Your Model, Your Walls

For teams who can't put sensitive material in someone else's cloud: an AI that runs inside your own walls and never phones home.

Précis. For a counselling service, a case-management team, a board handling commercial-sensitive matters, a hauora provider, or a security-minded unit inside a larger organisation, "it's in the cloud" is the problem stated politely. Most AI in business software reads your text by sending it to a vendor's model, on infrastructure you don't control, under a foreign jurisdiction. The Village keeps the model where the data already is: each community runs a modest model of its own, served from sovereign infrastructure on New Zealand and EU ground, and nothing leaves to train anyone else's product. In June we removed the last code path by which a Village could call an outside model at all. This essay is what *your model, your walls* means in practice.

Some organisations have found the whole AI conversation beside the point, because it starts with a disqualifying assumption: that to use AI you send your material to someone else's computer. A counselling note. A child-protection file. A board's strategy before it is public. A patient record. For the people responsible for these, the question was never "which AI assistant is best." It was "how do I get help with this without it leaving the room." From the major vendors, the answer has been: you don't.

The cause is the architecture itself. A cloud AI feature reads your text by transmitting it to the vendor's model, where it is processed under a policy you must take on trust and is, at the moment of use, visible to a party that is not you. You can read the data-handling page, tick the enterprise box, believe every word, and the fact is unchanged: the sensitive thing left the building so the machine could read it. For ordinary material that risk is manageable. For sensitive material it is the whole problem, and no toggle fixes it, because the toggle only governs what a vendor does with data it already holds.

Where "the cloud" actually takes your data

"In the cloud," for this kind of work, means three things:

- the record sits on infrastructure owned by one of a few very large companies, most reachable under foreign law wherever you and your members are;
- the AI reading it is the vendor's model, on the vendor's hardware;
- your text may be retained, may be reviewed, may shape the next version of the model, and is at least legible to the vendor at the moment of use.

None of that is malice. It is the shape of the deal: the intelligence belongs to them, runs on their ground, and your material has to travel to it.

For a security-minded team, that shape is disqualifying. Their threat model is one question: can I give an account, to a regulator or a court or the person whose file this is, of every place this information has been? Once the answer includes "and then it went to a model in another country," the account is broken. A better promise about data that has already left is worth nothing to them. They need the data not to leave.

A model inside your walls

The Village keeps the model where the data already is. Each kind of community runs a modest model of its own, an open model of around fourteen billion parameters fine-tuned for the work that community actually

does, on sovereign infrastructure in New Zealand and the EU, never a hyperscaler's cloud. An organisation that needs to can have the Village deployed on infrastructure of its own. There is no call to OpenAI, Google, or Anthropic anywhere in the path. A small model still does the useful things: answer a question, summarise a long thread, draft a note for a human to approve, help a newcomer find their footing. All of it happens inside the walls. How each model is fine-tuned, served, and kept off any outside cloud is set out in the Village AI explainer at mysovereignty.digital/village-ai.html.

In June we shut the last door. We removed `external_ai` from the purposes a member can consent to, so the audit log now records only a Village's own engine, or a local simulation, as the source of any AI output. There is no longer a code path by which a Village reaches an outside model. "The data doesn't leave" is no longer a privacy policy you have to believe; it is a property of the software, because there is nowhere for the data to go.

A word of care, because the wrong version of this is a lie. Our fourteen-billion-parameter model does not out-think the largest frontier systems, and the corpus has always said so. Those models are extraordinary; a community model is modest beside them. The offer here is a different one: intelligence that never leaves your custody. For genuinely sensitive work, that beats a cleverer model you have to mail your files to. A brilliant answer that needed your case file to leave the building is the wrong answer.

Situated, not averaged

A model that lives inside one community has a quieter advantage, and it cuts against the assumption that bigger is always better. A frontier model is built to be everyone's at once, which makes it, in the particulars, no one's. It knows the average of how the whole internet discusses a thing, not how your parish, your practice, or your case team actually works. A Village model is fine-tuned for the kind of community it serves and runs under that community's own rules. Seeing only your kind of work, in your context, it gets right the particulars a global average smooths away: the local vocabulary, the specific obligations, the way your people phrase things. And none of those particulars become training data for someone else's product. Situated, not averaged.

How far into your walls you can pull it

How much separation a community needs is its own call:

- most run on shared infrastructure, each community's data partitioned in software from every other, which is enough for the great majority;
- those who want more can take a database of their own, dedicated to them alone;
- those who will not put their data on anyone else's machine at all can have the Village deployed on infrastructure of their own.

Across all three, the constant holds: the data stays on sovereign ground, off every hyperscaler's cloud, and the model reading it is the community's own. The platform serves the knitting circle and the case-management unit alike; what differs is only how much separation you have asked for.

The walls also run inside a community, where the work demands it. For counselling, case management, mentoring, and recruitment, the Village has a mode where ordinary members reach only the moderators helping them, not one another. The person seeking help is never exposed to a roomful of strangers in the same situation; the support relationship stays private. The same instinct, applied to the shape of the community rather than its infrastructure. This hub-and-spoke mode runs in the recruitment demo at recruitment-demo.mysovereignty.digital.

And because it is yours, it looks like yours: your logo, your colours, the vocabulary your people actually use rather than someone's product jargon, your own domain, your own mail.

Who needs walls this thick

Most communities will be fine at the lighter end. For the ones who need the thick walls, the need is absolute, and they know who they are:

- the security-minded team inside a larger organisation, who have already concluded "it's in the cloud" is no answer for what they hold;
- the counselling and hauora services, whose duty of confidence is the basis of the work;
- the board handling matters that are commercial-sensitive or legally privileged;
- the unit answerable to a regulator who will, one day, ask for a full account of where a piece of information has been.

For all of them the offer is narrow and provable: the model runs inside your walls, your material does not leave them, and there is no longer a door through which it could.

This is the data-custody half of the series. Earlier essays argued the AI should be yours and should know its place; this one is what *yours* means when the material is too sensitive to mean anything less. The next pieces turn outward: running and branding your own networked Village, federating with others on terms you set and can revoke. They all rest on this floor. The walls come first.

The Village is a running system, not a brochure — see it at mysovereignty.digital. The Village runs its own models on self-hosted inference, with no external-AI code path; "the data doesn't leave" is a property of the architecture, not a policy promise. — John G. Stroh, My Digital Sovereignty Ltd., June 2026.