_source=chatgpt.com to the links it hands out and frequently sends no Referer at all, so the panel could read zero on a shop that was being sent real traffic. The screen now says how many visits arrived with each signal, because the two are not equally strong evidence and you should be able to see the split.

The delivery answer was the one place in the plugin that stated something the page did not support. It joined up to three settings with full stops, so a shop with "delivery time varies by product" ticked published "the delivery time is stated on each product page" on pages that stated nothing. It is now one sentence built from what you filled in, and when the time varies and the product carries no delivery value of its own, the delivery question is left out entirely. Business type became a list of real schema.org types instead of a free text box that published whatever was typed as a type; delivery time became two day figures and delivery cost a free tick or a figure with your currency, because both were prose being published as structured values. Anything you had typed before is read where it can be read, kept where it cannot, and shown back in the field so you can enter it properly. Nothing is guessed.

The plain text mirror, the extracted attributes and the extra structured data are now cached per page until that page is edited. The mirror was the only surface with no cache, and it is the one a serious crawl asks for hundreds of times. The visible "See the plain text version" link now also renders on block themes such as Twenty Twenty-Five, where the WooCommerce hook it used never fires; it appears exactly once whichever path places it. Every crawler counted on the AI Traffic screen can now also be allowed or blocked on the AI Crawlers screen, DeepSeekBot and DuckAssistBot included, and the two robots.txt control tokens are labelled rather than presented as crawlers.

The Title Match section was removed from the Status tab. It compared a page title against questions this plugin generated from that same title, so the only words it could ever report as missing were the words in the question templates. The vendor master key was rotated: the previous one had been committed to this project's own repository, and although that repository was never published, it is treated as exposed.

1.7.7 Every AI crawler request is now verified before it is counted. A user agent string is a claim, and anyone can send "GPTBot" from a laptop with one line of curl; until this version, that request was stored and shown to you as a visit from OpenAI. The plugin now asks the address what its hostname is, then asks that hostname what its address is, and only counts the request as genuine when the answer comes back to the same place. That second half is controlled by the crawler's operator, so a forged record fails it. Requests that fail are still served and still counted, but shown separately as not verified, because losing a real crawler would cost far more than counting a fake one. One lookup per address per day is cached, so a crawl of a thousand pages costs a single check, and on a host where lookups are unavailable the answer is "unknown" rather than a false "verified". Requests counted before this version carry no verdict and are shown as not checked. Google-Extended and Applebot-Extended are deliberately absent from the list of verifiable crawlers: they are robots.txt settings, not crawlers, so a request arriving under those names is by definition not what it claims. This release also restores automatic updates: while your licence is active, new versions appear in the WordPress updates screen like any other plugin, and lifetime keys are included.

1.7.6 Fixed a fault present in every earlier version: wherever HTML markup was the only thing separating two words, stripping the markup fused them. A heading broken over two lines produced "storereadable", and two figures in adjacent table cells produced "5AI crawlers". This corrupted the plain text version of a page, and more seriously it corrupted attribute extraction, because a term fused to its neighbour no longer matches the dictionary and the fact was silently missed. A price written across two cells could also be read as its second half only, so a product priced 1800 could publish as 800. Text containing ">" inside an HTML attribute, which page builders produce routinely, is no longer leaked into the published page. The "See the plain text version" link added in 1.7.5 now has its switch on the Modules tab; before this it could be stored but not reached.

1.7.5 Dimensions written as "180 by 90 by 75 cm" are recognised and normalised to the same "x" form as the rest. Diagnostics now flags WooCommerce "Coming soon" mode and the WordPress "Discourage search engines" setting, either of which hides the whole shop from every AI crawler while every other check still passes. A new optional link, off by default, lets a visitor open the plain text version of the page an assistant reads.

1.7.4 The AI Traffic tab: one screen showing whether AI read your shop, whether it sent you visitors, whether your catalogue is ready, and where else AI activity appears. Visits arriving from ChatGPT, Perplexity, Claude, Copilot and Gemini are now counted on your own server, with no cookie, no IP address and no per-visit record. A date range from 7 to 180 days drives every figure, and the selected range downloads as a spreadsheet. History extended

from 30 to 180 days, building up from the day you update. Every crawler now states who owns it and what it does, with Google-Extended and Applebot-Extended marked as robots.txt settings rather than crawlers. The whole admin redesigned for one consistent look.

1.7.3 The Google "Add as Preferred Source" button, as a plain link with no Google script. A visitor who taps it and confirms sees your content more likely badged "preferred" in AI Mode and AI Overviews - that visitor only. Off by default, because it is the first element the plugin renders for the human visitor rather than for machines. Enable it under Features, place it with the [ai_visibility_preferred_source] shortcode or automatically in the footer, and record the eligibility check in Diagnostics. Expect nothing in the first weeks: the mechanism depends on returning visitors choosing you.

1.7.2 llms.txt now links to the plain text version of each page instead of the normal page, which is what the specification asks for. Every page also points at llms.txt with the standard describedby relation. Both address forms for the text version keep working.

1.7.1 The crawler check no longer inflates the visit counter. Its own requests are marked and excluded, so the counter only ever shows genuine crawler traffic.

1.7.0 Official profiles published as sameAs, in llms.txt and on the markdown mirrors. AI crawler visits counted and shown in the Status tab, including the pages read most. Shared images skipped in alt autofill. Filename-like image titles never used as alt text. "L 130, l 80" reads as length and width, not twice the same letter. The overwrite checkbox now says exactly what it replaces.

1.6.9 Commercial questions are now only asked about things you actually sell. On a site that also runs a blog, an article could be given a delivery time and a return window, which mean nothing for an article.

1.6.8 Catalogues built on a custom post type now get buyer questions too. They were being skipped entirely. If your shop is a custom post type, run the bulk pass again with overwriting on.

1.6.7 Buyer questions, ratings, options and variations, delivery and returns, related items, factual comparison, and canonical attribute keys. Stock changes made by an order now refresh the text mirror.

1.6.6 Attribute extraction rewritten. A term followed by a verb or a connector is now skipped instead of published, so descriptions written as normal prose stop producing fragments like "table has the top". Two new forms are available in the dictionary: `#size` reads measurements out of the text directly, so "l 80, L 130, h 79 cm" is captured as written, and `* term` reads a count written before the word, as in "6 chairs" or "4 people". Terms that differ only by diacritics are listed once. Price is no longer repeated as an attribute, since it already appears in the offer.

If your dictionary lists dimensions as `length *, width *`, replace that line with `Dimensions: #size`. If it lists counts as `seats *`, write `* seats` instead.

1.6.5 Comparable attributes now reach the Product node even when Yoast or Rank Math builds it. Previously, on sites where an SEO plugin owned the product schema, WooCommerce never ran its own structured data filter, so `additionalProperty`, brand, GTIN and MPN were left out. Anything your SEO plugin already sets is still left untouched. If you run Yoast or Rank Math, update and re-check one product in validator.schema.org.

1.6.4 The dictionary test now reads the text in the box rather than the last saved version, so loading a preset and pressing Test gives a report on what you are actually looking at.

1.6.3 Interface and presets fully in English.

1.6.2 The dictionary test works before a licence is activated.

1.6.1 Guarded a call that could fail on hosts without the mbstring extension.