




INDITION AI ASSISTANT

Feature List

Every capability in the AI Assistant module — assistants, knowledge collections, chat apps, the inbox, feedback, quotas, and reporting.

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 — Feature List

Indition AI Assistant is a customer-managed AI assistant that gives grounded, on-brand answers drawn only from your own approved content. You build knowledge collections from files and web pages, attach an assistant, embed it on your site as a chat app, and review every conversation, citation, and piece of feedback from one workspace.

This document is a structured catalogue of what the module does, grouped by area. Use it as a coverage map, a requirements checklist, or an index into the User Manual and Technical Reference. AI generation, providers, models, cost, and logging are supplied by the **Indition AI Core** infrastructure layer (documented separately).

At a glance

Area	What it delivers
Assistants	Create, configure, and test grounded AI assistants with tone, model, and behaviour controls.
Knowledge	Collections of resources (files + URLs), chunked and indexed for retrieval with citations.
Chat apps	Embeddable widgets powered by one or more collections, with a public key.
Answer quality	Citations, clarification prompts, fallbacks, and pre-configured Q&A.
Inbox & feedback	Read every conversation transcript; capture thumbs-up/down and comments.
Governance	Per-customer quotas, cost tracking, security rules, and model-parameter control.
Reporting	Seven reports across cost, adoption, satisfaction, knowledge gaps, and security.
Platform	Provider-agnostic generation, embeddings, and request logging via Indition AI Core.

1. Assistants

CREATE & CONFIGURE

- **Assistant creator wizard** — a guided four-step flow: Setup → Sources → Starter Questions → Review & Create.
- **Identity & behaviour**: name, purpose, tone (e.g. professional / friendly / concise), language, welcome message, and fallback message.
- **Citations** toggle — require the assistant to cite the sources behind each answer.
- **Status** — activate or deactivate an assistant without deleting it.
- **AI model settings** — choose the model and tune parameters (e.g. temperature); gated by permission.

WORK WITH AN ASSISTANT

- **Assistant workspace** — one place to edit setup, training sources, Q&A, starter questions, email, model, security, feedback, and the embed widget.
- **Inline test chat** — ask the assistant questions live while you build it.
- **Starter questions** — pre-drafted prompts shown before the first message (assistant-answered or static replies).
- **Chat history** — per-assistant inbox, conversations, messages, and feedback history.

PRE-CONFIGURED Q&A

- **Global configured responses** — a central library of canned question/answer pairs deployable to any assistant.
- **Per-assistant overrides** — override, disable, or reset a global response, or add assistant-specific Q&A.
- Used to give fast, exact answers to known questions before the model is consulted.

2. Knowledge: collections & resources

COLLECTIONS (KNOWLEDGE BASES)

- A **collection** groups resources and holds assistant defaults (tone, fallback, citations, model).
- **Collection builder** — a single page to manage sources, defaults, and an inline test chat.
- Health metrics: total vs. indexed resources, chunk counts, and indexing status.

RESOURCES (SOURCES)

- Add knowledge as **files** (PDF, DOCX, TXT) or **URLs**.
- **Bulk import drawer** — add many files and links in one action, with per-item titles.
- **Chunked upload** with progress and checksum for large files.
- **Processing modes** — background, partial-inline, or inline indexing with clear status.
- Per-resource actions: reprocess/reindex, delete, and last-error visibility.

INDEXING & RETRIEVAL

- Resources are split into **chunks** and embedded for search (via AI Core's embedding service).
- Answers are grounded in retrieved chunks and return **citations** back to the source.
- **Clarification** — when a question is ambiguous, the assistant can offer follow-up options.
- **Fallback** — a configurable message when no grounded answer is available.

3. Chat apps & embedding

- **Chat app builder** — assemble an embeddable chat experience powered by one or more collections.
- Configure name, status, tone, welcome/fallback messages, citations, and clarification mode.
- Model selection (where permitted) and an inline test chat after saving.
- **Embeddable widget** — drop the assistant onto a public website; copy the embed code from widget settings.
- **Public widget key** — issue and **regenerate** the key that authorizes the widget.
- **Markdown answers** — responses render with lists, bold/italic, code, and safe links (http/https only; no raw HTML).

4. Conversations, inbox & feedback

- **Global inbox** — every conversation across all collections and assistants in one reader.
- **Transcript reader** — full message timeline with timestamps and response-time durations.
- Filter by search text, outcome (answered / fallback / blocked), channel, and feedback.
- **Feedback** — visitors rate answers (thumbs up/down) with optional comments.
- Turn a real conversation answer into a **configured response** for next time.
- **Email transcripts** — forward conversation transcripts to configured recipients.
- Captured context per conversation: visitor info, channel, intent, citations, and latency.

5. Governance: quotas, cost & security

PER-CUSTOMER QUOTAS

- **Structural limits:** max assistants, resources, chunks, and storage (MB).
- **Daily limits:** messages, AI requests, tokens, and cost (USD).
- **Monthly limits:** messages, AI requests, tokens, cost, and ingestion jobs.
- **Public messages:** custom "approaching limit" and "blocked" messages shown to end users.
- **Default quota settings** — platform-wide defaults applied to new customers, editable without disturbing existing quotas.
- Live usage tracking (daily/monthly buckets) with blocked-attempt counts.

COST & MODEL CONTROL

- Per-message and aggregated token/cost attribution (rates come from AI Core models).
- **Model parameter settings** — platform defaults for model choice and parameters.

SECURITY

- **Security rules** per collection — access restrictions and blocked-content policies.
- Usage & security events are surfaced in reporting and the request log (AI Core).

6. Reporting & analytics

- **Performance Overview** — citation coverage, fallback rate, response time, satisfaction.
- **Cost & Usage** — AI request volume, tokens, and cost by model/assistant over time.
- **Engagement & Adoption** — conversation volume, active collections, adoption trends.
- **Feedback & Satisfaction** — thumbs ratio, sentiment, and recurring themes.
- **Knowledge Gaps** — questions the assistant could not answer well.
- **Knowledge Source Effectiveness** — which sources actually drive cited answers.
- **Usage & Security Events** — access, blocked questions, and rule triggers.
- Reports support saved views, scheduled delivery, and export.

7. Platform & integration

- **Indition AI Core** dependency — generation, embeddings, document extraction, request logging, cost calculation, and provider/model management.
- **Provider-agnostic** — runs on the providers configured in AI Core (e.g. OpenAI, Google/Gemini).
- **Dashboard** — recommended actions, answer-quality snapshot, and "assistants needing attention".
- Role-aware access: assistant builders, report/cost admins, and quota/system admins.