




INDITION AI CORE

FAQ & Troubleshooting

Common questions and fixes for platform administrators integrating and operating AI Core, with a symptom-based troubleshooting index.

AI Core 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 Core — FAQ & Troubleshooting

Answers and fixes for platform administrators who integrate and operate AI Core — the shared layer behind every AI feature. A symptom-based troubleshooting index is at the end.

AI Core is infrastructure: there are no end users. When a product module reports an AI problem, start here.

General

WHAT IS AI CORE, AND HOW IS IT DIFFERENT FROM AI ASSISTANT?

AI Core is the platform's **AI infrastructure**: providers, models, credentials, request types, generation/embedding calls, logging, and cost. **AI Assistant** is a product that *uses* AI Core. Other features can use it too. You configure AI Core once; every AI feature benefits.

WHICH PROVIDERS AND MODELS ARE SUPPORTED?

AI Core is provider-agnostic. It supports the providers you register (detected integrations include OpenAI and Google/Gemini). You add models under each provider with their pricing and capabilities.

WHERE ARE API KEYS STORED?

Encrypted at rest (AES) as ciphertext in the credential record. Raw keys are never displayed after saving, and sensitive fields are redacted from stored request/response payloads.

DOES AI CORE STORE EVERYTHING SENT TO PROVIDERS?

If logging is enabled, it records each call's metadata (provider, model, tokens, latency, cost, status) and redacted payloads, attributed to the feature and business object that made the call. Redaction fields and retention are configurable.

Setup & operations

NOTHING AI-RELATED WORKS ANYWHERE ON THE PLATFORM.

Check the fundamentals, in order:

- **Master switch:** `AI_CORE_ENABLED` must be on.
- **Provider:** at least one *active* provider with a correct base URL.
- **Credential:** an active, valid API key for that provider.
- **Model:** at least one active model under the provider.
- **Sharing:** the provider/model is shared (or default-shared) to the customer using it.

CALLS FAIL WITH AUTHENTICATION ERRORS.

The credential is wrong, expired, or revoked at the provider. Add a new credential, mark it default, and deactivate the old one. Confirm the key has access to the models you registered.

REQUESTS TIME OUT OR FAIL INTERMITTENTLY.

Increase `AI_REQUEST_TIMEOUT_SECONDS` and/or `AI_REQUEST_MAX_RETRIES`. Check the request log for the provider's HTTP status and error message, and the provider's own status page.

COSTS LOOK WRONG IN REPORTS.

Cost is computed from each model's `cost_input` / `cost_output` rates. If figures are off, correct the model pricing in Manage AI Models — all downstream cost (including AI Assistant quotas/reports) recalculates from those rates.

A DOCUMENT WON'T PROCESS / EXTRACT.

Confirm the file type is in `AI_ALLOWED_INFO_DOC_TYPES` (pdf/docx/txt by default) and within `AI_MAX_FILE_SIZE_MB`. Scanned PDFs with no embedded text yield nothing to extract. Extraction is also capped by `AI_MAX_EXTRACTED_TEXT_CHARS`.

VISION/IMAGE INPUTS FAIL.

Ensure the chosen model has the **vision** capability flag, the image type is allowed, and you're within `AI_MAX_IMAGE_DOCS`. If a provider rejects URL images, leave `AI_VISION_FORCE_BASE64` on so images are sent inline.

AI EMAIL GENERATION RETURNS INVALID OUTPUT.

The model occasionally returns malformed JSON. Raise `AI_EMAIL_INVALID_RESPONSE_MAX_RETRIES` (0–3). If images aren't wanted, turn off `AI_EMAIL_GENERATION_IMAGES_ENABLED`; otherwise keep within `AI_EMAIL_GENERATION_MAX_IMAGES`.

HOW DO I GIVE ONE CUSTOMER ACCESS TO A MODEL BUT NOT OTHERS?

Use the sharing grid to share specific providers/models/request types with that customer, and leave **default share** off so it isn't exposed to everyone.

LOGS ARE GROWING TOO LARGE / CONTAIN SENSITIVE DATA.

Lower `AI_LOG_RETENTION_DAYS` to prune sooner, expand `AI_LOG_REDACTION_FIELDS` to mask more fields, or disable logging with `AI_LOG_ENABLED` (you'll lose cost/audit detail).

Troubleshooting index

Symptom	Likely cause & fix
No AI anywhere	<code>AI_CORE_ENABLED</code> off, or no active provider/credential/model.
Auth/permission errors	Invalid/expired key, or key lacks access to the model — rotate the credential.
Feature can't see a model	Model inactive or not shared to that customer — enable/share it.
Timeouts / flaky calls	Raise timeout/retries; check provider status and request-log errors.
Wrong cost figures	Incorrect model input/output rates — fix pricing in Manage AI Models.
Document not extracted	Disallowed type, oversized, scanned-image PDF, or over the char cap.
Vision input rejected	Model lacks vision flag, image type disallowed, or over image-doc limit; keep force-base64 on.
Email generation malformed	Raise invalid-response retries; adjust image settings.
Logs too big / sensitive	Lower retention, widen redaction fields, or disable logging.
Customer over-exposed to capabilities	Turn off default share; share explicitly via the sharing grid.

Diagnosing via the request log

The request log is the single best diagnostic. For a failing feature, filter by the source product/operation or the business object (assistant, conversation, message) and inspect status, HTTP status, error code/message, tokens, latency, and cost. Attribution and context rows tell you exactly which interaction and object produced the call.

Still stuck?

Capture the request log entry (or correlation id), the provider/model involved, and the exact error. Confirm the provider's own service status, then contact your Indition administrator or support contact.