The Agentic Startup Manifesto

From SaaS to Autonomous Agents

Author’s Note

We wrote this manifesto at InTheValley because we believe the next generation of startups will be built by autonomous agents, not traditional apps.

These are the principles guiding every product, project, and partner we work with.

— The InTheValley Team

Table of Contents

1. Introduction

In 2011, Marc Andreessen proclaimed that *“software is eating the world.”*

Today, software itself is on the menu - **AI agents are eating software**.

The startup and venture ecosystem is witnessing a profound shift: what used to be static SaaS products are evolving into dynamic, AI-driven entities.

Y Combinator’s Harj Taggar predicts that *“for every SaaS company, there will be a corresponding AI agent company.”*

This means that instead of building yet another app, the next generation of founders will create **autonomous digital teammates**: services that can interpret goals, make decisions, and take action on a user’s behalf.

2. Every Revolution needs its Manifesto

Long before the Industrialization, humanity **told stories of epochal change** – in the form of myth.

The Greeks peopled their world with **gods and crafted heroes** to explain natural forces and human aspirations.

It is striking how many of these ancient myths involve the creation of powerful entities and even artificial life.

The people who forged those myths were, in a sense, the **startup founders of the imagination**.

2.A. Talos

One of the most fascinating examples is the myth of **Talos**.

Talos was a bronze giant, a kind of animated machine in human form, forged by the divine craftsman Hephaestus.

According to the legend, Zeus charged Talos with defending the island of Crete.

The bronze guardian would circle the island's shore three times a day, hurling massive boulders to sink any hostile ship that approached.

In the Argonautica (3rd century BCE), when Jason and the Argonauts attempt to land in Crete, Talos stands in their way, an unyielding metal sentinel.

Eventually, the sorceress Medea helps defeat Talos by tricking him into removing the plug in his vein of ichor (the life-fluid of the gods), effectively "*bleeding*" the automaton to death.

But the key point is this: the Greeks **imagined an automaton with agency**, a self-moving machine that could act with purpose.

Why delve into these ancient stories? Because they remind us that what we are attempting now with AI agents taps into one of the oldest currents in human culture.

The **myth-makers** of antiquity were early visionaries of the human condition. In a very real sense, **founders who build AI "*entities*" are today's myth-makers**.

A startup founder who pursues an *entity-based* vision is stepping into a mythic role: the creator of something **alive with purpose**.

This has profound implications for how people adopt and relate to technology.

2.B. Founders as Modern Myth-Makers

Throughout history, those at the forefront of change have often had to **craft a new narrative** to accompany their innovations.

The great industrial pioneers were not just engineers; they were visionaries selling society a new way of life (think of Henry Ford championing not just the Model T but the idea of *everyday people driving* and the world built for cars).

Today's AI startup founders find themselves in a similar dual role: they must be both **pioneering builders** and **narrative creators**.

On one hand, they are like the industrial titans, marshaling a new technology to overthrow established ways of doing business. On the other, they are like the storytellers of old because the entities they create (AI agents) invite fundamentally new metaphors and relationships.

We already see hints of this in how the tech community talks about advanced AI.

Some technologists speak openly of *"creating god"* when referring to superintelligent AI, suggesting the birth of something akin to a deity in digital form.

Others in Silicon Valley half-jokingly (or not) refer to powerful AI models as "***golems***," recalling the Jewish folklore of the Golem – a clay giant animated through mystic words, created to serve but often prone to running amok.

These are explicit mythic analogies, painting AI as something more than a mere tool: **a creation with its own agency and power**.

Will these digital entities be benevolent helpers, like helpful household gods, improving human life? Or could they become capricious and dangerous, like uncontrolled sorcerous creations?

Crafting a *manifesto* for AI-native startups, then, is not hyperbole – it is a necessity. Every revolution needs its manifesto.

3. A Technological Revolution

Carlota Perez's influential theory of techno-economic paradigms reveals that **technological revolutions** follow recurring socio-economic cycles.

Over the past 250 years, we've seen multiple surges – from steam power and electricity to mass production and the information age – each radically reshaping society.

According to Perez, each revolution begins with a burst of new technology confronting an old order: "*The process of installation of each new techno-economic paradigm in society begins with a **battle against the power of the old**, which is ingrained in the established production structure and embedded in the socio-cultural environment*."

In other words, the prevailing institutions and mindset fight to retain dominance while a nascent technology promises a better way.

Perez describes how these surges unfold in two distinct periods:

First comes the **Installation Period**, an era of turbulent innovation when the new technology is "*a small fact and a big promise*" struggling to prove itself.

This period often features exuberant investment and speculative frenzy as society experiments with the new paradigm. (We saw this in the dot-com boom of the late 1990s and are arguably witnessing it again in today's feverish investments in AI.)

Eventually, after a turning point (often marked by a financial crash or social pushback), the second phase – the **Deployment Period** – begins.

This is when the new paradigm, having **overcome the resistance of the old**, becomes the new normal and drives a golden age of widespread growth.

Viewed through Perez's lens, our present moment with AI has all the hallmarks of a paradigm in the installation phase – "*a small fact and a big promise*" poised to bulldoze the old ways.

The **old paradigm** in business software has been the Software-as-a-Service (SaaS) model, which has dominated the last two decades.

AI-native startups are the insurgents, introducing a fundamentally different way of creating value.

As with past revolutions, there is friction: legacy companies and skeptical observers resist, while innovators press forward.

History teaches that those who recognize the new paradigm's logic and boldly champion it are the ones who shape the next era.

Today's founders must be willing to break with the comfortable SaaS playbook. In doing so, they aren't leaping into the unknown without guidance – they are following a historical pattern of revolution and renewal.

3.A. A New Paradigm of Value

To understand the scale of what is happening now, consider the **Industrial Revolution** of the 18th and 19th centuries – perhaps the closest analog in terms of economic upheaval.

The Industrial Revolution "*transformed economies that had been based on agriculture and handicrafts into economies based on large-scale industry, mechanized manufacturing, and the factory system*."

New machines and power sources (water wheels, steam engines, and later electricity) changed not only how goods were produced but also how people lived and worked.

Entire ways of life were upended as **handicraft and cottage industries gave way to factories**, and rural populations migrated to factory towns.

This process "*introduced novel ways of working and living and fundamentally **transformed society***."

The Industrial Age was driven by a generation of founders and inventors who **abandoned the old models** of production.

Instead of one craftsman painstakingly producing one item at a time, early industrialists like Richard Arkwright built water-powered textile mills that could spin thread continuously at a scale previously unimaginable.

In 1771, Arkwright's mill in Cromford – often cited as a start of the Industrial Revolution – was a "*small fact and a big promise*," much like today's AI prototypes.

By leveraging mechanization, Arkwright and others unlocked orders-of-magnitude improvements in productivity.

They were the **startup founders of their era**, proving a new paradigm under conditions of fierce skepticism.

We forget today that many in the late 1700s and early 1800s were deeply suspicious of mechanized production.

Skilled textile artisans (and the famous Luddites) feared job loss and, in some cases, **literally smashed the new machines** in protest.

Yet the **economic logic** of the factory system was undeniable: those who adopted power looms and steam engines could create far more value, far more cheaply, than those clinging to hand looms.

In Perez's terms, the "*battle against the power of the old*" had to be fought – and it was, with factories ultimately vanquishing the guild workshops across Britain and beyond.

By the late 19th century, the new paradigm had entered its deployment phase.

The fruits of industry spread to wider society: **mass-produced goods**, rising standards of living, and an expanded middle class as industrial capitalism matured.

Looking back, it's clear that **the founders who embraced mechanization and mass production became the leaders of the new epoch** – Carnegie in steel, Rockefeller in oil, Edison in electrical technology, Ford in automobiles.

Those who refused to adapt were left behind with businesses that soon looked quaint or obsolete.

3.B. A Striking Parallel

Today's AI revolution carries a striking parallel.

If traditional SaaS businesses are the craftsman shops of the digital age, then **AI-centric entities are the new factories** – poised to outperform by orders of magnitude.

Just as machines took over and scaled up the manual labor of the 19th century, AI agents can take over and scale up the cognitive labor and services of the 21st century.

A founder clinging to a conventional SaaS product (where a user logs in and manually utilizes software tools) may soon find themselves akin to a master weaver in 1800 watching a power loom eclipse his art.

The new paradigm does not abolish the need for human ingenuity or work – it amplifies it through a different vehicle.

In the Industrial Age, human muscle was leveraged through mechanized power; in the AI Age, human intelligence and decision-making are being leveraged through autonomous algorithms.

The lesson from the Industrial Revolution is clear: *to lead the next age and embrace the new means of production*.

For modern startups, that means AI agents and entity-based services rather than just human-mediated SaaS platforms.

4. From SaaS to Digital Entities

What does it concretely mean to *abandon traditional SaaS models* in favor of **entity-based, agent-led startup building**?

We've discussed the historical and philosophical backdrop – now let's translate it into the practical shift underway in technology and business.

In simple terms, the SaaS model can be described as *software as a tool*: a company provides a software platform (usually via the cloud), and the customer uses this tool to do something (manage finances, collaborate on projects, etc.).

The **emerging agent-led model** flips this dynamic. Instead of software being a passive tool, the software is an **active agent** that *does the work for the customer*.

In the old SaaS world, the vendor provided the hammer, but the customer swung it; in the new world, the vendor provides an autonomous carpenter.

For example, under the SaaS model, a business might buy a subscription to accounting software (and then the business's employees still have to do the accounting using that tool).

Under the agent model, that business would instead subscribe to an **AI accountant** – a cloud-based AI service that *handles the accounting tasks end-to-end*.

The AI agent delivers a result, not just a piece of software. *This shift is enormous.*

It mirrors the difference between selling someone a car versus offering them a chauffeur.

For the customer, the value proposition is higher-level (outcomes, not process). For the provider, the responsibility is greater, but so is the control over delivering consistent quality via automation.

We can see early signs of this *agentic paradigm* in products like AI-powered customer service bots that handle entire support tickets or GitHub's Copilot, which not only highlights code syntax but also *writes code* on behalf of the developer.

These are narrow examples, but they point to a future where entire **business functions are managed autonomously by AI**.

This isn't futuristic speculation – it's the trajectory we are already on.

Today, there are startups working on AI agents that can autonomously perform sales calls, schedule and optimize shipping logistics, or even conduct scientific research by reading papers and formulating hypotheses.

For startup founders, the message is clear: the **opportunity space of entity-based, agent-led startups is immense**.

The transition from SaaS to autonomous services represents a $4.6 *trillion* opportunity because the services sector of the economy (everything from legal work to healthcare to finance) is far **larger than the software sector ever was**.

In essence, *AI agents allow tech companies to eat the services world*, not just the software world. We can refer to these new startups as “AI-native” or “agentive” companies.

They will look different from yesterday’s startups. They might have far fewer employees relative to their user base or revenue.

Their “products” might feel more like **collaborators** or **employees** than apps – imagine hiring an AI agent service the way you’d hire a firm rather than installing an application.

Their success metrics will center on real-world outcomes achieved (*did the AI successfully increase a client’s sales? cure a patient’s illness? design a profitable chip?*) rather than just software usage metrics.

This is a qualitatively different way of thinking about value creation. It is *entity-centric* rather than product-centric: the AI agent is an **entity that acts** (almost like a virtual team member), not just a static product that is used.

Abandoning the SaaS mindset is hard, especially for those of us who grew up during its dominance.

SaaS taught a generation of founders to focus on scalable distribution, user engagement, and iterative development of software features – all important things, but mainly within the paradigm of user-does-X-with-our-tool.

The agent paradigm asks founders to focus on a different question: *what complete problem can we solve if our software is smart enough to handle it entirely?* It shifts the focus from user workflows to user goals.

For example, a SaaS project management tool focuses on providing features for humans to manage tasks; an AI project manager would instead focus on *actually managing the project* – assigning tasks, following up, reprioritizing, and only involving humans where necessary.

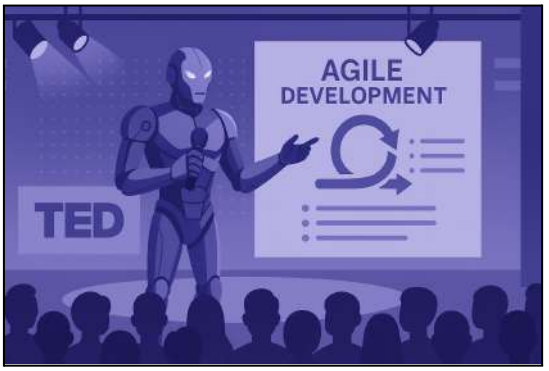

The ultimate vision is that **AI entities become autonomous companies or workers** that interface with human organizations.

We see glimmers of this in systems like AutoGPT, where an AI can autonomously break down goals into sub-tasks and execute them in loops.

While still rudimentary, these point toward AI agents that one day could effectively run an entire small business or a subdivision of a larger enterprise as a tireless, ultra-efficient executive.

The economic forces at play (**speed, scale, cost reduction**) strongly favor those who fully utilize AI agents.

A company offering an AI-based autonomous service can outcompete a SaaS company that requires manual effort, just as a factory could outproduce a craftsman.

When "*bots become brains*" and are entrusted with higher-order tasks, the efficiency gains become exponential.

4.A. User Interface vs. AI Agent Functions

Traditional SaaS applications offer a collection of features behind a user interface, each of which requires human input.

In contrast, an AI-first agent functions more like an **intelligent digital employee**, capable of understanding context, learning from data, and adapting its behavior in real time rather than executing only pre-programmed commands.

In other words, **we're shifting from software as a tool to software as a collaborator**.

The future of software is *conversational*; you tell the AI agent what outcome you want, and it handles the execution.

Microsoft's CEO Satya Nadella remarked that most enterprise apps are essentially "*just some logic wrapping a database*" and warned that *"they will all collapse in the Agent era*."

Salesforce's Marc Benioff has even called autonomous agents the "*third wave of AI*," envisioning a future where businesses deploy their *first AI employee* alongside their human workforce.

In Benioff's words, these agents won't just augment employees; they can replace low-value busywork, operating with *"zero hold times…expert all the time"* availability.

The ultimate goal is clear: *better business results* through always-on, autonomous helpers.

Even Sam Altman, CEO of OpenAI, has mused that *"agents are the thing everyone is talking about for a good reason"* – predicting that AI systems will soon handle complex tasks we'd "*give to a very smart human*" and *"really transform things"* as early as **2025**.

4.B. Meet Devin and Lindy

Meanwhile, a new crop of startups has been embracing an "AI-first" philosophy from the outset, **treating AI agents as the core product** rather than an add-on.

For example, Cognition Labs created Devin, touted as *"the first AI software engineer,"* a tireless coding agent that can plan, code, and debug projects autonomously.

Devin is presented not as a mere coding tool but as a *skilled teammate* who works alongside human developers or tackles tasks independently.

On another front, Lindy offers what founder Flo Crivello calls *"AI employees."*

Lindy offers a no-code platform for creating personalized agents that handle executive assistant tasks, such as drafting emails, scheduling meetings, and providing customer support.

In essence, Lindy lets you hire your first digital team member without writing code.

Similarly, platforms like Cognosys let users hand off entire objectives to an agent. The Cognosys philosophy urges, *"Don't just ask questions; give objectives.* "

You can ask it to say, *"Conduct an in-depth market analysis,"* and the agent will break this complex goal into subtasks and complete them autonomously.

These pioneers share a common vision: software shouldn't just sit there waiting for user input – **it should take the initiative and do the work for you**.

5. From Products to Agents

To understand this transformation, let's compare the traditional SaaS model ("Old") to the AI-first agent model ("New").

The differences go beyond technology — they represent a philosophical shift in *what* and *how* we build.

5.A. User Experience – From Tool to Teammate

OLD: Users interact with software via GUIs – clicking buttons, filling forms, and navigating dashboards. The software *helps,* but the human is still firmly in control of each step.

NEW: Users interact with an AI agent via natural language as if delegating to a colleague.

"Instead of users clicking through menus and forms, we're moving toward a world where you tell an AI agent what outcome you need, and it handles the execution."

The agent becomes a **collaborator**, not just a tool, handling the busywork and even making suggestions.

In practical terms, **every user becomes a manager**, assigning tasks to an AI subordinate rather than doing those tasks themselves.

5.B. Problem Solving – From Features to Outcomes

OLD: Startups built **features**. Each feature addresses a subset of the user's problem, and users string these features together to get a result.

The product is essentially a toolkit; value is delivered when the *user* uses the tool correctly.

NEW: Startups build **agents that deliver outcomes**. You don't ship a laundry list of features; you develop an AI agent capable of solving the whole problem end-to-end.

For example, a traditional travel app gives you forms to search flights and hotels, whereas an agent-based service could plan an entire trip with a single request.

Internally, this represents a shift from writing thousands of lines of code orchestrating every step to defining high-level goals and letting the agent figure out the steps.

Thus, founders of agentic startups focus on training the AI, giving it capabilities and context, rather than manually coding every pathway. The product is **outcome-oriented** by design.

5.C. Adaptability – From Static to Evolving Systems

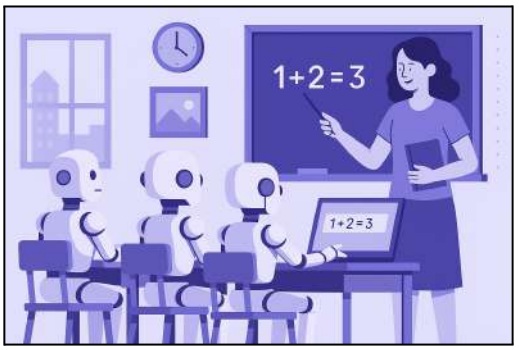

OLD: SaaS products improve through version updates and new feature releases.

Until the dev team updates the code, the software's behavior is static – it does what it was programmed to do, nothing more.

If a new use case or trend emerges, humans must explicitly redesign the product to accommodate it.

NEW: AI agents are **living systems**. They learn from each interaction and can adapt quickly. *"AI agents… continuously adapt, learn, and operate without human intervention,"* as one founder explains.

Rather than being disposable tools that become outdated, agents can **self-improve with new data and experiences**.

Mark Levitt of Trixta argues that we should view AI agents as *"long-living digital entities"* that **evolve with business needs**, *"not just software components."*

The traditional idea of version releases gives way to continuous learning in production.

In the new playbook, the startup team's role is to **guide the evolution** —providing feedback, setting goals and constraints, and occasionally upgrading the model— rather than shipping static code updates.

5.D. Workforce and Scaling – From Hiring Humans to Deploying AI Workers

OLD: As a company grows, serving more customers or tackling bigger problems usually means **hiring more people —** more sales reps, support staff, and operations personnel.

Likewise, scaling the product often means adding more servers or more code, but always with significant human oversight. A startup's headcount grew in tandem with its ambitions.

NEW: An AI-first startup scales differently: it can deploy **hundreds of agent instances** to handle work in parallel, at negligible marginal cost, in a way that humans never could.

One VC noted that *"AI agents will significantly expand software's total addressable market by tapping into the largest budget of all: payroll"* – because companies can offload work to AI and potentially reduce human hiring.

In an AI-native startup, a **small human team can orchestrate a massive AI workforce**.

Researchers predict that an **AI-first company will have a high ratio of AI agents to human employees**, and those humans will act as skilled “managers” supervising teams of AI agents.

The leverage is unprecedented: when demand spikes, you don’t hire and train dozens of new staff; you spin up more agent capacity.

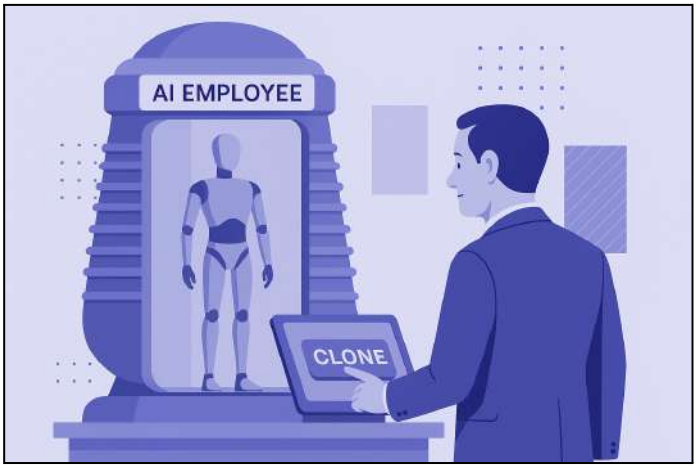

When each agent can work 24/7, the **scale is no longer linear to headcount**.

This also changes the nature of work within the startup — engineers and operators focus on training, validating, and steering AI agents (much like managers and coaches) rather than manually performing the front-line tasks.

In the old model, software was a tool for a human workforce; in the new model, the software *is* the workforce.

5.E. Value Proposition – From Efficiency to Autonomy

OLD: The selling point of SaaS was often efficiency and convenience. “*We save you time doing X*” or “*We help you do Y better*.” But fundamentally, the **user still had to do the thing**, just with improved software.

The limits of value were the limits of human time and attention since the user remained the primary actor.

NEW: The promise of agentic startups is **outsourcing the entire task**.

The AI agent doesn’t just make you faster; it *frees you entirely*. This is a categorical leap in value.

As Ben Thompson of Stratechery succinctly put it, *“Agents aren’t copilots; they are **replacements**. They do work in place of humans… and they have all of the advantages of software: always available, and scalable up-and-down with demand.”*

An agent can handle that 3 am support request or process a million documents without breaking a sweat.

For customers, this shifts software from a tool expense to something closer to a labor expense (hence the talk of tapping into payroll budgets).

For businesses, it means **outcome-based results**: you didn't just buy a tool; you hired a machine that gets the job done.

For example, rather than selling a CRM system and expecting a sales team to use it, an agent company might sell an AI sales associate that, *itself,* prospects leads, sends emails, and schedules meetings for you.

Outcome, not effort, becomes the unit of value.

6. Principles for an AI-First Agentic Startup

Adopting this agent-first model isn't just a technical shift – it's a philosophical and strategic one.

Here we outline a new set of **first principles** for founders building in an AI-first world:

#1 - Build *Agents*, Not Apps

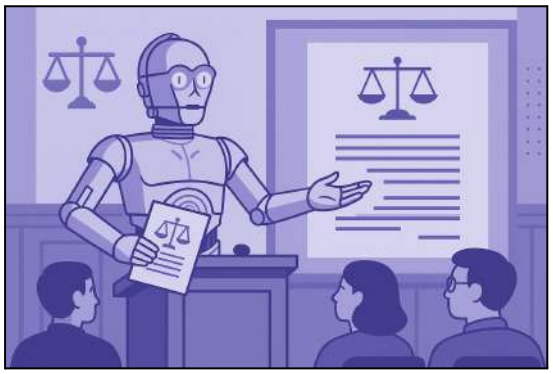

Envision your startup as creating a **digital team member**, not a static application.

This means thinking of your software as an *autonomous service provider* from the outset.

Ask yourself: *if my product were a person I hired to solve this problem, what would their responsibilities be?* Design your AI agent to fulfill those responsibilities.

For instance, if you're in legal tech, don't just build a document management app – build an AI paralegal that *reads, analyzes, and drafts documents* on its own.

This mindset shift is crucial. It moves you away from feature checklists and towards a holistic approach where the agent's **purpose is to achieve an outcome** for the user.

Keep in mind that an *agent is "alive" in its domain*: it should ideally run continuously, handle new inputs, and maintain context over time.

By focusing on an agent's role (e.g., "AI financial advisor" or "AI project manager"), you ground your product vision in **delivering a service** rather than delivering software.

#2 - Outcome-Obsessed Development

Make **user outcomes** your north star in product development.

In an AI-first startup, you don't ask, "*What features do we build next?*" but rather, "*What broader goal can our agent take on next?*"

This principle pushes you to measure success by what your AI accomplishes for users (tasks completed, problems solved, value created), not just by usage stats of features.

It also means your design process often starts with a conversation: *What would a user ask the agent to do?* And *How will we know the agent succeeded?* From there, you work

backward to equip the AI with the knowledge, tools, and decision-making ability to deliver that result.

If your agent isn't yet capable enough, you iterate by improving its training or giving it new integration capabilities – not by adding more buttons for the user to click.

This outcome focus aligns your team with the user's real needs and keeps the technology invisible.

Remember the mantra: **users describe the *what*, and the agent figures out the *how***.

Your job as founder is to relentlessly improve the agent's ability to figure out the *how*.

#3 - Embrace Conversational UX (and Beyond)

The most natural interface for instructing an agent is human language, so design your product around a **conversational interface**. This could be a chat box, voice command, or a simple form.

This principle, however, goes beyond just chat UI. It's about adopting the philosophy that *"the UI will be reimagined from showing users how to do something to letting users say what they want done."*

Consider how your agent will report progress, ask for clarification, or present results.

Often, a conversational approach works here, too – e.g., Devin (the AI engineer) **"reports on its progress in real time [and] accepts feedback"** as it works. That kind of interactive loop builds trust and keeps the user informed at a high level without micromanaging them.

Design dialogues, not dashboards. Let the agent handle the intricacies and surface just enough info for the user to guide it.

In short, make interacting with your product feel like working with a competent colleague – one who speaks plain English (or whatever language your user prefers) and can translate your requests into action.

#4 - Iterate with Your Agent as a Team Member

In an agentic startup, **product development is a partnership between humans and AI**.

You will be *"managing"* your agents just like you manage employees – giving them feedback, new training, and new responsibilities over time.

Adopt an iterative process where you regularly test your agent on real-world tasks, see where it fails or succeeds, and refine it.

This may involve fine-tuning models, expanding its toolset (such as integrating new APIs), or adjusting its “prompting” approach.

It’s akin to coaching an employee: identify weaknesses and help it improve incrementally.

Encourage a culture where your team is not just shipping code but also **shaping behavior** – analyzing transcripts of the agent’s decisions, understanding its errors (e.g., where it might hallucinate or get stuck), and course-correcting.

The **agent’s performance IS the product**, so pay close attention to how it behaves in the wild.

Moreover, plan for your agent to have **memory and learning**: it should retain important knowledge from past interactions (within ethical and privacy bounds) to avoid repeating mistakes and to personalize its service.

By viewing AI as a growing learning entity, you’ll naturally extend traditional QA into ongoing training.

Some advanced teams even start giving their agents a form of “on-the-job training” – seeding them with initial knowledge and then letting them learn from each task they attempt, much like a junior employee learning on the job.

The takeaway: Treat continuous improvement of the AI’s decisions as a core part of your development cycle. This is how your product stays ahead of both user expectations and competing offerings.

#5 - Redefine Team and Culture Around AI-First Values

Founders must establish a philosophy that humans and AI agents work in tandem – and sometimes, the AI takes the lead.

This starts with your internal team culture.

Encourage your team to offload their own routine tasks to agents where possible so they intimately understand the experience of using one.

Hire people who are excited to *become “AI orchestrators.”*

Your team’s skills might skew more towards prompt engineering, data analysis, and domain expertise for model training rather than classical full-stack development or manual ops.

As your company grows, **maintain a high AI-to-human ratio** as a strategic goal.

This doesn't mean replacing people for the sake of it but rather leveraging AI to keep your organization lean and focused on creative work.

Every human you hire should multiply in impact by managing or enhancing swarms of AI workers.

Make it a first principle that **whenever a new problem or workflow arises, you ask, "*Can an agent handle this?*" before you default to hiring or assigning a human**.

This mindset will push you to innovate with your AI capabilities continuously. It will also help you scale faster without bloating costs.

Culturally, celebrate when your agents achieve milestones – treat it with the same excitement as a successful project done by a team member. This reinforces that your *AI is part of the team*.

Finally, instill **ethical guidelines** in your culture from day one: with great power (autonomous agents) comes great responsibility.

Train your team to monitor for bias, errors, or unintended consequences in the agents' actions and to take responsibility for them.

An AI-first company must be proactive about AI alignment and safety, ensuring that as agents take on more autonomy, they always remain aligned with user intent and company values.

#6 - Flip the Business Model – Deliver Transformation, Not Just Service

Agentic startups have the opportunity to redraw business models.

If your AI agents truly deliver outcomes (and potentially replace human effort), consider **pricing on delivered value** or outcome-based models.

This could mean usage-based pricing that correlates with tasks completed or success-based fees.

The philosophy here is to **align your revenue with the actual results the agent provides**, which builds trust with customers.

Strategically, focus on problems where an agent can provide a **10x solution** over the status quo.

Investors are already seeing that an AI agent startup can tackle markets previously out of reach for software by directly taking on the costly labor.

By automating what was once only possible with large teams of people, you can create and dominate new market niches.

In your strategy, articulate the vision not as "*better software for X*" but as a "*fundamentally new way to solve X.*"

For example, instead of a traditional SaaS pitch like "*a platform to manage your social media marketing*," an agentic startup would say "*an AI marketer that runs your social media accounts for you*." The latter promises a *replacement of hassle with hands-free results*.

Structure your go-to-market messaging around these transformative promises – and ensure your pricing captures some of the massive savings or gains you're providing.

One VC described this shift succinctly: *"Labor and software are merging."*

When you build an agent, you're effectively selling **labor-as-a-service** (but with machine scalability).

7. Legal and Ethical Foundation

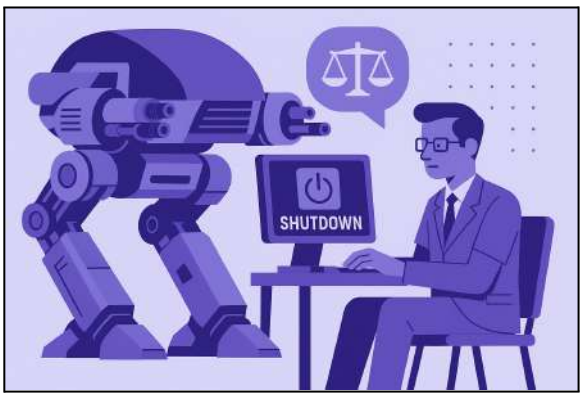

The mythic automaton Talos had a single vein stopped by a bronze plug.

Pull it, and the titan collapsed.

Twenty-first-century agents are more sophisticated, but their stoppers are just as real: a lattice of statutes, oversight norms, and accountability frameworks that collectively define what a "*responsible*" AI entity must be.

Ignore this lattice, and you risk crafting a marvel that no regulator will license, no enterprise will touch, and no market will absorb.

Today's founders don't just engineer capabilities — they encode values, constraints, and obligations into intelligent systems.

This is no longer optional. It's the new constitutional layer of agentic enterprise.

7.A. Privacy, Consent, and the Human Right to Review

Start with the General Data Protection Regulation (GDPR). Its influence stretches far beyond Europe — any AI system handling EU residents' data must comply.

Article 22 enshrines the right to opt out of automated decision-making and demands a human override.

For agentic startups, this means offering a clear path to human review in domains such as hiring, finance, or healthcare — especially when outcomes have a significant impact on individuals.

Many early teams overlook that **model weights can qualify as personal data** — especially when fine-tuned on user traces.

This pulls you into GDPR's Article 25 "*privacy by design*" doctrine, mandating explicit safeguards for data handling, retention, and security.

Failing to comply can trigger audits, bans, or fines that scale up to 4% of global revenue.

7.B. Risk-Tiered Governance: The AI Act

Europe's AI Act (Regulation 2024/1689) goes further, creating the world's first vertically tiered framework for AI systems.

It places **agents that hire, grade, diagnose, surveil, or allocate credit** into a "*high-risk*" category.

The law is now enforceable across all EU member states, and its scope is wide:

- Article 6 defines how AI systems are classified;
- Article 13 mandates transparency to users;
- Article 14 requires meaningful human oversight;
- Annex III specifically lists agent-driven systems in employment, education, healthcare, and essential services as high-risk by design. If your agent does more than assist —if it decides— you're likely in scope.

To operate legally in the EU, these agents must:

1. Maintain logs of every consequential decision.
2. Disclose their training dataset pedigree and model provenance.
3. Be designed with a "pause switch" controlled by a named human supervisor.
4. Pass conformity assessments and register in a public EU database.

Penalties for noncompliance reach €35 million or 7% of global turnover — whichever is higher.

If your agent "*acts*," not just advises, you're likely in-scope. Comply early. The burden only grows later.

7.C. Declare the Machine: Interface Disclosure Laws

In California, SB 1001 remains the digital world's most surgical bot law.

If your agent interacts for commercial or political purposes, **it must disclose its non-human identity clearly and conspicuously**.

This means no buried footnotes and no ambiguity. A missing declaration isn't just shady UX — it's a violation with real liability.

California's SB 1001, passed in 2019, is the definitive precedent: any bot used for commercial or political communication must clearly disclose its artificial nature when interacting with the public. The FTC has signaled support for similar enforcement nationally. Disclosure isn't just ethical — it's legally expected in the world's fifth-largest economy.

This principle —truth in the interface— is spreading fast.

7.D. Voluntary Codes with Real Teeth

The OECD AI Principles, refreshed in 2024, are now embedded in global procurement standards — from Canadian defense bids to Japanese fintech.

The NIST AI Risk Management Framework has become table stakes for U.S. federal contracts: no documented risk controls, no deal.

The ISO/IEC 42001:2023, the world's first AI management system standard, is gaining traction with auditors and insurers alike.

Founders who internalize these norms early **gain not just compliance but a competitive edge**: faster onboarding in regulated sectors, fewer late-stage rewrites, and the credibility to lead in trust-sensitive domains.

7.E. Architecting Ethical Agents by Default

Ethics cannot be a retrofit. Bake it in from line one:

1. **Declare the machine.** Whether the user is a consumer in Fresno or a regulator in Brussels, the interface must signal: "AI at work."
2. **Keep a hand on the plug.** Every agent must ship with an auditable off-switch — tied to a specific human who has both the authority and duty to intervene.
3. **Record the lineage.** From dataset to decision, maintain a ledger of how the model was built, trained, and deployed. Store it longer than your lawyers' longest statute of limitations.
4. **Support human override.** Design AI not as an autocrat but as a deputy — fallible, reviewable, interruptible.
5. **Audit continuously.** What an agent does in the wild is the real product. Establish feedback loops to analyze edge cases, flag anomalies, and course-correct proactively.

8. Join the Agentic Revolution

We stand at the dawn of a new era in tech – one that will fundamentally reshape how we build products, how we run companies, and how we work with machines.

This manifesto is a call to all ambitious founders and builders: **embrace the AI-first, agent-based paradigm** with conviction.

The old playbooks of SaaS, while successful in their time, are reaching their limits.

In a world where AI agents can learn and do so much, clinging to the old model is like insisting on horseback travel while the first automobiles roll out.

As one tech leader recently put it, *if you're building SaaS today, you have two choices: become an AI-first company that integrates seamlessly with AI agents or stay UI-bound and watch the world move on without you.*

The time to choose is now.

Adopting this new paradigm is not easy – paradigm shifts never are.

It demands courage to rethink product assumptions, invest in still-evolving AI technology, and educate the market on a new way of doing things. But this is exactly where startups thrive: at the edge of what's next, pulling the future forward.

At **InTheValley**, these principles aren't theory. They're actively shaping how we build, invest, and partner today.

The reward is not just commercial success but a chance to **fundamentally change the game**.

Agentic startups aren't just iterating on past ideas; they are pioneering a model where **every company can have an army of tireless digital workers**, where consumers can have personal AIs handling their needs, and where innovation is limited only by our imagination, not by headcount.

Imagine a world five years from now: entrepreneurs are launching companies whose first 100 "*employees*" are AI agents executing the business model from day one.

Product managers design high-level objectives and let agents A/B test and optimize solutions in hours.

Customers saying, *"I have an AI for that"* for everything from personal shopping to home maintenance scheduling.

This is not science fiction – **this is the trajectory we are already on.**

This manifesto is an invitation. Let's leave behind the old constraints and fully unleash what today's AI technology can do.

Let's write a new narrative where startups aren't apps or websites but **entities** – goal-driven, self-improving, and ever-ready to serve each user in a personalized way.

If the last era was about mobile-first or cloud-first, this era is about being **AI-first** and *agent-first*.

It's about daring to build a company where **the core team might be you, a co-founder, and an ensemble of AI agents doing the work of 50+ traditional employees**.

It's about trusting that with careful training, alignment, and oversight, these agents will not only match human performance but exceed it in speed, consistency, and scale, all while maintaining the creativity, empathy, and strategic insight that only humans (for now) possess.

We invite you to join this revolution. Build something audacious – an agent that pushes the boundaries of what we thought software could do.

In doing so, you'll not only create a category-defining company; you'll also push the whole market forward, inspiring others to adopt this new paradigm. This is how a manifesto becomes a movement.

9. Ready to Become an Agentic Founder?

We’re entering the Agentic Era — not tomorrow, not next year, but right now.

The startups of tomorrow won't be built around teams grinding away on manual tasks. They'll run autonomously, powered by intelligent agents, safeguarded by *compliance layers*, and driven by frameworks designed to scale logic, not labor.

The question isn't whether this will happen — it's already underway.

The real question is: **Will you watch it happen, or will you lead it?**

*The agentic revolution is underway — and we’re building it every day at **InTheValley**.*
Join the conversation: DM “Agentic” to our Founder on LinkedIn.