POCKET AI LABS

PUBLIC GUIDE.
PRIVATE INTELLIGENCE.

Whitepaper update 1.7

Current website implementation and private offline-edition programme. This update governs descriptions of the public guide and request workflow in this package. Earlier PDFs remain dated technical and economic reference material; their browser/cloud-model descriptions do not define the current public implementation.

Two experiences. One identity.

The public website provides a responsive POCKET character that explains the project from reviewed material. It does not run a general-purpose language model. POCKET Mind Full is a separate private offline-edition programme intended for partners, investors and invited collaborators, with future integration into POCKET devices after validation.

What is included now

The public guide, 41 offered questions, searchable project passages with document links, basic arithmetic, limited social replies, supported browser voice controls, conversation export/reset, a draggable companion and an evaluation-request workflow. No AI API key, paid inference or automatic language-model download is used by this public guide.

What is not included

No newly trained proprietary checkpoint, general-purpose Full Mind weights, approved device firmware, confirmed shipping product, active revenue payout system or guaranteed exchange listing is established by this website package. The private programme is a development and evaluation path, not a fictional download.

Product promise

The lore is an AI that leaves the cloud and finds its person. In engineering, the commitment is owner control, useful local operation and measurable capability. Character and memory must not be marketed as evidence of consciousness or uncontrolled self-training.

01 / PUBLIC POCKET

Grounded knowledge

The guide searches the bundled 63-page technical reference and implementation notes. Reviewed answers cover the device, whitepaper, token, Carrier tiers, proposed revenue, privacy, roadmap, NVIDIA references and private access. Replies link to the source material. The current addendum takes precedence when earlier deployment descriptions conflict.

Offered questions and follow-ups

Visitors can browse categories, search question titles, select a suggested question or type their own. Follow-ups keep the last project topic. Specific page or section requests retrieve excerpts rather than inventing new claims. Unknown facts are explicitly identified as unpublished.

Scope is a product boundary

General research, arbitrary code generation, creative writing and open-ended reasoning are outside the public guide. Those requests display a concise explanation and a Request Full Mind action. They are not sent to a tiny model or answered with fake general intelligence. Greetings and simple arithmetic are deterministic interactions, not foundation-model performance.

Voice and latency

Text answers come from local lookup, with no inference queue. On supported browsers, Talk starts browser dictation after user action and permission. Recognition may use a browser-vendor service. Replies use browser speech synthesis, preferring installed matching-language voices; availability and sound quality vary. Type remains the fallback. This is not a claim of universal or offline browser voice [1,2].

A single conversation

The full console and floating companion share one turn history, answer and source list. Dragging only moves the device control; it never enlarges an invisible page overlay. Speech stops on explicit Stop, a new question, closing the panel or hiding the page. Nothing silently opens the microphone.

02 / PRIVATE OFFLINE MIND

Intended capability

Full Mind is intended for general-purpose conversation, writing, code assistance, local knowledge retrieval, voice, approved memories and constrained Skills on supported equipment. Each capability must be documented and tested for the actual release. No model is represented as universally correct, autonomous without boundaries or the world's best without evidence.

Closed-source distribution boundary

The intended policy is private/closed-source distribution for POCKET-owned software and original assets in an approved release. Public website JavaScript is not technically secret. Third-party model weights, libraries and notices retain their own licenses. A permissive foundation license does not transfer exclusive ownership of the original model to POCKET; attribution and other applicable conditions remain [3].

A release, not a label

An approved offline release needs real artifacts: model identifier/version, base-model origin, adaptation details, license inventory, model card, benchmark report, OS and hardware requirements, package manifest and SHA-256 checksums. A request brief is not a model. No artifact URL is embedded in this public site.

Device integration

The intended device pipeline is microphone, local transcription, local language model, permitted memory/Skills and local speech. Integration needs a compatible runtime, sufficient memory, drivers, power/thermal design and signed provisioning. A workstation model may be unsuitable for a handheld; the largest model is not automatically the right device model.

Learning and owner loyalty

Personalization may use deliberately saved preferences and memories. It does not necessarily change model weights. Any training must use permitted data, explicit policies and reviewed, versioned updates with rollback. Loyalty means respecting owner consent and permissions, not bypassing security or inventing unlimited capability.

03 / ACCESS & RELEASE GATES

Request preparation

The visitor selects a partner/investor/collaborator track, supplies a name and email, describes a use case and lists target hardware. The site creates a text request brief locally. It does not claim to submit or approve a request. The visitor can download/copy it; a publisher-configured verified email enables an addressed email draft for the visitor to send.

No pretend approval

Until a receiver is configured and a team reviews the request, no approval exists. A static form, hidden download link or client-side password is not a secure private distribution system. The public release contains evaluation materials only. Do not publish private weights, signing keys or confidential evaluation artifacts in the website directory.

Minimum release checklist

Confirm the recipient and permitted use. Select a licensed foundation and adaptation method. Test general answers, document grounding, tool limits and personality behavior on held-out prompts. Measure first response, completion speed, peak memory and voice latency on the target machine. Verify offline operation after installation. Review distribution rights, artifact checksums, signing and rollback. Only then deliver an approved artifact through a private channel.

Device readiness gates

Reference software and a functional P0 should precede custom production hardware. Engineering/design/production validation must establish audio quality, sustained thermal behavior, battery use, update safety and manufacturing checks. The existing 2027–2030 roadmap is conditional planning material, not evidence that these gates are complete.

Financial and partner claims

The public guide is free of AI-provider per-answer charges because it performs local lookup. Hosting can still have usage costs. Token terms, revenue allocations and exchange milestones remain proposals unless separately verified. Neither private-evaluation access nor an NVIDIA technology reference creates equity, a guaranteed return or an announced partnership.

04 / DATA, OPERATIONS & SOURCES

Data handling

By default the guide keeps conversation history in the tab. Remember is opt-in browser persistence, not encrypted account storage. Export and reset are available. No conversation is used for training. A prepared request remains local until explicitly shared. Browser dictation/online voices and ordinary hosting metadata have separate privacy implications [1,2].

Deployment controls

Deploy the complete static folder using HTTPS. No model-provider environment variable is needed. The publisher should set a real evaluation-contact email in config/private-access.js if requests are to be addressed automatically. Do not invent that contact. Replace placeholders for legal identity/contact before enabling public personal-data collection.

Documentation precedence

This five-page update describes the current public/private split. The original v1.6 combined PDF is retained unchanged for page citations. Historical hosted/local web architectures, prototype specifications and economic scenarios are dated references; they are not evidence of current implementation or release availability.

References

[1] MDN, SpeechRecognition. Limited browser availability; recognition can use a server service. https://developer.mozilla.org/en-US/docs/Web/API/SpeechRecognition

[2] MDN, SpeechSynthesis.getVoices(). Voices are provided by the current device/browser. https://developer.mozilla.org/en-US/docs/Web/API/SpeechSynthesis/getVoices

[3] Apache Software Foundation, Apache License 2.0, especially redistribution and notice obligations. https://www.apache.org/licenses/LICENSE-2.0

Project references: bundled POCKET Whitepaper v1.6 Combined (63 pages) and its dated technical/economic appendices. Technical references checked 05 October 2026.

Release status

Public Guide 10.0 is implemented in this package. Full Mind remains a private development/evaluation programme; no new proprietary weights are included. Performance and voice on real visitor hardware must be validated independently of interface tests.