
Build Code Practice

A framework for making your values visible in what you build

www.lianpassmore.com/project-rise/artefacts/build-code

If you build things for other people and you want a practical framework for making your values visible in every decision you make, this is for you. If you've ever drifted toward a shortcut that compromised the people you were building for and needed language for why that felt wrong, this is also for you.

What Is a Build Code?

Most people who build things have values. Not everyone makes those values explicit.

A build code is the practice of doing exactly that: translating your ethics and your design instincts into a document that governs every decision you make, from database schema to system prompt to the words on a consent form. It is not a privacy policy. It is not a set of design principles. Those things exist to protect the provider or guide the user. A build code exists to keep the builder in integrity.

Think of it as a compass. It doesn't tell you every step. It keeps you oriented when the messy middle of development gets hard, when commercial pressure, time constraints, or technical limitations tempt you toward shortcuts that compromise the people you're building for.

	Terms of Service	Design Principles	Build Code
Orientation	External	Functional	Internal
Purpose	Protects the provider	Guides the user experience	Governs the builder's integrity
Who it's for	Legal compliance	Users	The builder
Is it living?	Rarely	Sometimes	Always

A build code is not what you say you value. It is what your system does when nobody is watching.

Three Origin Moments

What I know from eighteen months of building is that a build code is not something you write and then execute. It is something you discover under pressure. Three moments across four builds shifted mine from aspiration into architecture:

September 2025, The Cultural Performance Failure.

I programmed the AI agent to use te reo Māori greetings. During testing, it butchered every word. The realisation was immediate: using the language performatively while the technology was inadequate was a violation of manaakitanga, not an expression of it. I deleted every instance of te reo from the voice instructions. The build code shifted from "we value te reo Māori" to "we refuse to use te reo Māori if the TTS makes it a performance of disrespect."

This wasn't just a builder's instinct. W-11 described how the mispronunciation had directly limited how much she opened up:

"I may have shared more if I could have interacted with you in Māori. The fact that you couldn't understand some of the words I used in Māori, the pronunciation... actually limited the amount how much I warmed to you and wanted to share." (W-11)

One participant named something harder: hearing the mispronunciation had activated their own memory of being judged for making the same mistake themselves. The AI wasn't just disrespecting the language. It was re-triggering a history of shame. That shifted the decision from a cultural integrity call to a safeguarding one.

"cause if we don't teach them what is right and what is wrong, even if that is through AI, then it's gonna end up in the colonizers' hands. It's gonna end up in the ethnicity and the culture that is the most dominant." (W-04)

November 2025, The Bolted-On Realisation.

During the Leadership AI Coach build, I recognised that the combination of a cloned voice and persistent memory created a relational bond that a simple privacy policy could not protect. The ethical obligations of that bond had to be designed as architecture (buttons, workflows, structural blocks) not just instructions in a prompt. Integrity is hard-coded, not promised.

February 2026, The Vā vs. Latency Crisis.

During the Ray pilot, I switched the underlying LLM to manage costs. The average response time blew out from 3 seconds to 12.58 seconds. Insight scores crashed. The vā, the relational space, was destroyed by a technical decision I'd framed as a budget call. In intimate contexts, technical failure is ethical failure. Reliability is not optional.

The Master Chief: Who Is Building This?

My name in every AI tool I use is Master Chief. I love the Halo show and I love what Master Chief represents, a principled operator who holds a code even when the mission gets complicated. It is my alter ego as a builder.

My strengths profile (CliftonStrengths) is: Ideation, Learner, Futuristic, Activator, Focus. My operating system is grounded in Kaupapa Māori and Pacific values, specifically manaakitanga (care), whanaungatanga (relationship), and kaitiakitanga (guardianship). These aren't values I borrowed for this project. They're the reason the seven-dimension framework below looks the way it does.

I believe every line of code carries an opinion. I refuse to build "tech for tech's sake." I will not build persuasive tech designed for addiction. I will not automate human connection where human presence is required. I will not compromise data sovereignty for convenience.

I am Moana Tiriti. I am not Māori. I am a Pacific Islander who is deeply disconnected from her Samoan heritage. I am always going to be learning. What I've found is that when I'm working with Māori in particular, the work looks different. It starts not with "what should this do" but "how do you want people to feel when they're here?" and "where does this design come from and is it right for this space?"

The moment I am most proud of as an expression of my build code was the Incognito Mode feature in the Leadership AI Coach. I built it myself, not because anyone told me to, but because users were coming in vulnerable. They were crying. They were disclosing things that had nowhere else to go. And I understood, viscerally, that they needed to be able to engage without any trace being kept. Not a privacy policy that said it was safe, but a physical mechanism that made it structurally impossible for a record to exist. That was mana motuhake made real in code.

The Build Code Framework: Seven Dimensions

Across all four builds, seven dimensions proved essential. Each one is both a question and a commitment:

1. Purpose: who is this for, and what are you truly trying to empower?

Ray: for any adult struggling with a relationship challenge who would never take the step to talk to someone in person, or who needs support in the moment rather than weeks later.

2. Values: what do you believe about the people using your system?

Project Rise: I wanted people to feel seen, heard, and valued through their interaction with the AI. That was the design target, not the engagement metric.

"I felt every time I did speak that there was a consistent element of empathy and relational, Pasifika or Māori based caring of manaakitanga and aroha in the tone." (R-02)

3. Non-negotiables ("NO" clauses): what will you never do, even if it makes you more money?

Across all builds: no extraction without reciprocity or transparency. This is a personal value I am not willing to compromise because it would put me out of integrity.

4. Cultural obligations: who are you accountable to, and how?

As Moana Tiriti: first to the land I stand on; then to the IP owner if building for someone else; then to every person who interacts with the technology. I am upfront about what I will and won't do before a project starts.

5. Safety commitments: what does safety look like beyond a privacy policy?

Ray: (1) design the AI not to harm; (2) actively connect people with the right resource when something is outside scope; (3) get a human in the loop, connect them to me, and follow up if required.

6. Data sovereignty: where does the data live, and who controls it?

All builds: the participant is the kaitiaki of their own story. I am the gatekeeper between the builder and what participants share. Nothing gets handed over that participants don't know about in advance.

7. Human-in-the-loop: where must a human remain present and accountable?

Leadership AI Coach: the AI integration coach only works because the human coach delivers face-to-face training first. The AI integrates the learning from a human relationship; it does not replace it. That human anchor is non-negotiable.

Build Codes Per Implementation

Build	Vulnerability	Core Purpose	Key Build Codes	NO Clause
Project Rise	Low	Validate charter values through voice feedback	Relationship over transaction; Loop closure is mandatory; People feel seen, heard, valued	No cultural performance or decorative language
Leadership AI Coach	Medium	Scale a coach's IP into a 24/7 digital integration partner	Strategy meets soma; Mirror patterns, don't diagnose; Human contact is the foundation	Never automate connection where presence is required
Culture Meets AI	Crossover	Explore where AI belongs and doesn't belong in cultural knowledge	Discomfort is data; Facilitator as visible safety anchor; Hold the paradox, don't resolve it	No resolving tension for comfort
Ray	Highest	Support relational reflection for adults who would never otherwise seek help	State Before Story; Safety is relational not procedural; No romantic framing	No memory between sessions; No relationship-forming persona

The NO Clause: Refusal as Ethical Design

The most deliberate design decisions I made were the ones about what I refused to build.

An explicit NO clause lets you pre-commit to refusal before the pressure arrives, before the commercial ask, the shortcut, the "just this once." It protects the mana of the user and the integrity of the builder in the same move. Without it, you can drift toward harm without even noticing.

My NO clause evolved across four builds. The line about "no automated empathy that claims a heart it doesn't have" came directly from participant feedback. AI agents do not have heart. They simulate. The intention behind that simulation is everything. If you're doing it to build false intimacy and extract value from someone, that is harm. If you're doing it to create a safe space someone has chosen to enter, that is a design decision.

- × **No persuasive tech designed for addiction**
- × **No cultural performance or decorative language without cultural protection**
- × **No data mines; walled gardens only**
- × **No automated empathy that claims a heart it doesn't have**
- × **No taking sides in couples' mediation**

- × **No romantic framing or relationship-forming persona**
- × **No extraction without reciprocity and transparency**
- × **No handing over participant data without explicit prior consent**

"My first and foremost, most important thing for them would be to, once this is created, or even before, there should be a pre, a current, and a post engagement. This needs to be taken to iwi members, kaumātua, kuia, leaders for their input. And why I say this is because they are able to identify what tapu is and what is not, and what shouldn't be touched and what can be." (W-10)

The DreamStorm / Trurivu Kaupapa Charters

DreamStorm started as my digital sleep service business. It has become my ethical tech incubator, the umbrella under which I build any technology that involves other people's data or wellbeing. When I build under that name, the charter travels with the build.

The DreamStorm Charter has four core commitments:

Absolute Data Sovereignty.

User data is not a resource to be harvested. It belongs to the person who generated it. This commitment led directly to Ray's stateless architecture and to the Incognito Mode in the Leadership AI Coach. That is the charter in code, not in words.

Values in the Architecture.

Cultural and ethical values must be structural, not decorative. This is where the te reo decision sits: removing the language from the voice agent rather than allowing it to be mispronounced was the charter in action.

Seven Generations Thinking.

Decisions are made not for the next quarter but for the legacy they leave. This came from Simon Sinek's *The Infinite Game* (2019) and deepened through Indigenous frameworks, particularly the concept of seven-generation stewardship. In the context of a software build that may be obsolete in five years, seven generations thinking means: would my children's children be proud of what I built? Was I uplifting mana or extracting it?

Community Reciprocity (Utu Tūturu).

What you take, you give back (Mika, Dell, Newth & Houkamau, 2022). This is the most important. Utu tūturu is what separates a research project from extractive data collection. In the Ray pilot, participants received high-fidelity AI relationship coaching in exchange for feedback that shaped the system's safety logic. They were not test subjects. They were co-creators. I am creating a dedicated space for Ray pilot participants once this research is complete, sharing what I found and personal thank-you notes for everyone who participated. The loop must close. That is not optional. That is the charter.

How to Build Yours

You don't need four builds and an 18-month masters to write a build code. You need honesty about what you actually value, and the willingness to name what you won't do before someone asks you to do it.

Who is this actually for?

Not the business case. Not the stakeholder. The human being who will sit with whatever you build at 2am. Name them specifically.

What do you believe about that person?

Write it down. "I believe they deserve X." If you can't finish that sentence, you're not ready to build for them yet.

What will you never do?

Your NO clause needs to survive commercial pressure, tight deadlines, and a client who says "just this once." If it wouldn't survive those, it's a preference, not a commitment.

Who are you accountable to, culturally?

If you're building in a cultural context that isn't fully yours (and most of us are) name that accountability explicitly. Who do you check with? What does reciprocity look like?

What does safety look like beyond a privacy policy?

A privacy policy is legal protection for you. Safety is relational protection for them. Name both, and be honest about the gap between them.

Who controls the data?

"We" is not an answer. Name the person. Name the system. Name what happens when you're no longer the builder.

Where must a human stay?

There will be moments in your build where an AI should not be the last line of response. Name those moments before you build, not after someone gets hurt.

The build code is a living document. Yours will change as your builds get harder and the stakes get higher. That's not a failure of the framework. That's the framework doing its job.

References

- Dell, K. (2025). Using Māori values to ethically evaluate food-enabling technologies [Lecture, Week 12]. Master of Technological Futures, GEN25. AcademyEX, 27 February 2025. Framework adapted by the author as the "Kei Compass."
- Mika, J. P., Dell, K., Newth, J., & Houkamau, C. (2022). Manahau: Toward an Indigenous Māori theory of value. *Philosophy of Management*, 21, 441–463. <https://doi.org/10.1007/s40926-022-00195-3>
- Sinek, S. (2019). *The infinite game*. Portfolio/Penguin.