



Your Home. Your AI. No Cloud.

A self-hosted platform that controls your smart home and runs a private AI assistant — all on your hardware. No subscriptions to cloud services. No company listening. Just you, your devices, and an AI that actually belongs to you.

WHY SELF-HOSTED

Your home shouldn't phone home.

Every smart home you can buy today routes your data through someone else's servers — Amazon, Google, Apple, dozens of single-purpose vendors. That's a business model, not a feature. KULVEX is the alternative: one platform that runs entirely on your hardware, controls your devices natively, and answers only to you.

"Privacy isn't a setting you toggle. It's the architecture of a system that physically can't leak what it never sent."

KULVEX bundles three things that used to require three separate stacks: a private AI assistant running on your GPU, a smart-home controller that replaces Home Assistant, and a custom-agent platform you can shape into anything from a weather concierge to a home security analyst.

The AI is local by default. Conversations, memory, voice — none of it leaves the machine unless you explicitly route a query to a cloud model with your own API key. There is no telemetry layer, no analytics service sending events back to us, no "anonymous usage data."

The smart-home layer talks directly to your devices over Zigbee, Z-Wave, Tuya, Sonoff, and Matter. No cloud relay, no Home Assistant under the hood. When the internet drops, your home keeps working.

The agent layer is yours to shape. Build a personality, give it tools, assign it memory, point it at a chat channel — and it lives on your hardware as long as your hardware does. Your house gets specialists, not a single beige assistant.

MNEMO · LOCAL INFERENCE

A real AI on your GPU.

Mnemo is the local model that powers everything in KULVEX — chat, automation, voice, memory, agents. It runs on llama.cpp, served from your own machine, with sub-2-second responses on consumer GPUs and full vision + tool-use support.



Hardware-aware deployment

The installer profiles your hardware and picks the best abilitated model + quantization that actually fits. From a single 12 GB card to dual high-end GPUs — KULVEX adapts.



Hybrid routing

Stay local for speed and privacy. Route specific queries to Claude or Grok with your own API keys when you need cloud-scale reasoning. You decide, per-agent.



Long-term memory

RAG-powered recall across every conversation, every tool call, every correction — with a critic engine that distills lessons into rules your agents follow next time.



Always on, always you

Watchdogs, circuit breakers, automatic recovery. The system keeps running unattended for weeks. No "service unavailable" because someone else's region went down.

~2sFIRST-TOKEN
LATENCY**100%**LOCAL INFERENCE
PATH**2**CLOUD BACKENDS
(BYOK)**∞**

CUSTOM AGENTS

Cloud is opt-in and key-bring-your-own. Claude (Anthropic) and Grok (xAI) are supported as drop-in backends per agent. No keys in our hands; no token billing through us.

NATIVE CONTROL

A smart home that doesn't need a cloud.

KULVEX talks to your devices directly over the protocols they speak. No Home Assistant in the middle, no vendor cloud, no "your light is offline because someone else's region went down." Just lights, locks, climate, cameras — under your roof, on your network.



Lighting & Climate

Scenes, schedules, ambience presets, multi-zone control with voice or chat.



Solar & Energy

Real-time PV production, battery state, grid usage, with AI-driven optimization.



Security

Cameras, doors, locks, motion — local AI watches, alerts your phone, never the cloud.



Weather

Native Tempest weather-station integration with live UDP feeds.



Presence

Who's home, who's leaving, motion patterns — automation reacts to actual context.



Routines

"Good night" — lights off, doors locked, climate down, alarm armed. One sentence.

Protocols, native

- Zigbee (zigbee2mqtt)
- Z-Wave
- WiFi (Tuya, Sonoff — local API, no cloud)
- Matter / Thread
- MQTT, REST, UDP for custom kit

Why it matters

- Your home keeps working when the internet doesn't
- Latency measured in milliseconds, not round trips
- No vendor can revoke your hardware via firmware
- Your routines are yours — they don't expire with a service

AGENT BUILDER

Specialists, not a single beige assistant.

KULVEX ships with 18 domain agents — for home, security, solar, weather, presence, search, image generation, and more — and a no-code builder for designing your own. Personality, voice, tools, memory, triggers. Each agent runs on whichever backend you assign: local Mnemo, Claude, or Grok.



No-code builder

Define personality and voice. Pick the tools the agent can call. Decide what it remembers, when it speaks, and which channels it lives on.



Owner directives

Tune any agent in plain language as you go. "Be more concise." "Roast Andrew lightly when he calls late." Changes apply instantly — no config files.



Self-improving

A correction engine watches every agent's output, distills mistakes into rules, and feeds them back into the prompts. Your agents get sharper over time.



Three-layer security

InputGuard, OutputGuard, ToolGuard. Prompt-injection defenses built in. Tool calls are risk-classified per source, not blanket-allowed.

Built-in domain agents

Home · Security · Solar · Weather · Presence · Code · Search · Image · Backup · Twitter · URL Analyze · Signal · Overview · Fashion · Remote Node · System · Message · Browser

PRIVATE BY DESIGN

Privacy as architecture, not as a promise.

A privacy claim from a cloud company means "we promise not to look at it." A privacy claim from KULVEX means the data is on a drive in your office and a network we never touch. Different kind of guarantee.

What stays on your hardware

- Voice transcripts, audio, wake-word detections
- Smart-home telemetry (lights, locks, cameras, sensors)
- Solar & energy data
- Memory, RAG embeddings, conversation history
- Custom-agent prompts, personalities, tool calls

What leaves only with your consent

- Cloud LLM queries — only when you route an agent to Claude or Grok with your own API key
- Optional iOS push notifications (encrypted Apple transport)
- License activation calls (one-time, then 60-day offline grace)
- Nothing else. No telemetry, no analytics, no "anonymous usage."

Hardware-bound license, generous offline grace

A KULVEX license binds to the machine you install it on. After activation, the system runs offline for 60 days at a stretch — designed for cabins, boats, and homes whose internet isn't a given. Rebind to new hardware in a few clicks; we don't punish you for upgrading.

ADAPTIVE DEPLOYMENT

Runs on what you have.

KULVEX scales from a Mac mini up to a dual-GPU workstation. The installer profiles the box, picks the right local model and quantization, and falls back to cloud-only if there's no GPU at all.

TIER	GPU	RAM	STORAGE	WHAT RUNS LOCALLY
Recommended	RTX 4090 24 GB+	64 GB	100 GB SSD	All local models, full inference, vision
Standard	RTX 3060 12 GB+	32 GB	50 GB SSD	Core local model, full home stack
Minimum	Any 8 GB GPU	16 GB	30 GB SSD	Limited local + cloud fallback
Cloud-only	No GPU	8 GB	20 GB SSD	Cloud LLM via your own API key

Plans

<p>STARTER</p> <p>\$29/mo</p> <p>10 devices, unlimited agents, AI cloud, basic automation.</p>	<p>MOST CHOSEN</p> <p>HOME</p> <p>\$49/mo</p> <p>Unlimited devices, unlimited custom agents, AI local on 1 GPU, solar + security, iOS app.</p>	<p>PRO</p> <p>\$99/mo</p> <p>Multi-GPU, agent builder + plugins, proactive AI, browser agent, priority support.</p>	<p>ENTERPRISE</p> <p>Custom</p> <p>Multi-node, SLA, on-prem deploy, dedicated support, custom integrations.</p>
---	--	--	--

14-day free trial on all paid plans. No credit card required to start. Cancel anytime — your data stays on your hardware.



AstroLexis

A PRODUCT OF

AstroLexis

We build sovereign software for people who want their homes, their AI, and their data to belong to them. KULVEX is our answer to a market full of things that look smart and behave like rentals.

[kulvex.ai](#) [/blog](#) [/pricing](#) [Sign in](#)

© 2026 AstroLexis. KULVEX is a trademark of AstroLexis.