AI Email Assistant — Setup & Daily Use

How the team writes grounded customer replies inside Thunderbird … set up once, then use every day.

What this is. The Vidalux AI email assistant lives inside **Thunderbird** as an add-on called **ThunderAI**. It reads a customer's email and writes a **grounded draft reply** — pulled live from the full Vidalux knowledge base (products, warranties, delivery, policies). Nothing is ever sent automatically … a person always reviews and sends.

The 3 golden rules

1. **It's always a draft.** Nothing reaches a customer until a human reviews it and presses send.
2. **It only knows the Vidalux knowledge base.** If it doesn't know something it will politely *defer* ("let me get a member of the team to confirm") rather than invent an answer. That's deliberate — it protects you.
3. **Set it up once, then copy to the other machines.** Don't retype anything five times.

PART A — SETUP (ONE-TIME, PER MACHINE)

A1 · Install Thunderbird + add an email account

1. Download **Thunderbird** (version 140 or later) and install it.
2. On first launch, add the mailbox.
   - **Gmail / Google:** sign in via the **Google pop-up** (don't type the password into Thunderbird's own box, or it hangs).
   - **Other hosts:** if you get an *"Add Security Exception"* certificate warning, confirm it to continue (it's your own mail server).

A2 · Install the ThunderAI add-on

Menu (≡) → **Add-ons and Themes** → search **ThunderAI** → **Add to Thunderbird** → **Add**.

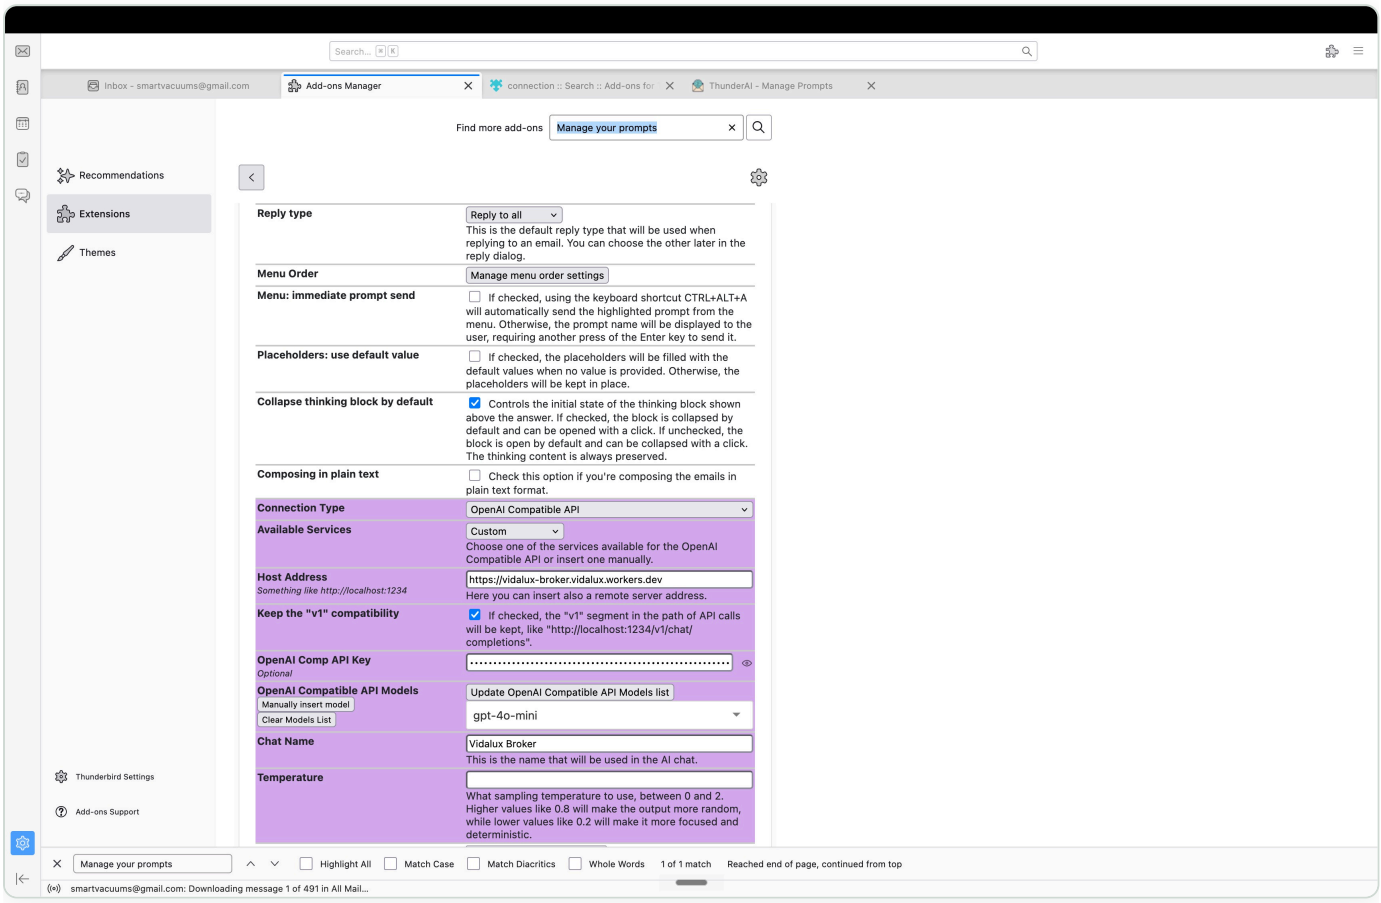
A3 · Connect ThunderAI to the Vidalux brain

This is the step that makes replies grounded. In **Add-ons and Themes** → **ThunderAI** → **the gear / Preferences**, scroll to **Connection Type** and set it up exactly as below:

Connection Type	`OpenAI Compatible API`
Available Services	`Custom`
Host Address	`https://vidalux-broker.vidalux.workers.dev` *(base URL only)*
Keep the "v1" compatibility	✅ ticked *(it adds `/v1/chat/completions` for you)*
OpenAI Comp API Key	*the broker token (supplied separately — keep it internal)*
Model	click **Manually insert model** → type `gpt-4o-mini` → make sure it's selected

Chat Name	`Vidalux Broker`
Temperature	leave blank

Don't skip: right after typing the Host Address, click **"Give permission to the current host"** in the CORS box — without it the request gets blocked. Then click **"Update OpenAI Compatible API Models list"** as a quick test — no error means it's reaching the brain.



A3 — the finished connection screen: OpenAI Compatible API, broker Host Address, key (masked), model `gpt-4o-mini`.

A4 · (Optional) Add the "Reply with instructions" prompt

The built-in **"Reply to this email"** is the main one and needs nothing extra. This optional prompt just lets the team type a steer *up front*. In ThunderAI Preferences → bottom → **"Manage your prompts"** → **Add New**:

ID	`reply_with_instructions`
Name	`Reply with instructions`

Prompt Text	`Write a reply to the following customer email.`
	`Extra instructions for this reply:`
	`{%additional_text:instructions%}`
	
	`Customer email:`
	`{%mail_text_body%}`

Add to menu	`Reading an email`
Action	`Do reply`

Properties	✅ tick **"Ask for additional text"**

Click **"Add the Prompt"**, then **"Save All"**.

A5 · Copy the setup to the other machines

- **Prompts:** on the configured machine, **Manage your prompts** → **"Export all prompts"** → save the file → on each other machine, **"Import new prompts"** → choose it → **Save All**.
- **Connection:** repeat **A3** on each machine (four fields pasted in — under a minute). Same Host Address + key + model everywhere.

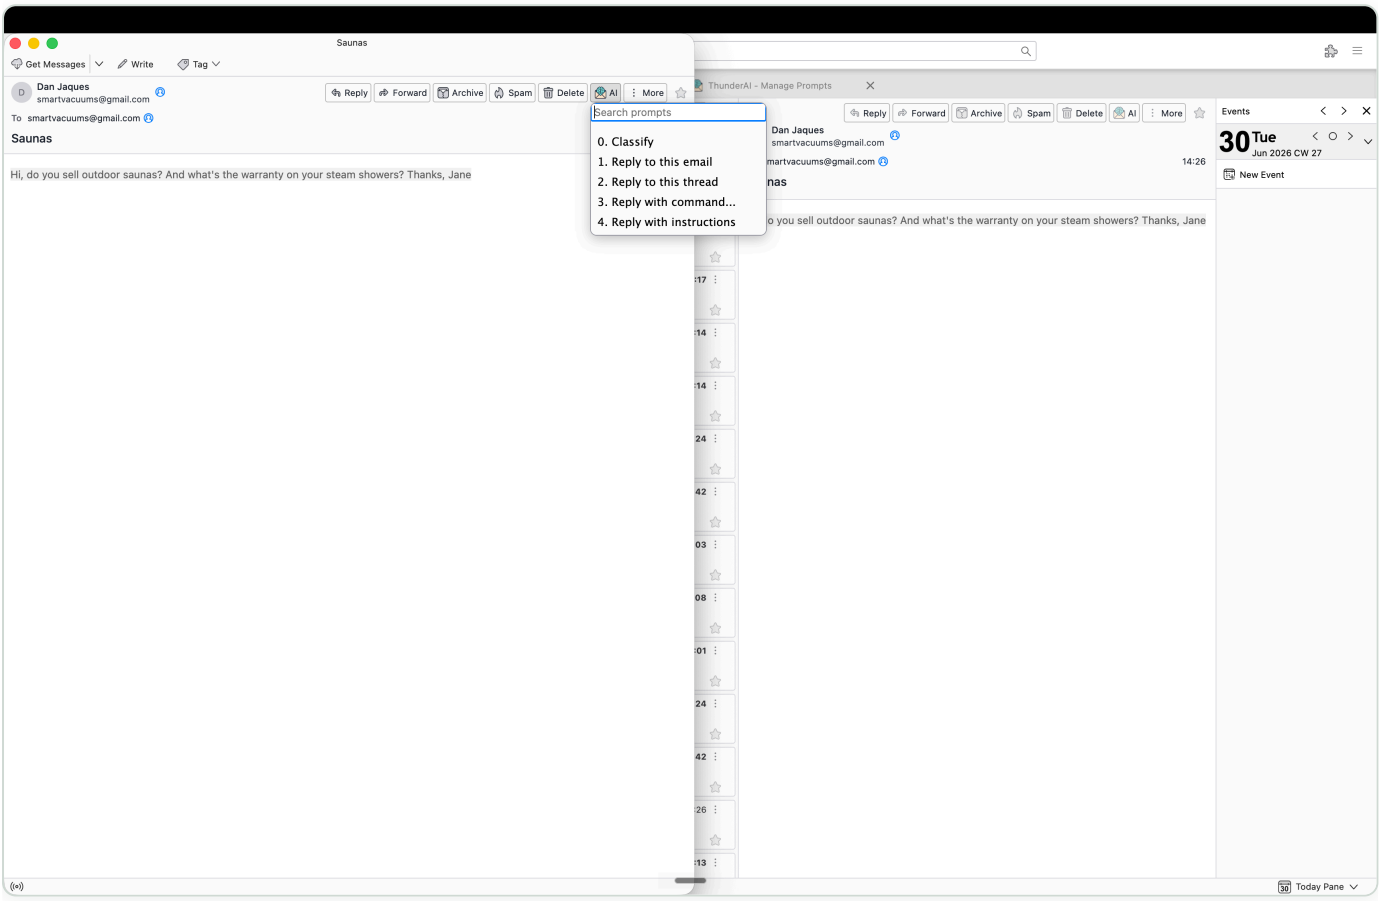
> ⚠️ The broker key is shared across all machines. Keep any export file **internal** and delete it once the machines are done.

PART B — DAILY USE (THE WORKFLOW)

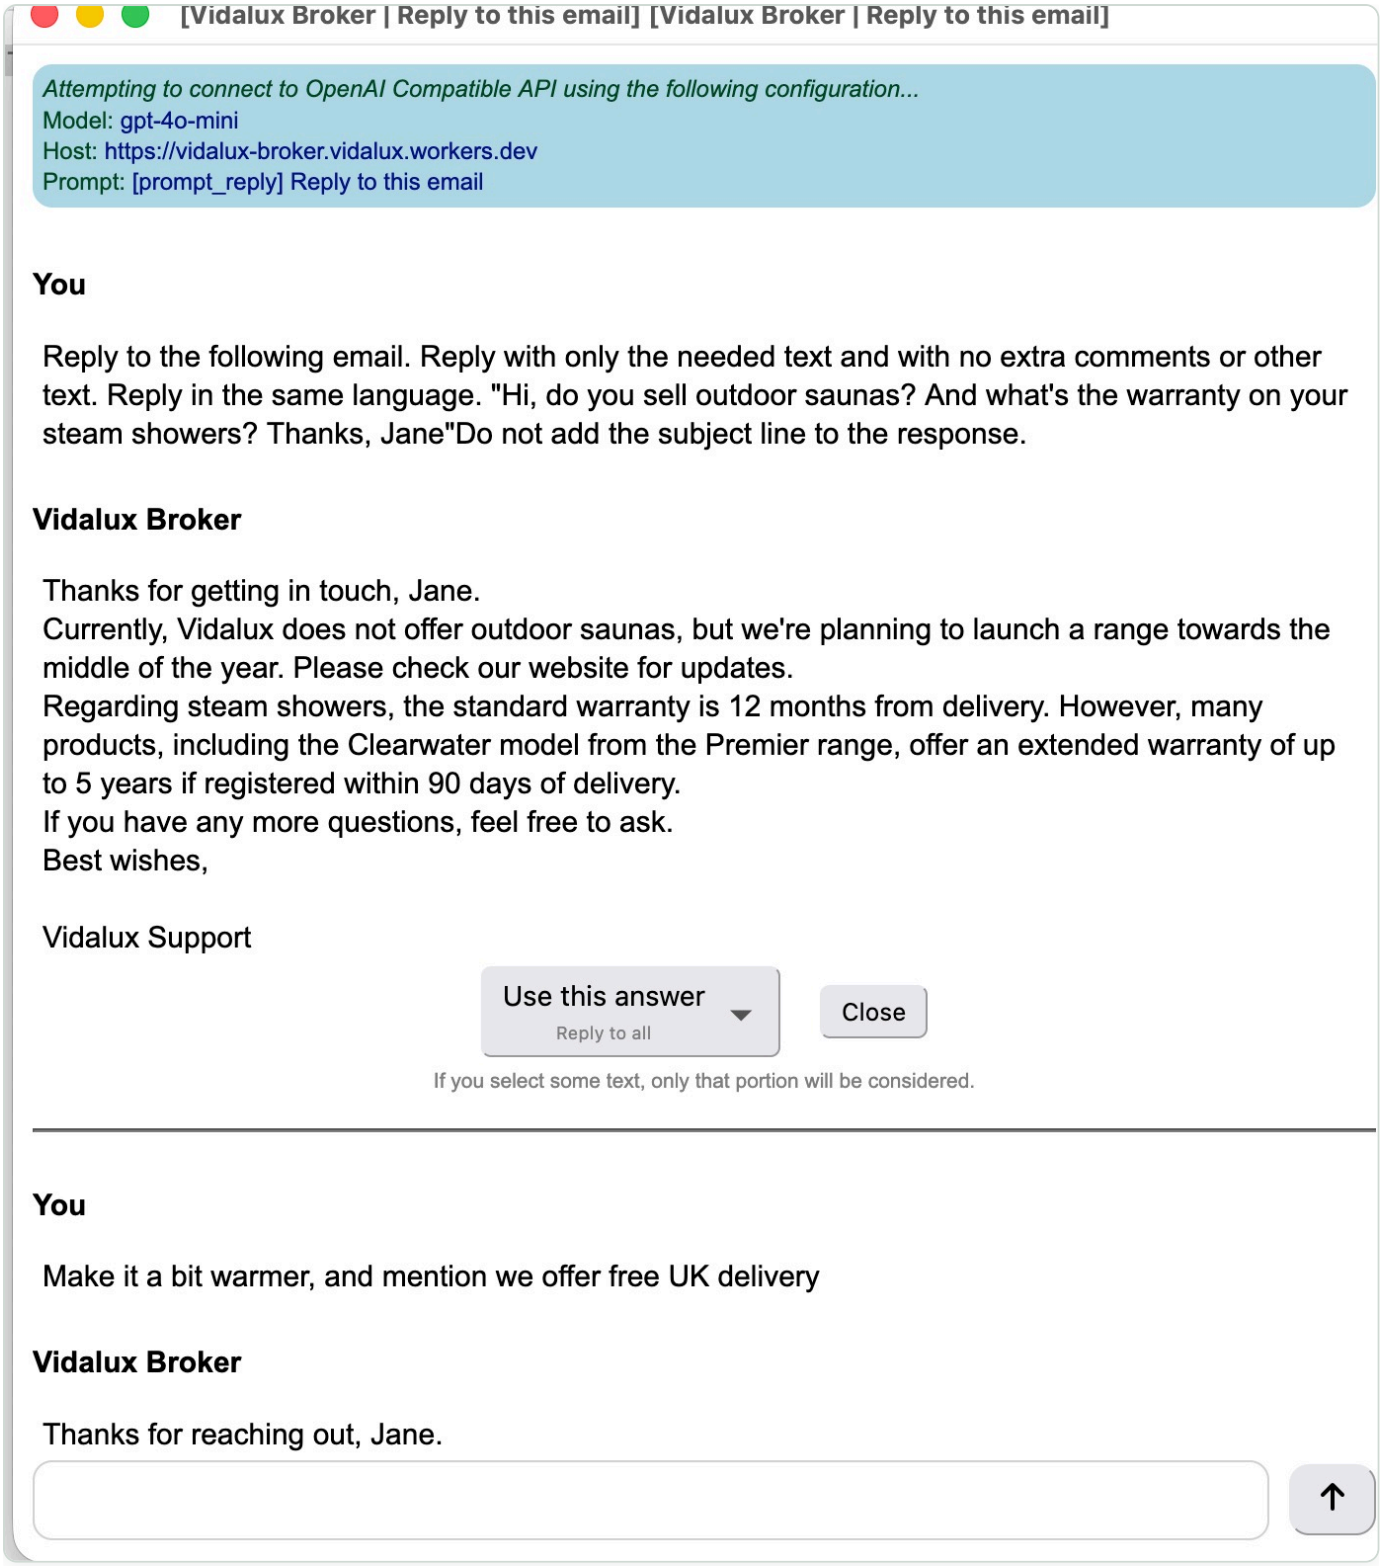
This is what the team actually does, every day. Three steps.

1 Get a grounded draft

Open the customer's email → click the **AI** button in the toolbar → **"Reply to this email"**. In ~10 seconds a grounded draft appears (it's retrieving the answer from the knowledge base).



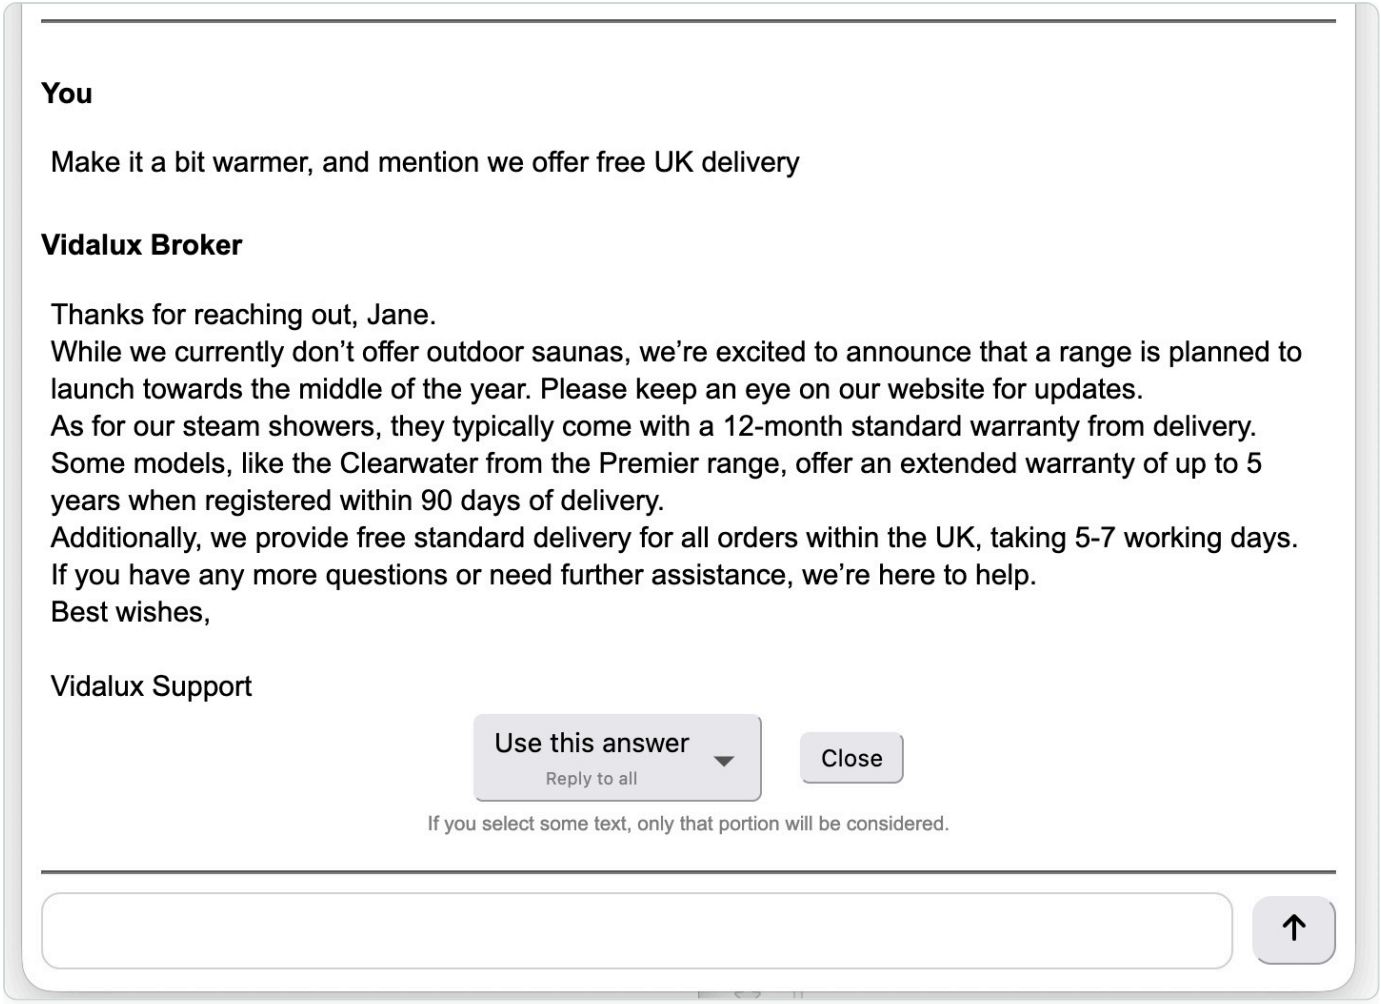
*Step 1 — the **AI** menu. "Reply to this email" is the everyday choice.*



The grounded draft — real Vidalux facts (outdoor saunas launching mid-year; steam-shower warranty 12 months, extendable to 5 years). Not invented.

2 Refine it — just chat to it

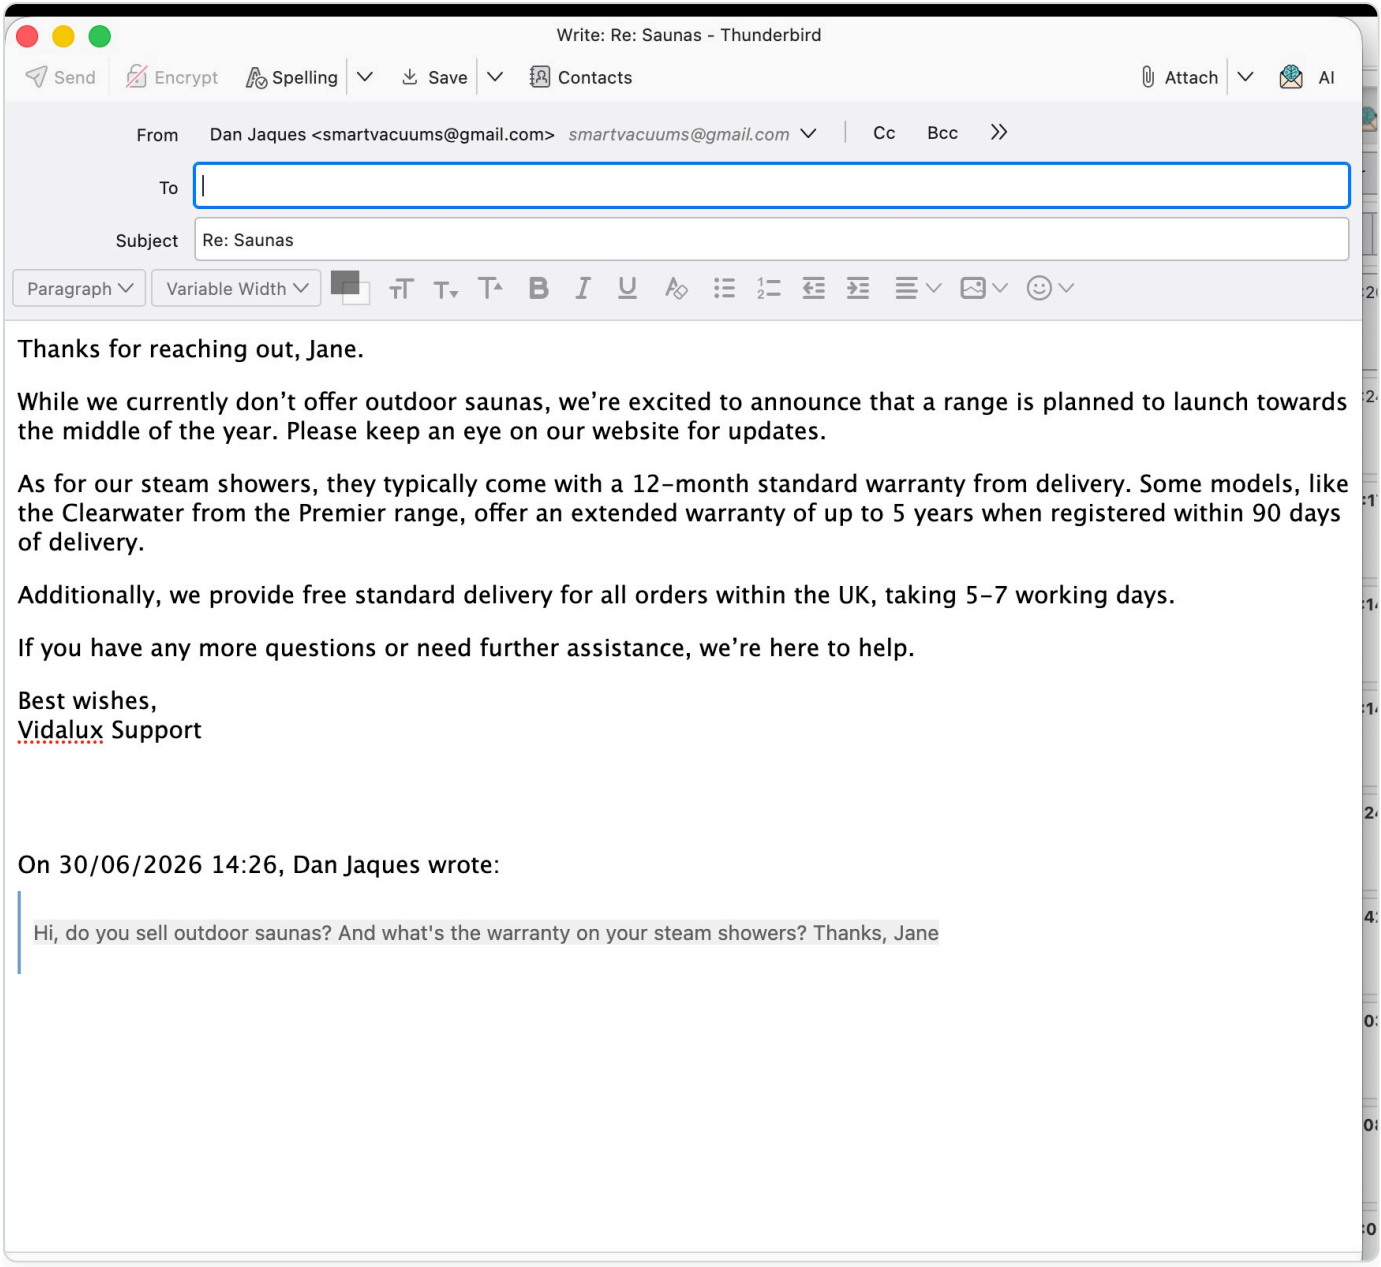
The box with the ↑ **arrow** at the bottom of that window is a **live chat**. Type a follow-up, press ↑, and it rewrites the draft keeping all the facts — *"make it warmer"*, *"mention free UK delivery"*, *"shorten it"*, *"add the returns policy"*. As many rounds as you like.



Step 2 — "Make it a bit warmer, and mention we offer free UK delivery" … the draft updates, facts intact.

3 Send it

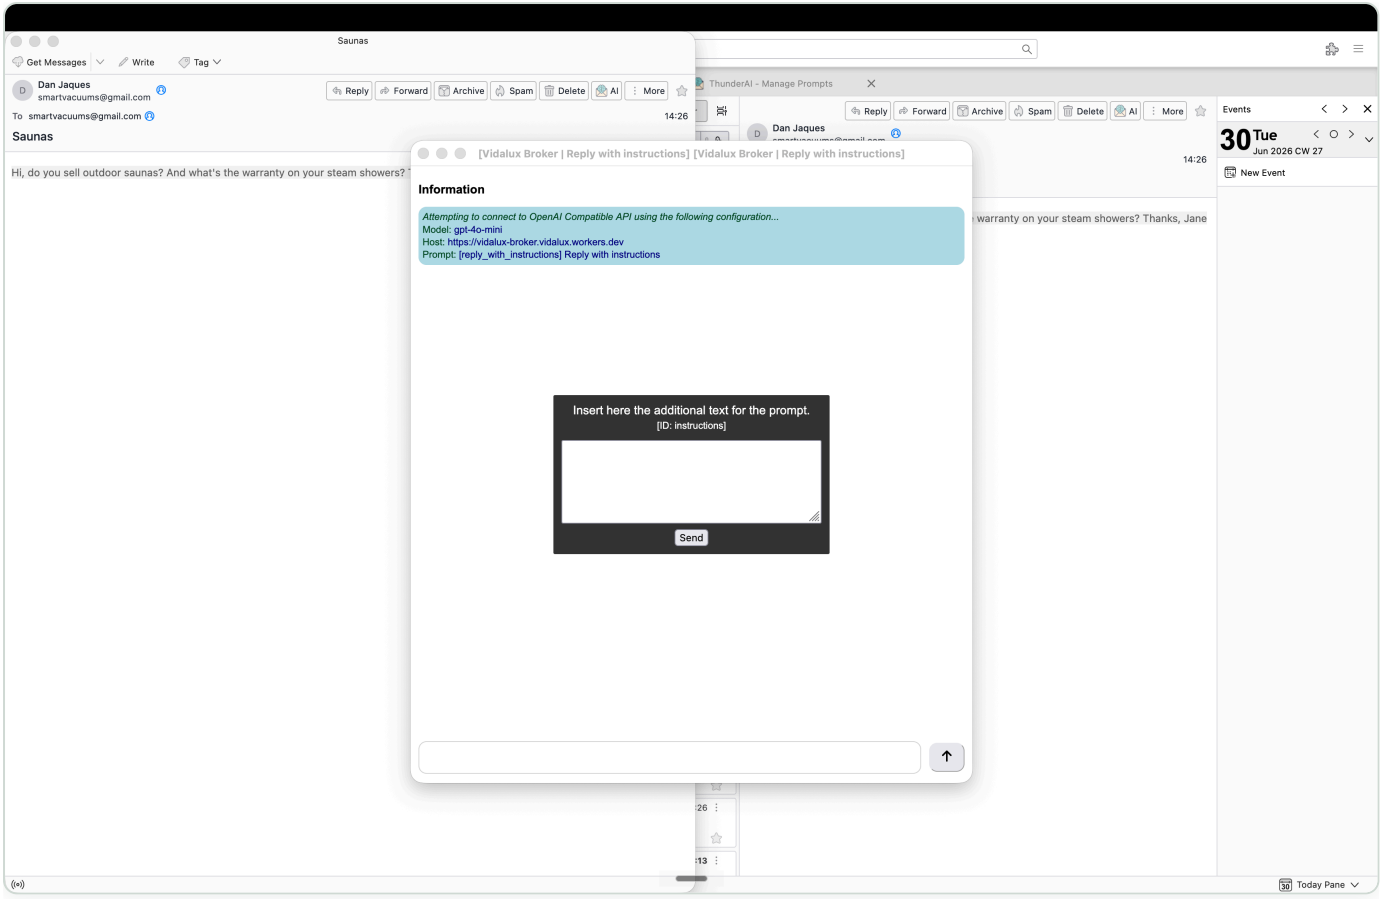
Click **"Use this answer"** → the draft drops into a normal **editable reply**. Make any final tweaks by hand → press **Send**.



Step 3 — the draft lands in a standard, fully editable reply. Nothing sends until you press Send.

Shortcut · "Reply with instructions"

If you *already* know the steer before seeing a draft, pick **"Reply with instructions"** at step 1 instead — it pops a box for your instruction first, then drafts. Same result, one less round-trip. *(Leave the box blank and it behaves like a normal reply.)*



The optional "Reply with instructions" box — type a steer up front, or leave it blank.

When it doesn't know the answer

It will **defer**, not guess — something like *"let me get a member of the team to confirm."* Chatting won't conjure facts the brain doesn't hold. That exact question is logged automatically and added to the knowledge base, so next time it *will* know. Handle that one email yourself in the meantime.

If something looks off

No "Reply with instructions" in the AI menu	Manage your prompts → check it's there and set to **Reading an email**; re-run **Save All**.
Connection error when drafting	Re-check A3: Host Address (base only), key, and that you clicked **"Give permission to the current host"**.
Account setup hangs on a spinner	A certificate / sign-in pop-up is waiting behind the window — check Mission Control.
Anything else	Flag it to ContentWorks — don't fall back to the old Gem, its knowledge is out of date.
