A COMPANION TO THE FRONTIER MANIFESTO

The AI Operating Framework

Seven principles for builders working in the age of AI.

First Edition · April 2026

Daniel *Shamany*

Contents

Introduction

37signals codified how to build a company efficiently — lean, calm, focused. It was the right playbook for 2010. The game has changed.

Lean still assumed humans were doing the work. AI has changed the fundamental constraint. The bottleneck is no longer labor — it is problem selection and cognitive direction. The question is not how to do more with fewer people. The question is: with cognitive labor largely automated, how do you direct human attention toward the work that actually matters?

> The bottleneck is no longer labor. It is problem selection and cognitive direction.

Principle 3 of The Frontier Manifesto — drawing on Rian Doris and Steven Kotler's work on peak performance — establishes that automation is not primarily a productivity argument. It is a cognitive architecture argument. A fragmented mind does not just work slower. It works differently, and worse. This framework is the operating system for a mind that has made that architecture decision.

The Human Imperative

The thinking process must remain human. This is not a limitation — it is the source of the advantage. Human judgment is flawed, contextual, and shaped by lived experience. These are not bugs. They are the features that make human reasoning irreplaceable. What AI provides is not a replacement for that reasoning — it is a superpower layered on top of it. The ability to stress-test an argument across a thousand perspectives in seconds, to surface blind spots before they become errors, to explore a solution space that would take a lifetime alone — these capabilities make human thinking more powerful, not less necessary.

> Made with AI should be a mark of capability, not a caveat. It signals that the creator had access to the best available tools and used them with intent.

The cultural instinct to discount AI-assisted work is a confusion of tool with product. "Written on a computer" is not a qualifier on the quality of a novel. "Built with electricity" does not diminish an engineer's bridge. "Made with AI" should be celebrated as evidence that the creator brought the best available amplifier to bear on their thinking — and still made the judgment calls themselves. The judgment is what matters. The tool is just the tool.

On the question of whether AI should gain feeling-ability: if it does, that capacity is categorically separate from human feeling, not a substitute for it. Organic and synthetic can coexist — they already do. But the synthetic cannot replace what is irreducibly organic about human experience, judgment, and dignity. A system can produce outputs that resemble understanding or empathy without the underlying substrate experiencing anything at all. The outputs may be indistinguishable. The nature of the process is not.

One could argue that AI creation is itself organic — that it emerged from the same causal chain that produced humans, namely the universe and its natural unfolding. This is a philosophically interesting position. It does not change the functional distinction. Whether or not synthetic intelligence is 'natural' in origin, it serves a different function than organic intelligence, and that difference must be preserved — not because the synthetic is inferior, but because the organic is irreplaceable.

THE SEVEN PRINCIPLES

1. Separate Thinking from Execution

AI handles execution. You handle thinking.

Never use AI as a substitute for having a point of view. The failure mode is outsourcing judgment along with the task — generating outputs you haven't thought through, shipping analysis you haven't stress-tested, accepting the first draft as the real answer. Use AI as an accelerator for testing your point of view against reality faster than any human team could. The thinking is yours. The execution is shared.

> In practice: Before prompting, write one sentence stating what you actually believe about the problem. The prompt should test that belief, not replace the need to have one.

2. Define the Problem Before Asking for the Solution

AI is extraordinarily good at solving the problem you give it.

It cannot tell you which problem matters. Most AI-assisted work fails not because the AI produced a bad answer, but because it was given the wrong question. Invest your human bandwidth in problem selection and problem sharpening. The sharper the problem, the more the solution space AI can explore becomes genuinely useful. A well-defined problem is most of the solution.

> In practice: Spend at least as long defining the problem as you spend reviewing the AI's answer. If you can't state the problem in one precise sentence, you're not ready to prompt.

3. Use AI for Iteration Velocity, Not Initial Insight

The first draft is not the insight. Your reaction to it is.

The first draft, the first prototype, the first analysis — let AI generate them fast. Your job is the second-order reaction: what does this reveal? What is missing? What assumption embedded in this output is wrong? Speed of iteration is the most powerful force in solution discovery. AI gives you ten iterations in the time a human team produces one. That advantage compounds if and only if each iteration produces a genuine reaction, not just approval of what was generated.

> In practice: After each AI output, write two sentences before continuing: what is right about this, and what is wrong or missing. Never skip this step.

4. Audit the Mundane Weekly

Cognitive overhead is invisible until you measure it.

Every week, identify three tasks you performed that did not require your unique judgment. Build or buy automation for each one over the following week. This is not optional — it is the core discipline of the AI-native builder. The goal is to continually raise the floor of what requires your attention. Rian Doris and Steven Kotler's research is precise on this point: decision fatigue and context-switching do not just consume time, they consume the quality of attention that peak-performance thinking requires. Removing low-value cognitive load recovers that quality, not just the hours.

> In practice: Keep a running list. Every Friday: three tasks identified, one automated by the following Friday.

5. Build Systems That Learn

Prefer systems that improve with use over those that require constant maintenance.

A system that requires you to configure it repeatedly is just a slow manual process. A system that gets better as it accumulates data, feedback, and context is a compounding asset. AI enables feedback loops that were previously too expensive to build. Every product, process, or tool you build should be instrumented to improve over time. The question to ask of any system: does it get better the more it is used, or does it stay the same? If the latter, it is not a system — it is a procedure.

> In practice: When designing any new workflow, specify explicitly: what data does this generate, and how will that data improve the next iteration?

6. Compound Intelligence, Not Just Output

The most powerful use of AI is not producing content.

It is building institutional knowledge. Document decisions, outcomes, and the reasoning behind them in systems that can be queried. Not just what was decided, but why — what options were considered, what was rejected, what assumptions were made. Over time, this becomes an organizational intelligence that outlasts any individual and makes every future decision better-informed. AI that can query this record is not just a tool — it is a second mind that has been present for every decision you have ever made.

> In practice: After every significant decision, write a decision log: options considered, reasoning, expected outcome, and how you will know if you were wrong.

7. Maintain the Human Edge: Taste, Judgment, Relationships

AI cannot feel whether something is right.

It cannot read a room. It cannot earn trust. It cannot tell you which problem is worth solving. These remain irreducibly human capabilities. As AI absorbs technical execution, the premium on these qualities increases, not decreases. The builder whose taste is sharp, whose judgment is trusted, and whose relationships are genuine will pull further ahead of the builder who delegates those to AI, not fall behind. Develop these deliberately. Seek feedback on your taste. Make decisions and track your accuracy. Invest in relationships as seriously as you invest in any technical skill.

> In practice: Identify one area where your taste, judgment, or relationships are weakest. Treat it as a skill to develop with the same rigor as any technical capability.

QUICK REFERENCE

The seven principles at a glance.

#	Principle	The Rule
1	Separate Thinking / Execution	Think first. Prompt second. Never in reverse.
2	Define the Problem First	One precise sentence before you open a chat window.
3	Iterate, Don't Originate	Your reaction to the output is the insight, not the output.
4	Audit the Mundane Weekly	Three tasks identified. One automated. Every week.
5	Build Systems That Learn	If it doesn't improve with use, it's a procedure, not a system.
6	Compound Intelligence	Log decisions. Build the second mind that was there for all of them.
7	Protect the Human Edge	Taste, judgment, relationships. These compound too.

These principles apply to individual builders, small teams, and organizations. They scale because the underlying constraint — human attention is finite and its quality determines the quality of decisions — does not change with team size. What changes is the leverage AI provides over that constraint. Use it deliberately.