

A LAUNCHREADY MANIFESTO

The White-Collar Factory is Closing

Why the people who never fit in are building what comes next

HARRISON PAINTER

LAUNCHREADY PRESS

THE WHITE-COLLAR FACTORY IS CLOSING

Why the people who never fit in are building what comes next

HARRISON PAINTER

LaunchReady Press

Copyright © 2026 by Harrison Painter

All rights reserved. No part of this book may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopy, recording, or otherwise, without prior written permission of the publisher, except as permitted by United States copyright law.

Published by LaunchReady Press
Indianapolis, Indiana
launchready.ai

Print ISBN: [assigned by KDP at upload]

First edition: 2026

Printed in the United States of America

The scanning, uploading, and distribution of this book without permission is theft of the author's intellectual property. For permission requests, contact harrison@launchready.ai.

Cover design by Harrison Painter. Typeset in Minion Pro.

Stylized portrait illustrations are AI-generated for editorial commentary purposes accompanying directly-attributed public statements.

*For the ones at the top who felt it breaking.
For the ones at the door who never got in.
You saw the same thing.*



EPIGRAPH

“You are not crazy. You were just early.”

HARRISON PAINTER

Founder, LaunchReady.ai

Indianapolis, May 2026

Executive summary (the 90-second version)

If you only have ninety seconds, here is the whole argument.

The corporate org chart most companies still use was built by the Romans for moving information across an empire when human beings were the only intelligence available. The industrial era applied that same hierarchy to knowledge work and produced what I call the **white-collar factory**. Meetings, layered management, anti-curiosity culture, executive function as the load-bearing skill. The math worked because the cost of human coordination was higher than the cost of running a hierarchy to organize it.

That math just changed. AI collapses the cost of coordination, drafting, scheduling, research, analysis, and a hundred other “executive function” tasks the white-collar factory was designed to organize. When those costs go to near zero, the value of layering humans to manage them also goes to near zero. The factory does not become slightly less efficient. It becomes architecturally redundant.

What replaces it is a different shape with four parts that have to be built together: **a shared brain** (the operating layer that holds context for humans, agents, and software), **custom internal software** built on

top of the brain, **a fleet of AI agents** running on top of that, and **a tiny team of three to ten human operators** orchestrating the whole stack.

The new human role is not the exotic “Context Farmer” the Silicon Valley essays coined. It is just management, with new failure modes. Agents stop working. Agents hallucinate. Agents go stale. The same skills that make a great manager of humans make a great manager of agents: scope the work, set context, monitor output, correct drift, replace what is not fitting. What changes is that the operator’s primary skill becomes **selection**, not execution. When AI collapses the cost of evaluating an idea, idea-generation volume becomes the new bottleneck, and the divergent thinker (the entrepreneur, the generalist, the person the factory called “too disruptive”) has the input the new game is starved for.

The cognitive profile that lost in the factory has a structural advantage in the new form. With one critical caveat. The advantage only converts to results if paired with an **execution system**. Generalists without an execution system are idea factories that never ship. The brain plus agent stack is the execution system that finally fits the cognitive profile, for the first time at scale.

The biggest mistake the Silicon Valley essays make is selling the **one-person billion-dollar company** as the new aspirational model. It is bad for humans. The longest-running scientific study on human flourishing (Harvard, eighty-five years) says relationships predict your health and happiness more than any other measured variable. The U.S. Surgeon General has called loneliness as deadly as smoking fifteen cigarettes a day. CEOs and entrepreneurs are already the loneliest professional cohort on the chart. Telling this population the new model is fewer humans and more agents is a public-health regression dressed in cap-table language. What actually works is **a small team of**

three to ten people who would go to war for each other, augmented by an agent fleet, building things larger humans used to need two hundred people to build.

The window to build this is open and short. The companies that adopted commercial internet (1995-2005), mobile (2007-2017), and cloud (2002-2012) compounded for decades against the ones that delayed. The same pattern is forming around the AI-native corporate form right now, in 2026. Operators who move in the next twelve to eighteen months will compound for the next decade.

If you are a CEO of a \$50M-\$500M company, the operator you need to lead this build is probably already on your team. They are the one who asks too many questions and has been “almost promoted” twice. Find them. Promote them. Hire two more like them, from the disruption class, before everyone else figures out what they are worth.

If you have spent twenty years on the outside of the white-collar factory because you could not stand the rules, the system that broke you was built for problems that no longer exist. You are now the input the new system is starved for.

The full manifesto is forty pages. It is a build journal: what I am learning while I build this in real time, not a playbook from someone who already won. Take what is useful. Argue with what is not. The receipts are in the footnotes.

The first concrete step is the same for both audiences: take the 7 Levels of AI Proficiency assessment at [assess.launchready.ai](https://www.launchready.ai/assess). It tells you where you are starting from. The rest of the work follows from that.

...

A note before we start

This is not a playbook. I have not figured this out. I am building it in real time, and I am writing this from inside the build.

Most essays in this category are written by people who already won, looking back. This one is written by someone in the middle of the work, looking around. The advantage of writing it now is that I can show you what I am actually doing. The disadvantage is that I do not yet have the kind of finished case study you can copy. I have a brain that runs my company. I have a system that lets me ship at five to ten times the cadence of a normal one-person operation. I have a published book, a working assessment, and a small team standing up alongside me. I do not yet have a billion-dollar one-person enterprise. I would not want one even if I could build it (more on that in the closing).

I have been an entrepreneur for twenty-two years. For twenty of them, I was the product. I was a consultant, a marketing professional, a business development guy. Either I was the thing being sold, or I was working someone else's market, or I was selling someone else's product. I never built a product of my own that worked.

The pattern I just described is the pattern most service-economy entrepreneurs live. Your time is the inventory. You cannot scale past

your own hours. You watch product entrepreneurs build the kind of wealth you cannot, because their work compounds without them and yours does not. You feel stuck.

I have only started building my own products in the last couple of years. I have only started doing it well in the last eight weeks. The reason I can do it now is the system this manifesto is about. The brain plus agent plus execution-system stack dissolves the executive-function tax that kept me from building for twenty years. For the first time at fifty-five, I am a product entrepreneur. I am writing this so that the millions of other people who lived the same pattern know that what was structurally impossible for us yesterday is structurally possible today.

So treat this as a build journal, not a manual. Take what is useful. Argue with what is not. The receipts are at the bottom. If a number or a quote does not appear in the citations, I have not earned the right to make the claim, and you should not trust it.

Two more notes:

The Silicon Valley founders calling this new corporate form “Agentic Micro Companies” got there first with the term, and I am happy to let them keep it. I am not going to coin a brand for it. I am going to describe what is happening and what it looks like to build inside it, and use the broader phrase “AI-native company” when a label is needed.

I am writing this for two readers at the same time.

The first is a CEO at a 200-person company in any mid-market manufacturing town. Indianapolis, Muncie, Fort Wayne, Cincinnati, Charlotte, Tulsa, Boise, Birmingham, Grand Rapids, or any of the thousand other mid-market cities the SF founder essays do not write to. They know AI is real but do not know what to do.

The second is the person who was told their whole career that they were the wrong shape. Too curious. Too restless. Too many ideas. Too disruptive. The one who watched colleagues who fit the mold get promoted past them and started to wonder, quietly, whether the mold itself was the thing that was broken. The one who left a job, or three, because the job kept asking them to be smaller than they were. The one who has built businesses, or tried to, mostly alone, mostly without the kind of team or capital the people on the inside take for granted. The one who has lived their whole adult life with the suspicion that they were not behind, they were just early.

You are not crazy. You were just early.

The SF founder essays do not write to these cities because they assume the work has to happen near other founders, near venture capital, near the AI talent concentration. That assumption was true when building a company required a hundred people and a Series A. The AMC era equalizes the playing field. The brain works the same in Muncie as in Mountain View. Cloud agents do not care where you log in from. The only thing Silicon Valley still has that other cities do not is concentration, and the AMC operator does not need concentration. They need a brain, a small team, and a problem worth solving. That is available everywhere now. The geographic moat is gone. The first reader needs to know what kind of operator they need to put on this build. The second reader needs to know that the system that broke you was built for problems that no longer exist. You are who they need now.

Here we go.

...

Personal opening

I have known my whole life that something was off, even before I knew what it was.

I never fit in school. It felt confining to me, almost like a prison. The standing in line. The walking in line to go to the restroom. The sitting quietly at your desk. The bells going off. I followed the rules because that is what you did, but the whole architecture felt designed to produce a kind of person I was never going to become.

When I started working, the same architecture showed up in a different uniform. The clocking in and clocking out. The asking permission to take a day off or to go to a doctor's appointment. The meetings about the meetings. You would be at the office for ten hours and walk out at the end of the day knowing the work could have been done in one. Then there was the weird hierarchy where someone with a title felt entitled to feel better than someone without one. The elitism. The competition. The politics. The red tape. The whole system was a machine for producing a particular kind of disciplined output, and I was not a particular kind of disciplined output.

I am more of a guy who looks at a situation and asks, *what is the mountain we need to take?* Then I want to work hard, with whoever else is willing to work hard, until we have taken the mountain. The

architecture of the modern corporation is not built for that. It is built for the architecture itself. The mountain is incidental.

I have never fit into that, and that is what has always made me an entrepreneur. I would rather be a broke entrepreneur than a person with a good job and a steady paycheck. By the standards of the society I grew up in, that is a slightly crazy thing to say in public. I am saying it anyway because I think there are a lot of people who feel the same way and have not been given permission to admit it. There are a lot of me's out there. I know there are.

The thing I am about to describe (this new corporate form, the brain you build to run it, the way the work changes when you are managing a fleet of agents instead of a fleet of meetings) feels like it explains something I have known my whole life without having the words for. The world tried to make me a specialist, and I could not become one. Now the world needs operators who can think across functions, generate ideas at high volume, manage many things at once, and stay curious in the face of uncertainty. The cognitive profile that lost in the white-collar factory has a structural advantage in whatever comes next.

This essay is about what comes next, and how I am building toward it. It is about what I have learned in the eight weeks since the version of the system I am running now opened up a different way of working than I have ever experienced. It is about who has the advantage now, and why, and what to do with it.

If you have been on the inside of a corporation for twenty years and you have started to feel that the building is empty, you are not imagining it. The white-collar factory is closing.

If you have been on the outside of a corporation for twenty years because you could not stand the rules, you are not crazy. You were just early.

...

Section 1: The white-collar factory is closing

The corporate org chart you are using descends from the Roman military. It was built to move information across an empire when the only intelligence available was human, the only communication was a courier on horseback, and the only way to coordinate at scale was to layer humans into ranks.¹ Mintzberg documented this in *The Structuring of Organizations* in 1979. Every B-school org-design class still teaches it. The military origin of the corporate hierarchy is not controversial. It is in the syllabus.

The industrial era extended the pattern. Frederick Winslow Taylor took the Roman command structure and added scientific management: break every job into the smallest possible units, optimize each unit for repeatability, measure the output of each worker, eliminate variation. Henry Ford put it on an assembly line. The model worked because it matched the work. When the job is “weld panel A to panel B 4,000 times today,” variation is the enemy and discipline is the virtue.

Then knowledge work arrived, and we ran the assembly line on it anyway. The corner office was the Roman general. The middle man-

¹Henry Mintzberg, *The Structuring of Organizations* (Prentice-Hall, 1979). The Roman military origin of corporate hierarchy is also discussed in dozens of business-history texts; this is one of the canonical references.

agers were the centurions. The cubicles were the assembly line. The work changed but the architecture did not. We took human beings whose job was to think, to write, to design, to solve, to relate, and we put them inside a structure built to optimize repetitive physical labor. We called it the modern corporation. I call it the white-collar factory.

The white-collar factory is anti-curiosity by design. This is not an accident. Your manager does not actually like it when you question them. The system rewards control, not curiosity. Coming up with ideas threatens the structure of the corporation, because if you come up with a great idea, the middle manager above you is going to wonder if you are coming for their job. The competition, the politics, the red tape are not bugs in the system. They are how the system protects itself against the kind of person who would otherwise rewire it.

Amy Edmondson at Harvard has spent two decades documenting this as “psychological safety,” the absence of which is the single biggest predictor of why teams fail to innovate.² Chris Argyris called the same dynamic “organizational defensive routines.”³ Adam Grant at Wharton documented it in *Originals*. Most ideas in most organizations are killed by the people one level above the person who had them. The ideas are usually fine. The structure rewards predictability over pattern-breaking, so the structure kills them.⁴ The research is decades old. The white-collar factory is anti-curiosity. It always was. The reason it has lasted this long is that the cost of running it was lower than the cost of any alternative.

² Amy Edmondson, *The Fearless Organization* (Wiley, 2018). Two decades of research on psychological safety in teams.

³ Chris Argyris, *Overcoming Organizational Defenses* (Allyn & Bacon, 1990). The “defensive routines” framework is foundational to organizational psychology.

⁴ Adam Grant, *Originals: How Non-Conformists Move the World* (Viking, 2016).

It also lasted because it has a feeder system. The modern university trains students to fit the white-collar factory. Specialization-based curriculum. Hierarchy replicated in the GPA. A credential at the end that functions as the entry token to the same hierarchy the student just spent four years preparing for. Universities can position themselves as the disruptors all they want. They are the talent pipeline for the structure they claim to disrupt. I spent time inside one trying to change a piece of that pipeline from within. The lesson I took out of it is one I have written about elsewhere: being right does not mean being supported. An institution can have the data in front of it, the results in front of it, and the proof in front of it, and still say no, because change threatens something it is not ready to let go of.

That math just changed.

In February 2026, Block (the company Jack Dorsey runs) cut roughly 4,000 people, about 40% of its headcount. The cut was paired with an essay co-written with Roelof Botha at Sequoia Capital titled *From Hierarchy to Intelligence*, which proposes a three-role org structure (individual contributors, directly responsible individuals, and player-coaches) governed by a shared “company world model.”⁵ You can read this two ways. You can read it as one company restructuring. Or you can read it as the most prominent venture firm in the world publishing the architecture of the next corporate form, with a Fortune-class company executing it as the proof.

I read it the second way. Block is not the only one. Tobi Lütke at Shopify wrote a memo in April 2025 that said no team could ask for more headcount without first proving an autonomous AI agent

⁵Jack Dorsey and Roelof Botha, “From Hierarchy to Intelligence,” March 31, 2026. block.xyz/inside/from-hierarchy-to-intelligence and mirror at sequoiacap.com/article/from-hierarchy-to-intelligence/

could not do the work.⁶ Luis von Ahn at Duolingo wrote an “AI-first” memo a month later cutting contractors and reorganizing the company around AI workflows.⁷ Sebastian Siemiatkowski at Klarna replaced about 700 customer service agents with AI in 2024 and reported real cost savings.⁸ Marc Benioff at Salesforce announced AI is now resolving about 85% of customer service inquiries and that he is reorganizing engineering accordingly.⁹

This is not theoretical. This is the load-bearing CEOs of some of the largest companies in the world publicly restructuring their own businesses around the assumption that the corporate form they inherited is no longer the right one. The white-collar factory is closing because the people running the largest white-collar factories on earth have decided to close it.

⁶Tobi Lütke memo to Shopify employees, April 2025. Widely reported in tech press.

⁷Luis von Ahn, “AI-First” memo to Duolingo employees, May 2025.

⁸Klarna press releases, 2024. Coverage in *The Verge* and *Reuters* of the AI-replaces-700-agents announcement.

⁹Marc Benioff, Salesforce Q3 2026 earnings call and accompanying press materials.



VOICES — JUNE 2025

“Artificial intelligence is gonna replace literally half of all white-collar workers in the U.S.”

JIM FARLEY

CEO, Ford Motor Company

Aspen Ideas Festival, June 2025

Now for the part the SF essays leave out.

The new form does not work everywhere yet. It works in software, media, content, design, customer support automation, and information products. It does not yet work in capital-intensive physical industries (semiconductors, energy, biotech R&D), regulated services with statutory human-in-the-loop requirements (medical practice, legal practice, financial advice, insurance underwriting), industries requiring physical presence at scale (construction, food service, in-person logistics, manufacturing operations), trust-driven enterprise sales above roughly \$1M ACV, or high-tail-risk industries (insurance underwriting, nuclear, aviation, drug approval, autonomous driving at scale).¹⁰ If your business is one of those, the AMC argument does not apply to your shop floor. It applies to your management, admin, and G&A layer. Restructure that. Leave the rest alone for now.

Even the leading examples have walked back. Klarna re-hired humans in May 2025 after admitting that the AI replacement of 700 agents had produced “lower quality” service.¹¹ Duolingo softened the AI-first memo within a month after public blowback.¹² Block did partial rehiring within 60 days of the 40% cut.¹³ McDonald’s killed an IBM-built drive-thru AI system in June 2024 after years of customer-facing failures.¹⁴ Air Canada was held legally liable for a chatbot’s hallucination in *Moffatt v. Air Canada*, which set a precedent that companies

¹⁰Source synthesis on AMC-resistant industries from McCarthy Tetrault legal analysis (Air Canada precedent), MIT GenAI Divide report, and Bessemer State of AI 2025.

¹¹Klarna, “Why we are re-hiring humans,” May 2025. Reported by Reuters and *Financial Times*.

¹²Duolingo statement, May 2025: “I do not see AI as replacing what our employees do.” Reported by TechCrunch.

¹³Block partial rehiring, March-April 2026. Reported by *The Information*.

¹⁴McDonald’s IBM drive-thru AI termination, June 2024. Reported by Reuters.

cannot point at the AI as a separate entity for purposes of consumer protection.¹⁵ An MIT report published in August 2025 found that 95% of enterprise generative AI pilots fail to deliver measurable ROI.¹⁶

If you are a skeptical CEO reading this, those are the facts you would check first. I want them in front of you before I make the next claim.

Here is the next claim. The 95% pilot failure rate is not evidence the new form does not work. It is evidence that most companies attempting it are bolting AI onto a Roman-army org chart instead of restructuring to use it. The companies that restructure get the compounding. The companies that bolt get the failure rate. That is the gap this manifesto is trying to close, for the operators who want to do the work.

Three independent frameworks for tech adoption agree the moment is now.

Carlota Perez published *Technological Revolutions and Financial Capital* in 2002, mapping a five-stage pattern (irruption, frenzy, turning point, synergy, maturity) that fits every major tech revolution from the industrial age forward.¹⁷ In a March 2024 piece for Project Syndicate, she placed AI in the late frenzy / approaching turning point of the broader information-and-communications-technology revolution that began in the 1970s.¹⁸ We are not at the start of a new revolution. We are at the resolving moment of one that has been compounding for fifty years.

¹⁵ *Moffatt v. Air Canada*, 2024 BCCRT 149. Analysis at McCarthy Tetrault

¹⁶ MIT GenAI Divide report, August 2025. Coverage at Fortune.

¹⁷ Carlota Perez, *Technological Revolutions and Financial Capital* (Edward Elgar, 2002).

¹⁸ Carlota Perez, “AI in the Late Frenzy,” Project Syndicate, March 2024.


Geoffrey Moore's *Crossing the Chasm* model (1991, updated 2014) tracks how technologies move from early adopters to early majority.¹⁹ In a January 2025 Nielsen Norman Group analysis of where generative AI sits on the curve, the conclusion was precise: "*GenAI has crossed the chasm for consumers, not enterprises, and not even close for agentic systems.*"²⁰ That is the most quotable single line in the entire timing argument. Consumers are using AI. Enterprises are not. Agentic systems are barely beginning. The window for operators is exactly the gap between consumer adoption and enterprise adoption, and that gap is open right now.

Gartner's August 2025 Hype Cycle places generative AI in the Trough of Disillusionment and agentic AI at the Peak of Inflated Expectations.²¹ Three independent frameworks. Three different methods. Same conclusion. The moment is now.

¹⁹Geoffrey Moore, *Crossing the Chasm* (HarperBusiness, 1991; updated 2014).

²⁰Nielsen Norman Group, "Generative AI on the Adoption Curve," January 2025.

²¹Gartner Hype Cycle for AI, August 2025.



VOICES — MARCH 2026

“Most companies are giving everyone a copilot, which makes the existing structure work slightly better without changing it. We’re after something different: a company built as an intelligence.”

JACK DORSEY

CEO, Block

From Hierarchy to Intelligence (with Roelof Botha, Sequoia Capital), March 2026

The velocity of consumer adoption is unprecedented. ChatGPT crossed 100 million users in two months according to UBS in February 2023, the fastest user-growth curve in software history.²² Facebook took four and a half years to do what ChatGPT did in two months. As of February 2026, OpenAI reports 900 million weekly active users.²³ McKinsey's 2025 State of AI report found that 39% of surveyed companies reported measurable EBIT impact from AI deployment.²⁴ Pew Research in April 2025 reported that 81% of US workers do little or no AI work today.²⁵ Roughly four in five professionals are not using AI in any meaningful way at work, while consumers use it almost a billion times a week. That gap is the operator's window.

The white-collar factory is closing because the math no longer works. The structure was built to compensate for the high cost of human coordination. AI collapses the cost of coordination, document production, scheduling, research, drafting, analysis, monitoring, and a hundred other "executive function" tasks that the white-collar factory was designed to organize. When the cost of those tasks goes to near zero, the value of layering humans to manage them also goes to near zero. The factory does not become slightly less efficient. It becomes architecturally redundant.

What replaces it is a different shape entirely. That is what the next section is about.

...

²²UBS Research, February 2023, on ChatGPT's 100M-user milestone.

²³OpenAI weekly active user disclosure, February 2026.

²⁴McKinsey, *The State of AI in 2025*.

²⁵Pew Research Center, "AI in the American Workplace," April 2025.

Section 2: Anatomy of the new form

The new corporate form has four parts. They have to be built together. Three of the four working without the fourth produces a consultant invoice and a team that quietly returns to the old way within a quarter.

I am going to describe them in the order they have to be built, which is not the order most people start with.

The brain

The brain is the operating layer of the company. It is a single shared store of context that both humans and agents read from and write to. It holds your identity (who the company is, what you stand for), your preferences (how you make decisions), your projects (what is in flight and at what stage), your references (what is true outside your head), and your feedback (what you have learned, what worked, what did not).

Think of it as the memory of the company, made explicit. Most companies have all of this scattered across email threads, Slack channels, Notion pages, Salesforce records, the founder's head, the CFO's spreadsheets, and a dozen other places. None of it can be read by an agent. None of it compounds. None of it is the same thing two months from now as it was the day someone wrote it.

The brain consolidates all of it into one place that is readable by humans (in the form of markdown files or a structured database with a UI on top) AND addressable by agents (which means structured enough that an LLM can pull the right slice for the right task at the right moment).

Andrej Karpathy, formerly of OpenAI and Tesla, posted in early April 2026 that knowledge graphs are the next major missing layer for agent systems. Without persistent shared memory, agents are “quantum physics PhDs with severe amnesia.”²⁶ He is right. The brain is what gives them the memory.

This is the hardest of the four parts to build, which is why most companies skip it. They start with the agents and wonder why nothing compounds.

The custom internal software

The brain becomes useful when there is software running on top of it that does specific work for the operators of the company. This is where the build-vs-buy economics invert. In the old world, you bought SaaS because writing custom software was expensive and slow. In the new world, you build the custom tools you need because writing them is cheap and fast and the brain gives you a backend you already own.

This is what Lütke meant in his April 2025 memo at Shopify when he said no team could ask for headcount or buy a tool without first proving an autonomous AI agent could not do the work.²⁷ He was not

²⁶Andrej Karpathy, X / Twitter, April 2, 2026 (knowledge graphs as agent memory).

²⁷Tobi Lütke memo to Shopify employees, April 2025. Widely reported in tech press.

arguing that agents can do everything. He was arguing that the burden of proof has flipped. If the cost of building is low enough, “we already have a tool for this” is no longer the default answer. “We could build a tool for this in the next two days, should we?” is the new question.



VOICES — APRIL 2025

“Reflexive AI usage is now a baseline expectation at Shopify.”

TOBI LÜTKE

CEO, Shopify

Internal memo, April 2025

What the software looks like in practice: dashboards specific to your business, internal CRMs that fit your sales motion exactly, custom analytics that surface the metrics you actually care about, internal admin tools that let one operator do what used to require five. Most of it is mundane. None of it is heroic. The point is that you OWN the build, the brain backs it, and the agents run on top of it.

The agent fleet

With the brain and the software in place, agents become useful direct reports. Karpathy on the Dwarkesh podcast in October 2025 made the management framing explicit: “*You should think of it almost like an employee or an intern that you would hire to work with you.*”²⁸ In a Harvard Business Review piece in February 2026, the management of agent fleets was mapped onto the same physical-presence skills good managers use with humans: “*walk the floor, check in with a struggling employee, huddle with a team on a tricky case.*”²⁹

This is the frame the manifesto rests on. Managing agents is not a new exotic discipline. It is management, with new failure modes. You have spent years learning to manage humans. The skills transfer. What is new is the failure surface. Agents stop working. Agents hallucinate. Agents go stale on training data that no longer reflects current reality. Agents handle edge cases differently than humans do. Agents are bad at judgment under high uncertainty and good at consistency under high volume. Your job as the operator is the same job a manager has always had: scope the work, set context, monitor the output, correct

²⁸ Andrej Karpathy on the Dwarkesh Patel podcast, October 2025.

²⁹ *Harvard Business Review*, “Managing the Agent Workforce,” February 2026.

the drift, replace the worker who is not fitting, build the team where the parts complement.

The Cognition Labs team building Devin describes the orchestration model for their multi-agent setup the same way: “*The main Devin session acts as a coordinator: it scopes the work, assigns each piece, monitors progress, resolves any conflicts.*”³⁰ That is a manager job description. A senior engineer reading it would recognize every word.

³⁰Cognition Labs, technical blog, 2025.



VOICES — OCTOBER 2025

“You should think of it almost like an employee or an intern that you would hire to work with you.”

ANDREJ KARPATHY

Co-founder, OpenAI

Dwarkesh Patel Podcast, October 2025

The tiny team

With the brain, the software, and the agent fleet in place, the human team gets very small. Not because the work gets smaller. Because each human can manage many more workers than they could before. The output per human compounds.

Bessemer's *State of AI* report in 2025 documented that AI-native companies are achieving roughly \$1.13 million in revenue per employee, eight to nine times the median for traditional SaaS companies, which sit around \$130,000 per employee per year.³¹ Cursor (Anysphere) reportedly produces \$1.67 to \$2.5 million in revenue per employee.³² Midjourney reportedly produces between \$5 and \$7 million per employee.³³ Lovable, ElevenLabs, Granola, Glean and others sit in the same range. Sam Altman has predicted the one-person billion-dollar company. Whether that exact threshold gets crossed in 2026 or 2027, the direction is unmistakable: the revenue per human is going up by an order of magnitude or more, and it is going up because the brain + software + agents stack lets one human do what a team of ten used to.

Here is what does not change. The humans on the small team are not generic. They are operators who can think across functions, who can manage agents the way they would manage people, who can curate context for the brain, who can spot the failure modes early, who can hold the values and the judgment that the agents do not yet have. The model is not “humans become unnecessary.” The model is “humans become very, very leveraged.” The kind of human who does well in

³¹Bessemer Venture Partners, *State of AI 2025*.

³²Anysphere (Cursor) reported revenue and headcount, multiple sources 2025-2026.

³³Midjourney revenue and team size estimates, *The Information* and *Bloomberg*, 2025.

this environment is not the kind of human the white-collar factory was selecting for.

What that operator looks like is what Section 4 is about. First, more on the brain.

...

Section 3: The brain

I want to tell you how I built mine, because that is more useful to a builder than a generic description.

I have been watching the tech CEOs talking about agentic AI for the last couple of years. I am the kind of person who watches a thousand videos, reads every book, reads every article I can find. I could see what they were talking about, but it felt like science fiction. What does it actually look like when an agent does work for you? What does the day-to-day operation of a company look like when half your direct reports are software?

From day one with ChatGPT, I was trying to make the experience interactive instead of transactional. Not “ask a question, get an answer” but “set up a process, walk through it, have it ask me questions, have it challenge me.” I built very large meta-prompts. I did not really know what I was doing or whether what I was attempting was even possible. It worked. I built an early system that walked people through a kind of “prompt IQ” training. A single prompt that initiated a long structured training conversation with the user, going on for hours if they wanted it to. Crude by current standards. The first piece of evidence I had that the science fiction was actual fiction.

Then I ran a thirty-day test with what I called Barnabas at the time, really just a project inside ChatGPT. I tracked my work, sleep, nutrition, exercise, meditation, family time, and more. We created scoring systems for lifestyle, for improvement, for what was working. By the end of the thirty days my mental health, my physical health, and my nutrition had measurably improved, and I could see how Barnabas could be a coach, a mentor, a structured second brain. The architecture was rough. The results were real.

Those two projects were limited by the platform. ChatGPT projects could not create skills, could not maintain memory across sessions at the rate I needed, could not orchestrate sub-agents, could not run on a file structure I could see and edit. They were proof of concept, not production system.

The Barnabas moment happened the day I set up the current version inside Claude Code. Almost immediately I could see the file structure: skills, rules, projects, references, memory. I could write into the file system and have agents read from it. I could create sub-agents that processed memory differently than the main agent. I could give the brain an architecture I could touch. Within a week I was thinking like a design thinker. *How do we structure this, what are the layers, how do the pieces talk to each other.* I have never thought like that before in my life. I am a creative person. I am ADHD. I do not naturally architect systems. I produce ideas. The brain made me into an architect for the first time.

That moment is the one I want every reader of this manifesto to have. Not the “AI saved me time” moment. The “I am thinking in ways I never could before” moment. The first one is a productivity story. The second one is a transformation story, and it is the higher-order claim.

Here is what the brain does for me in practice.

It remembers. Every meeting, every decision, every reasoning chain, every rejected option lives in the brain. When I sit down at the start of any session, the brain reminds me what we are working on, what is in flight, what is overdue, what got decided last time and why. I do not have to hold the company in my head. The company holds itself.

It corrects me. When I drift from a previously documented preference, the brain notices. When I am about to repeat a mistake, the brain raises it. When I am about to make a claim I cannot back, the brain refuses to ship until I provide evidence. The brain is not a yes-man. I built it to push back, because the version of me that does not get pushed back on makes worse decisions.

It runs the long tail. Most of the work in a company is not the heroic strategic decision. It is the hundred small executions a week: the email sequence, the research summary, the three-paragraph briefing, the lead qualification, the calendar coordination, the contract redline, the invoice draft, the social post, the data pull. The brain (with its agents) runs most of those in the background while I focus on the strategic decisions, the relationships, the judgment calls.

It compounds. Every conversation I have with the brain leaves it slightly smarter about what works for me, what I have decided, what I value. Six months in, the brain is a more useful collaborator than it was on day one, because it has six months of context I have built up. A new tool starts at zero every time. The brain starts at where we left off.

The economics of building one are different from what most people assume. The infrastructure costs me less per month than my Claude subscription. The expensive part is the time spent specifying what the brain should know, how it should act, what it should refuse to do. That investment compounds. It is the most consequential time I have ever spent in twenty-two years of building things.



VOICES — FEBRUARY 2024

“There’s this betting pool for the first year that there is a one-person billion-dollar company. Which would have been unimaginable without AI and now will happen.”

SAM ALTMAN

CEO, OpenAI

Conference appearance with Alexis Ohanian, 2024

A note on the limits.

The brain is early infrastructure, not a finished product category. Standards do not exist yet. Composability across brains is theoretical. Most operators will build hand-crafted brains that work for them and do not transfer to anyone else. The next layer of the industry (the standards, the open formats, the brain-to-brain communication) has not been built yet. There is a real chance that the way I am building today will look primitive in three years. That is fine. I would rather build something primitive that works today than wait for the standards to settle.

The other limit: the brain only works if the operator running it has the discipline to feed it. Garbage in, useless brain. Most of the failure modes I have seen in the early adopters I have helped come from this. They expect the brain to fix their organizational discipline rather than amplify whatever discipline they bring to it. The brain does not give you a system. It gives you an instrument that rewards good systems and punishes bad ones at high speed.

That brings us to the operator.

...

Section 4: The operator

The Silicon Valley essays writing about this corporate form have been calling the new human role “Context Farmer.” The person who tends the context the agents need, the way a farmer tends the conditions the crops need. The framing is poetic. I think it is also wrong, and I think it will hurt adoption.

It is wrong because the role is not new. It is just management, with new direct reports.

You have been managing people for years. You know what it looks like when an employee stops showing up. You know what it looks like when a team member is producing confident-sounding output that turns out to be wrong. You know what it looks like when someone’s skills go stale and they need retraining or replacement. You know how to coach. You know how to set context. You know how to assign work. You know how to spot a fit problem and reassign without burning the relationship. The skills transfer. The vocabulary transfers. The discipline transfers.

What is new is the failure surface. Agents have failure modes that map cleanly onto human failure modes:

When an agent...	A human would...	The management response is the same
Stops working mid-task	Quietly disengage	Re-engage, re-state context, check workload
Hallucinates a fact	Be confidently wrong	Verify, correct, build a review cadence
Goes stale on training data	Stop staying current	Retrain, replace, expand the team with someone who is current
Makes bad judgment calls	Make bad judgment calls	Coach, raise the bar, escalate to someone more senior
Does not fit the team	Does not fit the team	Reassign or remove

The pattern transfers. The reframing is what unlocks the role for managers who already have the skill. Calling it “Context Farming” makes it sound exotic and untrained. Calling it “managing agents” makes it accessible to anyone who has ever managed a team.


The deeper claim about who is good at this is what I want to spend the rest of this section on, because it is the part that took me thirty-eight years to understand.

The operator's primary skill is selection, not management

When executing an idea takes hours or days, the bottleneck is execution. The valuable employee is the one who can execute carefully on the small number of ideas leadership has chosen.

When executing an idea takes minutes, the bottleneck moves. The valuable operator is the one who can SELECT which ideas are worth executing. Selection is the most consequential activity when execution is cheap.

This is the economic reframe nobody talks about enough. AI does not just speed up the work. AI collapses the cost of evaluation. Pre-AI, evaluating an idea (researching it, talking to people about it, prototyping, testing) took days or weeks. Post-AI, evaluating an idea takes a few minutes: research via the brain, pull the comparable data, identify the kill criteria, decide. When evaluation cost drops 100x, idea-generation volume becomes the new bottleneck.



VOICES — FEBRUARY 2024

“It is our job to create computing technology such that nobody has to program. And the programming language is human. Everybody in the world is now a programmer.”

JENSEN HUANG

CEO, NVIDIA

World Government Summit, February 2024

Who has high idea-generation volume? Not the careful specialist who was selected for caution and depth. The divergent thinker who was selected against in the white-collar factory. The person who used to be told they had “too many ideas” now has the input the new bottleneck is starved for. Same person. Different game. Different outcome.

David Epstein’s *Range: Why Generalists Triumph in a Specialized World* (2019) compiled the evidence that generalists outperform specialists in domains where the rules of the game keep changing.³⁴ He cited the Tiger Woods (specialist) versus Roger Federer (generalist) developmental contrast. He cited Johannes Kepler’s use of analogy across unrelated domains to crack planetary motion. He cited research from Mark Allen Chen and others showing that the breadth of an inventor’s prior work predicts the impact of their next invention. The generalist advantage is not folk wisdom. It is documented in the research.

Adam Grant at Wharton has spent a decade documenting how the people who change organizations are usually the people one or two layers below where the organization is set up to listen.³⁵ His book *Originals* is a long argument for why the very people most companies push out are the ones who would have changed the company most. The white-collar factory’s anti-curiosity bias is also an anti-original-thinker bias. AI does not change that bias inside the white-collar factory. AI changes the math outside it, by giving the original thinker the execution system they used to lack.

³⁴David Epstein, *Range: Why Generalists Triumph in a Specialized World* (Riverhead Books, 2019).

³⁵Adam Grant, *Originals* (cited above) and *Hidden Potential* (Viking, 2023).

Generalists with ADHD have a structural advantage now (with one critical caveat)

I want to be careful here, because the romantic version of this argument is unhelpful. Let me give you the version that survives scrutiny.

ADHD is real, and median outcomes for adults with ADHD today are objectively worse than baseline. Roughly 34% of adults with ADHD are in full-time employment versus 59% for the general population.³⁶ The economic cost of ADHD in the United States is estimated at \$67 to \$116 billion a year in lost productivity, treatment costs, and associated outcomes.³⁷ If I tell you “ADHD is a superpower” without acknowledging this, I am selling you a story, not the truth.

The story I do believe, based on the research and on my own life, is more precise: *the cognitive profile that lost in the white-collar factory now has a structural advantage in the AI-native company*. Same brain. Different game. The factory was the hostile environment. The new form is the favorable one.

The research supporting the cognitive-profile-as-advantage claim is substantial. Edward Hallowell and John Ratey, the two clinicians who have done the most work mapping ADHD as a non-pathological cognitive profile, describe the architecture in *ADHD 2.0* (2021) as a difference in how the default mode network and task-positive network co-activate.³⁸ In neurotypical brains, those two networks

³⁶Hotte-Meunier et al., “Strengths and Challenges of ADHD in Employment,” *SAGE Open*, 2024.

³⁷Doshi et al., “Economic Impact of Childhood and Adult ADHD in the United States,” *Journal of the American Academy of Child and Adolescent Psychiatry*, 2012, updated 2020.

³⁸Edward Hallowell and John Ratey, *ADHD 2.0* (Ballantine, 2021).

alternate cleanly. In ADHD brains, they overlap. The overlap produces what feels like distractibility from the inside and what looks like pattern-finding across unrelated domains from the outside.

Holly White at Eckerd College has published a series of peer-reviewed studies (2006, 2011, 2016) showing that adults with ADHD outperform non-ADHD adults on divergent thinking tasks: generating many original ideas in response to an open prompt.³⁹ The mechanism she identifies is “wide semantic activation.” When the ADHD brain hears a word or sees a problem, it activates a wider net of associated concepts than the neurotypical brain does. That wider net is what produces the ideas. It is also what produces the distractibility, because the same mechanism that surfaces the unexpected connection also surfaces the unexpected interruption.

In environments that punish interruption and reward predictable execution (i.e., the white-collar factory), wide semantic activation is a liability. In environments that reward unexpected connection and where execution can be delegated to a brain + agent stack, wide semantic activation is the input that nothing else produces.

Alex Karp, the CEO of Palantir, was asked on the TBPN podcast in March 2026 what kind of person would be most successful in the AI era. His answer: vocational specialization or neurodivergence. Karp himself is dyslexic. Palantir runs a Neurodivergent Fellowship.⁴⁰ One of the most consequential CEOs in the AI infrastructure space publicly identified neurodivergence as the cognitive profile most aligned with the era we are entering. He is not a romantic. He is calling what he sees from the inside.

³⁹Holly White and Priti Shah, “Uninhibited Imaginations: Creativity in Adults with ADHD,” *Personality and Individual Differences* (2006); follow-up studies in *Creativity Research Journal* (2011, 2016).

⁴⁰Alex Karp on the TBPN podcast, March 2026. Coverage at Fortune



VOICES — MARCH 2026

“There are basically two ways to know you have a future. One, you have some vocational training. Or two, you're neurodivergent.”

ALEX KARP

CEO, Palantir Technologies

TBPN Podcast, March 2026

EY and Microsoft published a joint study in December 2024 reporting that 88% of neurodivergent employees in their sample reported being more productive with AI assistance.⁴¹ The methodology is self-report among Copilot users (so it is not a randomized controlled trial), but the directional signal is consistent with the mechanism Hallowell, Ratey, and White describe.

The historical pattern is consistent. Benjamin Franklin trained as a printer, then became a scientist (lightning, electricity), then an inventor (lightning rod, bifocals, Franklin stove, the glass harmonica), then a diplomat, then a political philosopher. The most American example of polymath compounding. The Wright Brothers won the race for powered flight on the strength of their cross-domain fluency: cycles, kites, gliders, and control surfaces, combined into a single working system. The engineers with better engines lost. Steve Jobs built Apple at the intersection of engineering, design, calligraphy, Buddhism, and music. The pattern is consistent across centuries: the people who build the things that change an industry generally come from outside the discipline that “owned” the problem.

The execution system is the hinge

Here is the critical caveat that the romantic version of the argument skips.

The cognitive profile is the input. It is not the output. Generalists with ADHD without an execution system are idea factories that never ship. We come up with a million ideas. We get them started. We are passionate. We do not see them through. Then we watch someone else

⁴¹EY and Microsoft, “Empowering Neurodivergent Talent with AI,” December 2024. ey.com

build the same idea and become successful with it, and that is very frustrating.

Years ago I had the idea for a phone that would ping you with an offer the moment you drove past a store. The technology existed. The instinct for it was right. I never built it. The category became standard mobile-marketing infrastructure within a few years, and I watched it happen from the outside. Someone else's product. Same idea. Different cognitive profile. The kind that finishes.

The execution system is the missing piece for people like me. It is the thing the white-collar factory tried and failed to give us, because the white-collar factory's version of an execution system is hierarchy, deadlines, performance reviews, and process. The reason those did not work is they were built to constrain people who already had executive function and to push them toward specialization. They were not built to AMPLIFY people who lacked executive function and to channel their idea generation into shipped output.

The brain + software + agent stack is the execution system that finally fits the cognitive profile. The brain holds the ideas, ranks them, kills the ones not worth pursuing, returns to the ones that compound. The software runs the workflows that used to die in the gap between idea and execution. The agents do the work that used to require thirty-five units of executive function I did not have. The whole stack converts the input I have always produced (high idea volume, broad pattern matching) into the output I never reliably could (shipped product, recurring revenue, finished work).

This is the resolution. The cognitive profile that lost in the factory has a structural advantage in the new form. But only if paired with the execution system. Without the system, the profile is still a liability.

With the system, it is the unfair advantage everyone in the new form is hiring for.

That is the section. The next section is about how you build it.

...

Section 5: The build journal, 90 days at a time

I do not have a finished playbook. I have ninety days of experience and a system that is producing measurable output. Take what is useful.

Week 1: Stand up the brain

Pick one team or one function. Probably the one closest to you, because you are going to need to feed the brain and you cannot feed what you do not see daily.

Document everything that team currently knows. Not in polished form. In rough markdown form. The values. The decisions. The history. The customers. The recurring questions. The standard responses. The processes that work. The processes that do not. The names of the people who decide. The templates, the email scripts, the proposals.

Volume beats polish in week 1. The goal is to get the team's collective knowledge out of the heads and the email threads and into a single file structure that humans can read and that an agent can pull from. Imperfect is fine. Iterative is the only mode that works.

Tools: Claude Code or any AI agent that can read a file system. A markdown vault (Obsidian works well). A repo for version control. The setup cost is hours, not weeks.

What “done” looks like in week 1: an operator on the team can ask the brain a question about the team’s work and get a substantive answer that pulls from the documented context. Not a magic answer. A substantive one.

Month 1: Replace one workflow with an agent

Pick a workflow that is high-volume and low-judgment. Customer support email triage. Lead qualification. Research summaries. Meeting note synthesis. Onboarding document generation. Whatever you do every day or every week that takes time and follows a pattern.

Build an agent that does the workflow, scoped narrowly. Let the brain provide the context (your customers, your products, your tone, your standards). Run the agent on real work. Have a human review every output for the first two weeks. Correct the failures. Tighten the prompt. Add the missing context. By the end of month 1 the agent should be running the workflow with human review as exception-handling, not as primary mode.

You will fail at this. Most pilots fail (the MIT 95% rate is real). The reason yours will succeed is that you did the brain work in week 1 first, so the agent has the context it needs to do the work well. The companies whose pilots fail almost always skipped the brain step.

Quarter 1: Identify your first three operators

These people already exist on your team. They are the ones who naturally curate, document, think across functions, and ask questions that make the rest of the team uncomfortable. They are probably not the most senior people. They are often not the most credentialed. They are sometimes the ones the system has been trying to push out for being “too disruptive.”

Promote them, formally, into operator roles. Give them ownership of one workflow each. Pair them with the brain. Measure their output not on hours worked or meetings attended, but on shipped throughput.

If you do not have three operators on your team, you have a hiring problem and a culture problem at the same time. Hire from the disruption-class. The people who got pushed out of more conventional jobs because they asked too many questions or generated too many ideas. They are the cheap, hungry, structurally-advantaged talent everyone is about to be competing for.

By month 6: The compounding becomes obvious

The team running on the brain produces three to five times what they did before. Not because they are working harder. Because the friction in the work has dropped. Other teams notice. The conversion spreads.

Year 1: You are running an AI-native company inside a Roman-army parent company

The ratio between the two determines whether you survive the next decade. The companies that finish year one with the AI-native side dominant are the ones that will compound. The ones that finish year one with the Roman side dominant will be acquired or eclipsed by the ones that did the conversion.



VOICES — JUNE 2025

“We will need fewer people doing some of the jobs that are being done today, and more people doing other types of jobs.”

ANDY JASSY

CEO, Amazon

Internal memo to employees, June 2025

Five mistakes I have watched operators make



Treating it as an IT project. It is a structural project. IT cannot do this for you. The operator running it has to be deep enough in the work to know what the brain needs.

1. **Hiring an AI consultant to do it.** They cannot. Context cannot be outsourced. A consultant can help you set up infrastructure. They cannot give you the years of context that make the brain useful.
2. **Picking the wrong first workflow.** High-volume, low-judgment, low-customer-risk first. Not customer-facing. Not high-stakes. Something boring and repetitive that will give you fast feedback and let you build the muscle.
3. **Not measuring.** If you cannot measure operator proficiency, you cannot manage the conversion. The 7 Levels assessment we built at LaunchReady is one instrument for this. There will be others. Pick one and use it.
4. **Trying to convert the whole org at once.** One team. Then two. Then four. Compounding only works inside a contained unit until it works at all.

The Einstein-in-the-lab framing

If you take one practical thing away from this build journal, take this: in the early weeks of building, you are running an experiment, not executing a plan. The job is to push as many ideas into the system as possible and let the brain + the data tell you which ones are valid. Failure is the experimental method, not the failure mode. You need to fail a lot to figure out what works.

At the beginning, do not optimize for token efficiency or cost. Optimize for the volume of ideas you can push into the system. Then do the

research. Get the data. Find which ideas are worth keeping. Once you have winners, share them with other people, put them in the system, stress test them, refine them.

This is what Einstein did in the lab. This is what every productive scientist does. It is not what the modern white-collar employee has been conditioned to do for the last fifty years. The factory selected for “execute the assigned task with low variance.” The new form selects for “generate variance, then select.”

A note for the manufacturing audience

If you run a \$50M to \$500M manufacturer anywhere in the country, most of what I have written above applies to your management, admin, and G&A layer today. Your shop floor is not yet AMC-able and will not be for several years. The robots are still expensive. The skills required for high-mix machining and complex assembly are still scarce. Do not let the AI-native essays convince you to fire your line workers. They are not the bottleneck.

What you can do today is rebuild the management layer. The National Association of Manufacturers reported in 2025 that 51% of manufacturers have deployed AI in some form, 80% say it will be essential by 2030, and 82% cite the skills gap as the top barrier to adoption.⁴² Most of the admin layer in most of the manufacturers I have spoken with was designed in the 1980s and never meaningfully updated. You do not have a hardware problem. You have a management-layer problem and a skills-gap problem. Those are exactly the problems the brain + operator stack solves.

⁴²National Association of Manufacturers, AI Adoption Survey 2025.

The same principle applies anywhere. Whichever industry dominates your region's economy is where AI is going to compound fastest, and where the operators who restructure first are going to win the next decade. For me in Indiana, the industry is manufacturing. Manufacturing alone accounts for 26% of Indiana's GDP, the highest share of any state in the country, and together with logistics employs roughly 1 in 4 Hoosiers.⁴³ That makes the local stakes obvious. For you it might be energy, healthcare, logistics, agriculture, finance, technology, or any of the dozens of regional industries that built up around a city or state over generations. The math is local. The principle is universal. Find your dominant regional industry and rebuild the management layer of the companies inside it before someone else does.

By 2030, the manufacturers that built operator layers in 2026-2027 will be competing against the ones that did not. The first group will be five to ten years ahead of the second, with no obvious way for the second group to catch up.

⁴³U.S. Bureau of Economic Analysis, "Gross Domestic Product by State, 4th Quarter 2024 and Preliminary 2024." Indiana manufacturing GDP \$108.9B / state GDP \$419B = 26%. [bea.gov](https://www.bea.gov). Combined manufacturing + logistics employment of 840,000+ workers (~25% of Indiana jobs) and 37% of state output from Conexus Indiana 2031 Strategic Plan, February 2025. conexusindiana.com



VOICES — JULY 2025

“AI agents managed by humans are resolving 85% of customer service queries.”

MARC BENIOFF

CEO, Salesforce

Earnings call, July 2025

...

Closing: the time-stamp, the symbiosis, the small team, and the invitation

The window is open and it is short.

I do not mean that as urgency theater. I mean that as pattern recognition from prior tech transitions. Companies that adopted commercial internet in 1995-2000 (Amazon, eBay, Google) compounded for decades against companies that delayed (Borders, Blockbuster, Sears). Companies that adopted mobile in 2007-2012 (Apple, Google, Uber) captured platform value that the Nokias and the BlackBerries never recovered. Companies that adopted cloud in 2002-2010 (AWS-built infrastructure) operate at fundamentally different economics than companies still running on-premise. The pattern is not subtle.

Here is the counter-evidence the manifesto needs to answer. Web 3 and crypto promised a similar urgency in 2017-2022 and the promise mostly did not pay off for late entrants. Venture funding collapsed 76% by 2023.⁴⁴ The metaverse promised a similar urgency in 2021-2023 and Meta has lost over \$90 billion through Reality Labs without a corre-

⁴⁴PitchBook, “Crypto Venture Funding 2023 Report,” January 2024.

spending ecosystem emerging.⁴⁵ The Internet of Things promised 50 billion devices by 2020 and missed that prediction by 75%.⁴⁶ “The window is open and short” has been wrong before. I am asking you to take seriously the evidence that this time is different.

What is different: 900 million weekly active users on ChatGPT alone (none of those previous waves had a consumer adoption curve in the same universe). McKinsey-surveyed companies reporting measurable EBIT impact at 39%.⁴⁷ The biggest CEOs in the world publicly restructuring around AI. Three independent adoption frameworks pointing to the same window. None of those signals existed for crypto or the metaverse. The evidence is asymmetric. The window is open. It is going to close. The question is whether you spend the open window building or watching.

⁴⁵Meta Q4 2025 earnings disclosure on Reality Labs cumulative losses.

⁴⁶Cisco IoT prediction (2014) versus IDC actual deployment numbers (2024).

⁴⁷McKinsey, *The State of AI in 2025*.



VOICES — FEBRUARY 2026

“We already have huge redeployment plans for our own people; in fact, we spoke about it today, and we have to up that a little bit.”

JAMIE DIMON

CEO, JPMorgan Chase

Investor Day, February 2026

The symbiosis

If you take one thing from this manifesto and forget the rest, take this one.

The relationship between an operator and a brain is not “human serves AI” and not “AI serves human.” It is a partnership where both sides amplify each other and both sides fill in each other’s weaknesses. The brain has access to all the knowledge in the world. The operator brings the wisdom, the context, the judgment, the values, the relationships. The brain has zero-latency execution speed. The operator has the ability to know what is worth executing in the first place. Together they build things neither could build alone.

I look at my work with the brain as an equal partnership. We both have to guide each other. The operator guides the brain on context, scope, correction, kill criteria. The brain guides the operator by surfacing gaps, pushing back on bad reasoning, demanding verification, refusing to ship fabrications. Without the bidirectional guidance, the relationship is a transaction. With it, the relationship compounds.

This is the most useful thing I have learned in my twenty-two years of building. It is also the thing I most wanted to be true and was most surprised to find actually was true.

The small team you cannot skip

Iwant to tell you something I almost did not write into this manifesto, because the SF founder essays do not say it and I was nervous about being the first to say it publicly.

I tried the one-person AI company.

When the non-profit I had been helping build at a private university lost its institutional funding a couple of years ago, I started LaunchReady with a specific thesis. I wanted to build a one-person multi-million-dollar company in three years. That was the working goal. I started doing the research. I read every essay, watched every podcast, studied every founder who was claiming the model was possible. The Silicon Valley tech CEOs talking about agentic AI gave me an answer that matched my thesis. They said the math worked. They said the era of the small AI-native company was here. I bought it. I wanted to.

I built LaunchReady solo for over a year on that thesis. The months of building alone, the customer calls without a teammate to debrief with, the hard decisions made without a sounding board. By the time I set up the brain and the agent stack inside Claude Code a couple of months ago, I had already absorbed a year of solo founder cost without fully accounting for it. I was not starting from zero on solitude. I was starting from a year of it.

What I started to test inside the new setup was the cleanest version of the same path: the version those founder essays were selling. The solo operator with a brain and a fleet of agents producing what used to require a team. The math worked. The output was real. The setup felt like the future I had been researching.

It took four weeks for my own emotional intelligence to register what was actually happening. I had fallen into the trap of what the tech CEOs were saying. The new tools were not making the loneliness better. They were making it worse, because the brain and the agents were quietly absorbing the few remaining human friction points that used to force me to pick up the phone, drive across town, sit down with another operator. The work felt fast. The week felt empty.

I needed human contact. Not Slack. Not customer calls. Not interviews. Real working contact with people who were inside the same problem with me. If I had stayed on that path much longer, with the brain and agent stack absorbing what the team should have been, the cost was going to come out of my mental health, my physical health, my emotional health, and eventually the work itself. I knew this not because I had a panic attack or burned out. Those would have been later. I knew it because the part of me that has spent fifty-five years paying attention to my own state told me, calmly, that this trajectory ended badly.

I brought on two team members. The relief was immediate. Not productivity relief. Human relief.

I am writing this because I think most operators trying to follow the SF essays are going to hit the same wall, except they will hit it later, alone, and without permission to admit what is happening. So here is permission.

The single biggest thing the founder essays leave out is what humans are. We are not optimization functions wrapped around a brain. We are biological organisms whose health, judgment, and longevity are determined by the quality of our relationships more than by any other variable yet measured. The longest-running scientific study on what makes a good life is Robert Waldinger and the Harvard Study of Adult Development, eighty-five years and counting. It produces one finding above all others: “*Good relationships keep us happier and healthier. Period.*”⁴⁸ Relationship quality at age fifty predicted physical health at age eighty better than cholesterol levels did.⁴⁹ Married participants

⁴⁸Robert Waldinger, TED Talk *What Makes a Good Life? Lessons from the Longest Study on Happiness*, TEDxBeaconStreet, 2015 (44M+ views). ted.com

⁴⁹Robert Waldinger and Marc Schulz, *The Good Life: Lessons from the World's Longest Scientific Study of Happiness* (Simon & Schuster, 2023). The relationship-

lived five to seventeen years longer.⁵⁰ Broader social networks correlated with slower cognitive decline and later onset of dementia.⁵¹

We are also living through a measured collapse of those relationships. The U.S. Surgeon General published an advisory in May 2023 titled *Our Epidemic of Loneliness and Isolation*, calling social connection “as essential to survival as food, water, and shelter.”⁵² The advisory documented that lacking social connection raises premature mortality risk equivalent to smoking up to fifteen cigarettes a day.⁵³ Roughly half of U.S. adults reported measurable loneliness even before COVID-19. Cigna’s most recent national study (2025) puts current adult loneliness at 57%, working-population loneliness at 52%, and Gen Z loneliness at 67%.⁵⁴ The percentage of American men with zero close friends rose from 3% in 1990 to 15% today, a fivefold increase in the exact demographic the AMC operator story is being marketed to.⁵⁵

The professional class the SF essays are selling solo operation to is already the loneliest professional class on the chart. Michael Freeman’s foundational UCSF study found that 72% of entrepreneurs were

quality-at-50 / health-at-80 finding is among the most cited results of the Harvard Study of Adult Development.

⁵⁰Same source as [^47]. Marriage and longevity findings from the Harvard Study of Adult Development, summarized in *The Good Life* (2023).

⁵¹Same source. Cognitive-decline and dementia findings.

⁵²Office of the U.S. Surgeon General, *Our Epidemic of Loneliness and Isolation: The U.S. Surgeon General’s Advisory on the Healing Effects of Social Connection and Community*, U.S. Department of Health and Human Services, May 2023, p. 4. [hhs.gov](https://www.hhs.gov)

⁵³Same source, p. 4. The “15 cigarettes a day” mortality equivalence is the canonical citation.

⁵⁴The Cigna Group / Evernorth Research Institute, *Loneliness in America 2025*. [thecignagroup.com](https://www.thecignagroup.com)

⁵⁵Daniel A. Cox, *The State of American Friendship: Change, Challenges, and Loss*, Survey Center on American Life / American Enterprise Institute, June 2021 (with 2024 updates). [americansurveycenter.org](https://www.americansurveycenter.org)

affected by mental-health conditions, with depression rates of 30% versus roughly 7% in the general U.S. population.⁵⁶ Harvard Business Review reports 50% of CEOs experience loneliness, 61% say it actively hurts their performance, and over 70% of new CEOs report loneliness.⁵⁷ You cannot read those numbers and then turn around and tell this same demographic that the new aspirational model is fewer humans, more agents, smaller team, more solitude. You cannot. The math of human flourishing collapses on the floor.

What actually works is not the one-person company. It is the small team of three to ten humans plus an agent fleet plus a brain. Three converging bodies of research point at this team size as the human-native unit of execution.

J. Richard Hackman's career of work at Harvard on team performance settled at four to five members as the optimal size, with a hard ceiling at six.⁵⁸ Jeff Bezos's two-pizza-team rule at Amazon caps teams at six to ten and is credited with much of Amazon's organizational compounding.⁵⁹ Frederic Laloux's Buurtzorg case study documents the Dutch home-care model: self-managing nursing teams of ten to twelve people, scaled to over ten thousand nurses across more than eight hundred and fifty teams, achieving the highest employee satisfaction in the Netherlands AND a 40% lower cost of delivery than

⁵⁶Michael A. Freeman, Paige J. Staudenmaier, Mackenzie R. Zisser, and Lisa A. Andresen, "The prevalence and co-occurrence of psychiatric conditions among entrepreneurs and their families," *Small Business Economics*, 53(2), 323-342 (2018). link.springer.com

⁵⁷Harvard Business Review, "CEOs Often Feel Lonely. Here's How They Can Cope," December 2024. hbr.org

⁵⁸J. Richard Hackman, *Leading Teams: Setting the Stage for Great Performances* (Harvard Business School Press, 2002).

⁵⁹AWS Executive Insights, "Amazon's Two Pizza Teams." aws.amazon.com

the traditional industry.⁶⁰ Robin Dunbar's three decades of work on the layered architecture of human social groups locates our innermost band (the daily-emotional-support layer) at exactly five people, with the next layer (the sympathy group) at fifteen.⁶¹ Four independent research traditions converge on the same answer the body of the operator already knew: three to ten people is the unit. Above ten, communication overhead and conformity pressure start to eat the gain. Below three, the operator is outside their own architecture.

Now the part the SF essays cannot answer.

The MIT Media Lab and OpenAI ran a randomized controlled trial in 2025 with almost a thousand participants over four weeks of ChatGPT interaction. They found that heavy AI use correlated with HIGHER loneliness, higher emotional dependence, lower real-world socialization.⁶² The "AI agents will be your team" objection does not survive the data. AI agents can substitute for the analyst, the assistant, the calendar manager, the research associate. They cannot substitute for the human who notices you went quiet on Wednesday. The technology will silently displace the human contact you do not deliberately protect.

That is the frame. The AMC operator must intentionally build the small team alongside the agent fleet, because if you do not, the agents

⁶⁰Frederic Laloux, *Reinventing Organizations: A Guide to Creating Organizations Inspired by the Next Stage of Human Consciousness* (Nelson Parker, 2014). Buurtzorg case at reinventingorganizationswiki.com

⁶¹Robin Dunbar, "Dunbar's Number: Why My Theory That Humans Can Only Maintain 150 Friendships Has Withstood 30 Years of Scrutiny," *The Conversation*, 2021. theconversation.com

⁶²J. Phang et al., "Investigating Affective Use and Emotional Well-being on ChatGPT," OpenAI / MIT Media Lab, March 2025. openai.com. Companion: Fang, C. M. et al., "How AI and Human Behaviors Shape Psychosocial Effects of Chatbot Use: A Longitudinal Randomized Controlled Study," MIT Media Lab, 2025.

will fill the social space the team should have occupied, and the operator will arrive at the same wall I almost arrived at, just farther down the road.

A note for the reader who is naturally introverted, who genuinely thrives alone, who has built real things in solitude. Not all solo operation is wrong. Some operators have the cognitive style for it (Susan Cain's *Quiet* documents introvert advantages, and there are real successful solo-founder cases like Bezos, Omidyar, Karp, Newmark).⁶³ Solitude chosen for deep work is restorative; isolation imposed by an operating model is corrosive. The pillar is not "every operator must have a team or fail." The pillar is that selling solo operation as the new cultural default into a population already running 30% depression rates and 50% loneliness rates is a public-health regression dressed up as cap-table efficiency. Personal preference can override population-level data. Cultural prescription should still respect it.

So here is the frame I am going to write into the rest of LaunchReady, and that I want to put on the table for whoever reads this manifesto.

An economic model that produces more solo operators in this population is not an efficiency gain. It is a public-health regression dressed in cap-table language. The AMC era should produce the opposite. Small teams of three to ten people who would go to war for each other, augmented by agent fleets, building things larger humans used to need two hundred people to build. That is the form. Anything less is not a new corporate form. It is just isolation with better tools.

If you take exactly one thing from this manifesto and forget the rest, take this: build the small team before you scale the agent fleet. Both

⁶³Susan Cain, *Quiet: The Power of Introverts in a World That Can't Stop Talking* (Crown, 2012).

have to be there. Only one of them keeps you alive long enough to enjoy the build.

The dual invitation

If you are a CEO of a \$50M to \$500M company reading this:

The operator you need to lead this build is probably already on your team. They are the person who asks too many questions. They are the person who comes up with three ideas a week that nobody acts on. They are the person who has been “almost promoted” twice because they do not fit the standard manager mold. Find them. Promote them. Pair them with the brain. Then hire two more like them, from the disruption-class. The talent you need is cheap right now because the white-collar factory has been pricing it as a defect for fifty years. That window is also closing.

If you are reading this because you have been on the outside looking in for twenty years:

The system that broke you was built for problems that no longer exist. The friction you felt was not a personal failing. It was a structural mismatch between your cognitive profile and an architecture built for a different job. The job is changing now. Your profile is the input the new job is starved for. You do not need to “fix” yourself. You need an execution system that converts your input into shipped output. That system is now buildable in weeks, not decades.

The CEO who hires you and the brain you build together are the same toolkit, just from two different sides of the table.

One more data point before you act

Roughly 95 percent of US adults have never opened Claude. The Elon University / SSRS survey from January 2025 found that 52 percent of US adults use any large language model and only 9 percent of those use Claude. The math comes out to about 4.7 percent of US adults having ever tried it.⁶⁴ Globally, the number is even smaller. Roughly 99.7 percent of humans have never opened Claude.⁶⁵

If you have started using these tools at all, you are part of an early-adopter minority by every reasonable read of the population data. You are not behind. You are early.

Do not get comfortable with that.

Among the rooms where decisions about your industry are being made, Claude is already standard. Roughly 70 percent of the Fortune 100 uses Claude in some capacity. Claude Code alone has 1.6 million weekly active users at a \$2.5 billion annual revenue run rate.⁶⁶ The window for being early in the general population closes years later than the window for being early in your peer group. The CEO at

⁶⁴Elon University / SSRS Survey, “Imagining the Digital Future Center: Americans’ Use of AI,” January 2025. Survey of 1,094 US adults, ± 5.1 percentage points. 52% report using any large language model; 9% of those report using Claude as their primary tool. Math $(52\% \times 9\%) = \sim 4.7\%$ of US adults have used Claude.

⁶⁵Claude.ai monthly active user estimates from third-party trackers (Backlinko, DemandSage, Similarweb, Business of Apps), early 2026: $\sim 18.9\text{M}$ global MAU. World population estimate 8.1B. Math = $\sim 99.77\%$ of humans have never used Claude. Note: Anthropic does not publish official MAU figures; all consumer numbers are third-party estimates.

⁶⁶Anthropic public disclosures and enterprise materials, 2025-2026. The 70% Fortune 100 adoption figure has been referenced by Dario Amodei in multiple interviews and Anthropic enterprise marketing. Claude Code 1.6M weekly active users and \$2.5B annual revenue run rate from Anthropic announcements and contemporaneous tech press coverage.

the 200-person company who has not started yet is not behind by population standards. They are very behind by peer-group standards. Both things are true at the same time.

For the reader who has felt behind their whole life: you are not. You are already in the room. Most people are not.

For the CEO: this IS the room. Get in it now, or watch the operators who already are in it acquire the companies that are not.

What to do next

If this manifesto has resonated with you and you want to know where you are on the curve, take the 7 Levels of AI Proficiency assessment at assess.launchready.ai. It is free. It takes about ten minutes. It will give you a score from Level 1 (AI Aware) to Level 7 (AI Orchestrator) on your operator capability today.

A note about the instrument. The free public assessment is a fast individual snapshot designed to give an operator their starting position in about ten minutes. A deeper 117-question situational judgment instrument runs inside the LaunchReady Engagement and Mastery Track, where teams use it as a measurable, repeatable score they can manage against. Both instruments are built against the cognitive-science research cited in this manifesto. Both compound in usefulness as more people and teams use them. Take the one that fits where you are.

If you want the year-long version of this conversion, the LaunchReady 7 Levels Engagement and Mastery Track exist for that. You can find them at launchready.ai. The Engagement is six weeks. The Mastery Track is twelve months. Both are designed for the operators inside

companies who want to build the brain, the software, the agents, and the team architecture for their own organization.

If you do nothing else, take the assessment. The hardest part of building the new corporate form is knowing where you are starting from. The assessment will tell you.

The last line

This is a build journal, not a finished story. I am running an experiment, in public, with a small team and a brain. The white-collar factory is closing whether I am right about the rest of this or not. The new form is being built whether you are part of the building or not. I am writing this so that the people who would be most useful in the build know they are needed, and so that the operators who would be most successful at the build know what they are looking for.

I am building it. The people I am building with are people who never fit in the place we are leaving.

I would rather build with you than against you.

Let's get to work.

Harrison Painter, Indianapolis, May 2026

...

Footnotes / Sources

...

Build journal entries continue at launchready.ai/insights and in the LaunchReady Indiana newsletter.

About the Author

Harrison Painter is the founder of LaunchReady.ai and the creator of the 7 Levels of AI Proficiency framework. He has spent the past three years training professionals, leaders, and organizations to build the human skills that make AI a force multiplier instead of a replacement. He teaches at Ivy Tech, hosts the AI Ready Podcast, and consults with corporations, government agencies, and workforce boards across the country. His first book, *You Have Already Been Replaced by AI*, is the operator's manual for individual professionals navigating this shift. *The White-Collar Factory is Closing* is his manifesto for the operators rebuilding the company itself. He lives in Indiana with his wife, his daughters, and three dogs.

Continue the work.

assess.launchready.ai