



When Software Runs Every Company

*A Founder's Thesis on AI, Organizations,
and the Next 18 Months*

Purusharth Sharma
Kami Labs
April 2026

This document outlines a founder's thesis on the most significant restructuring of work in human history, the company being built in response, and the plan for the next 18 months.

Section 1: The Era

It is 2027. A mid-sized software company wakes up before its founders do. By 6am, three AI agents have reviewed overnight customer signals, flagged a churn risk in the enterprise segment, drafted a response playbook, and queued it for human review. A fourth has negotiated a revised delivery timeline with a vendor — autonomously, within pre-approved parameters. A fifth is running a content campaign across four channels, optimizing in real time against the previous hour's engagement data. The founders arrive at 9am to a company that has already been thinking, deciding, and moving for three hours without them. They spend their morning on the work that actually requires them — strategy, relationships, the calls that demand human judgment. The rest runs itself.

Three hundred kilometers away, a manufacturing business making precision components for the automotive industry is running on similar logic. The promoter — a second-generation founder who took over from his father — is not buried in WhatsApp threads chasing production updates. An intelligence layer sits above his ERP, his inventory system, and his dispatch records, synthesizing everything into a single morning briefing. Overdue vendor payments have been flagged. A production bottleneck that emerged overnight has been identified, escalated, and assigned. A customer who has not reordered in 45 days has already received a personalised follow-up. He walks the factory floor with full visibility — not because he spent two hours pulling reports, but because the system did that before he arrived. His plant runs with the operational discipline of a company three times its size.

Both companies share the same underlying reality. Small teams producing what forty-person organizations would have delivered three years ago. No head of operations — an operations layer instead, a system of interconnected agents that monitors, flags, follows up, and executes across every function of the business. No content department — a content engine, trained on the company's voice and strategy,

deploying without waiting for a brief. Finance, customer success, business development — each function with an intelligence layer underneath it, handling the mechanical work autonomously, surfacing only what genuinely requires a human decision. These companies do not feel small. They feel precise. Lean by design, not by limitation.

This is not science fiction. It is the logical destination of a shift that is already underway. The companies building toward it now are creating distance that will not be easy to close.

For most of human history, thinking was personal. Problem solving lived inside a single mind, or passed between minds slowly. We have never, in the entire history of organized human activity, had access to another entity that could think alongside us. Until now.

That changes everything.

To understand why this moment is genuinely different, it helps to trace how thinking and work have been organized across history. In agrarian societies, value was physical — the land you worked, the crop you harvested, the labour you could sustain. The mind mattered at the margins. Decisions were simple, hierarchies were fixed, and the pace of change was generational.

The industrial revolution changed the relationship between human effort and output. Machines could do physical work at scale. But the machines were dumb — they could press, cut, assemble, and move, but they could not think. Human cognition remained irreplaceable. The factory floor needed engineers, managers, clerks, planners, and

strategists. Entire industries of knowledge work were born. The mind became the engine of economic value.

The information age accelerated this further. Software became the medium through which knowledge workers operated. Computers could store, retrieve, calculate, and communicate at speeds no human could match. But they still could not reason. They could not write a proposal, diagnose a problem, design a system, or make a judgment call. Those remained human capabilities. The knowledge worker remained indispensable.

For the first time in history, we have an entity capable of writing code, generating content, reasoning through complexity, and executing tasks — autonomously, at scale, and at a cost that approaches zero.

That is no longer true.

The emergence of large language models — and more specifically, the agentic AI systems being built on top of them — represents something categorically different from every previous technological shift. Not because it automates physical labour, as the industrial revolution did. Not because it stores and retrieves information faster, as the computing revolution did. But because, for the first time, it performs cognitive work. It reasons. It writes. It codes. It designs. It plans. It executes.

This is not a marginal improvement on existing tools. It is the arrival of a new category of entity in the history of work — one that can think alongside us, and increasingly, instead of us.

What This Means for Organizations

Most conversations about AI focus on the individual knowledge worker — the developer whose code gets written faster, the marketer whose copy gets generated in seconds, the analyst whose data gets processed in minutes. These are real and significant. But they miss the more profound implication.

Organizations are, at their core, coordination machines. They exist to organize human thinking and effort toward a shared objective. The org chart, the department structure, the tool stack, the meeting cadence, the reporting hierarchy — all of it is infrastructure designed to coordinate cognition across people. Every workflow, every process, every SOP is essentially an answer to the question: how do we get the right thinking to happen, in the right order, at the right time?

When the entity doing the thinking changes — when it is no longer exclusively human — the entire coordination machine has to be redesigned. Not tweaked. Redesigned.

The org chart, the department structure, the tool stack — all of it is up for redesign. This is not a software upgrade. It is an organizational restructuring event.

The businesses that understand this are not asking how to make their existing workflows faster using AI. They are asking a more fundamental question: if AI can think, reason, write, and execute — what is the right structure for an organization built around that capability from the ground up?

That is the question Kami Labs exists to answer.

The Reinvention of Thinking, Creativity, and Productive Work

Beyond organizations, this shift has profound implications for the nature of thinking and creativity itself. When a human being solves a problem, they draw on experience, intuition, pattern recognition, emotional intelligence, and domain knowledge accumulated over years. These capabilities are deeply personal and deeply contextual. They are also slow to develop, difficult to transfer, and impossible to replicate at scale.

AI changes the economics of cognition. Ideas that previously required a specialist to generate can now be produced in seconds. Documents that took hours to draft can be assembled in minutes. Code that required a team of engineers can be written by a single founder. Research that demanded weeks of reading can be synthesized in an afternoon.

This does not make human thinking less valuable. It makes a specific kind of human thinking dramatically more valuable — the kind that cannot be automated. Abstract reasoning under ambiguity. Leadership through complexity. The ability to ask the right question, not just answer the one in front of you. Creativity that draws on genuine lived experience and emotional depth. The capacity to build trust with other human beings.

The productive work of the future will be defined by how well individuals and organizations combine human and artificial intelligence. Not one or the other. Both, together, in a deliberately designed system.

Systems That Are Finally Alive

There is one more implication of this shift that deserves particular attention, because it is the least understood and the most consequential.

Until now, organizations were essentially mechanical. They moved when humans pushed them. Energy in, output out. The moment human attention shifted — when the founder got sick, when the key employee resigned, when the team got distracted — things slowed down or stopped entirely. Boring work and interesting work were

indistinguishable from the organization's perspective. Both required a human to do them. Both competed for the same finite pool of human attention.

AI-native organizations are different in kind, not just degree. The mechanical work — the follow-ups, the reporting, the data entry, the routine communication, the status tracking — runs itself. The organization has its own momentum. It does not stop when human attention shifts, because the systems sustaining it are autonomous. Feedback loops are faster, so the organization learns and corrects in near real-time. It becomes genuinely alive — not metaphorically, but functionally.

Current organizations are mechanical. They move only when humans push them. AI-native organizations are organisms. They have their own momentum, their own intelligence, their own capacity to learn and correct.

This also changes what work feels like for the humans inside these organizations. When the boring, repetitive, mechanical work is handled autonomously, what remains is the genuinely interesting work — the exploration, the strategy, the creativity, the relationship-building, the decisions that require judgment. Work becomes more alive too. It becomes, in a real sense, more like a game — with faster feedback, clearer progress, and more meaningful choices at every turn.

This is not a distant future. It is available now, to any organization willing to redesign itself around the new reality. The gap between those who do and those who don't is already opening. In the next 18 months, it will become very difficult to close.

The following section outlines the specific thesis points that Kami Labs is betting on — the claims about the world that, if true, define the opportunity and the obligation to move now.

Section 1 — Key Takeaways

- For the first time in history, an entity exists that can think, reason, write, code, and execute — autonomously, at scale, and at near-zero cost. This is not an upgrade to existing tools. It is a new category of entity in the history of work.
- Every previous technological shift left human cognition irreplaceable. This shift does not. It performs cognitive work directly, which changes everything downstream.
- Organizations are coordination machines built to organize human thinking toward shared objectives. When the entity doing the thinking changes, the entire machine must be redesigned — not tweaked.
- Human thinking does not become less valuable — a specific kind of human thinking becomes dramatically more valuable: abstract reasoning, leadership under ambiguity, genuine creativity, and the ability to build trust.
- Current organizations are mechanical — they move only when humans push them. AI-native organizations are organisms — they have their own momentum, their own intelligence, and their own capacity to learn and self-correct.
- When mechanical work runs autonomously, what remains for humans is the genuinely interesting work — exploration, strategy, creativity, judgment. Work becomes more alive, more gamified, more meaningful.
- The gap between organizations that redesign themselves around this reality and those that don't is already opening. In the next 18 months, it will become very difficult to close.

Section 2: The Thesis

There are nine things we believe to be true about what is happening in the world right now. Not trends. Not predictions. Bets. Specific, falsifiable claims about the direction of technology, organizations, and work — claims that, if correct, define both the opportunity and the urgency of what Kami Labs is building.

WHAT IS HAPPENING

1. AI will eat software. Software will run organizations.

Marc Andreessen said software will eat the world. That turned out to be correct. Every industry that could be reorganized around software was. The companies that understood this early — and built accordingly — became the most valuable companies in history.

The next version of that claim is this: AI will eat software. The software products that accumulate data about a business — the ERP, the CRM, the project management tool, the communication platform — will not remain passive repositories of information. They will become active. They will start interpreting that data, making recommendations, and eventually making decisions. The line between a tool that assists a business and a system that runs a business will dissolve. This is not a feature update. It is a fundamental shift in what software is for.

The line between a tool that assists a business and a system that runs a business will dissolve. This is not a feature update. It is a fundamental shift in what software is for.

There is a second dimension to this shift that deserves equal attention: the economics of building software itself are being restructured from the ground up. For decades, the build vs buy vs lease decision was shaped by one constant — writing custom software was expensive, slow, and risky. Most businesses bought packaged solutions because building their own was out of reach. That assumption is dissolving.

Coding agents — AI systems that can write, test, debug, and ship production-grade code — are collapsing the cost and time required to build usable software. A workflow that previously required a team of engineers and months of development can now be prototyped in days and shipped in weeks. This does not mean every business will build everything from scratch. It means the decision of whether to build, buy, or lease is now genuinely open in a way it has never been before. Custom software — tailored precisely to a business's specific workflows, data, and context — is becoming accessible to companies that could never have considered it.

This does not spell the end of software developers. It spells the end of the developer as a bottleneck. The best developers — those who can think architecturally, ship fast, and maintain complex systems — will be more valuable than ever. They will simply build more, faster, with smaller teams. The developer who once managed a feature backlog measured in months will ship products measured in weeks. The constraint shifts from capacity to judgment.

2. Knowledge work is being completely reinvented.

The way people work, think, communicate, create, and coordinate inside organizations is changing at a speed that most businesses and most people are not prepared for. The workflows that were designed around human cognitive limitations — the approval chains, the status meetings, the weekly reports, the email threads — were never efficient. They were the best available option given what humans could do. That constraint is gone.

Every workflow in every organization is now up for redesign. Not optimization. Redesign. The question is no longer how to make the existing process faster. It is whether the existing process needs to exist at all in its current form, or whether it can be restructured entirely around the capabilities that AI now provides.

3. The 18-month window is now.

This is not a ten-year thesis. The window is now. Companies like Spotify and others have already stopped writing code manually. Entire software products are being built, iterated, and killed in months. The gap between AI-native organizations and traditional ones is not theoretical — it is already measurable in productivity, speed, and output quality.

The businesses that move in the next 18 months will create distance that cannot be closed by simply adopting the same tools later. First-mover advantage in organizational redesign compounds. The learning, the systems, the institutional knowledge built during this period will be proprietary. Latecomers will be able to buy the tools. They will not be able to buy the experience of having built with them from the beginning.

Latecomers will be able to buy the tools. They will not be able to buy the experience of having built with them from the beginning.

4. For the first time, systems can be truly alive and autonomous.

Every organization in history has been, at its core, a machine that stops when humans stop pushing it. Momentum was always human-dependent. The moment attention shifted, things slowed. The moment a key person left, institutional knowledge walked out with them. The moment the founder was unavailable, decisions waited.

AI-native systems break this dependency. For the first time, the infrastructure of an organization can have genuine autonomy — monitoring, flagging, following up, executing, and learning without requiring a human to initiate every action.

Organizations can develop their own momentum. They can become, in a functional sense, alive. Not metaphorically. The feedback loops are real. The self-correction is real. The compounding intelligence is real. This changes what it means to build a company — and what it means to run one.

HOW TO THINK ABOUT IT

5. Business fundamentals remain the same. The game just got more complex.

None of this means the fundamentals of business have changed. Companies still need to find customers, deliver value, manage cash, build teams, and earn trust. The market still rewards those who solve real problems better than anyone else. Relationships still matter. Reputation still matters. Execution still matters.

What has changed is the complexity of the game. The number of variables a business must manage has increased. The speed at which the environment changes has increased. The tools available to navigate that complexity have also increased — but only for those who learn to use them. AI does not simplify business. It raises the ceiling on what is possible and raises the floor on what is required to compete.

6. Founder-led companies will move the fastest.

The single greatest predictor of how quickly an organization will adapt to this shift is the agency of the person at the top. Founders — particularly those who are deeply technical or deeply curious — are the highest-agency workers in any economy. They make decisions faster, tolerate more ambiguity, and are more willing to restructure everything in pursuit of a better answer.

This is why founder-led companies will lead the AI-native transition. Not because they have more resources — most don't. But because the bottleneck to adoption is not

budget. It is willingness to change. And founders, by definition, have built their identity around changing things.

The bottleneck to AI-native adoption is not budget. It is willingness to change. Founders, by definition, have built their identity around changing things.

7. The human premium rises. A new mindset is required.

As AI handles more of the mechanical, repetitive, and even analytical work, the capabilities that remain distinctly human become disproportionately valuable. Communication. Leadership. Abstract thinking. The ability to build. The capacity to hold complexity without losing clarity. These are not soft skills. They are the hard skills of the AI era.

But possessing these capabilities in isolation is not enough. The new mindset required to thrive in this economy is one that holds all of them together simultaneously — that can move between strategy and execution, between human relationship and system design, between creative vision and operational discipline. The leaders who will define this era are not specialists. They are integrators. People who can think at the level of systems while remaining deeply connected to the human realities those systems serve.

WHAT TO BET ON

8. The entire technology stack of businesses will be disrupted and augmented.

Every layer of the technology stack that businesses currently run on is being rebuilt. Not replaced wholesale — augmented, extended, and in many cases fundamentally reimaged. ERPs will develop intelligence layers. CRMs will become relationship engines. Communication tools will become coordination systems. Project management platforms will become execution environments.

The businesses that thrive will not be those that find the one AI tool that solves everything. They will be those that develop a coherent view of how AI fits across their entire stack — and build the workflows, the integrations, and the institutional knowledge to make that stack operate as a unified system. The winners will be systems thinkers, not tool collectors.

9. We are betting on specific technologies.

Not all AI platforms are equal. Not all tools will compound over time. We have made deliberate bets on the technologies we believe will define the infrastructure of AI-native organizations over the next decade.

Claude — for reasoning, writing, coding, and agentic execution. Notion — for knowledge management, workflow design, and organizational memory. And a select set of other tools that, together, form the stack we are building on, building with, and building for.

These are not vendor preferences. They are architectural decisions. The depth of mastery we develop in these specific tools is itself a competitive advantage — one that will compound as the tools evolve and as our experience with them deepens.

Section 2 — Key Takeaways

- AI will eat software. Software will run organizations. The line between a tool that assists a business and a system that runs it will dissolve.
- The economics of building software are being restructured. Coding agents are collapsing the cost and time to build. The build vs buy vs lease decision is now genuinely open. The best developers will build more, faster — the constraint shifts from capacity to judgment.

- Knowledge work is being completely reinvented — every workflow is up for redesign, not just optimization.
- The 18-month window is now. First-mover advantage in organizational redesign compounds in ways that cannot be bought later.
- For the first time, systems can be truly alive and autonomous — organizations can develop genuine momentum independent of human attention.
- Business fundamentals remain the same. The game just got more complex. AI raises both the ceiling on what is possible and the floor on what is required to compete.
- Founder-led companies will move the fastest. The bottleneck to adoption is not budget — it is willingness to change.
- The human premium rises. Communication, leadership, abstract thinking, and building become the hard skills of the AI era. A new mindset — one that integrates all of them — is required to thrive.
- The entire technology stack of businesses will be disrupted and augmented. Winners will be systems thinkers, not tool collectors.
- We are betting on Claude, Notion, and select others — as architectural decisions, not tool preferences. Mastery in these compounds over time.

Section 3: The Journey

The thesis in the previous section is not something I arrived at by reading about AI. It is something I earned through a specific sequence of decisions and experiments over the last six months.

The Starting Point

By August 2025, I had something most early-stage founders spend years chasing: a functioning business with proven product-market fit.

I was working as a fractional CMO for Golden Earth, a technology-rooted company with a genuine appreciation for documentation and systems thinking. The agency was generating \$5,000 per month. The team was lean and remote. The work was genuinely enjoyable. I had found the rhythm that most people describe as success at this stage — stable revenue, good clients, interesting problems, room to grow.

Most people in that position stay. I didn't.

The Decision

Two things happened simultaneously that made staying feel like the wrong choice.

The first was technical. AI had gotten good enough — specifically, agentic AI had gotten good enough — that the barrier to building real products had collapsed. I no longer needed to wait for a co-founder with a specific technical skill, or a funding round to hire engineers, or months of development to test an idea. I could build, test, and iterate on

solutions myself, faster than I had ever been able to before. The excuse to wait had disappeared.

The second was personal. I am, at my core, a technologist and a nerd. And I was watching the most significant technological shift in my lifetime unfold in real time — a shift that would fundamentally change how human beings work, communicate, and think. The idea of spending that moment creating marketing content for other people's businesses felt deeply wrong. Not because the work was bad. Because the moment was too important to be on the sidelines of.

So I made the decision. I scaled down the agency. I redirected everything toward building.

The barrier to building had collapsed. And the most important technological shift of a generation was underway. Staying on the sidelines was not an option.

The Exploration

What followed was a period of intense, sometimes chaotic exploration. I started working on products for knowledge workers and organizations — tools that could make the kind of AI-native workflows I was imagining accessible to real businesses. In time, two ideas crystallized: OrgOS, a framework for building and operating AI-native organizations, and Kirrin, an AI Chief of Staff — an agentic intelligence layer that sits above a company's existing systems and gives founders and operators the visibility and follow-through they had always lacked.

I hired a small team. I built early versions. I pushed myself technically in ways I had not done before. The work was hard and energizing in equal measure.

Then I decided I needed to go deeper.

Bangalore

I traveled to Bangalore and lived with the CPO of a blockchain company — a friend who was operating at a technical depth I wanted to reach. I learned Claude Code hands-on, building directly, making mistakes in real time, developing the kind of embodied understanding of agentic development that you cannot get from documentation or tutorials. I built what I could. I broke things. I rebuilt them.

This was not a detour. It was the most important investment I made in this entire period. The difference between understanding agentic AI conceptually and understanding it through direct building is the difference between knowing a language and thinking in it. I came back from Bangalore thinking in it.

The difference between understanding agentic AI conceptually and understanding it through direct building is the difference between knowing a language and thinking in it.

The ICO

Around the same time, I attempted something unconventional. I ran an ICO — an initial coin offering — for Kirrin, piggybacking on a friend's platform and its existing leverage in the crypto community. The timing was unfortunate. The broader crypto market was in a significant downturn, and the raise did not come together in the way I had hoped.

But the experience was not a failure. It was an education. The process of putting Kirrin in front of a community of investors and technologists — of having to explain it, defend it, and refine it under real scrutiny — sharpened the thinking considerably. More importantly, it opened a network. The conversations I had during that period, and the relationships I built, remain assets. In this business, the network you unlock through a failed attempt is often more valuable than the capital you were trying to raise.

Now

We are living through the most significant restructuring of work in human history. The technology is here. The window is open. And there is nowhere I would rather be than building inside it.

Section 3 — Key Takeaways

- By August 2025, the agency had proven PMF at \$5K/month — a stable, enjoyable business that most founders would have stayed in.
- The decision to pivot was driven by two simultaneous realizations: the barrier to building had collapsed, and the most important technological shift of a generation was underway. Staying on the sidelines was not an option.
- OrgOS and Kirrin emerged from a period of intense exploration — a framework for AI-native organizations and an agentic intelligence layer for founders and operators.
- The Bangalore chapter was the most important technical investment of the period — learning Claude Code hands-on, developing embodied understanding through direct building rather than documentation.
- The Kirrin ICO did not raise capital, but sharpened the thesis through real scrutiny and unlocked a network that remains a genuine asset.
- The technology is here. The window is open. There is nowhere else to be.

Section 4: What Kami Labs Is

Kami Labs is an AI-native company builder. It exists at the intersection of two realities: the most significant restructuring of work in human history is underway, and most businesses — from early-stage startups to established SMEs — do not yet have the frameworks, the systems, or the infrastructure to navigate it.

We are building those frameworks, those systems, and that infrastructure. Not as consultants who advise and leave. As builders who deploy, operate, and iterate alongside the businesses we work with.

Kami Labs operates across four functions. Each one is deliberately designed. Each one feeds the others.

Earn — Mugen Engines

Mugen Engines is Kami Labs' revenue engine. It is a content and distribution infrastructure business that builds and deploys AI-native content and email marketing systems for technology companies, software businesses, and remote-first organizations globally.

The core philosophy of Mugen Engines is systems over services. We do not deliver traditional agency work — campaigns created, delivered, and forgotten. We deploy infrastructure. Engines that generate, distribute, and optimize content autonomously, compounding in value over time rather than requiring constant human intervention to sustain.

Mugen Engines serves two purposes simultaneously. It generates the revenue that funds Kami Labs' broader product development. And it is the laboratory where we test, validate, and refine the AI-native workflows that eventually become the foundation of our products. Every client engagement is both a commercial relationship and a research opportunity.

The immediate objective for Mugen Engines is to identify and scale a \$20,000 per month agentic service — a productized offering that can be delivered consistently, at quality, without being constrained by the hours in a day.

Every client engagement is both a commercial relationship and a research opportunity. Mugen Engines is the laboratory where the thesis gets tested in the real world.

Build — Kirrin and OrgOS

Kirrin and OrgOS are the product bets at the heart of Kami Labs. They represent our long-term view of what AI-native organizational infrastructure looks like — and our conviction that the businesses which build on that infrastructure early will operate at a fundamentally different level than those that don't.

OrgOS is an open-source framework for building and mapping AI-native organizations. It is the operating system for a new kind of company — one where workflows are documented, decision logic is captured, and the systems running the business are visible, auditable, and continuously improving. OrgOS is being dogfooded first, deployed across Kami Labs' own entities before being offered externally.

Kirrin is the intelligence layer that sits above a company's existing systems — the ERP, the CRM, the communication stack — and makes the organization genuinely responsive. It provides founders and operators with daily visibility into what is happening across their business, ensures nothing falls through the cracks, and over time develops into a

full agentic Chief of Staff. The entry point is operational visibility and follow-through. The destination is an organization that thinks.

Kirrin's near-term GTM focus is Gujarat-based MSMEs — promoter and founder-led businesses that are scaling but still running on the founder's personal attention. The pitch is not abstract: we sit above your existing systems, we don't replace them, and we give you back the visibility and follow-through you have been trying to maintain manually.

The entry point is operational visibility and follow-through. The destination is an organization that thinks.

Learn — R&D and Training

The third function of Kami Labs is where proprietary knowledge gets built. We run active R&D engagements with a select group of companies — including Stone Sapphire, Golden Earth, Acros, Smoothmix, Atliq, and Dalnex, among others from our network — using these relationships to develop, test, and refine the agentic frameworks that feed back into our products.

Alongside this, we run training programs — intensive, hands-on engagements that teach founders, operators, and teams how to work with AI natively. Not conceptually. Practically. The workshop format we are developing brings cohorts of six to ten high-agency founders through a weekend of building with Claude Code, learning by doing rather than by listening.

This function exists for two reasons. It generates additional revenue. And it accelerates learning at a rate that no amount of internal R&D alone could match. Every company we work with teaches us something about how AI-native systems behave in the real world, under real constraints, with real people. That knowledge compounds.

Distribute — Personal Brand and Media

The fourth function is distribution. Building the right products is not enough. Reaching the right people — consistently, credibly, at scale — requires owning an audience rather than renting one.

Kami Labs is building its distribution infrastructure around a personal brand anchored by PXSH — a writing and content presence that operates at the intersection of AI, organizational design, and the future of work. The voice is measured, architectural, and direct. The content strategy is built around the same philosophy as everything else we do: systems over one-off outputs. A content engine, not a content calendar.

The distribution function serves the entire Kami Labs ecosystem. It brings clients to Mugen Engines. It builds the audience for Kirrin and OrgOS. It attracts the collaborators, partners, and talent that make everything else possible. And it ensures that when Kami Labs has something important to say, it has a platform to say it from.

The Internal Logic

These four functions are not parallel tracks. They are a deliberately designed system.

Mugen Engines earns the revenue that funds the build. The R&D work generates the knowledge that sharpens the products. The products generate the proof points and the stories that fuel the distribution. The distribution brings in better clients for Mugen Engines. Each function strengthens every other.

This is what it means to build AI-natively — not just to use AI tools, but to design the organization itself as a system where every part compounds every other part.

Section 4 — Key Takeaways

- Kami Labs is an AI-native company builder operating across four functions: Earn, Build, Learn, Distribute.
- Mugen Engines is the revenue engine — a content and distribution infrastructure business deploying AI-native systems for technology companies globally. Immediate target: \$20K/month productized service.
- Kirrin is an agentic intelligence layer for founders and operators. OrgOS is the open-source framework for building AI-native organizations. Both are being dogfooded internally before being sold externally.
- The R&D and training function builds proprietary knowledge through active engagements with select companies and intensive hands-on workshops with cohorts of high-agency founders.
- Distribution is built around PXSH — a personal brand and content presence that owns audience rather than renting it, serving the entire Kami Labs ecosystem.
- The four functions are a system, not parallel tracks. Each one compounds every other. This is what it means to build AI-natively.

Section 5: The Market

Kami Labs is not trying to serve everyone. The businesses that will benefit most from what we are building share a specific profile — and identifying them clearly is as important as building the right products for them.

We are operating across two distinct markets simultaneously, with different motions, different objectives, and different time horizons.

The Global Market — Mugen Engines

The primary commercial focus of the next 90 days is the global market for technology companies, software businesses, and remote-first organizations.

These are businesses that already understand the value of content and distribution. They are not being convinced that marketing matters — they know it does. What they lack is the infrastructure to do it consistently, at quality, without it consuming disproportionate amounts of founder or team time. They are the right first clients for Mugen Engines because they move fast, they understand systems thinking, and they are willing to pay for outcomes rather than hours.

The ICP for Mugen Engines is being refined through the first 90 days of active sales and client work. We have strong directional clarity — technology companies, remote businesses, founder-led software products — but the specific profile of the best client will sharpen through direct engagement. This is deliberate. We are not pretending to have certainty we don't yet have. The 90-day sprint is partly about earning and partly about discovering exactly who benefits most from what we build.

What we are confident about is the offer architecture. AI-native content and distribution infrastructure — deployed as a system, not delivered as a service. Productized, scalable, and designed to compound in value over time for the client.

We are not pretending to have certainty we don't yet have. The 90-day sprint is partly about earning and partly about discovering exactly who benefits most from what we build.

The India Market — Kirrin and OrgOS

The second market is closer to home and longer in time horizon. Gujarat-based MSMEs — promoter and founder-led businesses with revenues between 25 crore and 500 crore — are the near-term GTM target for Kirrin.

These businesses share a specific and well-understood problem. They are run by high-agency founders who have built significant companies largely on personal attention, relationships, and hustle. As they scale, the founder becomes the bottleneck. Things fall through the cracks not because the business lacks capability, but because the founder lacks visibility. The ERP captures transactions. The CRM captures contacts. But nothing captures the full operational picture and surfaces it in a way the founder can act on.

Kirrin is the answer to that problem. It sits above the existing systems — not replacing them, not requiring a rip-and-replace — and provides the intelligence layer that makes the organization genuinely responsive. The pitch to a Gujarat MSME promoter is not about AI in the abstract. It is about getting back the operational visibility and follow-through that currently lives only in the founder's head.

The sales motion for this market bypasses traditional enterprise channels entirely. We are going directly to promoters and founders through personal networks — Parshva's

relationships, our existing client base, and the trust built through the R&D engagements already underway with companies in this segment.

The pitch is not about AI in the abstract. It is about getting back the operational visibility and follow-through that currently lives only in the founder's head.

The Agency Network — A Ready Testing Ground

There is a third segment worth highlighting separately: digital and technology agencies. Kami Labs has direct access to a significant network of agency founders and operators — built through years of working in and around the agency ecosystem. These businesses are structurally well-suited to early Kirrin and OrgOS adoption. They run on knowledge work, they are founder-led, they struggle with the same operational visibility and follow-through problems as any scaling SME, and they are typically more technically literate than traditional businesses — which means shorter sales cycles and faster deployment.

This network is not a future aspiration. It is an existing asset, and it gives Kami Labs a ready testing ground for OrgOS and Kirrin that most early-stage product companies would spend years trying to build. The agency network accelerates validation, generates real-world proof points, and creates a pipeline of early adopters who are predisposed to move fast and share what they learn.

Why These Two Markets Together

The combination is deliberate and mutually reinforcing.

The global technology market through Mugen Engines generates the revenue and the AI-native workflow knowledge that funds and informs the product build. The Gujarat MSME market through Kirrin generates the deep organizational intelligence that shapes

the product into something genuinely differentiated. The agency network accelerates validation across both.

One market funds the company. The other market builds the product. The network proves it. Together, they create a business that is commercially sustainable and product-led simultaneously — without requiring us to choose between revenue and building.

Section 5 — Key Takeaways

- Kami Labs is operating across two distinct markets with different motions and different objectives.
- The global market — technology companies, software businesses, remote-first organizations — is the primary commercial focus of the next 90 days through Mugen Engines. ICP is being refined through active sales and client work.
- The India market — Gujarat-based MSMEs, promoter and founder-led businesses — is the near-term GTM target for Kirrin. Sales motion goes directly to founders through personal networks, bypassing enterprise channels entirely.
- Kirrin's pitch in this market is not abstract: sit above existing systems, provide operational visibility and follow-through, give founders back the clarity they have been maintaining manually.
- The agency network is an existing asset — technically literate, founder-led businesses with shorter sales cycles that give Kami Labs a ready testing ground for OrgOS and Kirrin.
- One market funds the company. The other builds the product. The network proves it. Together they create a business that is commercially sustainable and product-led simultaneously.

Section 6: The Plan

The thesis is clear. The architecture is designed. What follows is how we execute against it — broken into three phases across 18 months, each one building on the last.

We operate in three-month sprints because that is the cadence at which reality is moving. Each sprint has a specific objective, a specific set of deliverables, and a specific definition of success. We do not plan in years and execute in quarters. We plan in sprints and learn in weeks.

Phase 1 — Now to 90 Days: Research, Development, and One Scalable Product

The objective of the first phase is not to scale. It is to discover, validate, and crystallize. By the end of 90 days, Kami Labs will have one clear, productized Mugen Engines offer — validated through real engagements, refined through real feedback, and ready to be sold at volume.

The R&D Engagements

We are currently in active conversation with a range of companies across our network who will serve as the R&D ground for this phase. These include Golden Earth, Stone Sapphire, Acros, Smoothmix, Atliq, Dalnex, Pranav Sir's company, and additional companies from Parshva's, Divya's, and Kshitij bhai's networks.

These are not sales conversations. They are structured learning partnerships. Each engagement is designed to surface a specific problem, test a specific solution, and generate the evidence that informs the productized offer. The Golden Earth engagement

is already underway — we are building a Claude Code-powered content engine for them, which serves as the live proof of concept for the broader Mugen Engines thesis.

We are also actively developing solutions for cybersecurity companies and low-code platform businesses — two verticals within the B2B technology space that have strong content and distribution needs and are structurally well-suited to the AI-native infrastructure we are building.

These are not sales conversations. They are structured learning partnerships — designed to surface specific problems, test specific solutions, and generate the evidence that informs the productized offer.

The Hypothesis

The working hypothesis for Mugen Engines is this: B2B technology companies — software businesses, agencies, cybersecurity firms, low-code platforms, remote-first organizations — need AI-native content and email infrastructure that runs as a system, not as a service. Not campaigns. Not retainers. Engines. Built once, compounding over time, generating pipeline without consuming founder attention.

The 90 days are about validating this hypothesis in enough real contexts to know exactly what the best version of this product looks like — and who it is built for.

The Productized Offer

By day 90, the objective is one clear, scalable Mugen Engines product. An AI-native content and email engine for B2B technology companies — designed, priced, and packaged for consistent delivery at quality. Not a bespoke engagement for each new client. A defined system that can be deployed repeatedly, with predictable outcomes, at a price point that works for both sides.

Kirrin and OrgOS — Internal Dogfooding

In parallel, Kirrin is being deployed internally across Kami Labs' own entities — Mugen Engines, the Kirrin product company, and the founder's office. The Notion-based prototype serves as the initial implementation. The coded Kirrin.ai product runs in parallel development. We will not ask clients to trust a system we have not trusted ourselves first.

Distribution

The PXSH personal brand launches in this phase. Writing begins. The content engine is built and starts running. The objective is not viral reach. It is consistent, credible presence — the kind that compounds over months rather than exploding and fading.

Partnerships

Three partnership tracks are active in this phase. Parshva is co-founder across Kami Labs and Mugen Engines — the execution partner on both the commercial and product sides. Divya is a development collaborator, with potential to deepen into a co-founder relationship for specific development-led initiatives. Siddharth Parwatay is a potential co-founder for the media and technology content side — relevant to the distribution function and the longer-term media vision.

Phase 2 — 90 Days to 6 Months: Scale, Deepen, and Ship

Phase 2 begins with a validated product and the first real commercial traction. The objective of this phase is to scale what is working, deepen the product development, and begin shipping Kirrin externally.

Mugen Engines

With a productized offer in hand, the commercial push begins in earnest. Target: \$20,000 to \$25,000 per month. The R&D partners from Phase 1 become the first wave of paying clients where fit is strong. The refined ICP drives outreach. Delivery is systematized so that growth does not require proportional increases in founder time.

Kirrin — First External Pilots

The internal dogfooding from Phase 1 generates the first external pilot opportunities. Select companies from the R&D network and the agency network become the first Kirrin pilots — structured engagements with clear success metrics, designed to validate the product in real organizational environments outside Kami Labs.

The agency network is a particular priority here. Digital and technology agencies in our existing network are technically literate, founder-led, and predisposed to move fast. They are the ideal first external Kirrin adopters — and their feedback will shape the product faster than any other segment.

OrgOS — Open Source Release

The OrgOS framework, refined through internal deployment and R&D engagements, begins moving toward an open source release. This builds community and credibility around the framework and creates a distribution channel for Kirrin.

Training and Workshops

The workshop program launches in this phase. Cohorts of six to ten high-agency founders. Weekend format. Hands-on building with Claude Code. Two workshops per month. Revenue, network, and pipeline for Mugen Engines and Kirrin simultaneously.

Phase 3 — 6 to 18 Months: Double Down and Build for Scale

Phase 3 is where Kami Labs makes its most consequential decisions. Real data — on Mugen Engines clients, Kirrin pilots, OrgOS traction — drives resource allocation.

The Double Down

If Mugen Engines is at or above \$25,000 per month and Kirrin pilots are generating strong signals, engineering investment accelerates. The Notion prototype gives way to the coded Kirrin.ai product as the primary deployment vehicle. GTM for Kirrin expands beyond the initial Gujarat MSME focus.

Fundraising Optionality

A scaled agency at \$25,000 per month combined with clear PMF signals from Kirrin pilots creates the conditions for a preseed or seed round from a position of strength. We are not planning to raise in Phase 1. We are building the evidence that makes a raise on good terms possible in Phase 3.

Distribution at Scale

By Phase 3, the PXSH content engine has been running for six months. The compounding effect of consistent quality content begins to show. Distribution becomes a genuine competitive advantage rather than an aspiration.

The Operating Principle

Every phase is designed around one principle: learn faster than we spend. The capital deployed in each phase is an investment in learning — about the market, the product, the right clients, and what works. The faster we learn, the faster we double down on what's working and stop what isn't.

This is why we operate in sprints. Not because 90 days is a magic number. Because 90 days is long enough to generate real signal and short enough to change direction before a bad bet compounds.

Section 6 — Key Takeaways

- Phase 1 is a research and development sprint, not a sales sprint. The objective is one validated, productized Mugen Engines offer by day 90.
- Active R&D conversations are underway with Golden Earth, Stone Sapphire, Acros, Smoothmix, Atliq, Dalnex, Pranav Sir's company, and networks from Parshva, Divya, and Kshitij bhai. The Golden Earth content engine on Claude Code is the live proof of concept.
- The Mugen Engines hypothesis: AI-native content and email infrastructure for B2B technology companies — with early traction in cybersecurity and low-code verticals.
- Three active partnership tracks: Parshva as co-founder, Divya as development collaborator, Siddharth Parwatay on the media side.
- Phase 2 target: \$20-25K/month Mugen Engines, first external Kirrin pilots, OrgOS open source release, workshop program launch.
- Phase 3: double down on product based on real data, build toward preseed or seed round from a position of strength, scale distribution.
- Operating principle: learn faster than we spend. Sprints generate signal fast enough to change direction before bad bets compound.

Section 7: The Ask

This is a direct capital deployment plan. What is needed, what it funds, what it delivers, and what comes after.

The Relationship

Firdosh Sir has supported this work before. This document is an update as much as it is a proposal — full transparency on where things stand and where the next capital goes.

The Ask

Tranche	Amount	Timing
Tranche 1	\$10,000 — \$15,000 USD	Immediate
Tranche 2	\$10,000 — \$15,000 USD	End of April 2026
Total	\$20,000 — \$30,000 USD	April 2026

Capital Deployment

Category	Amount	Purpose
Hardware and tooling	\$5,000	Development infrastructure for agentic systems
R&D and development	\$5,000	Claude Code builds, Kirrin and OrgOS prototypes

Mugen Engines sales	\$5,000	LinkedIn sales program, ads, outreach tool stack
Server and operational costs	\$5,000	Infrastructure, hosting, agentic system runtime
Travel and miscellaneous	\$5,000	Client meetings, R&D engagements, April – May
Runway and selectivity	\$5,000	Time to find right first clients without revenue pressure

90-Day Deliverables

Deliverable	Definition of Success
Mugen Engines productized offer	One validated, repeatable AI-native content and email engine for B2B tech companies
Internal Kirrin deployment	OrgOS running across all three Kami Labs entities on Notion prototype
Kirrin.ai parallel development	Coded product in active development alongside Notion prototype
PXSH distribution	Content engine live, writing published consistently, audience building
R&D engagements	Active structured partnerships with Golden Earth, Stone Sapphire, Acros, Smoothmix, Atliq, Dalnex, Pranav Sir's company and network
Foundation for Phase 2	Refined ICP, first client results, evidence base for preseed or seed round

Scenario Planning — Day 90

Scenario	Condition	Response
Double down	Mugen Engines has validated product and early traction. Kirrin internal deployment live. Strong signals.	Scale Mugen Engines to \$25K/month. Accelerate Kirrin engineering. Pursue preseed or seed round from strength.
Adjust and continue	Hypothesis partially validated. Some elements need refinement.	Deploy next sprint against sharpened thesis. Capital bought the learning. Direction remains the same.

The worst outcome is not a failed hypothesis. It is not moving while the window is open.

The Commitment

Full transparency through this phase. Regular updates. The work will speak for itself.

Section 7 — Key Takeaways

- Immediate ask: \$10,000 to \$15,000 USD. Second tranche of similar size by end of April 2026. Total: \$20,000 to \$30,000 USD.
- Capital deployed across six categories: hardware, R&D, sales infrastructure, server costs, travel, and runway for client selectivity.
- Six defined deliverables by day 90 — each with a clear definition of success.
- Two scenarios at day 90 — both productive. The only failure mode is inaction.
- This is a continuation of an existing relationship. Full transparency and regular updates are the commitment.

