AI Navigator

Administrator Guide

The complete reference for configuring, operating, and maintaining AI Navigator

v2.3.0 | speak2web | 2026 | speak2web.com

ABOUT THIS GUIDE

This guide covers every setting, feature, and maintenance task in AI Navigator v2.3.0. It is written for the WordPress administrator responsible for running the plugin on a live site. If you are installing AI Navigator for the first time, start with the Quick Start Guide — it gets you live in five required steps — and come back here when you want the full picture.

New in v2.3.0: WooCommerce product indexing (Section 3), indexing that completes even on hosts without WP-Cron (Section 3), ask-first navigation (Section 10), per-visitor usage limits that protect your OpenAI credits (Section 11), and a conversation History tab in Analytics (Section 14). Appendix C has the full list.

Section	What it covers
1. Installation	Requirements, uploading, activating, what activation creates
2. API Key Setup	Entering your OpenAI key — the widget does not appear without it
3. Content Indexing	How indexing works, WooCommerce products, chunking, cron-less hosts, free tier
4. Index Viewer	The AI Navigator → Index page — search, summaries, per-page re-index
5. Page Roles	Role overrides in the post editor — the highest-leverage accuracy tool
6. AI Models (Pro)	Embedding, chat, and voice models and what they cost
7. Pinecone (Pro)	Vector database setup, wipe-before-sync, dimension rules
8. Widget Appearance & Engagement	Position, colors, greeting bubble, starter questions
9. Voice & Language	TTS, STT, Mic Mode, multilingual behavior, WPML/Polylang
10. Navigation Behavior	Ask-first navigation, the aggressiveness slider, page suggestions
11. Usage Limits & Abuse Protection	Per-visitor rate limits and input caps on AI calls
12. Data Retention & Privacy	Retention policy, right to erasure, PII scrubbing, cookies

13. License (Pro)	Activating and managing your Pro license
14. Analytics (Pro)	Four-tab dashboard including the new History tab
15. Maintenance & Updates	Upgrade checklist, re-index triggers, shared updater
16. Uninstall & Data Cleanup	Exactly what deletion removes
17. Troubleshooting	Symptoms, causes, and fixes

1. INSTALLATION

1.1 System requirements

Requirement	Details
WordPress	6.0 or later
PHP	8.0 or later
MySQL	5.7 or later
OpenAI account	API key required — see Section 2. The widget does not appear without a valid key.
Browser (admin)	Any modern browser. Keep the admin tab open during a full index — see Section 3.4.
Browser (visitors)	All modern browsers, including iOS Safari and Android Chrome.
Pinecone (optional)	Pro only. Recommended for sites with 200+ pages or when semantic chunking is needed.
WooCommerce (optional)	Any active WooCommerce install unlocks the product indexing option (Section 3.2).
Companion plugins	AI Concierge and AI Gateway are optional. If installed, all S2W plugins must run the same infrastructure version (S2W_INFRA_VERSION 1_2 in this release) — see Section 15.1.

1.2 Installation steps

- Log in to your WordPress admin dashboard.
- Go to Plugins → Add New → Upload Plugin.
- Click Choose File and select the s2w-ai-navigator.zip file.
- Click Install Now and wait for the upload to complete.
- Click Activate Plugin.
- AI Navigator appears as a top-level menu in the admin sidebar with four subpages: Settings, Index, Privacy, and Analytics (the Analytics page appears once a Pro license is active).

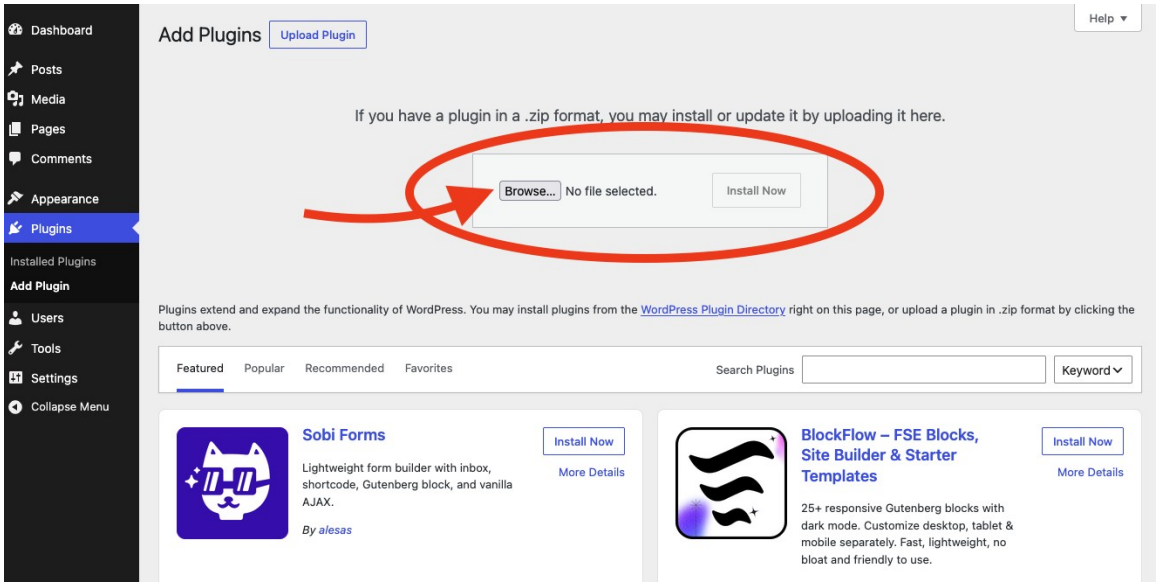

Plugins → Add New → Upload Plugin

After activation The plugin creates three database tables: s2w_ain_index (page vectors and metadata), s2w_ain_chunks (chunk-level vectors for Pinecone, Pro), and s2w_ain_conversations (conversation log). Sites upgrading from a legacy install are migrated automatically, provided the old plugin is v2.1.17 or later — an admin notice explains what to do if it is older.

2. API KEY SETUP REQUIRED

No API key = no widget AI Navigator needs a valid OpenAI API key to function. Until one is saved, the chat widget does not appear anywhere on your site — no button, no panel — and a persistent red notice on every admin page reminds you to add the key.

2.1 Getting a key

- Go to platform.openai.com and sign in or create an account.
- Open API Keys in the left sidebar.
- Click Create new secret key and copy the key (it starts with sk-).
- Store it securely — OpenAI shows it only once.

API key cost AI Navigator uses your own OpenAI account. Typical cost for a small site (20–50 pages, moderate traffic) is $1–$3 per month. OpenAI bills you directly — speak2web adds no markup. Per-visitor usage limits (Section 11) keep a single visitor from driving the bill up.

2.2 Entering and verifying the key

Go to AI Navigator → Settings → General tab. Paste the key into the OpenAI API Key field, click Verify Key to confirm it works, and click Save Settings. The Index All Pages Now button stays disabled until a key is

saved — the tooltip says so if you hover it. One key handles everything: indexing, answering, and optional voice.

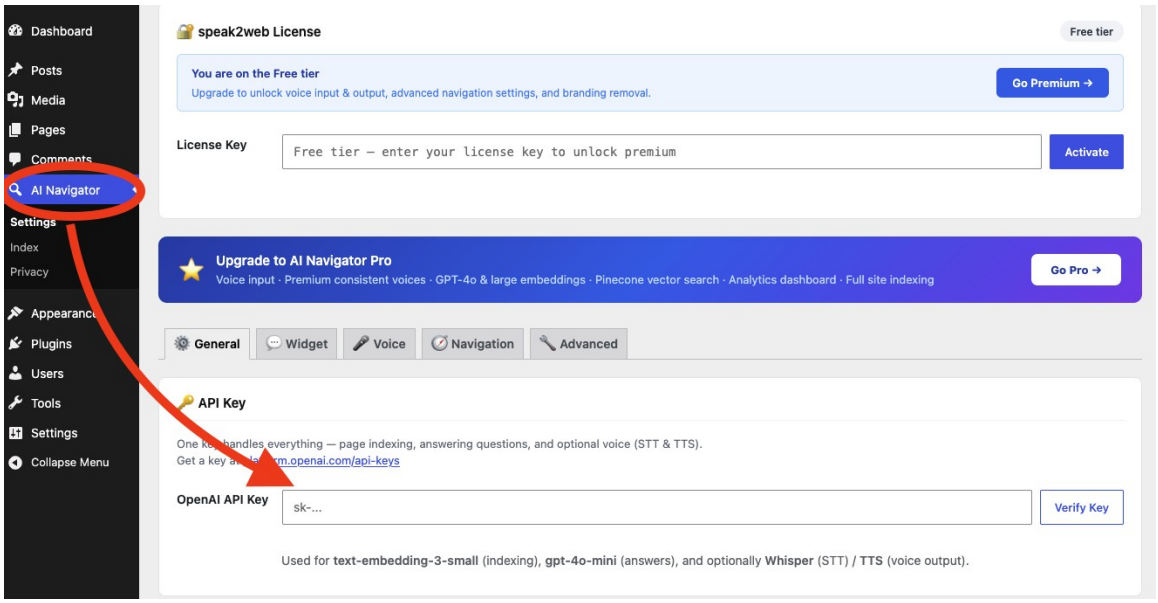

Settings → General tab — API Key field with verification

3. CONTENT INDEXING

3.1 How indexing works

When you click Index All Pages Now, AI Navigator reads every published item of the content types you selected. For each page it:

- Extracts the title and content, converting block HTML to plain text while preserving paragraph boundaries.
- Classifies the page role (contact, product, pricing, and so on) using structural rules, with a lightweight LLM call for ambiguous pages.
- Sends the content to the OpenAI Embedding API and stores the returned vector, role, authority score, and metadata in the s2w_ain_index table.
- For long pages on Pro with Pinecone: splits the content into overlapping semantic chunks and embeds each chunk separately (Section 3.3).
- If Pinecone is configured, syncs vectors to Pinecone after the full index completes — a full re-index wipes Pinecone first, so no stale vectors survive.

Automatic re-indexing After the first manual index you rarely need the button again. Every time a post or page is saved, AI Navigator re-indexes that page in the background, and unpublished or trashed pages leave the index immediately. The index stays current on its own.

Settings changes lock the button The Index All Pages Now button disables itself the moment you edit any settings field, with the note "Save your changes before indexing." Indexing always runs against saved settings, never the live form — click Save Settings first.

3.2 Choosing content types

Under Content Indexing on the General tab, check the types AI Navigator should read. Pages and Posts are the defaults; every public custom post type appears here too, and your selection persists exactly as saved. Enable only the types visitors would want to be navigated to.

Index WooCommerce Products. When WooCommerce is active, a dedicated checkbox adds the product catalogue to the index — off by default. Each product is embedded with its price and currency, sale price, stock status, categories, SKU, attributes, and a variations summary, so the widget can answer questions like "do you have this in blue?" and send the visitor to the right product page. Products are automatically classified with the product role and a small authority boost so they outrank blog posts that merely mention them. Run a re-index after enabling.

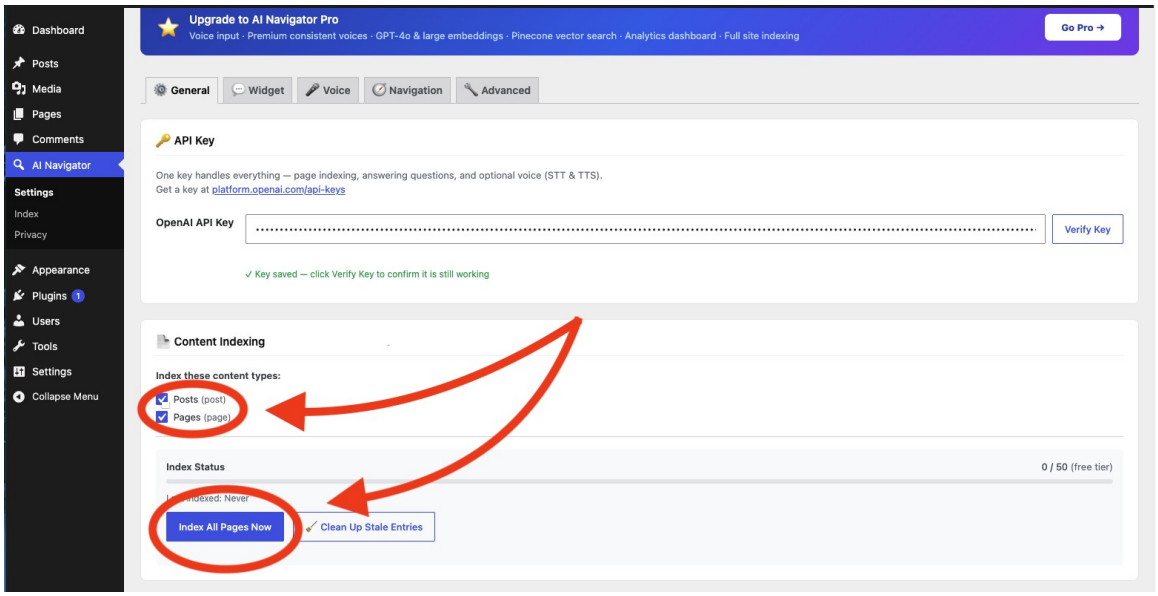

Settings → General tab — Content Indexing and Index Status

3.3 Semantic chunking for long pages (Pro + Pinecone)

A single vector represents a whole page well when the page is short. On long pages the vector captures the general topic but loses the details buried further down — a visitor asking about something in paragraph 8 of a 2,000-word page can get a low-confidence answer even though the answer is right there. Chunking fixes this: with a Pro license and Pinecone configured, long pages are split into overlapping chunks and each chunk is embedded separately.

Parameter	Value
Target chunk size	350 words
Overlap	50 words — the end of one chunk repeats at the start of the next
Minimum chunk	50 words — short trailing pieces merge into the previous chunk
Split points	Paragraph boundaries only. Never mid-sentence.
Short pages	Pages under 350 words keep a single page-level embedding — no extra API call
Storage	Chunk rows live in s2w_ain_chunks; each has its own embedding and Pinecone

	vector

Chunking requires Pinecone MySQL search (free tier, and Pro without Pinecone) always uses the page-level embedding. Chunking benefits Pinecone-backed search only.

3.4 Indexing on hosts without WP-Cron

Some hosts disable WordPress cron (DISABLE_WP_CRON), which used to leave background indexing stranded. In v2.3.0 the admin page itself drives the index job: while indexing runs, your browser advances the queue batch by batch alongside whatever server cron exists — on a healthy host the two cooperate, and on a cron-less host your tab completes the job alone. The status area shows "Keep this tab open until indexing finishes," and on cron-less hosts it adds that pages are being indexed right in your browser. If you try to leave mid-job, the browser warns you; leaving pauses the job, and re-opening the page resumes it. No configuration is needed either way.

3.5 Stale entry cleanup

A stale entry is an index row for a page that no longer exists or is no longer published. AI Navigator removes stale entries automatically when a page is trashed, drafted, or made private, and again at the end of every full index run. The Clean Up Stale Entries button next to Index All Pages Now removes orphans on demand without a full re-index.

3.6 Free tier limit

The free tier indexes up to 50 pages, and the counter reads "X / 50 (free tier)" in the Index Status block. Pages skipped over the limit are reported in a banner. Automatic re-indexing of pages already in the index never counts against it. To index more, upgrade to Pro.

4. INDEX VIEWER

AI Navigator → Index lists every indexed page with full detail, 25 rows per page, with a search box for filtering. Each row offers Re-index this page (refresh one page on demand) and View full summary (the stored answer text the widget draws on).

Column	What it shows
Title / URL	The page title (linked for editing) and its public URL.
Role	The detected or overridden page role (Section 5).
Words	Word count of the indexed content.
Chunks	Number of chunk vectors. 1 = single-vector page; more = chunked (Pro + Pinecone). Long pages showing 1 chunk were indexed before chunking — re-index them to benefit.
Last Indexed	When the page was last embedded.
Pinecone	Whether the page's vectors are synced to Pinecone.

Stored Summary	An excerpt of the stored summary; View full summary shows all of it.

5. PAGE ROLES

The highest-leverage five minutes in this guide AI Navigator classifies every page's role automatically, and for most pages it gets it right. When it doesn't, the override lives in the post editor sidebar — not in the plugin settings — and many administrators never find it unaided. If the widget sends visitors to the wrong page, fix the role first.

5.1 Overriding a role

- Open the page or post in the WordPress editor.
- Find the "AI Navigator — Page Role" box in the right-hand sidebar. It shows the detected role and how it was detected (rule or LLM).
- Pick the correct role from the Override role dropdown.
- Check Mark as canonical destination if this is the definitive page for its role — your main Contact page rather than a blog post that mentions contact details.
- Save the post. The change applies immediately; no re-index needed.

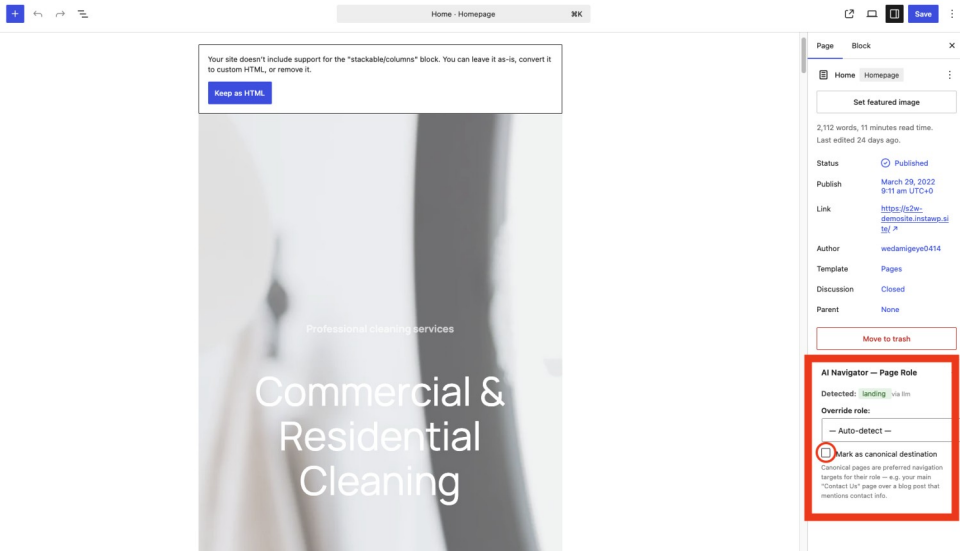

Post editor sidebar — AI Navigator Page Role box

5.2 Why roles matter

Before searching, AI Navigator classifies what the visitor wants — contact, pricing, a product. Ranking then blends semantic similarity (the dominant factor) with how well each page's role matches that intent, plus a small authority weight. Since v2.3.0, the role boost only applies when the page is also semantically related to the question, so a Contact page no longer outranks genuinely relevant content on unrelated queries. Setting correct roles on your five to ten most important pages is still the single most effective accuracy improvement available. The full role list is in Appendix B.

6. AI MODELS PRO

Setting	Details
Embedding model	Free: text-embedding-3-small (1536 dims). Pro: text-embedding-3-large (3072 dims) for higher-quality matching. Changing it requires a full re-index.
Chat model	Free: gpt-4o-mini. Pro: gpt-4o for more nuanced answers and better multilingual output. No re-index needed.
Intent & translation	Always gpt-4o-mini — fast and accurate enough for classification. Not configurable.
TTS (Pro)	OpenAI tts-1 with nine voices: Alloy (default), Echo, Fable, Onyx, Nova, Shimmer, Ash, Coral, Sage. Free tier uses browser TTS.
STT (Pro)	OpenAI whisper-1, 50+ languages. Not configurable.

On the free tier the model dropdowns are locked, and the lock is enforced server-side — the free models apply even if the form is manipulated.

7. PINECONE PRO

7.1 When to use it

Default search runs on MySQL cosine similarity — reliable, no external account. Pinecone becomes worthwhile for sites with 200+ indexed pages or heavy concurrent traffic, where it answers in milliseconds, and it is required for semantic chunking (Section 3.3). The Advanced tab nudges you when your index size crosses 200 and again at 1,000 pages.

7.2 Setup

- Create a Pinecone account at pinecone.io — your account, your data.
- Create an index with dimension 1536 (small embedding model) or 3072 (large model), metric cosine. The Advanced tab suggests an index name based on your domain.
- Enter your Pinecone API key on Settings → Advanced and click Test Connection.
- Click the Sync to Pinecone button. The sync wipes existing Pinecone vectors first, then pushes everything current from MySQL.
- Check AI Navigator → Index afterwards — long pages should show more than one chunk.

Dimension must match the embedding model Switching from the small model (1536) to the large model (3072) means creating a new Pinecone index with the right dimension and running a full re-index. Mismatched dimensions fail on every sync.

7.3 Wipe-before-sync

Every full re-index that reaches Pinecone deletes all vectors there first, then rebuilds from MySQL — stale vectors cannot survive, even for pages that were rewritten rather than deleted. MySQL remains the source of truth; Pinecone is an acceleration layer. If a wipe fails mid-run (network error, rate limit), matching vectors

are still overwritten, the event is logged under WP_DEBUG, and a manual sync resolves any leftover inconsistency. If Pinecone is unreachable at query time, search falls back to MySQL automatically.

8. WIDGET APPEARANCE & ENGAGEMENT

Everything visual lives on Settings → Widget.

Setting	Details
Chat Title	Header text in the chat panel. Default: "Simon AI assistance" — change it to match your brand.
Input Placeholder	Text in the empty input field. Default: "Type your question..."
Accent Color	Button, header, and link color. Any hex value. Default: #2563EB.
Widget Position	Bottom Right (default), Bottom Left, Bottom Center, or Top Center.
Powered-by credit	"Show “Powered by speak2web” in chat footer" — off by default.
Proactive greeting	A speech bubble above the widget button after a configurable delay (default 3 seconds). Default text: "Hi! How can I help you today?" Dismissed the moment the visitor opens the widget or taps the mic.
Starter questions	Preset questions visitors can tap to fill the input — on by default with three examples. They translate automatically in multilingual mode.
Mic pulse animation	A pulse on the microphone icon to signal voice input. On by default.

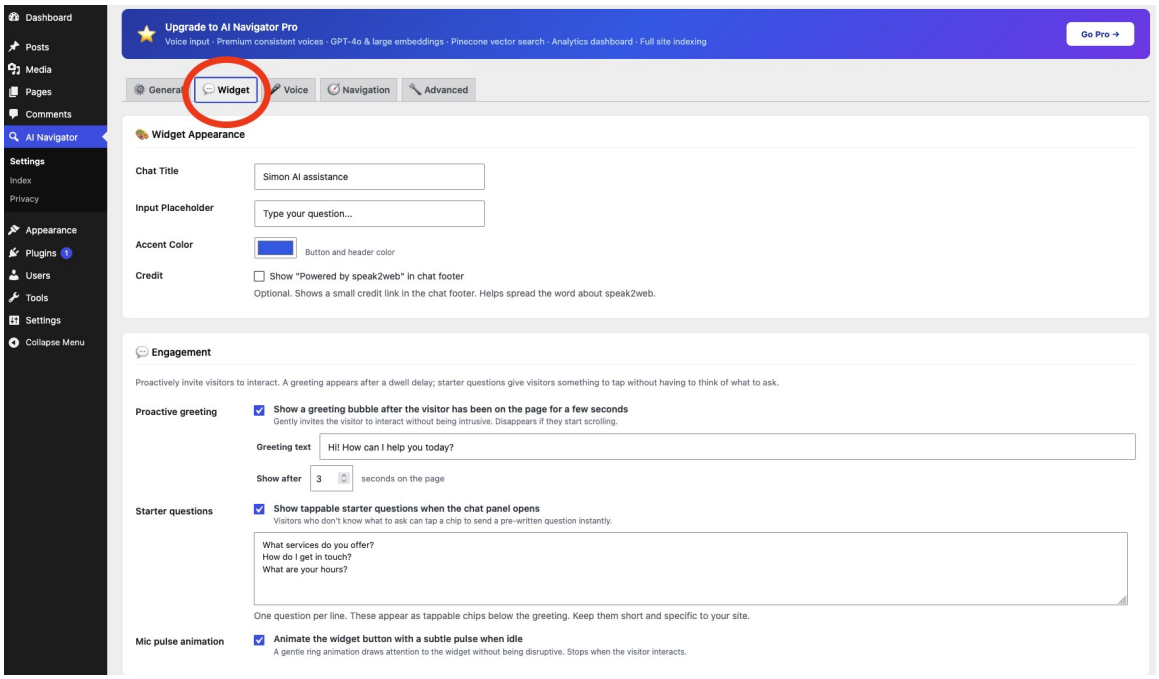

Settings → Widget tab — appearance and engagement

9. VOICE & LANGUAGE

9.1 Text-to-speech

Free tier: the browser's built-in speech synthesis reads answers aloud, with a gender preference (Female default, Male, Neutral) that picks the best available voice in the visitor's language. No API cost; quality varies by browser. Pro: OpenAI TTS with the nine voices listed in Section 6 — consistent quality everywhere, billed to your OpenAI key at roughly $0.015 per 1,000 characters. Pro sites can also prefer browser TTS when available to save API cost, optionally naming a preferred browser voice.

9.2 Speech-to-text (Pro)

The microphone button records the visitor, transcribes the audio with Whisper in the visitor's selected language, and submits the transcript as a query. Recordings stop automatically at the Max voice recording length (5–60 seconds, default 15). Uploads are capped at 5 MB and checked against an audio-type allowlist before they reach OpenAI (Section 11). Mic Mode replaces the chat icon with a mic-plus-pencil pair for hands-free use — tapping the mic starts listening in one step, which is a meaningful accessibility win on mobile.

9.3 Multilingual mode

On by default and server-side — this is not browser translation. The visitor's language is detected from their browser and any hreflang tags (WPML and Polylang emit these, so both are picked up automatically). The greeting and starter questions are translated before display, a flag button in the widget header lets the visitor switch languages, and every answer — and voice transcription — follows the selected language. To lock the widget to one language, uncheck Multilingual mode on the Voice tab and set the Language dropdown.

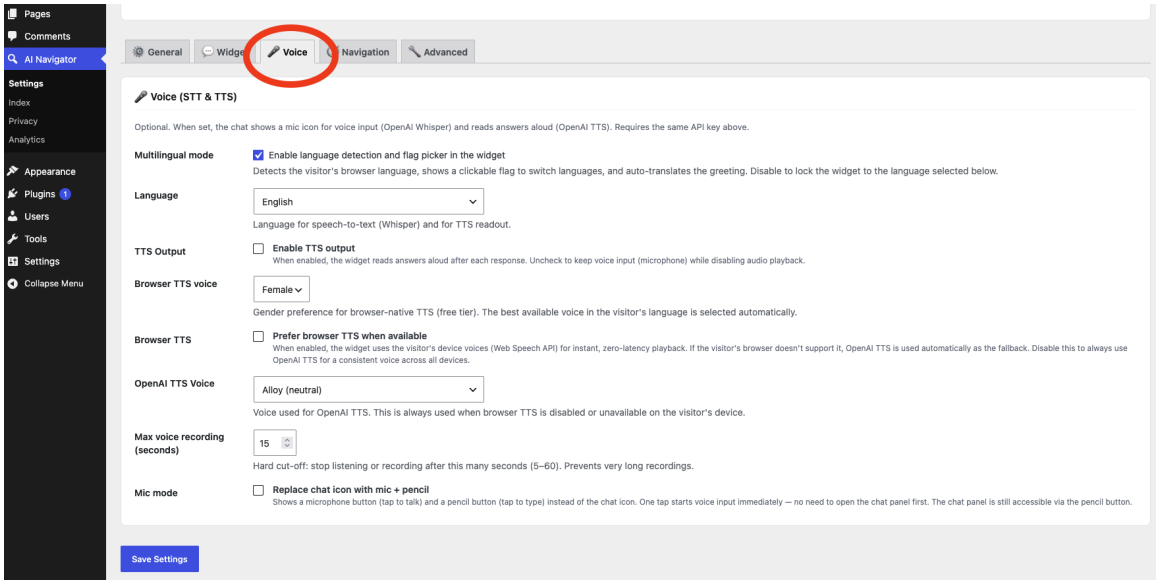

Settings → Voice tab — language, TTS, and Mic Mode

10. NAVIGATION BEHAVIOR

Navigation has its own settings tab, and v2.3.0 changes the philosophy: **navigation is now ask-first**. A strong content match no longer moves the visitor on its own. The widget navigates automatically only when the visitor asked to be taken somewhere, or when you have explicitly opted into auto-navigation with the

aggressiveness slider — otherwise it offers the page and waits for a confirming tap. A casual on-topic question never ejects the visitor from the conversation.

10.1 What happens on a match

Match	Visitor asked to go there	Visitor just asked a question
Strong match	Navigates automatically (after a short pause to read the answer).	Aggressiveness 3+: navigates automatically. Below 3: offers the page with a confirm button.
Moderate match	Shows a "Take me there" button.	Offers the page with a confirm button.
Weak / no match	Answers in chat only.	Answers in chat only.

10.2 The controls

Setting	Details
Navigation aggressiveness	A 1–5 slider from Cautious to Aggressive, default 3. Levels 3 and up allow auto-navigation on strong matches without an explicit ask; levels 1–2 always ask first. Level 1 never navigates automatically at all.
Page suggestions	When to show related page links under answers: Never · Low confidence only · Medium & low confidence (default, recommended) · Always.

Navigation always happens in the visitor's current tab. Where the browser supports it, AI Navigator appends a text fragment so the destination scrolls to and highlights the relevant passage.

11. USAGE LIMITS & ABUSE PROTECTION NEW IN 2.3.0

Every AI-backed endpoint the widget exposes to visitors is rate-limited per visitor (by IP address) in fixed five-minute windows, with hard caps on input sizes. This stops a single visitor — or a bot hammering your site — from driving up your OpenAI usage. Administrators are exempt, nothing needs configuring, and legitimate visitors will not notice: the limits sit far above normal conversational use.

Endpoint	Limit per visitor	Input caps
Chat	30 requests / 5 minutes	Messages up to 4,000 characters; conversation history up to 6 turns
Speech-to-text (Pro)	20 requests / 5 minutes	Audio up to 5 MB, common audio formats only
Text-to-speech (Pro)	30 requests / 5 minutes	Text up to 4,096 characters
Greeting translation	20 requests / 5 minutes	Greeting up to 2,000 characters; up to 12 starter questions

A visitor who exceeds a window gets a polite "too many requests" response and can continue once the window resets. v2.3.0 also hardens the widget itself: it validates the origin of browser postMessage events and sanitizes the configured accent color before rendering.

12. DATA RETENTION & PRIVACY

Conversation data is collected on all tiers — free sites accumulate history that becomes visible in Analytics the moment they upgrade. The Privacy page (AI Navigator → Privacy) is available on every tier and is the hub for compliance work.

Control	Details
Retention policy	"Auto-delete conversations older than": 30, 60, 90 days (default, recommended), 6 months, 1 year, or Never. A daily job prunes older records. Set it on the Privacy page or Settings → Advanced.
What is stored	The "What AI Navigator Stores" card inventories every data point, including the note that queries pass through OpenAI under its Data Processing Agreement.
PII scrubbing	Emails, phone numbers, SSNs, card numbers, UK NI numbers, SINs, and passport numbers are redacted from queries before they are written to the log.
Right to erasure	Erase by session ID — enter the visitor's session ID and click Erase This Session to honor a GDPR Article 17 request without touching other records.
Bulk purge	Delete Records Before This Date removes everything older than a chosen date. Irreversible.
Disclosure template	A ready-to-paste privacy policy paragraph describing what the plugin stores.

Conversations are keyed by an anonymous session ID held in a 30-day cookie (`s2w_ain_session`) — a random token with no personal data and no link to WordPress accounts. Include it in your cookie disclosure; the Privacy page template helps.

13. LICENSE PRO

The speak2web License card sits at the top of the Settings page, above the tabs. Enter your license key and click Activate — the key is validated against the speak2web license server and Pro features unlock immediately. The badge beside the field reports the state: Active, Expired, Limit reached, Invalid, or Free tier. A weekly background check keeps the status current, and Deactivate releases the key for use elsewhere. For a license transfer to a new domain, contact speak2web support.

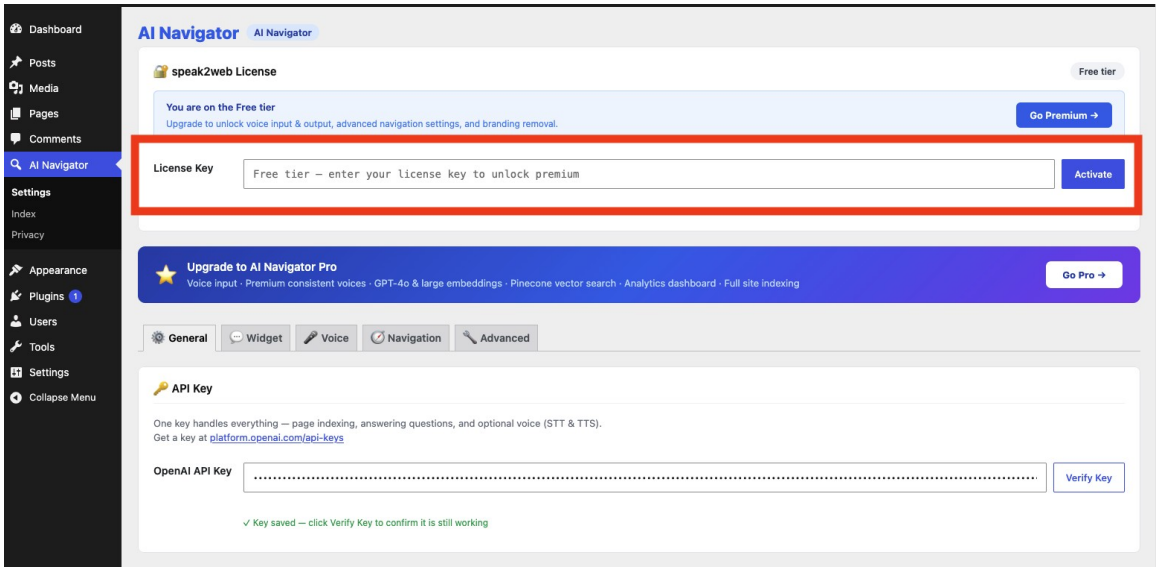

Settings — the License card above the tabs

14. ANALYTICS PRO

AI Navigator → Analytics opens with six at-a-glance cards — Conversations, Sessions, Nav Success Rate, No-Match Rate, Avg Latency, and Estimated API Cost — a date-range selector, and four tabs.

Tab	What it answers
Conversations	What are visitors asking? Daily volume, the confidence distribution, top queries with their matched pages, and top navigated pages. Every row of the top-queries table is a visitor telling you what they need.
Performance	Is the system healthy? Latency and its trend, high-confidence navigation rate, no-match rate, token usage and estimated cost, and the active configuration. Latency above 4 seconds meaningfully hurts the experience.
Gaps & Insights	What is your site missing? Queries with no confident match sorted by frequency — the questions your content cannot answer — plus repeated questions and medium/low-confidence answers that point at pages needing work.
History	New in v2.3.0: the raw conversation log. Each turn shows date, session, query, confidence, the page navigated to, and the response — click Show to expand the full answer. Filter by confidence level; 25 turns per page. Queries appear PII-scrubbed because scrubbing happens at write time.

AI Optimization Export

The Export for AI Analysis button builds a structured plain-text report of all analytics data with a pre-written analysis prompt. Copy it to the clipboard or download it, paste it into Claude, ChatGPT, or Gemini, and you get back a prioritized content roadmap grounded in what visitors actually asked.

15. MAINTENANCE & UPDATES

15.1 The shared updater

AI Navigator, AI Concierge, and AI Gateway share versioned update infrastructure (S2W_INFRA_VERSION 1_2 in this release) covering update checks, license verification, and telemetry. When the infrastructure version bumps, all installed S2W plugins must update in the same release window — mixing old and new infra in one PHP request causes "cannot redeclare class" fatals. Updates delivered through speak2web handle this automatically; the one rule is not to hold back one plugin while updating the others. A Check for updates link on the Plugins screen forces a fresh check on demand.

15.2 Upgrade checklist

- Read the changelog in readme.txt before upgrading.
- Back up the database — the tables that matter are s2w_ain_index, s2w_ain_chunks, and s2w_ain_conversations.
- Upload and activate the new version. Database migrations run on the first admin page load.
- Open AI Navigator → Index and re-index any long pages still showing one chunk.
- If Pinecone is configured, run a sync so chunk vectors are current.
- Clear your browser cache and any server-side page cache.
- Ask the widget three to five real visitor questions on the front end.

15.3 When to run a full re-index

- After upgrading the plugin, especially across a schema change.
- After changing the embedding model — every page must be re-embedded.
- After a bulk content migration or import.
- After enabling chunking, WooCommerce products, or a new content type for the first time.
- When the widget serves stale answers for pages that changed significantly.

16. UNINSTALL & DATA CLEANUP

Deleting AI Navigator from Plugins → Delete runs a complete sweep:

- Drops all three plugin tables: s2w_ain_index, s2w_ain_chunks, s2w_ain_conversations.
- Deletes every plugin option — license state, Pinecone credentials, retention policy, and index progress.
- Sweeps the daily telemetry counter options shared across S2W plugins.
- Clears all six scheduled jobs, from conversation pruning to the weekly license check.
- Removes the page-role and canonical post meta from every post.

Safe to reinstall After uninstall the site is fully clean. Reinstalling starts from empty tables — no orphaned counters, no stale license state, no cron hooks firing into dead handlers.

17. TROUBLESHOOTING

Symptom	Cause and fix
Widget does not appear at all.	No OpenAI API key saved. The red admin notice confirms it. Enter a key

	on Settings → General, click Verify Key, then Save Settings.
Widget appears but does not answer.	The key is invalid, out of credits, or rate-limited by OpenAI. Re-run Verify Key and check the PHP error log for OpenAI error codes.
Index All Pages Now is grayed out.	Either no API key is saved, or you have unsaved settings changes — the tooltip says which. Save Settings re-enables it.
Indexing stalls partway.	Keep the admin tab open — the browser drives the job (Section 3.4). If your host disables WP-Cron the tab is doing all the work, and closing it pauses the job until you return.
Widget answers but rarely navigates.	Working as designed since v2.3.0 — navigation is ask-first (Section 10). Raise the aggressiveness slider to 3+ to auto-navigate strong matches, or teach visitors they can tap the offered page button.
Widget navigates to the wrong page.	A role mismatch. Set the correct role on the page it should have chosen, and mark it canonical if appropriate (Section 5).
Long pages get low-confidence answers.	Chunking is not active. Confirm Pro + Pinecone, then re-index pages showing one chunk in the Index Viewer.
Pinecone returns stale results.	Run a full re-index (wipe-before-sync clears everything stale) or a manual Pinecone sync from the Advanced tab.
Greeting and starter questions not translating.	Multilingual mode is off, or the visitor's browser language matches the site default. Check the Voice tab.
A visitor reports "too many requests."	They hit a per-visitor usage window (Section 11). It resets within five minutes; no admin action needed. Frequent reports from real visitors would be unusual — check for bots hitting the site.
Analytics is empty after upgrading to Pro.	Data is collected on all tiers and should appear immediately. Check the retention setting has not pruned a young install, and widen the date range.
"Cannot redeclare class S2W_Updater…" fatal.	One S2W plugin was updated while its siblings stayed on older infrastructure. Update AI Navigator, AI Concierge, and AI Gateway together (Section 15.1).

APPENDIX A — FREE VS. PRO

Feature	Free	Pro
Smart semantic navigation	✓	✓
Greeting bubble + starter questions	✓	✓
Browser text-to-speech	✓	✓
Multilingual detection & language picker	✓	✓
Page Role override in post editor	✓	✓

Index Viewer	✓	✓
WooCommerce product indexing	✓	✓
Usage limits & abuse protection	✓	✓
Conversation data collection	✓	✓
Privacy tools: erasure, purge, retention	✓	✓
Pages indexed	Up to 50	Unlimited
Voice input (OpenAI Whisper STT)	—	✓
Premium OpenAI TTS voices	—	✓
GPT-4o + large embedding model	—	✓
Pinecone vector database	—	✓
Semantic chunking for long pages	—	✓
Analytics dashboard incl. History tab	—	✓

APPENDIX B — PAGE ROLE REFERENCE

Role	Use for
contact	Contact pages — phone, email, form, office locations
pricing	Pricing pages — plan comparison, cost, tier details
product	Product or service detail pages (WooCommerce products get this automatically)
faq	FAQ, help center, knowledge base pages
about	About us, team, history, culture pages
blog	Blog index or category pages
home	Homepage — use sparingly, only for the actual homepage
legal	Terms of service, privacy policy, compliance pages
careers	Jobs, careers, open positions pages
partner	Partner, reseller, affiliate pages

APPENDIX C — WHAT'S NEW IN V2.3.0

Change	Details
WooCommerce product	Products are indexed alongside posts and pages — enriched with price,

indexing	stock, categories, SKU, attributes, and variations — so the assistant can answer questions about what you sell (Section 3.2).
Ask-first navigation	A matching page no longer redirects the visitor on its own. The widget navigates on an explicit ask, or offers the page with a confirm button; auto-navigation requires aggressiveness 3 or higher (Section 10).
Smarter role ranking	Pages are no longer boosted by their page role when they are semantically unrelated to the question (Section 5.2).
Browser-driven indexing	Index All completes on hosts where WP-Cron does not fire reliably — the admin tab drives the job (Section 3.4).
Analytics History tab	A filterable, expandable log of every conversation turn (Section 14).
Usage limits on AI calls	Anonymous chat, speech-to-text, text-to-speech, and greeting-translation requests are rate-limited per visitor, with caps on message length, history size, and audio size and type (Section 11).
Widget security hardening	The widget validates the origin of postMessage events and sanitizes the configured accent color.
Safer indexing UX	Index All Pages Now disables itself while settings are unsaved, and non-default post type selections persist correctly.

First-time setup? Start with the **AI Navigator Quick Start Guide**. Questions? Visit **speak2web.com**.