



CELLULE-PRO DOCUMENTATION

User Guide

Cellule-PRO employee manual

Contents

Contents	2
0.1 Cellule-PRO — User Guide	3
Contents	3
1. Welcome to Cellule-PRO	3
2. First sign-in	4
3. Interface tour	5
4. Starting a conversation	7
5. Continuing a conversation	7
6. Attaching documents to a conversation	8
7. Debate mode — 3 agents thinking together	9
8. Long-term memory — [MEMORIZE: ...]	10
9. Managing your conversations — export, delete	11
10. Project mode — collaborative spaces	12
11. Your rights — GDPR, export, deletion	13
12. FAQ and troubleshooting	14

0.1 Cellule-PRO — User Guide

Audience: any employee using Cellule-PRO day-to-day from their browser (lawyer, physician, engineer, analyst, operator, support function).

Reader prerequisites: none. Being able to open a browser and type a question is enough.

What this guide is not: it does not cover IT-side installation (see [ADMIN_GUIDE.en.md](#)), nor command-line / API usage (see [API_REFERENCE.en.md](#)). It addresses you, the end user, in the standard web interface.

Contents

1. Welcome to Cellule-PRO
2. First sign-in
3. Interface tour
4. Starting a conversation
5. Continuing a conversation
6. Attaching documents to a conversation
7. Debate mode — 3 agents thinking together
8. Long-term memory — [\[MEMORIZE: ...\]](#)
9. Managing your conversations — export, delete
10. Project mode — collaborative spaces
11. Your rights — GDPR, export, deletion
12. FAQ and troubleshooting

1. Welcome to Cellule-PRO

Cellule-PRO is an intelligent assistant **deployed on your organization's servers**. No data leaves your network: not your questions, not the documents you share with it, not the answers it produces. You can therefore use it for your most sensitive professional matters — contracts, patient records, internal notes, litigation, confidential projects.

1.1 What you can do with it

- **Ask open questions** in a field you master, to save drafting time or structure your thinking.
- **Have a document analyzed** by attaching it (contract, study PDF, meeting notes) and get a summary, key points, follow-up questions.
- **Compare multiple viewpoints** on an ambiguous topic with **debate mode**: three agents (builder, skeptic, pragmatist) take positions before a synthesis is produced.
- **Build a long-term memory**: facts you want to retrieve in all your future conversations ([\[MEMORIZE: ...\]](#)).
- **Collaborate as a team** on a shared file (project mode): shared documents, shared conversations, shared memory.

1.2 What you cannot (or should not) do

- **No personal use** outside the professional scope — your organization logs requests for GDPR audit (article 30).
 - **No self-administered medical question** or legal advice on your own situation: this is neither a doctor nor a lawyer, it's an assistant that summarizes text.
 - **No absolute truth**: the assistant can be wrong, especially on exact figures, precise dates, fine legal references. Always verify what's important.
-

2. First sign-in

Your IT administrator has provided you with:

1. **A URL** — for example <https://cellule.company.local> (or a hostname ending in `.local`, `.lan`, or an internal subdomain).
2. **A login ID** — typically your work email or an LDAP identifier.
3. **An initial password** — to be changed on first sign-in (the system will ask).

2.1 Opening Cellule-PRO

In your browser (Edge, Chrome, Firefox, Safari — all supported), enter the provided URL. You land on a sober login page, optionally with your company logo.

No public signup. If you have not received credentials, contact your IT admin.

2.2 First login

- Enter your ID and initial password.
- You are prompted to choose a personal password.
- You land on the main screen — the **empty conversation** ready for your first question.

2.3 API token (optional)

If you also use a command-line tool ([opencode](#), [Continue](#), [Cursor](#), etc.) that calls Cellule-PRO via API, your IT admin provided a **token** formatted as `sk-cellule-...`. This token is only needed for programmatic use — for the web UI, your login suffices.

2.4 First prompt with opencode (30-second setup)

[opencode](#) is an OpenAI-compatible CLI agent that integrates natively with Cellule-PRO. If your IT admin gave you the API token (§2.3), here is the minimal setup on your workstation:

Step 1 — Installation (Linux / macOS)

```
| curl -fsSL https://opencode.ai/install | bash
```

(Windows: follow the official [opencode.ai docs](#) — install via PowerShell or WSL2.)

Step 2 — Configuration

Create the file `~/.config/opencode/opencode.json`:

```
{
  "$schema": "https://opencode.ai/config.json",
  "model": "cellule/CELLULE",
  "provider": {
    "cellule": {
      "npm": "@ai-sdk/openai-compatible",
      "name": "Cellule-PRO",
      "options": {
        "baseUrl": "http://<POOL_HOST>:18095/v1",
        "apiKey": "sk-cellule-<YOUR_TOKEN>"
      },
    },
    "models": {
      "CELLULE": {
        "name": "CELLULE",
        "limit": { "context": 131072, "output": 131072 }
      }
    }
  }
}
```

Replace `<POOL_HOST>` with the IP or hostname of your Cellule-PRO pool (your IT admin gave it to you — often the cluster VIP) and `<YOUR_TOKEN>` with the token you received.

Step 3 — First prompt

Launch `opencode` in interactive mode (TUI) from any project folder:

```
| opencode
```

Type your question (e.g. *“Refactor this file to ...”*) + Enter, `opencode` calls Cellule-PRO, the LLM answers and may use standard tools ([Read](#), [Edit](#), [Glob](#), [Bash](#), etc.) on your local tree.

Non-interactive mode (one-shot, scripts):

```
| opencode run "Write a Python function that ..."
```

Important — long-term memory: your `sk-cellule-*` token opens an encrypted memory on the pool side (per-user RAG). Anything learned through `opencode` is retrievable from the web UI (and vice-versa) as long as you use the **same token**. See §8 for the `[MEMORIZE:]` pattern.

3. Interface tour

Cellule-PRO is laid out as a web app with three zones:

TOPBAR (session status + language + sign-out)

CONV	MAIN ZONE	(context
RAIL	(active conversation OR	brief rail
(~280px)	projects page / API / etc.)	if project)
+ New		
- Conv 1	messages...	
- Conv 2		
- Conv 3		
	[composer..]	

3.1 The topbar (top)

- **Company logo** — if configured by your IT admin.
- **Language selector** — FR / EN (the AI replies in the UI language, unless you explicitly ask in another).
- **Status indicator** — green = pool responds, gray = you are offline or pool in maintenance.
- **“Sign-out” button** — closes your browser session.

3.2 The conversation rail (left)

Always visible, ~280px. Style close to ChatGPT.

- **+ New conversation** — starts a fresh thread.
- **Conversation list** — sorted by recent activity, most recent on top. Each line shows the title (inferred from the first message) + the message count + the date.
- **Delete a conversation** — × button right of each line (with confirmation).
- **Click a conversation** — switches the main zone to it and reloads the full history.

3.3 The main zone

- **In conversation mode:** message history (you on the right, assistant on the left) + composer at the bottom.
- **In project mode:** project dashboard (documents, members, shared conversations).
- **In API mode:** ready-to-copy code snippets to use Cellule-PRO from your tools.

3.4 The composer (input area, bottom)

Above the input field:

- **Paperclip** — attach a document to this conversation.
- **Theater mask** — enable debate mode for the next question (3 agents).

The input field supports:

- **Enter** = send the question.
- **Shift+Enter** = new line (line break without sending).

While the AI is thinking, the “Send” button becomes inactive and a “thinking...” indicator appears.

4. Starting a conversation

4.1 Your first question

On the home screen, you see an empty composer. Type your question, for example:

« Draft me a 10-line summary of DPO obligations under the GDPR. »

Press **Enter** (or click “Send”). The assistant starts answering in streaming — words appear as they are generated. You can read before it’s done.

Once the first message is sent, **the conversation is created server-side** and appears in the left rail with an auto-generated title.

4.2 What to do for better answers

- **Be specific about context:** “for an IT audience”, “with a formal tone”, “as a numbered list”.
- **State the goal:** “to help a client understand”, “to prepare a board meeting”.
- **Indicate expected length:** “in 5 lines max”, “in one paragraph”, “as a Markdown table”.
- **Iterate freely:** if the first answer doesn’t fit, say “rewrite shorter”, “add a concrete example”, “rephrase for a non-lawyer audience”.

4.3 If you change your mind

Click + **New conversation** in the left rail to start a fresh thread, without losing the current one.

5. Continuing a conversation

All your conversations are saved server-side (encrypted by your token). You can close your browser, log back in tomorrow from another machine — they are still there, in the left rail.

5.1 Pick up where you left off

Click a conversation in the rail. The main zone loads the full history. Type your new question in the composer at the bottom — the AI benefits from **all the prior context** to answer.

5.2 Context limit

A conversation can grow long, but past a certain volume of text, the AI starts forgetting the earliest exchanges (a technical limit of language models, not a bug).


If you handle a long, complex topic, two strategies:

- **Summarize then new conv:** ask « summarize our exchange in 10 points » at the end, copy the answer, paste it at the start of a new conversation.
 - **Use project mode** (cf. §10): memory is shared and persistent at project level.
-

6. Attaching documents to a conversation

One of the most useful features: making the AI work on **your** document.

6.1 How to attach

1. Click the  icon in the composer.
2. Select one or more files from your disk.
3. A “pill” appears above the composer for each document, with the name + a status (**indexed**, **in progress**, **failed**).
4. Once the status is **indexed**, ask your question.

6.2 Supported formats

- **PDF** — any readable PDF (not unreadable image scans).
- **DOCX** — recent Word documents.
- **TXT** — plain text.
- **MD** — Markdown.
- **ODT** — LibreOffice / OpenOffice.

Max file size: **50 MB** (your IT admin can adjust).

6.3 What happens behind the scenes

The AI does not “read” the whole document on every question (too slow). Instead:

- The document is sliced into small passages (“chunks”).
- Each passage is turned into a semantic vector.
- For each question, the 5-10 **most relevant** passages for your question are injected into the AI’s context.
- The AI replies explicitly grounded in these excerpts.

Consequence: your questions must be **targeted enough** to allow semantic search. Prefer:

« What are the termination conditions in this contract? »

over:


« What do you think of this document? »

6.4 Detach a document

Click the × right of the pill. The document is removed from the conversation (but remains accessible in the “My documents” tab if your IT admin enabled it).

6.5 Tip: prepare a conversation before the 1st message


You can **create an empty conversation** then attach documents **before** asking the first question. Useful to set up a work session:


1. Click “+ New conversation”
2. Click  and attach all your PDFs
3. Wait for all to be **indexed**
4. Ask your first question — the AI already has all the material.

7. Debate mode — 3 agents thinking together

An advanced feature, especially useful for **ambiguous, strategic or sensitive questions** where a single answer risks being superficial.

7.1 How to enable

In the composer, click the  icon (theater mask). A “Debate mode enabled” hint appears. Type your question then send.

The mode is **opt-in** and **disabled after each send**. If you want to debate several questions, click  each time. Reason: debate mode costs more time (×2 to ×3 vs simple answer) — opt-in avoids paying for it by accident.

7.2 What happens live

During the debate, you see live:

- **Round 1** — each agent (builder, skeptic, pragmatist) takes a position in parallel.
- **Convergence score** — between 0 and 1, indicates how close the 3 positions are. If 0.92, debate stops (agents agree).
- **Round 2-3** — each agent sees others’ positions and refines its own.
- **Final synthesis** — a fourth, more powerful model (30B+) takes the 3 final positions and produces the definitive answer.

7.3 The 3 characters

- **Builder** — proposes, opens paths, sees potential. Tends toward optimism and action.
- **Skeptic** — raises counter-arguments, identifies risks, questions assumptions. Anti-runaway safeguard.
- **Pragmatist** — balances, contextualizes, brings back to concrete feasibility. Mediator between extremes.

7.4 When to use it

Good case: - « Should we accept this 24-month non-compete clause? » - « What evaluation strategy for this ambiguous record? » - « Which trade-off between option A and option B on this project? »

Bad case (simple answer suffices): - « Draft me a polite follow-up email. » - « What's the definition of DPO? » - « Rephrase this paragraph shorter. »

7.5 Review the debate after the fact

Once synthesis is rendered, below the assistant message you'll see something like:

Debate in 2 rounds, convergence 0.94 — **view debate**

Clicking **view debate** unfolds the detail: for each round, each agent's position, the model used, elapsed time. Useful to understand **why** the synthesis reached its conclusion, or to spot an interesting point of disagreement.

8. Long-term memory — [MEMORIZE: ...]

The AI forgets everything between two conversations... **except** what you explicitly ask it to remember.

8.1 The pattern

At the end of **any question**, add:

[MEMORIZE: the fact to remember]

Examples:

What are the DPO obligations under GDPR?

[MEMORIZE: our internal DPO is Sophie Martin, reachable at sophie.martin@company.com]

Help me structure my note for the committee.

[MEMORIZE: the executive committee meets every 1st Tuesday of the month, chaired by Jean Dupont]

8.2 What happens

- The pool **intercepts** the `[MEMORIZE: ...]` tag.
- It **records** the fact in your personal memory (encrypted by your token).
- It **removes** the tag from your question before forwarding it to the AI (so you won't get a "noted!" reply).
- The AI answers your question normally.

8.3 How the fact resurfaces later

In any future conversation, if you ask a semantically close question, the fact is automatically re-injected into the AI's context. Example:

```
Session 1 (Monday): ... [MEMORIZE: parse_log() bug line 42 when log empty, current_section not init]
```

```
Session 2 (Friday, other machine): « What bug was I supposed to fix this week? » →  
The AI answers: « UnboundLocalError line 42 in parse_log() when the log is empty.  
Probable cause: current_section not initialized. »
```

8.4 Best practices

- **Concise:** one fact per MEMORIZE, phrased in one clear sentence.
- **Specific:** « DPO = Sophie Martin » is more useful than « we have a DPO ».
- **Avoid redundancy:** don't replant the same fact in every conversation.
- **Confidential by default:** your memory is isolated, nobody else in the company can read it.

8.5 Also works in API mode

If you use Cellule-PRO via `opencode`, `Continue`, or any other API client, the `[MEMORIZE:]` pattern works identically — your memory is shared between the web UI and your command-line tools, as long as they use **the same token**.

9. Managing your conversations — export, delete

9.1 Export a conversation

In the main zone, above an open conversation, two buttons:

- **Export JSON** — structured file, ideal for archiving or re-injecting elsewhere.
- **Export MD** — readable Markdown, ideal for pasting in a report, a wiki, or an email.

The export includes all messages, metadata (date, model used, debate scores if any), but **not** the attachments.

9.2 Delete a conversation

Two places:

- **In the left rail** — small × right of each line.
- **In the open conversation** — “Delete” button on top.

Deletion is **immediate and final** server-side (with GDPR audit). You are asked to confirm.

9.3 Encrypted conversations

If your conversation is marked `encrypted`, it has been encrypted by a personal key (rare, generally enabled for ultra-sensitive topics). You can only read it from the session where it was created. If you change machine / reset your password, the encrypted history remains unreadable.

10. Project mode — collaborative spaces

Project mode lets you **work as a team** on a shared file. Each project has:

- A **member list** (who can access).
- A **document space** (shared RAG — all members see the same documents).
- A **shared memory** (shared MEMORIZE facts).
- **Associated conversations** (visible to all members).

10.1 Create a project

1. Click “My projects” in the top navigation.
2. “New project” → name and description.
3. You are automatically the **owner**.

10.2 Invite colleagues

In the project page, “Members” section, enter your colleague’s ID (the same they use to sign in to Cellule-PRO) then “Invite”. They will have access at their next sign-in.

10.3 Add documents to the project

“Documents” tab → “Add” → drag-and-drop or selection. Same formats and limits as in conversation (PDF, DOCX, TXT, MD, ODT, 50 MB max).

All project documents are indexed once, then available **to all members** in all their conversations tied to the project.

10.4 Project conversations

In the project page, “New project conversation” launches a thread that **automatically benefits** from the project’s documents + memory. Ideal for: multi-person analysis of a complex file, shared synthesis note, meeting prep.

10.5 Leave / delete a project

- **Leave** (non-owner member) — you lose access, your past conversations remain.
 - **Delete** (owner only) — deletes documents, memory, conversations, **definitively**. Confirmation required.
-

11. Your rights — GDPR, export, deletion

In compliance with GDPR, at any time you have:

11.1 Right of access — export all your data

“My settings” page → GDPR section → “**Export my data**”. You receive a ZIP archive containing:

- All your conversations (JSON + MD).
- All documents you have attached.
- All facts you have planted ([MEMORIZE](#)).
- Connection audit history.

11.2 Right to erasure — delete all your data

Same page → “**Delete my account and all my data**”. Strong confirmation required (enter your password + tick a box).

Effect:

- All your conversations, documents, memories are erased.
- Your account is disabled.
- You can no longer sign in.
- A GDPR audit keeps trace of the deletion (anonymized), without the content of the erased data.

11.3 Your activity audit

You can consult your own activity log (who signed in when, which conversations modified, which documents added). Useful to verify nobody signed in on your behalf.

11.4 What your administrator sees

Your IT admin has access to **aggregate statistics** (number of conversations, token volume, capacity used) and a **security audit** (who signed in when, from which IP).

Your IT admin does not have access to the content of your conversations, your memory, your personal documents (encrypted by your token). Project conversations, however, are visible to **project members**.

12. FAQ and troubleshooting

12.1 The AI is slow to answer

Normal for a first question (~3-10 seconds of “thinking” before streaming). If more than 30 seconds without any character, check:

- That the status bar on top is **green**.
- That you attached documents before asking (indexing can take 10-30 s).
- That you are not in `debate` mode (counts 30-90 s).

If nothing happens, **refresh the page** (F5) and retry.

12.2 « I don't have access to the content » while I attached a PDF

- Check the pill status is **indexed** (not `failed` nor `in progress`).
- If `failed`: the PDF may be an unreadable image scan. Try OCR beforehand.
- If `indexed`: rephrase your question more **specifically**: cite a precise term that appears in the document.

12.3 The answer is invented / wrong

The AI can **hallucinate**, especially on figures, dates, names, legal references. Rule: anything that must be exact (number, name, date, citation) **must be verified** by you before professional use. Tip: re-ask « cite the exact source in the document » — the AI will cite or admit it doesn't know.

12.4 I lost my conversation

Check:

- The **left rail**: it may have slid down the list.
- The **current tab**: you may be on “Projects” rather than “Conversations”.
- If it was **deleted** by mistake: unfortunately final (no trash). Ask your IT admin if a backup restore is possible (per organization policy).

12.5 I forgot my password

Click “forgot password” on the sign-in screen. Per your organization’s policy, a reset email is sent, or you have to contact your IT admin.

12.6 Can I use Cellule-PRO from my phone?

Yes, the interface is responsive and works on tablet and mobile (Safari iOS, Chrome Android). The conversation rail is hidden automatically below 900px width, accessible via a menu.

12.7 Difference from ChatGPT, Claude.ai, etc.

Aspect	Cellule-PRO	ChatGPT (public)
Data stored	On your servers	On OpenAI servers
Confidentiality	LAN-guaranteed	Subject to ToS
Documents	Locally indexed	Sent to cloud
GDPR compliance	Native	Complex
Debate mode	Yes (3 agents)	No
Long-term memory	Yes (cross-session)	Limited

12.8 Who to contact

- **Technical issue** (can’t sign in, 500 error): your IT admin.
- **Usage question** (« how do I... »): this guide, or your functional referent.
- **Improvement suggestion**: forward to your IT admin who centralizes feedback for the vendor.

Enjoy Cellule-PRO. The tool is designed to fade behind your work — you remain the pilot, the AI a silent copilot.