

## Chapter 1

# Escaping the Iteration Trap

COPYRIGHTED MATERIAL




Escaping the  
Iteration Trap



**“It’s not a faster  
machine, a better  
database ... AI is a  
fundamental game  
changer: a new kind  
of intelligence.”**



# Escaping the Iteration Trap



**A**ccording to recent research by IBM and Oracle, 78% of C-suite executives recognize that achieving maximum benefit from agentic AI requires a new operating model. Yet more than three-quarters also admit their AI investment focuses on improving existing processes rather than developing net-new capabilities [29]. In short, most leaders know they need or aspire to achieve business reinvention, but they're largely funding AI iteration and optimization.

That's exactly why AI initially didn't deliver the ROI that so many predicted. The pervasive trend in the years following the onset of generative AI, AI agents, and agentic AI was to launch siloed pilots making incremental updates, extending the familiar to create a new era of "AI business as usual."

If there's one truth business history teaches us, it's this: every disruptive technology eventually forces companies to choose between preserving the past and inventing the future. The steam engine, electricity, the internet, mobile and social media...all of them offered leaders a choice: cling to old models or design entirely new ones. And each time, those who chose preservation ultimately ceded ground to those who dared to reimagine.

For centuries, our progress narrative has been one of linear evolution. The Industrial Revolution taught us that automation and mechanization could scale productivity far beyond human limits. That's been our working model ever since. In the so-called Digital Revolution, each wave brought extraordinary gains, but also a subtle

trap: the belief that the next era was just a more advanced version of the last.

That's why so many companies misunderstood "digital transformation" and are misunderstanding the AI revolution. They don't appreciate how different AI is. It's not a faster machine, a better database, or a more connected network. It's a fundamental game changer because it's a new kind of intelligence.

The intelligence revolution challenges every assumption about how we think, how we work, how we measure success, and even what it means to be a company at all.

Even though digital transformation was widely touted as the great leap forward, once the movement started, most companies were not motivated to become digital, let alone digital-first. There was very little true *transformation* in digital transformation. Most companies took the analog playbook and put it online. They built systems of record around silos: HR, sales, marketing, operations. Then they optimized each in isolation, because that's just how things were done. The result? Fragmented data, disconnected workflows, and marginal gains in productivity. Companies failed to design for a connected enterprise.

Even in creating cloud-based systems, organizations didn't truly reimagine HR, sales, marketing, or finance, and they didn't connect those systems and data. They adopted Workday for HR, Salesforce for sales, Adobe or HubSpot for marketing, and SAP for enterprise resource planning...tools that optimized rather than revolutionized how a department operated. That's why, despite trillions of dollars invested, most businesses never unlocked transformative performance gains.

MIT Sloan School of Management researcher George Westerman famously said [26] that when digital transformation is done right, it's like "a caterpillar turning into a butterfly," but when it's done wrong, "all you have is a really fast caterpillar."

The same is true for AI-driven business transformation.



## Iteration Versus Innovation

With the digital transformation, we mistook iteration for innovation.

What's the difference?

It's easy to conflate the two, especially when we're using new, innovative technologies. But the differences are important. Iteration is doing what we did yesterday better tomorrow. We become faster, more efficient, more scalable, all while reducing costs and resources. Innovation is doing something we didn't do yesterday, to create new value. Both are necessary.

Then there is transformation: a fundamental break from a previous state.

AI's ability to exceed human capacity in many ways provides an unparalleled opportunity for companies to actually transform. AI-driven business transformation can help fix the failures and missed opportunities of digital transformation.

## The Constraint of Cost Cutting

In the early days of digital (so-called) transformation, companies pressured technology leaders to use digital to not just modernize, but to cut costs. "Your goal is to take out 10% each year!" was a common refrain, and one that CIOs around the world would fret over throughout the year. It was the bane of their careers. Technology was viewed as a cost center, and as such, every investment in tech was viewed through a different lens than investments in parts of the organization that were viewed as growth drivers. Business executives looked to maximize cost cutting on the one side, and on the other side to grow the business through targeted investments based on historical performance and use cases.

Over time, this approach hinders a company's competitiveness. If yesterday's work is automated to the point of total efficiency, and headcounts are cut to realize ROI, the company ends up with limited capacity to create new value.

One of the most instructive examples comes from General Motors in the 1980s. We know, it's a long time ago. But it's still a cautionary tale that every business leader should heed.

In an effort to stay ahead of Japanese competitors, GM's CEO Roger Smith launched a highly ambitious automation strategy known as the "factory of the future" [31]. The company reportedly made over \$45 billion in capital investment in the early 1980s, intending to drive efficiency and productivity through cutting-edge technology. But the results were disappointing. GM's global market share barely budged, rising only 1%. Meanwhile, critics inside the company noted that the same investment could have allowed GM to acquire major competitors like Toyota or Nissan, instantly increasing market share by a far greater margin [32].

So what happened?

Automation became the strategy; it became the end rather than the means. While the robotic factories delivered efficiency gains, they absorbed capital and attention, leaving little room for exploring new vehicle designs, customer experiences, or alternative business models. Innovation pipelines shrank. Resources were tied up in automation projects, with fewer dedicated to research and development, conceptual ventures, or emerging technologies like electric vehicles or autonomous systems. What's more, culture shifted. As efficiency became the dominant

metric, risk taking and experimentation, noted hallmarks of innovation, were sidelined.

Other car manufacturers watched, learned, and acted, and over time, competitors (both domestic and overseas) were able



**Iteration is  
doing what we  
did yesterday  
better tomorrow.  
Innovation is doing  
something we didn't  
do yesterday.**



to leap ahead because they retained the bandwidth to explore what came next.

The bottom line: without deliberate investment in innovation and forward-looking capabilities, organizations risk automating themselves into irrelevance.

This is what we call the Iteration Trap. With the advent of AI, companies are falling into the trap again. We're primarily seeing cost cutting scaled by automating work and reducing headcounts.

Infinite companies see cost cutting as the status quo. It's a given. Leaders of these companies are looking for much more, asking "How can AI challenge our assumptions and models and enable us to create new value for the market, employees, and stakeholders?"

## The Combined Forces of Inertia

In 2024, *The Economist* ran a series of articles that explored the reasons for companies reporting limited results from investments in AI. One headline read, "What Happened to the Artificial-Intelligence Revolution?" [33].

The magazine reported that only about 10% of firms were using AI meaningfully, despite executives publicly praising its potential, and that the boost to the stock performance of established companies expected to benefit most from AI was failing to materialize, signaling slow practical uptake. It laid out a set of potential reasons for the limited results, including:

- **Economic and organizational friction.** Most companies have siloed infrastructure and data, but integrated datasets and cloud systems to integrate workflows are required for optimizing with AI, let alone for innovating.
- **Public choice dynamics hampering progress.** Like many governments, most firms suffer from self-interested actors (especially middle managers) who hold authority and can

delay or sabotage adoption. This serves, albeit counterintuitively, as a form of protectionism, to keep power, and people in power, in place.

- **Resistance to expected job losses.** Historical and modern studies show workers and managers generally resist productivity-enhancing technologies when they fear job losses or disruptions to their roles. Middle managers often block AI adoption because automating lower-tier tasks threatens their own positions.
- **Legal and regulatory barriers.** In-house lawyers and compliance teams focus on liability, data privacy, and discrimination risks, creating significant slowdowns and a risk-averse culture. In reality, the lack of established legal precedents does add uncertainty and encourage caution. But establishing new governance protocols and effective guardrails is eminently doable. You just have to switch from a mindset of “We can’t” to “How might we...?”
- **HR-driven job protectionism.** HR departments are naturally worried about the employment impacts of AI, with good reason. This may drive them to erect roadblocks to certain AI projects in order to protect their roles and the roles of employees at large. They may also champion initiatives that prioritize iteration to give employees time to reskill and upskill. This is perfectly normal, but it impedes the process of creating new types of jobs. The irony is that HR should be leading in envisioning new jobs and reskilling employees to move into them.
- **Market forces eventually drive adoption, but slowly.** As with past technologies, from tractors to smartphones, competitive pressure will eventually push firms toward adoption. They will finally heed the old mantra of “adapt or die!” The challenge is whether adoption will be enough to allow them



to compete against AI natives and infinite companies whose mantra is “innovate or die!”

While these are all understandable reasons why companies fall into the Iteration Trap, they are not excuses. Leaders must first acknowledge that these forces are at play in their companies, then actively combat them.

## **Building New Tracks**

One day, the two of us were talking about horrible holiday experiences of traffic on highways and why so many of us in the US drive rather than taking a train. One of us had recently been stuck in highway traffic for over seven hours traveling the 300 miles from Santa Barbara to San Diego. The idea of taking the train for the next trip was appealing, except it turned out it would actually take even longer: over nine hours. The journey required multiple transfers, each involving a wait time, with zero high-speed trains available.

Why is the American passenger train system so far behind Europe and Asia in introducing high-speed trains? Part of the answer lies the same problem many executives run into when they try to “AI” their way out of legacy operations: you can optimize a system right up to the edges of what it was built to do. But you cannot optimize your way into a future your infrastructure was never designed to reach.

This brings to mind the saying often attributed to business professor and author Oren Harari, “The electric light did not come from the continuous improvement of candles.”

If your track system can’t accommodate high-speed trains, there is no point in investing in faster trains. You’ll never run them at full speed. The same is true for organizations incorporating AI. An optimized workflow will allow you to get maximum benefit from the technology and infrastructure on which it runs. But to run a bullet train to its full potential, you need to lay new track.

Which leads to the question, why hasn't new rail track been laid in the US while it has been in Europe and Asia? A key reason is that the use of the rails we have is profitable, but for freight companies rather than passenger service. Union Pacific explains: in the US, you will see a freight rail network that is "the envy of the world" [34]. American freight carriers, which are private companies, own the tracks, and they are hugely profitable, some with profit margins of 40% or more in recent years [35].

Similarly, many legacy companies are "the envy of the world" at moving today's work through yesterday's systems. They're not failing. They're succeeding at exactly what they were designed to do.

The success of the freight carriers has impeded US passenger railway operators like Amtrak from innovating. And this is what it feels like inside many legacy enterprises trying to modernize with AI. Your workflows run on "host tracks," systems of record you don't want to disrupt, data structures you inherited, and process logic that was optimized for control, not learning velocity. You can absolutely add AI and see meaningful gains. You can reduce friction, automate steps, generate insights, and compress cycle times. But you will eventually hit an efficiency ceiling because the track geometry, meaning the workflow itself, was laid for a different era.

For a while, the dashboards still glow green, the margins hold, and customers still show up. But one day, all that will stop. It probably won't happen abruptly, though sometimes it does. You might not notice the downturn immediately. What's more, once you do, while you might focus resources on turning things around, you might miss the root cause of why this is happening.

The simple, uncomfortable truth is that many companies don't see, or feel, the need to radically change until the market has already changed around them. They become finite companies not by choice, but by drift.

If you apply AI to an old process, you'll see progress, but not transformation. And meanwhile, AI-native companies are doing



what Europe and Asia did: building high-speed lanes from scratch.

Of course, the obvious objection is that you can't reinvent the train system while trains are still running on the tracks. It's the railway equivalent of the familiar adage, "you can't build the plane while flying it."

True.

Just as you can't rip up the freight lines that keep an economy running, you can't shut down the systems of record that run payroll, billing, claims, patient care, or global supply chains. But you can build parallel tracks.

That is exactly how an enterprise becomes infinite without breaking itself in the process.

## The Nature of True Reinvention

In the early 2000s and 2010s, a few companies refused to start with the old playbook. These were the true digital transformers: Amazon, Netflix, PayPal, Uber, Airbnb, Tesla. They began with a blank page, or should we say, a blank screen.

Amazon not only moved retail online, it redefined commerce as an entirely digital-native experience. The company's core business wasn't "online retail"; it was a logistics, data, and recommendation engine that happened to sell products, and then services. Uber was also essentially a logistics company, which happened to offer rides, and then delivery services. Netflix didn't just shift from mailing DVDs to streaming movies; it reinvented how the world consumes stories. The company developed into a global content platform powered by algorithms, personalization, and direct-to-consumer storytelling, coming to know viewers so well that it could also successfully develop original content.

All of these companies were built from the outset with digital-native thinking. They rejected inherited business logic and asked a different set of questions: What if the constraints everyone

else accepts simply don't exist? This mindshift was their true differentiator.

Another thing that set these companies apart was that they thought in platforms, not products. They built ecosystems that allowed innovation to compound: marketplaces, APIs, partner networks. They also designed from the customer backward, eradicating friction, removing trade-offs, and delivering experiences so intuitive they felt inevitable in hindsight. To achieve this, they rewired their cultures to move at the speed of change. Decision-making became data-driven and decentralized. Experimentation became the operating system. Technology was the catalyst of innovation, not a cost center.

Now, we stand at a similar inflection point. While AI is the next chapter of digital transformation, it's also the last chapter of that book. It sets the stage for a completely new book, and a new era of business: digital transformation connected, captured, and automated. The difference is that AI perceives, learns, reasons, and acts. And it will keep learning; it is designed to constantly self-improve. Whereas with digital, humans had to do all the programming, AI can take charge of that, creating the capacity for itself to perform a wide range of functions.



**Automation became the end  
rather than the means.**

One of us had a conversation with a leadership team that was dismissing the power of AI because you still need control, governance, and monitoring of agents. For that, the company would need to invest in outside companies that specialize in the work, a cost the team were not receptive to. They didn't appreciate that they could use GenAI to build a solution that controls, governs, and monitors all



of their agents. When that was pointed out, the response was interesting. People said, “You can’t just do that!” But you can! The infinite company will use AI to build the solutions it doesn’t have.

AI is redefining what business can be.

The companies that will thrive in this new era will be the ones that, like the digital pioneers before them, refuse to simply retrofit yesterday’s logic. They will build for a world where human imagination and machine intelligence cocreate the future.

## **Stop Thinking in Silos, Start Thinking in Systems**

Most organizations are still structured around functional silos (HR, finance, sales, operations), each with its own technology stack, its own data, and its own tools and metrics. Each function is also stacked with its own people, processes, politics, and beliefs. This has meant that most AI deployment happens department by department; marketing uses one tool, HR another, IT something else. But work shouldn’t stay inside silos, and AI natives already know this.

The next stage of AI business evolution is what we call an *AI Workforce Platform*: a shared, cross-functional layer of intelligent agents available on demand, like internal talent. To build an infinite enterprise, intelligence must flow throughout the organization. Customer service data should inform marketing strategy. Sales patterns should guide supply chain optimization. Workflows should be connected across domains and become intelligent and adaptive.

Imagine a central repository where business units can “hire” AI agents for specific projects. That might be a compliance reviewer for finance this week, a product analyst for R&D next week. These agents don’t belong to departments; they belong to the enterprise. In effect, this creates a fractional workforce, where agents act as flexible capacity that scales up or down on demand.

Infinite companies will build platforms that continuously learn, adapt, and orchestrate work across human and AI agents. They will dissolve silo boundaries and reveal new ways of working.

In this new model, work is orchestrated by intelligent agents that assemble the right capabilities, human and machine, in the right moment. The job description as we know it becomes fluid. It's shaped by context, opportunity, and evolving capability. Infinite companies seek to create human–AI collaborations that accomplish what neither could do alone. They aim to spot opportunities no one else sees, solve problems before they surface, and create experiences so intuitive they feel like magic. When data, insights, processes, and people connect seamlessly, AI doesn't just automate. It amplifies curiosity and imagination. It becomes a force multiplier for invention.

## **Be Truly Bold...and Wildly Curious**

Infinite companies don't limit their experimentation and growth to the standard use cases vetted by their peers. To them, proven use cases serve as the status quo. They're important, but they don't redefine the boundaries of possibility. Infinite companies venture toward the edge to analyze emerging trends shaping the future. Their goal is to understand the next waves of business transformation as they take shape. This requires vision, boldness, openness to seeing things differently, curiosity to ask different questions, and a culture of empowerment to experiment and learn. You need a culture that serves as the foundation for a new type of company...an infinite culture.

An infinite culture is empowering and voraciously curious. The finite culture is risk-averse, cautious, but experienced in operating within the dynamic of business as usual. Neither has the answers to what the future world of living with AI will evolve into, but one accepts that massive change is underway.



Infinite companies understand that the transformation ahead should involve reimagining everything, from:

- **Disconnected work to augmented collaboration**, where human and machine creativity compound each other's strengths in real time.
- **Siloed processes to intelligent orchestration**, where data, insights, and actions flow freely across the enterprise.
- **Incremental efficiency to exponential value creation**, where entirely new products, services, and business models emerge that no human alone could have imagined.
- **Limited hours to 24/7**, where work is no longer limited by human capacity or time. AI extends work hours, humming along even when the lights are out and the doors are closed.

In this new model:

- **AI becomes a teammate, not just a tool**. It works proactively, anticipating needs and offering solutions before problems surface.
- **Agents and digital workers operate continuously**, learning and improving with every interaction.
- **Leadership becomes distributed and accelerated**, with decisions informed by machine foresight and many executed at the edge of the organization.
- **Workflows evolve dynamically**, reshaping themselves to meet shifting customer demands, market conditions, or strategic opportunities.
- **Systems become transparent and audited**, with decision criteria and logic being exposed and explained.

Imagine an operating system that displays a dashboard for tracking activity across the organization. It's a living, adaptive system for

orchestrating intelligence across every human and AI node in the enterprise. Now imagine this dashboard dynamically evolving as humans and AI collaborate to predict areas of opportunity!

This is the foundation of the infinite company: an organization that scales without friction, learns without limits, and uses AI not just to improve what is, but to create what has never been.

Infinite companies are not just organizations, they are living systems. They grow smarter with every decision. They redefine not just how work is done, but what work is worth doing. Instead of asking, “How can we use AI to improve what we already do?” they ask, “What could we do now that was impossible before?” And they act on the answers, even when those answers demand rewriting the rules entirely.

This is the leader’s moment to architect the unknown; to look beyond automation and see augmentation as the starting point, not the end goal; to build companies that don’t just respond to the future but *foretell* the future. The lesson of the digital era is clear and humbling: the greatest opportunities belong to those who abandon the safety of the familiar and expected, freeing themselves from the restraints of iterative and linear thinking.

## The Operating System of Infinite Companies

In sharing his view on the future of enterprise software, Aaron Levie, CEO of Box, offered a clear articulation of where we’re headed. His premise is simple yet profound: in the coming years, AI-native software will be the backbone of how organizations operate, not as a feature layered on top, but as the default way work gets done [36]. In the infinite company, this vision comes to life in three major ways:

### 1. Core platforms for deterministic work

For the most critical, high-stakes, and repeatable workflows (think payroll, enterprise resource planning, customer



relationship management, IT service management, customer support, and document management), enterprises will still rely on core platforms. These will be built from the ground up for machine interaction. The interface is built for people of course, but also tuned for AI agents and digital employees to navigate, interpret, and execute tasks.

Humans will still have “seats” in these systems, but the majority of transactions, updates, and decisions will come from AI agents operating at machine speed. Incumbent software players will either adapt to this model or disappear, as entirely new competitors emerge with AI-native architectures.

## **2. The rise of agent-only companies for nondeterministic work**

Beyond those core systems, we’ll see an explosion of agent-only platforms, purpose-built, to handle complex, less predictable, and highly specialized work. Imagine AI-native companies dedicated to coding, compliance audits, bug detection, or financial analysis, each with its own “fleet” of agents. These won’t be add-ons to existing software. They’ll be stand-alone ecosystems. Infinite companies will thrive here because there are few entrenched incumbents, and the speed of innovation will favor those unburdened by legacy infrastructure.

## **3. A seamless web of agents, platforms, and people**

The real magic happens as these systems connect: AI agents are embedded in core platforms, enabling them to integrate with other agents across multiple platforms. Here, you don’t just open an app to do a task. You call upon an ecosystem of intelligence that assembles the right capabilities in the moment you need them, whether that’s inside a platform, through a horizontal agent layer, or via a conversational interface.

For power users and niche needs, employees will be able to generate micro-apps on the fly. These instant micro-apps are

lightweight, purpose-built tools spun up in minutes to solve a problem, integrate multiple systems, or create an entirely new workflow.

The implications for leadership are massive and transformational. Software doesn't disappear; it becomes more important. But the way we interact with it changes completely. Work becomes fluid, adaptive, and machine-assisted by default. Your competitive advantage comes from orchestration and imagination.

The winners will be those who design an enterprise-wide intelligence fabric, connecting deterministic core systems, specialized agent platforms, and human creativity into one seamless operating model. Just as we now move fluidly between desktop and mobile and generative AI without thinking (sometimes too literally), the infinite company will move effortlessly between human-led, agent-led, and cocreated work.

Meanwhile, AI will live in the background, continuously sensing, learning, and acting.

## Take Heed of the 5%

MIT's *State of AI in Business 2025* report, published in July 2025, contained a surprising finding: despite tens of billions in enterprise investment, 95% of organizations are getting no measurable return [37]. Why share this one report that captured a moment in the past? Because it caused the market to tank. Writing for *Fortune*, Jeremy Kahn described it perfectly [38]. Traders, he observed, "chose to only read the headlines....As for that MIT report, the market chose to read it as an indictment of AI as a whole and head for the exits, even though that's not exactly what the research said."

The report revealed that the real challenge wasn't to do with the limitations of AI itself, as many executives assumed, but what



the researchers called the “GenAI Divide,” referring to a profound “learning gap” about AI they uncovered. They found that in most organizations, people didn’t yet know how to effectively integrate AI into their work. Why? Because they were thinking about it through the lens of yesterday’s problems and solutions.

Their research revealed that only 5% of integrated AI pilots were extracting meaningful value, while the majority stalled before creating real profit and loss (P&L) impact. If nearly all pilots underperform, the instinct is to pull back. But stepping away was and is the wrong lesson. The right lesson is to conduct your pilots well, learning from the small cohort who are doing so.

The MIT researchers also found that the organizations seeing measurable returns from AI pilots took a radically inventive approach to them: that of end-to-end redesign, creating workflow systems that learn, remember, and act. Additionally, those pilots were not disconnected in silos; they were decentralized, with ownership by line leaders.

The result: faster time to deployment, higher employee usage, lower total cost, and durable moats built from proprietary workflows and learning loops.

McKinsey’s 2025 *State of AI* research echoed MIT’s finding about this gap, reporting that a small cohort of high performers have separated from the pack by wiring AI into their operating model rather than sprinkling it on top [39].

In other words, adoption ≠ transformation. You have to build new tracks.

## **The Office of AI Business Reinvention**

Most companies are trying to meet an exponential moment with a finite vision and structure. They stand up a Center of Excellence (CoE), run a few pilots, publish a policy, create a knowledge article, host a training session, establish guardrails as governance, and then

wonder why adoption stalls, why risk teams say “not yet,” and why pilots don’t deliver noteworthy ROI.

This approach is wholly inadequate. This moment doesn’t call for another committee. It calls for a fully functioning, one-stop shop that makes AI real in the only place it matters: inside the work, empowered and resourced to make the work happen.

The moment demands a function dedicated to AI transformation, innovation, and scale; something like an “AI Resources Office,” or what we call the “Office of AI Business Reinvention.” This is a place that houses the business logic for AI, the engine that turns ideas into outcomes and outcomes into redesigned workflows that can safely run at scale.

AI agents, and eventually physical AI, change who and what does the work, how decisions are made, where accountability lives, and what “safe” looks like at speed. If you leave that to scattered teams, you get the predictable outcome: fragmented tools, inconsistent governance, shadow AI, irregular results, greater risk exposure, and a growing gap between what leaders hope AI will do and what it actually does.

This office exists to do what most organizations are missing: question the strategy, align it to company strategy, and then ship the new way work flows. The mandate begins the moment someone says, “We want to deploy AI for reason X.”

Then, the office goes to work helping that person answer key questions:

- What are you trying to achieve?
- What outcome changes do you want for customers, employees, risk, revenue, cost-to-serve, or resilience?
- If you “save” time or headcount in one area, where does that capacity go next so it drives growth?
- Which workflows will you redesign end to end so the business actually feels the difference?



That questioning is how you prevent AI from becoming a thousand disconnected experiments and instead turn it into a compounding advantage.

This is also why the Office of AI Business Reinvention naturally complements a role or concept we'll introduce later, the chief workflow officer (CWO). No single leader can be informed enough to redesign everything. The advantage comes from an empowered, cross-functional group that can see the chasms between the vision and the goals, teams, and systems needed to achieve it. This group can identify and prioritize lighthouse workflows that matter, and standardize what should be reusable so every new build does not start from zero. The CWO is accountable for flow and outcomes. The AI Office makes flow redesign repeatable, governable, and scalable across the enterprise.

## **Not a Center of Excellence, an Operating Desk That Ships**

A traditional CoE advises. It recommends. It hosts best practices.

An Office of AI Business Reinvention ships...*reinvention*. It operates. It governs. It enables.

It is the place where AI becomes a managed capability, and where reinvention is operationalized. This means the office helps teams use models while also helping them redesign the workflows the models touch. The goal then is to instrument those workflows so they can improve every time they run.

The office becomes the keeper of learning loops, to compound that learning into growth and reinvention at scale.

## **Treat Agents and Physical AI Like a Workforce, Not a Feature**

If AI is going to run inside the work, then agents and AI-powered robots and devices must be treated with the discipline you apply to anything that can take action across the enterprise.

Every production agent and robot should have an owner and a role. A job description. A work chart. Clear authority boundaries and permissions. Escalation paths. Circuit breakers. And when agents or robots are retrained, reassigned, or retired, that must be done with a lifecycle model so you don't end up with "orphaned" automations acting without accountability.

Reinvention now becomes a new operating model: a hybrid workforce where humans move up the stack and agents do the heavy lifting, continuously, under policy.

## **Governance That Travels with the Flow**

Governance cannot live in a PDF or knowledge center if agents and robots live in the workflow.

Governance must travel with the flow. As agents act, you need policy, observability, and intervention rights embedded into the workflow itself. Who can access, trigger, supervise, escalate, and shut down is not paperwork. It is product design. And it is how speed and safety become one motion.

This is also where the office earns its keep. It standardizes "safe patterns" and makes them reusable: templates, connectors, prompts, evaluation harnesses, policy envelopes, and control patterns that compress cycle time and reduce risk assessment fatigue. That is how you scale without multiplying governance overhead.

## **The AI Control Tower**

Later, we'll introduce the AI Control Tower, a single pane of glass for orchestrating AI, AI agents, and physical AI. The AI Control Tower acts as a central hub to discover, manage, monitor, and secure all of the company's AI agents, models, and workflows across the entire lifecycle, from ideation to deployment and retirement.

The Office of AI Business Reinvention is not the only team that would use an AI Control Tower. It becomes a shared capability across



the business, from IT to HR, legal, and risk, because everyone is accountable for outcomes and assurance.

In practice, IT is responsible for the agent platform, observability, security, and the Control Tower's capabilities; HR redesigns roles and career paths for human-agent work; business leaders own outcomes and decide where human-agent leverage is applied. The office orchestrates the system, so the Control Tower becomes instrumental in managing day-to-day activities and performance.

## How the Office Actually Runs

Finite organizations move in calendar time. Infinite companies learn to move in event time. Signals trigger directives; directives become bounded actions; actions generate telemetry; telemetry refines the system (Figure 1.1).

That loop is small enough to fit inside a minute and sturdy enough to run all night, 24/7, even in dark offices and factories. The Office of AI Business Reinvention is the team that makes the loop real for the enterprise.

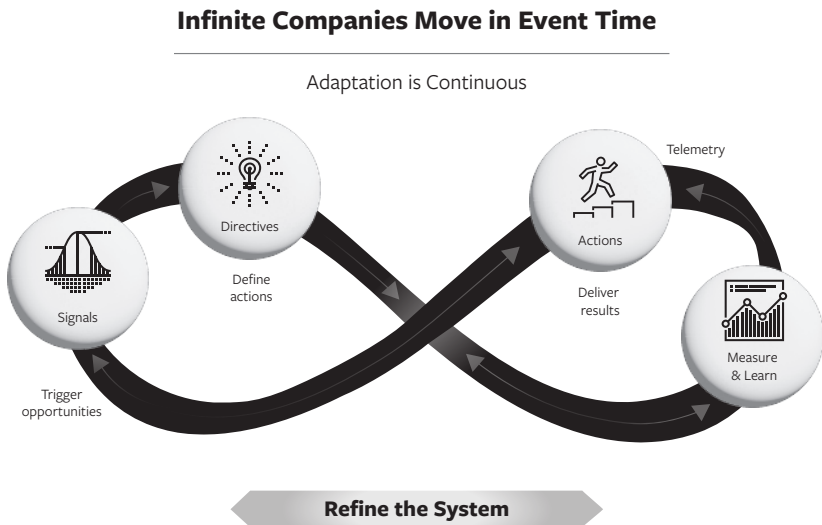


Figure 1.1. Infinite companies move in event time.

It is also how you avoid “pilot purgatory.” The office picks a small number of lighthouse workflows, redesigns them end to end with the CWO, instruments them with the Control Tower, and then scales what works by franchising the pattern across the next workflow. Vertical reinvention before horizontal rollout. One journey at a time, until the organization starts to gain exponential momentum.

## **You’ll Know It’s Working When...**

You’ll know the Office of AI Business Reinvention is gaining ground when:

- Every AI effort begins with an outcome brief, not a tool request.
- The office is allowed, expected, and equipped to challenge the “why” before anyone builds the “how.”
- Every production agent has an owner, privileges, escalation paths, and lifecycle discipline.
- Governance is embedded in the workflows, and intervention rights are explicit.
- The AI Control Tower is used to steer the vital signs of work, not just report activity.
- Each run teaches the next, so the organization exponentially compounds learning rather than repeating effort.

This is how a company stops “adding AI” and starts becoming an infinite company: by building the operating model that makes reinvention inevitable.

## **Compete with AI Natives or Become a Finite Company**

To cross the divide, legacy enterprises must stop benchmarking against other incumbents and pretend they’re competing in Silicon



Valley, because they are. Those AI natives are building their companies with AI in their DNA, becoming a new species of organization.

To understand what you're competing with in Silicon Valley, and the "Silicon Valleys" around the world, you must become familiar with the ways organizations that are winning with AI are reinventing with it. In the chapters ahead, we'll take you on a tour of their processes of reinvention.

