




INDITION AI ASSISTANT

User Manual

A practical guide to building and running AI assistants —
sources, settings, the inbox, feedback, and analytics.

AI Assistant Module · Version 1.0 · June 2026

 © 2026 Indition Corp. All rights reserved. Indition™ and the Indition logo are trademarks of Indition Corp. All other product names, logos, and brands are the property of their respective owners.

Indition AI Assistant — User Manual

A practical guide for the people who **build and run** AI assistants — from adding your first source to reviewing conversations and reading the reports. It explains what each screen does and how to get work done, not how the system is built.

Indition AI Assistant lets you create AI assistants that answer questions using **only your own approved content**. You gather knowledge into a collection, attach an assistant, embed it on your site, and then monitor and improve it from one workspace.

Key terms

| Term | Meaning |
|---------------------|--|
| Collection | A knowledge base — a set of sources plus default assistant settings. |
| Resource | One source in a collection: an uploaded file or a web page. |
| Assistant | The AI chatbot that answers using a collection's knowledge. |
| Chat app | An embeddable widget powered by one or more collections. |
| Citation | A reference in an answer pointing back to the source it came from. |
| Configured response | A pre-written answer to a known question (global or assistant-specific). |
| Starter question | A suggested prompt shown before the visitor's first message. |
| Fallback | The message shown when the assistant has no grounded answer. |

Get started in five steps

1. **Create a collection** and add a few sources (files or URLs).
2. Wait for the sources to finish **indexing**.
3. **Create an assistant** with the creator wizard and point it at your content.
4. **Test chat** with it, and add pre-configured Q&A for anything it gets wrong.
5. **Embed** the chat app on your site and watch the inbox and reports.

1. Building your knowledge base

Create a collection

Open the **Collection Builder** and give the collection a name and description. A collection holds your sources and the default behaviour (tone, fallback, citations, model) for assistants built on it.

Add sources (resources)

- **Upload files** — PDF, DOCX, or TXT. Large files upload in chunks with a progress bar.
- **Add URLs** — point at web pages to pull their content in.
- **Bulk import** — use the bulk drawer to add many files and links at once, then edit their titles before saving.

Indexing

After a source is added it is split into **chunks** and indexed so the assistant can search it. Sources show a status — *pending*, *indexed*, or *failed*. Depending on size, indexing runs in the background or inline.

If a source fails: open it to see the last error, then use **Reprocess** to try again. Common causes are unsupported file types, files over the size limit, or a URL that couldn't be fetched.

Test as you go

Both the collection builder and the assistant workspace include an **inline test chat**. Ask real questions to confirm the assistant answers from your content and cites the right sources before you publish.

2. Creating an assistant

The creator wizard

The **Assistant Workspace** opens in a four-step wizard when you create a new assistant:

1. **Setup** — name, purpose, tone, language, welcome and fallback messages, and whether to require citations. *(Required.)*
2. **Sources** — optionally upload files or add URLs now. *(Optional — you can do this later.)*
3. **Starter questions** — optionally pre-draft suggested questions. *(Optional.)*
4. **Review & Create** — confirm the summary and create the assistant.

Nothing is created until the final step. The wizard holds your choices locally and only saves when you click **Create Assistant** on Review & Create.

Editing an assistant

Once created, the workspace gives you a left navigation grouped into sections:

- **Assistant** — inline test chat, status toggle, and description.
- **Training** → **Sources** — add/reprocess/delete sources.
- **Settings** → **Setup** — name, tone, messages, status, citations.
- **Settings** → **Pre-configured Q&A** — assistant-specific or overridden global answers.
- **Settings** → **Starter Questions** — add, edit, reorder, or remove prompts.
- **Settings** → **Email Setup** — who receives conversation transcripts.
- **Settings** → **AI Model** — model choice and parameters **PERMISSION-GATED**.
- **Settings** → **Usage & Security** — access rules and blocked content.
- **Settings** → **Feedback** — feedback options shown to visitors.
- **Settings** → **Widget** — embed code and public key.
- **Chat History** — Inbox, Conversations, Messages, and Feedback History.

3. Improving answers

Pre-configured Q&A

For questions you want answered exactly the same way every time, use configured responses:

- **Global Pre-configured Q&A** — a central library of question/answer pairs you can apply to any assistant.
- **Per-assistant** — override a global answer, disable it for one assistant, reset it back to the global version, or add an assistant-only answer.

You can also turn a good answer from a real conversation into a configured response straight from the inbox.

Citations, clarification & fallback

- **Citations** show which sources an answer drew on — keep them on to build trust and spot weak sources.
- **Clarification** lets the assistant offer follow-up options when a question is ambiguous.
- **Fallback** is the message shown when there's no grounded answer — make it helpful (e.g. offer to contact a human).

Starter questions

Starter questions appear before a visitor types. Use them to showcase what the assistant can do. They can be answered live by the assistant or given a fixed static reply.

4. Publishing a chat app

- Use the **Chat App Builder** to create an embeddable experience powered by one or more collections.
- Set the name, tone, welcome/fallback messages, citations, and clarification mode; pick a model if allowed.
- Open **Settings** → **Widget** to copy the embed code and place it on your website.
- Use **Regenerate key** if the public widget key is ever exposed or needs rotating.

Before you embed: make sure the underlying collection's sources are indexed and you've test-chatted the most important questions. An assistant with no indexed sources can only fall back.

5. The inbox & feedback

- The **Inbox** (per assistant or the **Global Inbox**) lists conversations with the latest question and answer.
- Open a conversation to read the full **transcript** with timestamps and response times.
- Filter by search text, outcome (answered / fallback / blocked), channel, and feedback.
- **Feedback** — visitors leave thumbs-up/down and optional comments; review them in Feedback History.
- **Email transcripts** — configure recipients under Email Setup to have transcripts forwarded automatically.

6. Reports

Reports help you measure and improve the assistant. All support saved views, scheduling, and export.

| Report | What it tells you |
|---------------------------------------|--|
| Performance Overview | Citation coverage, fallback rate, response time, and satisfaction at a glance. |
| Cost & Usage | AI requests, tokens, and cost by model and assistant over time. |
| Engagement & Adoption | How much the assistant is used and whether adoption is growing. |
| Feedback & Satisfaction | Thumbs ratio, sentiment, and recurring feedback themes. |
| Knowledge Gaps | Questions the assistant couldn't answer — your content to-do list. |
| Knowledge Source Effectiveness | Which sources actually drive cited answers (and which are dead weight). |
| Usage & Security Events | Access patterns, blocked questions, and security-rule triggers. |

Dashboard

The **Dashboard** summarizes health: recommended actions ranked by impact, an answer-quality snapshot (citation coverage, fallback rate, negative feedback), and a list of **assistants needing attention** with quick actions (open workspace, add sources, review feedback, test chat).

A healthy weekly routine

1. Check the dashboard's recommended actions and "needs attention" list.
2. Read **Knowledge Gaps** and add or fix sources for unanswered questions.
3. Skim **Feedback & Satisfaction** and convert good answers into configured responses.
4. Watch **Cost & Usage** against your quota.

7. Administration

Some areas are for administrators rather than everyday builders.

Quotas ADMIN

- **Manage AI Assistant Quotas** — set per-customer limits: assistants, resources, chunks, and storage; daily and monthly messages, AI requests, tokens, and cost; and monthly ingestion jobs.
- **Public messages** — customize what end users see when they approach or hit a limit.
- **Default AI Assistant Quotas** — set the defaults applied to new customers.
- The quota grid shows live usage and blocked-attempt counts so you can see who is near a limit.

Model parameters ADMIN

AI Assistant Model Parameters sets the default model and parameters (e.g. temperature) used by assistants. Per-assistant model settings are available to permitted users in the workspace.

Where AI generation comes from

The assistant's intelligence — generation, embeddings, document text extraction, cost, and request logging — is provided by the **Indition AI Core** layer, which administrators configure with providers, models, and credentials. See the AI Core documentation for that setup.