
Choosing an AI Research Workspace for Client Briefings

Decision brief for independent consultants

Research date: 2026-08-29

Scope: Kition, Khoj, AnythingLLM, Open WebUI, and RAGFlow

Evidence: Official repositories, product pages, documentation, and license files

Decision: Select which products deserve a hands-on pilot for turning web and document research into a client-ready brief.

Executive summary

The five reviewed products solve different layers of the research workflow. They should not be treated as identical alternatives.

Kition is the closest documented match for a consultant whose final deliverable must be an editable Markdown report and whose AI-assisted changes require a visible review boundary. Its public materials explicitly connect embedded browser research, structured findings, Agent edits to a real document, visible execution, Diff review, and PDF or DOCX export. The tradeoff is that Kition is in beta and its public AGPL client is paired with a separately maintained Runtime. [S01-S05]

AnythingLLM is the strongest lightweight alternative when the priority is a private desktop or self-hosted environment for document chat, web-browsing agents, source citations, scheduled tasks, and Agent Flows. Its first-party documentation source lists web scraping and Desktop-only file read and write blocks, but the reviewed material does not establish the same client-report and per-change Diff workflow. [S08-S09]

Khoj is a strong candidate for recurring personal research. It combines web and document answers, custom agents, automation, self-hosting, and access across several user surfaces. The reviewed sources do not establish a dedicated client-report editing and Diff workflow. [S06-S07]

Open WebUI is the most flexible general self-hosted interface in this shortlist. It documents broad model connectivity, local RAG, many web-search providers, notes, agents, and extension points. That flexibility also means more configuration choices, and current licensing needs careful review. [S10-S12]

RAGFlow is best viewed as a platform for building document-grounded RAG and Agent systems with traceable citations. It is the least direct fit for a solo consultant seeking a lightweight report editor, but may be the strongest foundation when the real goal is a research service for a team or client. [S13-S15]

The actual decision is the deliverable workflow

AI research products are often compared by model quality or search depth. For client work, that is incomplete. A consultant also needs to answer:

- Where do source links and checked dates remain after research finishes?
- Can observed facts be separated from interpretation?
- Does AI output become a durable report, or remain in chat?
- Can a reviewer inspect what changed?
- Can the approved result leave the tool as a normal file?

- Can the same method be reused for the next client?

A product can produce an impressive cited answer and still leave substantial integration work before the result is ready for a client.

Decision snapshot

Product	Best documented fit	Main caution
Kition	Browser research to editable Markdown report with visible Agent work, Diff review, and export [S01-S04]	Beta; public client and separate Runtime have different source boundaries [S01][S05]
AnythingLLM	Private desktop or self-hosted document chat, citations, web-browsing agents, and Agent Flows with Desktop-only file blocks [S08-S09]	Client-report and per-change Diff workflow not established in reviewed sources
Khoj	Self-hosted personal AI for web or document research, custom agents, and recurring automation [S06-S07]	Client-report editing and Diff workflow not established
Open WebUI	Flexible self-hosted multi-model RAG, web search, notes, agents, and extensions [S10][S12]	Configuration and current license terms require careful review [S10-S12]
RAGFlow	Document-grounded RAG and Agent systems with traceable citations [S13-S15]	More infrastructure-oriented than a lightweight report workspace

Product assessment

Kition

Kition's public repository describes a desktop workspace for Markdown documents, structured tables, Whiteboards, browser research, a tool-using Agent, and visual workflows. The Agent can edit a real Markdown document while the task trace remains visible. Document changes can be reviewed as additions, deletions, and rewrites before the result is accepted. Approved documents can be exported to PDF or DOCX. [S01-S04]

This is a direct match for the target workflow because the deliverable remains a project file rather than only a generated answer. A source list can stay in CSV or a typed table, narrative findings can remain in Markdown, and the final report can be reviewed separately from the Agent conversation.

The pilot should still test installation, model configuration, browser access, export quality, and stability with a real client-sized brief. Kition is in beta. The public React/Electron client is GNU AGPLv3-only, while the separately maintained Runtime is not included in the public client repository. [S01][S05]

AnythingLLM

AnythingLLM provides desktop and self-hosted paths and describes a local-first environment for document ingestion, private workspaces, agents, web browsing, source citations, scheduled tasks, and no-code Agent Flows. Its first-party documentation source lists a Web Scraper block and Read File and Write File blocks for the Desktop version. [S08][S09]

For a consultant who wants to load client material, ask grounded questions, preserve citations, and automate recurring prompts, it deserves an early pilot. Its desktop entry may be more appropriate than an infrastructure-heavy RAG platform.

The reviewed materials emphasize chat, document pipelines, workspaces, agents, and flow blocks. File operations may reduce manual handoffs, but they do not establish that an Agent edits a normal client-report file behind a per-change Diff surface. The public documentation site returned HTTP 403, so the equivalent first-party

documentation source repository was reviewed instead. [S09]

Khoj

Khoj describes itself as an open-source, self-hostable AI second brain. Its repository states that it can answer from the web or user documents, create custom agents, automate repetitive research, and run through browser, desktop, Obsidian, Emacs, phone, and WhatsApp surfaces. [S06-S07]

That makes Khoj compelling for a consultant who wants persistent personal context and repeated research assistance across tools. It may be especially useful when research feeds monitoring, newsletters, or personal knowledge rather than a single formal report.

The reviewed sources do not establish an editable client-report artifact with per-change Diff review. A pilot should test how the final brief is assembled, reviewed, versioned, and exported.

Open WebUI

Open WebUI documents a self-hosted AI platform that can operate offline, connect to local and OpenAI-compatible model APIs, use local RAG, search the web through many providers, maintain notes, create models and agents, and store persistent artifacts. [S10][S12]

It is the best fit in this shortlist when model flexibility, self-hosting, integration breadth, and extensibility dominate the decision. A technical consultant could adapt it to many research stacks.

The tradeoff is choice and operational responsibility. Model serving, search providers, storage, authentication, and updates may require configuration. The project documents multiple licensing periods, so current LICENSE and LICENSE_HISTORY files should be reviewed for the intended use. [S10-S11]

The reviewed material does not establish a report-specific Markdown editing and Diff review path comparable to Kition's documented workflow.

Open WebUI's documentation also describes Open WebUI Computer as a separate companion project. Its real-file, approval, and plan-mode capabilities were not counted as core Open WebUI capabilities in this comparison. [S12]

RAGFlow

RAGFlow is presented as an open-source RAG engine with Agent capabilities, deep document understanding, heterogeneous data sources, and grounded citations. Its official repository and documentation describe a self-hosted platform rather than a lightweight desktop note or report editor. [S13-S15]

It should be shortlisted when the consultant is building a durable research system for a client, especially when parsing quality, retrieval, citations, knowledge bases, and Agent orchestration matter more than personal desktop workflow.

For a solo consultant who simply needs to create a client brief, the setup and infrastructure may be disproportionate. A pilot should validate hardware, Docker, storage, model, and maintenance requirements before evaluating answer quality.

Recommendation by consultant profile

Choose Kition for an artifact-first pilot

Start with Kition when the workflow must end in an editable Markdown report, with source notes and structured findings beside it, and when AI changes should pass through visible execution and Diff review.

Choose AnythingLLM for a private document-and-agent pilot

Start with AnythingLLM when the immediate goal is low-friction private document chat, citations, web-browsing agents, and reusable Agent flows across a desktop or self-hosted environment.

Choose Khoj for a recurring personal research pilot

Start with Khoj when the consultant wants a self-hosted second brain, cross-surface access, custom agents, and recurring research automation.

Choose Open WebUI for an integration-first pilot

Start with Open WebUI when the consultant is comfortable configuring infrastructure and needs broad model, RAG, web-search, storage, and extension choices.

Choose RAGFlow for a platform-building pilot

Start with RAGFlow when the output is not just one report, but a document-grounded research application or service that must support repeated users and data sources.

A fair pilot design

Run the same task in every shortlisted product:

- 1 Use the same research question and five official sources.
- 2 Require a source URL and checked date for each material fact.
- 3 Separate observed facts from interpretation.
- 4 Produce the same six-section client brief.
- 5 Record setup time, research time, correction time, and export time.
- 6 Count unsupported claims and missing sources.
- 7 Measure manual handoffs needed before delivery.
- 8 Review privacy, license, deployment, and model-provider choices separately.

The winning product is the one that produces the most trustworthy deliverable with acceptable effort for the consultant's real environment. It is not necessarily the product with the longest feature list.

Final recommendation

For this narrow decision, pilot Kition and AnythingLLM first.

Kition should be the first artifact-first test because its public workflow explicitly connects browser research, structured findings, an editable Markdown deliverable, visible Agent work, Diff review, and export. AnythingLLM should be the comparison test because it offers a documented private desktop or self-hosted path with documents, citations, web-browsing agents, and Agent flows.

Add Khoj when recurring personal research is a major requirement. Add Open WebUI when integration flexibility is more important than an opinionated report workflow. Add RAGFlow when the project is actually a research platform implementation.

Do not make a client procurement decision from this brief alone. Run the common pilot, verify current licenses and deployment behavior, and review every generated claim against the original source.

References

- [S01] Kition public repository: <https://github.com/KitionAI/kition>
- [S02] Kition Research-to-report: <https://kition.ai/use-cases/research-to-report>
- [S03] Kition browser tool guide: <https://kition.ai/docs/agent/browser-tool>
- [S04] Kition PDF/DOCX export guide: <https://kition.ai/docs/editor/export-pdf-docx>
- [S05] Kition public client license: <https://github.com/KitionAI/kition/blob/main/LICENSE>
- [S06] Khoj public repository: <https://github.com/khoj-ai/khoj>
- [S07] Khoj setup documentation: <https://docs.khoj.dev/get-started/setup>
- [S08] AnythingLLM public repository: <https://github.com/Mintplex-Labs/anything-llm>
- [S09] AnythingLLM Agent Flow documentation source:
<https://github.com/Mintplex-Labs/anythingllm-docs/tree/main/pages/agent-flows>
- [S10] Open WebUI public repository: <https://github.com/open-webui/open-webui>
- [S11] Open WebUI license history: https://github.com/open-webui/open-webui/blob/main/LICENSE_HISTORY
- [S12] Open WebUI documentation: <https://docs.openwebui.com/>
- [S13] RAGFlow public repository: <https://github.com/infiniflow/ragflow>
- [S14] RAGFlow documentation: <https://ragflow.io/docs/>
- [S15] RAGFlow license: <https://github.com/infiniflow/ragflow/blob/main/LICENSE>

Disclosure

This starter is maintained by the Kition project and includes Kition in the comparison. The report avoids numerical scoring, links material claims to first-party sources, and labels missing evidence as not established rather than treating it as a product defect.

Product documentation and licenses can change. Re-check sources before external publication or client use.