

The Off-Switch Is Not Enough

Why “switch off the AI” is the wrong thing to ask for — and what the third option looks like

Précis. A countermovement against everywhere-by-default AI has reached the mainstream, and its demand is an off-switch: a toggle, an opt-out, a premium tier whose feature is *absence*. That demand is right about the discomfort and wrong about the remedy. An off-switch quietly concedes that “AI” means the kind built by a handful of large companies, hosted on their infrastructure, trained on your behaviour, and that your only freedom is to decline it. There is a third option between that AI and none: AI that belongs to the community that runs it, on infrastructure it controls, under rules it sets, with nothing leaving the building. We have spent two years building it. This essay is why the off-switch is a poverty of ambition, and what the alternative is. It opens a short series on that alternative.

A movement is forming, and the tech press has named it. In *Android Authority* this June, under the headline *Why AI-free tech is growing*, the case is that the absence of AI has become a selling point, the way ad-free once was. The reader poll is not close: **64% want fewer AI features, 23% want better controls, just 2% want more.** The products it points to share a single move, switching the AI down:

- a search engine that leaves it optional;
- a browser that adds it with restraint;
- a note app that keeps it to a plugin;
- a chatbot that won’t touch your data unless you paste it in.

The pattern is familiar. Ad-bloat produced ad-free subscriptions; surveillance produced privacy-first products; AI-everywhere is producing its counter-market.

The discomfort is real and the diagnosis sound. The remedy is what deserves a harder look, because every product on that list offers the same one: a way to turn it off.

What the off-switch concedes

Reach for an off-switch and you have agreed to two things without noticing. First, that AI is something done *to* you: a feature that arrives switched on, that you can at best decline. Second, that “AI” means one specific thing — large models, built by a few large companies, running on their cloud, improved by watching what you do. With those in place, the only move left is abstention. The off-switch is the dignified version of having none.

That is a poverty of ambition dressed as a principled stand. It treats a structural problem (who owns the intelligence, where it runs, what it learns from, whom it answers to) as a matter of personal preference, soluble by a toggle. And it leaves the structure intact. The company still owns the model. It still runs on their infrastructure. The toggle means only that they are not running it on *you* this week. You have declined a serving; you have not gained control. Next year the toggle moves, defaults back on, or becomes the paid tier, and the negotiation restarts from the same weak position, because the thing was never yours.

The demand stops at the off-switch for a reason. For most people abstention is the only option on the menu, because the menu is written by the companies the movement objects to. You cannot ask for AI you own from a vendor whose business is owning it for you. So the want shrinks to the one thing a vendor can grant: a setting. The movement has the problem right. It has been handed too small a vocabulary to name it.

The third option

The choice was never *AI or no AI*. The real questions are *whose* AI, on *whose* infrastructure, trained on *whose* terms, answering to *whom*. Ask those and a third option appears, one the off-switch crowd mostly doesn't know exists: AI that belongs to the community that uses it.

This is what we build. A Village is a platform where a community (a parish, a whānau, a society, a club, a board, a team inside a larger organisation) runs its own AI, on infrastructure it controls. The model is not a line to someone's cloud; it runs in-house, and the community's content never leaves to train anyone else's product. It is also *situated* rather than averaged. A Big Tech model is built to be everyone's, which makes it no one's in particular: it knows the global average, not your parish, your paddock, your boardroom, or your department. A model that sits inside one community is tuned to it, and knows the context because the context is all it is for. There is no toggle to protect you from the vendor, because there is no vendor in the relationship. You are running your own.

The difference, put plainly, is custody. With the off-switch the intelligence belongs to the company, and you negotiate how much of it touches you. With the Village it belongs to the community, and the only question is what it should do. Moving custody from vendor to community dissolves most of what the AI-free movement is reacting to:

- no surveillance to opt out of, because no data leaves;
- no creeping default, because the community sets the defaults;
- no premium for absence, because absence was never the product.

You do not have to choose between intelligence and control, which is the choice the off-switch forces. You keep the AI and make it yours. A small model that belongs to the community can answer questions, draft notes, summarise a long thread, help a newcomer find their footing, without any of it being a window the vendor watches through. The off-switch gives up the usefulness to escape the watching; the Village keeps the usefulness and removes the vendor.

Why you wouldn't have to switch it off

A fair objection: an AI you own is still an AI, and the unease is not only about who profits. It is about systems that decide things they shouldn't, that act before anyone has thought, that flatten judgement into a default. Owning the model does not fix that on its own. So a second principle sits alongside custody, and together they make the off-switch beside the point.

We build the AI to know its place. Some things a machine may do; others it must not, and the line is not a matter of taste. Values cannot be automated, only verified by the people who hold them. Agency cannot be simulated, only respected. Where a question cannot be reduced to a rule, it returns to human judgement rather than a confident machine. The architecture enforces this; it is not a hope printed on a values page. The AI in a Village proposes and assists. It does not decide what matters, and it cannot act on what matters without a person in the loop. The premium here is intelligence that cannot quietly take authority it was never granted.

The off-switch can't offer that. A toggle is a confession that the AI cannot be trusted, a way to hold it at arm's length. An AI built to know its place can be trusted closer, because its reach is fixed in the building rather than by your vigilance. You do not have to switch it off to be safe from it.

Who this is for

Three kinds of people will recognise themselves here.

- **Community groups, small businesses, and the bodies that govern them** — the club, the trust, the parish, the practice, the board that meets monthly, the regional office. They were never going to build an AI department; until now the choice was a Big Tech product that monetises them, or no help at all. Governance feels it most sharply: a board carries duties and a paper trail, and “we pasted it into a chatbot” is not a defensible record.

- **Security-minded teams inside larger organisations**, who already know “it’s in the cloud” is not an answer for sensitive material, and need intelligence that cannot leave the room.
- **The principled** — the growing segment the AI-free movement is made of, reaching for the off-switch because it is the only lever they have been handed. To them the message is simplest: there is a bigger lever, and it is already built.

That segment will keep growing, because the conditions producing it are not going away. As long as the dominant model of AI is *theirs, watching you*, the want for an alternative compounds. The off-switch satisfies some of it for a while. But a demand for absence is unstable: what people want is not less intelligence, but more control over it, and absence is only the shape that want takes when control is off the menu. Put control back on the menu and the want finds its true form.

What comes next

This essay makes a claim and owes you the proof. The rest of the series gives it, one piece at a time:

- what it means to build an AI that knows its place, and why the boundaries are the premium;
- *your model, your walls*, for teams whose data cannot leave;
- running, branding, and federating your own Village on terms you can revoke;
- the control tower we built for operators: oversight that reports health and never reads members’ content;
- and the words themselves, made hard to say, and worth taking back.

For now, one idea to carry out of this: the off-switch is not enough, because it leaves the thing it objects to intact and asks only to be spared. The alternative is not less AI. It is AI you own, built to know its place, that never had to be switched off because it was yours all along.

You can walk through it rather than take it on description. The Village runs as live demo communities at mysovereignty.digital/demos.html, and the free Governance course sets out the case in full. If the whole idea is new, the Your Community, Your AI article series opens with what AI actually is.

*The Village is a running system, not a brochure — see it at mysovereignty.digital. Source for the AI-free movement and poll figures: “Why AI-free tech is growing,” *Android Authority*, June 2026. — John G. Stroh, *My Digital Sovereignty Ltd.*, June 2026.*