

AI Assist Handbook

A shared, advisory-only AI layer the whole platform draws on — drafting proposal and contract sections today, with a metered, budget-controlled gateway that keeps every request gated, audited and firmly in the author's hands.

Version 1.0 · ixlcore.com

Reference

AI Assist is not a separate app. It is a shared capability that lives at the centre of IXL CORE and that individual modules reach into when a task benefits from a first draft or a suggestion. Rather than each module building its own connection to an AI provider, every request flows through a single platform **AI gateway**. The gateway owns the things that matter across the whole platform — who is allowed to ask, what context is attached, what may and may not be done with the answer, how the request is priced, and how it is recorded — so that AI behaves consistently wherever it appears.

The most important rule to understand up front is that AI Assist is **advisory only**. It drafts, summarises and suggests. It never approves, posts, notifies, changes a record or executes a workflow on your behalf. A person always reads, edits and decides. This guide describes IXL CORE **version 1.0**.

Overview

Where AI shows up today falls into two confirmed places in the product, plus the shared machinery underneath them.

- **CRM — proposal and contract section drafting.** When you are writing a proposal or a contract, you can ask AI to draft or rewrite a single narrative section. This is a live, metered call through the gateway.
- **Operations — assist seam across the service desk and project work.** Operations routes a set of assists (ticket summaries, suggested replies, request classification, duplicate ranking, project-health explanations, service-review narratives, meeting-action extraction and more) through the same gateway. Because Operations is designed to keep working with the AI layer switched off, each of these degrades gracefully to a deterministic, offline result when no live provider is configured.

Underneath both sits the gateway itself, which enforces permission gates, applies guardrails, attaches only the context you are entitled to see, chooses the provider, meters the cost and writes an audit record and usage log for every single request.

It is worth being explicit about what is **not** AI-generated. Operations **status reports** and **service reviews** produce their numeric snapshots by straightforward aggregation and are flagged `ai_generated = false`. A narrative can be requested for a service review through the assist seam, but the review's figures themselves are not AI output. There is also a **memory** facility referenced in the request envelope; in version 1.0 it is readiness-only and disabled.

Proposal and contract drafting (CRM)

In CRM, both proposals and contracts expose an AI draft action for a single section at a time. You give the section a title, optionally paste the current body and add a short instruction, and choose whether you want a fresh **draft** or a **rewrite** of what is there. The gateway assembles a prompt from the proposal or contract (client, company, deliverables and so on), asks the provider for professional prose, and returns advisory content for you to edit before it goes anywhere.

These are the calls in the platform that request a live provider by default (the mid-tier model), so when a real key is configured they run against it and are metered; when no key is present they fall back to the internal advisory provider. The response you get back always carries an `advisory_only` flag and its own usage summary, so it is clear that the text is a starting point, not a finished document. Access is gated on the same permission you need to edit that proposal or contract, and a person's edit is always the final step.

Operations status reports and service reviews

Operations treats AI as an assist seam rather than a dependency. A dedicated service maps each Operations assist to a registered AI capability and asks the gateway for an advisory completion. Crucially, it is built to **degrade gracefully**: when no AI provider, capability or model is active, the gateway declines and Operations returns a structured `ai_unavailable` result carrying a deterministic fallback computed offline — so the interface never breaks and you always have something usable, with no network required.

This is why, in version 1.0, the Operations assists are best described as **wired but readiness-dependent**. The seam is real and the routes exist (ticket summary and reply, classify, duplicate-rank, project health, service-review narrative, meeting actions), but whether you receive an AI answer or the deterministic fallback depends entirely on whether a provider has been switched on for your organisation. The status report and service review figures, as noted above, are deterministic aggregations regardless of provider state.

Credits, budgets and control

Because AI has a real cost, the gateway meters it — but only on one specific path, and only when your organisation is set up to be metered.

Provider selection follows a **hybrid** rule for each request, in order:

- 1. Your own key (BYO).** If your organisation has configured and enabled its own Anthropic credential, the request runs on that key and is **not** metered — your key, your cost.
- 2. The platform key.** If the platform operator's key is configured and enabled, the request runs on it and **is** metered, billed to your prepaid AI credits.
- 3. The internal advisory provider.** If neither is available, the request falls back to an offline stub. No real spend, no failure. This is the default state, and the safety fallback whenever a live call cannot be made. A master switch can also hard-disable both live paths, forcing the stub everywhere.

On the metered (platform-key) path only, cost is controlled at two points. **Before** the provider is ever

called, the gateway makes a conservative estimate of the maximum cost and blocks the request if your balance is too low or a monthly budget cap would be breached — you are told your budget is exhausted and asked to top up. **After** a call actually completes, it debits the true cost from the returned token counts, once and atomically, so a retry can never double-charge and the wallet can never be driven negative. A failed or degraded call is never charged.

Credits are prepaid, denominated as one credit to one Kenyan shilling, and topped up in **Starter**, **Standard** and **Bulk** packs. There is also a base monthly bundle of included AI credits. Budget caps, when set, are enforced per calendar month.

How AI Assist connects

- **Access.** Every request passes the same permission gates as the rest of the platform. If you are not entitled to the underlying record or capability, the AI request is denied.
- **Context.** The gateway attaches only the context you are allowed to see, and records a summary of what was included and excluded on the usage log.
- **Guardrails.** Requests that ask AI to take an action — approve, reject, post, notify, create a document, execute a workflow or change a record — are blocked. Attempts to override the authenticated actor are blocked. AI output stays advisory.
- **Metering.** Cost flows through the shared tenant wallet, the same wallet that meters SMS, so budgets and top-ups are managed in one place.
- **Audit.** Every request — whether completed, blocked by a guardrail or metered — writes an audit record and a usage log, with a correlation identifier, so AI activity is as traceable as any other event on the platform.

In short, AI Assist gives your team faster first drafts and useful summaries where they help most, while the platform keeps a firm hand on permissions, cost and accountability — and keeps working exactly as before when AI is switched off.

How-to guides

Draft a proposal or contract section with AI

Ask AI Assist to draft or rewrite a single narrative section of a proposal or contract, then edit the advisory result before it goes anywhere.

When you are writing a proposal or a contract in CRM, AI Assist can draft a fresh section for you, or rewrite one you have already started — professional, concise prose that you own, edit and approve. This flow is **live** today, and it is the one place AI is metered when your organisation runs on the platform key.

Before you start

- You need permission to edit the record: **crm.proposals.update** for a proposal, or **crm.contracts.update** for a contract. If you cannot edit it, you cannot draft it.
- Have the proposal or contract open. The prompt draws on its client, company and deliverables automatically, so the more complete the record, the sharper the draft.
- Understand that the result is **advisory only**. AI never posts, sends or finalises anything — a person always reads and edits before it counts.
- If your organisation is metered and out of AI credits, the action is blocked with an "AI budget exhausted" message. See Manage AI credits and budgets.

Steps

1. Open the proposal or contract and go to the section you want to write.
2. Select **Draft with AI** on the section. [screenshot: proposal section with the Draft with AI action]
3. Enter a **Section title** (required) — for example "Scope of work" or "Terms and conditions".
4. Choose the **mode: Draft** for fresh prose, or **Rewrite** to improve the current text.
5. If you chose Rewrite, the section's **current body** is passed in so AI works from what is there.
6. Optionally add a short **Instruction** (up to 500 characters) — for example "keep it to three short paragraphs" or "emphasise the support SLA". [screenshot: AI draft panel with Section title, mode toggle and Instruction]
7. Select **Generate**. The request flows through the platform AI gateway, which applies guardrails, picks the provider and meters the call where applicable.
8. Review the returned text, marked **advisory only**, then **edit it in place** and save the section. [screenshot: advisory draft returned into the section editor]

Result

The section is populated with an AI-drafted starting point that you have reviewed and edited. Where the call ran on the platform key, its cost is debited from your AI credits once and shown in the response's usage summary; on your own key or the offline advisory provider, nothing is charged. The record itself only changes when you save your edits — AI never commits anything on your behalf.

Related

- [AI Assist reference](#)
- [Manage AI credits and budgets](#)
- [Draft and e-sign a proposal](#)

Manage AI credits and budgets

See how AI Assist is metered, top up prepaid credit packs, set a monthly budget cap, or bring your own key to run unmetered.

AI Assist has a real cost, so the platform meters it — but only on one path, and only when your organisation is set up to be metered. This guide explains the included monthly bundle, prepaid top-up packs, budget caps, and the bring-your-own-key option that makes AI unmetered.

Before you start

- Metering runs through the shared tenant wallet — the same wallet that meters SMS — so credits and caps are managed in one place.
- AI credits are prepaid and denominated **one credit to one Kenyan shilling (1 credit = KES 1)**.
- Only the **platform-key** path is metered. On your **own key** or the offline advisory provider, requests are **not** charged.
- You need the relevant billing or administration permission to view the wallet, buy packs and set caps.

Steps

1. Open your organisation's **AI credits** (wallet) area. [screenshot: AI credits wallet with current balance]
2. Check your **AI credit balance**. Your Base subscription includes a **monthly AI bundle** that is granted once per calendar month.
3. To add more, buy a prepaid top-up pack — **Starter**, **Standard** or **Bulk**. Each credit buys KES 1 of platform-key AI usage. [screenshot: Starter / Standard / Bulk top-up packs]
4. Optionally set a **monthly budget cap**. When a cap is set, the gateway blocks a request before any provider is called if it would push this calendar month's AI spend over the cap.
5. To run AI **unmetered**, add your own Anthropic key: open **Integrations**, configure the **Anthropic** credential and enable it. [screenshot: Anthropic integration credential, enabled]
6. Save. Provider selection is **hybrid** and automatic per request — your enabled key first (unmetered), then the platform key (metered), then the offline advisory provider (free).

Result

Your AI usage is controlled and predictable. When metered, the gateway makes a conservative cost estimate **before** each call and blocks it if your balance is too low or the monthly cap would be breached; **after** the call it debits the true cost once, so a retry never double-charges and the wallet can never go negative. A failed or offline call is never charged. With your own enabled key, live AI runs on

your credential at your cost — no platform credits are used.

Related

- [AI Assist reference](#)
- [Draft a proposal or contract section with AI](#)
- [Draft and e-sign a proposal](#)

