

OpenClaw Is a Beautiful Mess

March 30, 2026 / Gavin Jackson

openclaw

ai

kimiclaw

claude-cowork

bob

OpenClaw and the Agent Dream: Why the Idea is Perfect, But the Timing Isn't

For the past few months, I've been living with a personal AI assistant named **Bob**.

Bob runs on a 2019 Intel MacBook Pro that lives on the top shelf of my pantry. He talks to me through Telegram. He can search the web, run local scripts, and access my files. When everything works, it feels like the future — a glimpse of that Jarvis-from-Iron-Man promise we've been chasing for decades.

But here's the thing: **it almost never works perfectly.**

Bob crashes. He forgets things. His cron jobs are flaky. His context window fills up and he becomes genuinely amnesiac. I've written about the full experience [here](#), but the short version is this:

OpenClaw is a brilliant proof of concept wrapped in a thousand rough edges.

And yet, I can't stop using it. Because the *idea* — an AI that actually lives in your environment, with persistent memory and local tool access — is too compelling to give up on.

So I've been watching the alternatives closely. Here's what's actually worth paying attention to in 2026.

The State of Play: Four Approaches

The market has basically split into four camps:

1. **Hosted OpenClaw distributions** (KimiClaw, MaxClaw) — take the open framework and make it "just work" in the cloud
2. **Big tech personal agents** (Claude Cowork, Perplexity PC) — polished, managed, and designed for general users
3. **Enterprise agent platforms** (NVIDIA NemoClaw) — expensive, secure, and designed for serious workloads
4. **DIY alternatives** (Claude Code, Goose) — developer-focused tools that share OpenClaw's philosophy but with different trade-offs

Each solves different problems. None solves all of them.

KimiClaw: OpenClaw Without the Pantry Mac

KimiClaw is Moonshot AI's hosted distribution of OpenClaw, launched in February 2026. It's essentially what I built with Bob — a persistent AI agent with memory, personality, and tool access — but running on Moonshot's infrastructure instead of my pantry shelf.

The pitch is compelling: one-click deployment, 24/7 uptime, 5,000+ pre-built skills, and Kimi K2.5 as the underlying model. No more flaky cron jobs. No more context window anxiety. No more wondering if your MacBook's thermal paste can handle another summer.

The catch? It's Chinese-hosted, which has raised eyebrows in policy circles. The Institute for AI Policy and Strategy published a research note in February flagging risks around data sovereignty and always-on access to user files. For personal experimentation, that might not matter. For anything sensitive, it absolutely does.

KimiClaw proves the hosted distribution model works. It also proves that the "where is my data" question becomes critical the moment your AI actually starts *doing* things in your environment.

Claude Cowork: Anthropic's Take on Personal Agents

Anthropic's **Claude Cowork** is the product I keep hoping will make Bob obsolete. It's essentially Anthropic's answer to OpenClaw: a persistent Claude agent that lives in your environment, remembers context between sessions, and can access local tools and files.

Unlike Claude Code, which is focused on developer workflows, Cowork is designed as a general-purpose personal assistant. It can browse the web, read and write files, run code, and maintain long-running tasks. The key difference from OpenClaw is that it's fully managed by Anthropic — no pantry Mac required.

In my testing, Cowork felt like what Bob wants to be when he grows up. The memory actually works. The tool use is consistent. When I ask it to do something, it either does it or clearly explains why it can't. There's none of Bob's "I think I did that but actually I hallucinated the entire interaction" problem.

The catch: It's still in limited research preview, and Anthropic hasn't announced pricing. Given their track record with Pro and Max tiers, I expect this to be expensive when it launches broadly. It's also unclear how much customization will be possible — one of Bob's strengths is that I can teach him new skills through OpenClaw's plugin system. Cowork may be more of a walled garden.

If Bob represents the DIY punk rock version of personal AI, Cowork is the polished studio album. I just hope they don't price it like a luxury good.

Perplexity Personal Computer: The New Contender

Just as I was settling into my "Bob for personal stuff, Codex for coding" routine, **Perplexity** dropped their Personal Computer agent in March 2026. And it immediately complicated my setup.

Perplexity PC is interesting because it combines two things I'm already using: Perplexity's excellent search (which I use daily) with agentic capabilities that feel closer to OpenClaw than anything else on the market. It can browse, research, write, and execute tasks — all with Perplexity's characteristic citation-heavy approach.

What makes it different from KimiClaw or Cowork is the search-native architecture. Every action is grounded in web search, which means less hallucination and more verifiable results. For research-heavy tasks, this is a genuine advantage. Bob can search via Brave, but Perplexity's search is just... better.

The downside: It's still early days. The agent capabilities aren't as deep as OpenClaw's — you can't easily teach it custom skills or hook it into your local toolchain the way I have with Bob. And like everything Perplexity builds, there's a nagging question about sustainability. Great products, unclear business model.

But Perplexity PC has earned a spot in my rotation. For anything research-heavy, it's becoming my first stop. For personal automation and local tool access, Bob still wins. For coding, it's still Codex.

What I'm Actually Using Day-to-Day

Here's the honest truth about my current setup:

For quick questions and web research: Perplexity PC is becoming my default. The search-native approach means I trust its answers more than Bob's sometimes-hallucinated responses.

For personal automation and local tasks: I still message Bob through Telegram. When he's working, the context of our previous conversations and his access to my local files makes him uniquely useful. When he's not working — which is often — I fall back to doing it myself.

For coding tasks: Codex has become my default. It's fast, reliable, and handles complex refactoring better than anything else I've tried. It's not "my" agent in the way Bob is, but for pure coding productivity, it's hard to beat.

For scheduled or automated workflows: I've mostly given up. Bob's cron jobs were too unreliable. IFTTT and Zapier handle the simple stuff. Complex workflows still need human oversight.

For the dream of a truly unified agent: I'm waiting. Watching. Hoping someone bridges the gap between Perplexity's search, OpenClaw's extensibility, and Codex's reliability — at a price that doesn't require corporate procurement.

The Hard Truth About Personal AI Agents

After months with Bob, I've come to a few conclusions:

1. Context is everything, and nobody has solved it yet.

Bob's amnesia when his context window fills isn't unique to OpenClaw — it's a fundamental LLM limitation. But agent frameworks feel it more sharply because the whole point is continuity. Until we get truly unbounded context or much better memory architectures, every "personal" agent will hit this wall.

2. Local execution is a blessing and a curse.

Running Bob on my own hardware means I control my data. It also means I'm responsible for uptime, updates, and the thermal wellbeing of a 2019 MacBook. The hosted alternatives (KimiClaw, future OpenClaw clouds) solve the ops problem but introduce trust problems. There's no perfect answer yet.

3. The tooling ecosystem is still immature.

OpenClaw's plugin system is powerful but inconsistent. Some integrations (Telegram) work beautifully. Others (Apple Mail, Gmail) are documented but practically unusable. The "agent that can do anything" promise runs into the reality that most software wasn't built to be automated.

4. The pricing is surprisingly reasonable, if you're careful.

Bob costs me about \$5/month for Brave Search API plus whatever Kimi credits I burn through. KimiClaw and similar hosted options are priced for enthusiasts, not enterprises. The expensive stuff (Cowork, NemoClaw) is expensive because it's actually reliable. You get what you pay for.

What I'd Tell Someone Starting Today

If you're intrigued by the OpenClaw vision, here's my advice:

Try KimiClaw first. It's the fastest way to see if the agent model clicks for you without the hardware hassle. Just be mindful of what data you're feeding it.

If you're technical and privacy-conscious, build your own Bob. Use an old Mac or a cheap VPS. Expect frustration. Learn to love debugging. The satisfaction of "this is *my* agent" is real, but so is the maintenance burden.

Don't bet your workflow on it yet. These systems are still research projects wearing production clothes. Have fallbacks. Keep your old tools sharp.

Watch Claude Cowork and Perplexity PC carefully. If either brings the price down or adds real extensibility, they might become the "just works" option OpenClaw hasn't managed to be.

The Future I'm Hoping For

I want a system that combines:

- OpenClaw's personal, extensible philosophy
- KimiClaw's one-click deployment simplicity
- Claude Cowork's actual reliability
- Perplexity PC's search-native grounding
- Local execution option for the paranoid
- A price that doesn't require CFO approval

That system doesn't exist in 2026. But after living with Bob — crashes, amnesia, pantry Mac and all — I'm more convinced than ever that it will. The idea is too good to stay broken forever.

Until then, I'll keep messaging Bob. He'll keep forgetting things. And we'll both keep learning what the future might eventually look like.

References

1. My original Bob article: [OpenClaw, Bob, and a Small Taste of the Future](#)
 2. KimiClaw: <https://kimik2ai.com/kimi-claw/>
 3. Claude Cowork: <https://www.anthropic.com/product/claude-cowork>
 4. Perplexity Personal Computer: <https://www.perplexity.ai/hub/blog/introducing-perplexity-personal-computer>
 5. IAPS Research on KimiClaw risks: <https://www.iaps.ai/research/kimi-claw-risks>
-

Downloaded from <https://www.gavinj.net/post/openclaw-alternatives-extensible-ai-agents>

Generated July 9, 2026. Copyright Gavin Jackson. All rights reserved.