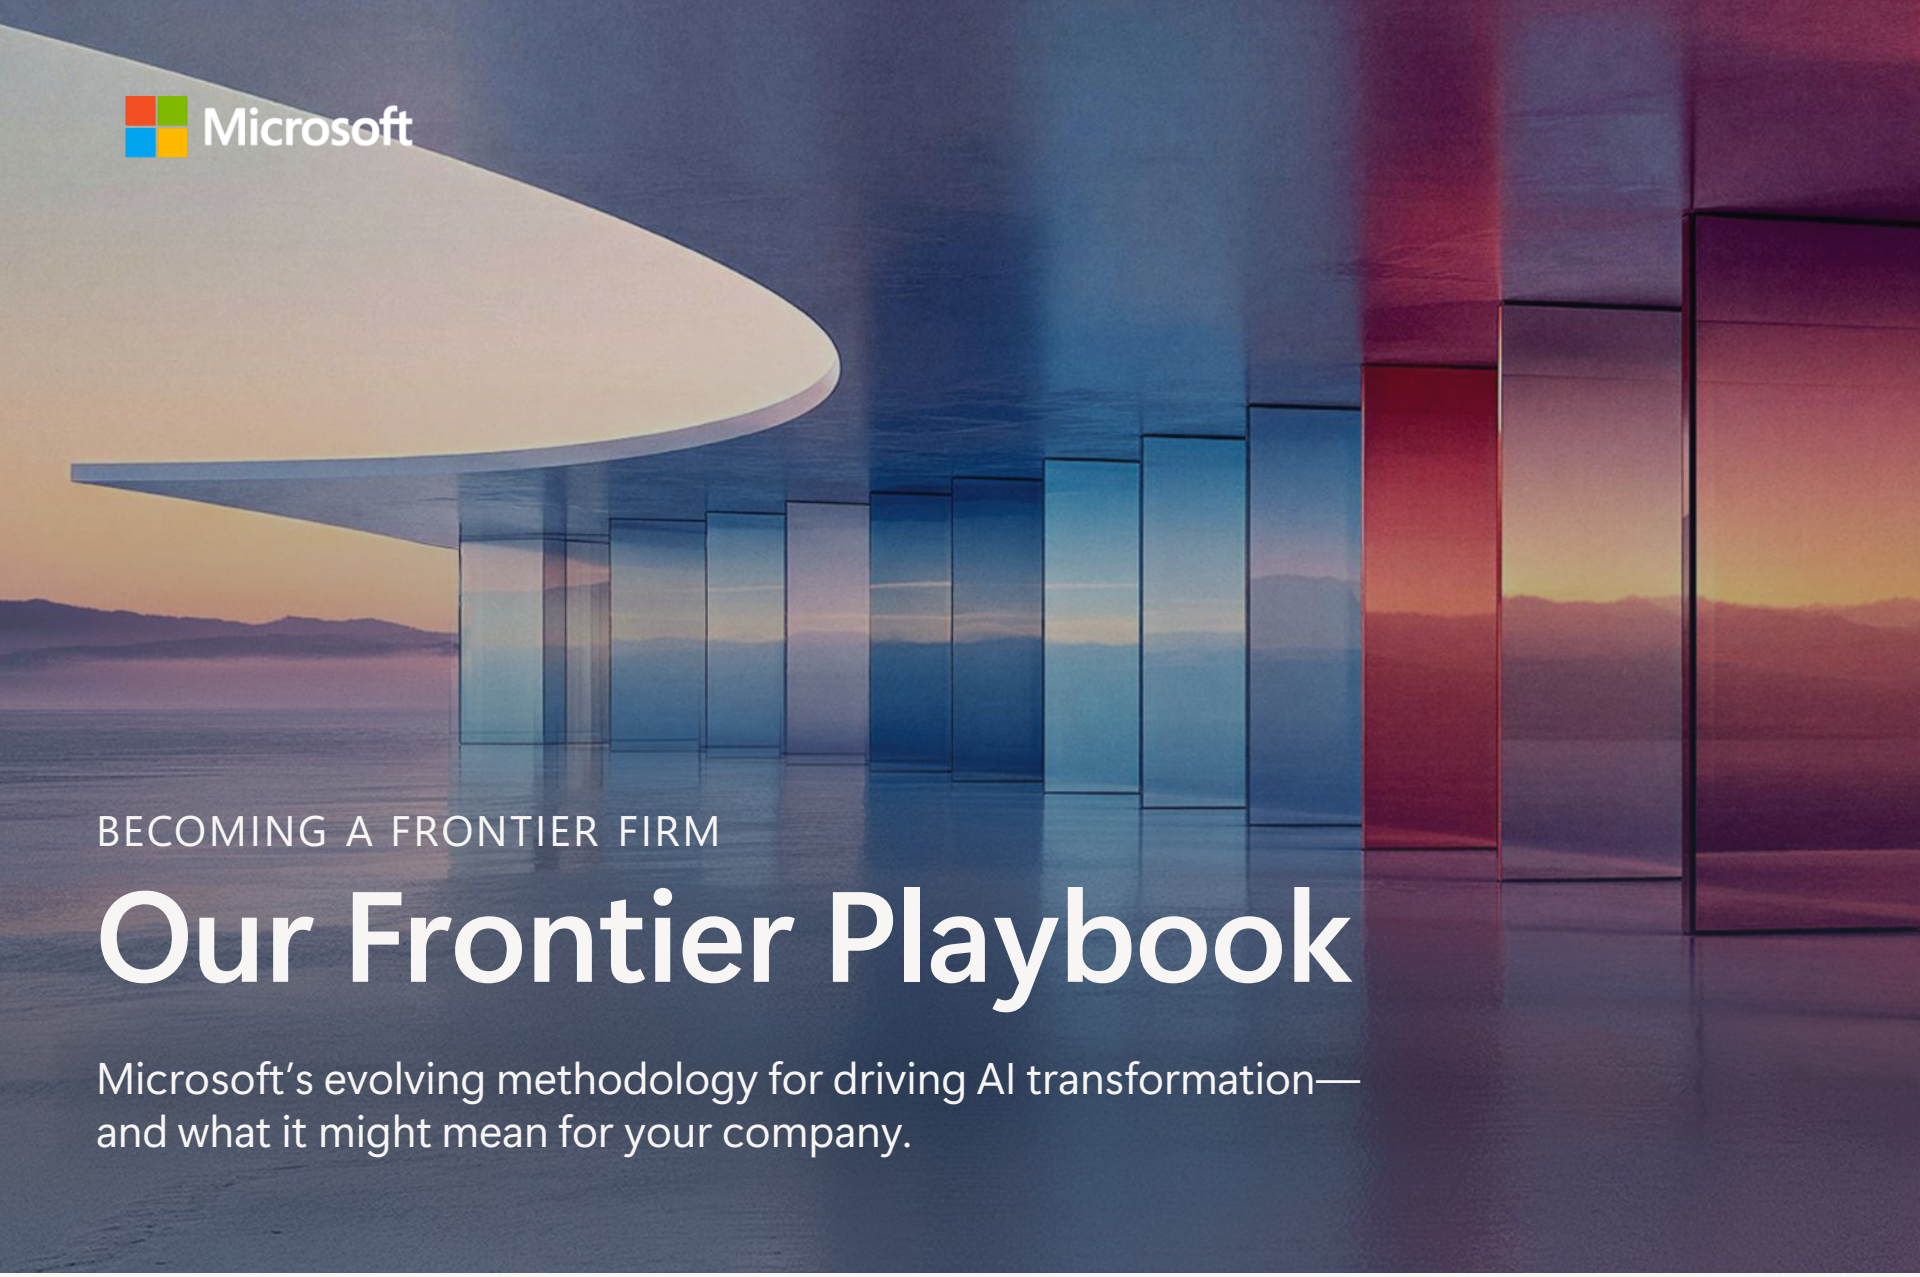

BECOMING A FRONTIER FIRM

Our Frontier Playbook

Microsoft's evolving methodology for driving AI transformation—and what it might mean for your company.

This Playbook reflects Microsoft's current experience and evolving perspective. It is intended to share our working approach and to advance thought leadership in AI transformation, not to prescribe a solution nor guarantee results for others.

Introduction to the Frontier Playbook

Microsoft is in the midst of its own AI transformation toward the frontier, not at the end of that journey.

What we are sharing here is a snapshot of that work: what we've tried, what has endured, and what continues to evolve. This Playbook reflects our current experience and perspective, knowing both will keep advancing.

The strategic questions we encounter are consistent across every boardroom:

- How do we anchor AI to real enterprise value?
- How do we move with speed without losing what makes us distinct?
- How do we bring people along as we redesign work?
- How do we compound value by investing in both human and token capital?
- How do we govern agents while remaining deeply human-centered?
- How do we protect our IP and sustain competitive advantage?
- How do we safeguard privacy responsibly?

In response, we've developed a working method—a set of repeatable patterns and a growing library of cases where this approach has delivered measurable impact.

Our goal is to help organizations use AI to advance their business strategies and position themselves at the new frontier—not to adopt AI for its own sake, but to make their business fundamentally stronger.

"Imagine a world where every person can get help from a researcher, a coder, or an analyst on demand. Not just information, but deep, contextual expertise paired with action. Or where every organization, no matter its size or sector, can reinvent employee experiences, reimagine customer engagement, reshape business processes, and bend the curve on innovation for their people, businesses, and industries.

This is the new frontier and how we will unlock the next level of productivity and growth for the world."

Satya Nadella, Chairman and CEO, Microsoft

Our Playbook is built around five elements that enable AI to deliver lasting business outcomes

A Set business ambition & goals

Anchor AI in service to your business strategy and creating enterprise value

B Build your diffusion engine

Build an AI operating model that delivers against your business goals at scale

C Invest in your people

Adapt roles, manager behaviors, and learning so AI expands human capability instead of eroding it

D Codify your advantage & controls

Codify what makes your company distinct into evals and a learning system—then keep that intelligence compounding inside your boundary

E Safeguard your security

Apply enterprise-grade governance to every agent action before you scale


A
Set business
ambition & goals

How is AI transforming your business results?

AI transformation is not a product launch, nor a tech project. It is truly a business transformation and therefore must be guided by clear business outcome objectives.

We have found that AI has the potential to drive performance relative to four key objectives.

1. Enrich employee experience

2. Reinvent customer engagement

3. Reshape business processes

4. Bend the curve on innovation

We're actively pursuing these goals ourselves, so we're not simply advising without taking our own advice. We anchor every AI initiative against four frontier goals—what we call the Frontier Firm Success Framework.

We have organized our own internal transformation to the frontier around this Success Framework. The leaders across each of these four areas ensure that there is clear focus and aspiration for each, guided by our business strategy.

"Organizations are activating human ambition—engaging customers more effectively, reshaping business processes, and accelerating innovation without adding operational complexity—turning gains into competitive advantage."

Judson Althoff, CEO,
Microsoft Commercial

The Frontier Firm Success Framework

Enrich
employee experiences

Reinvent
customer engagement

Reshape
business processes

Bend
the curve on innovation

The Frontier Firm Success Framework

Enrich employee experiences

Reinvent customer engagement

Reshape business processes

Bend the curve on innovation

What we are doing

We seek to improve our employees' ability to focus on their highest value activities, reducing the "toil" from routine work and increasing both their productivity and experience. AI is embedded in the daily flow of work—drafting, learning, deciding, collaborating—and removes friction that drains energy and crowds out the work that develops people. In addition, AI learning prompts are now embedded in many of our day-to-day routines so that we improve human skill-building in the flow of work.

We monitor the impact of AI usage on our employee engagement scores to assess whether we are having the impact intended. We are delighted to see that colleagues who use AI the most also report the highest "Thrive" scores, based on doing energized, empowered, and meaningful work.

Example outcome metrics

Engagement scores ("Thriving", for us)

Hours reclaimed per employee per week

Accuracy of self-service rate

Time to proficiency

Skill acquisition rate

The Frontier Firm Success framework

Enrich employee experiences

Reinvent customer engagement

Reshape business processes

Bend the curve on innovation

What we are doing

Our AI approach is designed to reflect and scale our unique customer engagement model. We have leveraged AI to magnify the work of our marketers by reducing manual drafting effort, accelerating content cycle times, and reallocating capacity toward higher-value work. We have developed a suite of AI solutions for our field personnel, including a custom Sales Chat tool, a Coaching agent, a Deal agent, and even a Sales agent that is able to help close business with small accounts. Our Customer Support team has been transformed by agents that connect the customer to the best human support engineers to address their issue and help deliver faster, more consistent resolution of routine inquiries.

We see AI's role as scaling the organization's best judgment into every moment: personalized, in-context, and consistent at enterprise scale.

Example outcome metrics

- Revenue per sales rep
- Qualified pipeline coverage
- Customer lifetime value
- Net revenue retention
- Conversion rate
- Net Promoter Score/CSAT
- Win rate

The Frontier Firm Success Framework

Enrich employee experiences

Reinvent customer engagement

Reshape business processes

Bend the curve on innovation

What we are doing

We are working across each of our corporate functions to reinvent end-to-end business processes. Our goals include increased internal customer satisfaction, speed, insight, improved quality, risk reduction, and efficiency. We convene a Customer Zero Council with leaders from every function to share insights and best practices.

Two watchouts from our own experience: first, we have had to push ourselves to work customer-back and end-to-end on the processes that matter most—not reinforcing the siloes of our org chart. Second, we have learned to redesign the work rather than simply automate it. AI should bring speed, accuracy, and efficiency, but also introduce capabilities only AI makes possible—prediction, scenario analysis, and action recommendation—that lift internal customer experience and reduce enterprise risk.

Example outcome metrics

End-to-end cycle time

Risk event reduction

Customer experience

Cost per transaction

Quality and error rate

Audit exceptions per period

The Frontier Firm Success Framework

Enrich employee experiences

Reinvent customer engagement

Reshape business processes

Bend the curve on innovation

What we are doing

Innovation is the heart of our company, and where we focus the most. Our goal is to leverage the power of AI to bring better products to market to support customer needs—not just to make our coding function faster. We are also applying AI across our data center operations to deliver against capacity goals.

Innovation is where AI shifts from efficiency to strategic advantage. Applied to the current roadmap, it delivers incremental speed; applied to the innovation system itself, it changes the slope. AI expands the range of hypotheses worth testing, enables continuous customer input, compresses the cycle from idea to evidence, and allows smaller teams to reach insights once reserved for scale. The result is a higher rate of product freshness and a product development portfolio weighted toward customer-visible value.

Example outcome metrics

Time from idea to market

Customer success

Developer productivity and experience

Experimentation velocity

Patents filed per R&D dollar

Orchestrating AI Transformation

Given our organization's scale and technical maturity, each of our functions and business areas has established their own AI Transformation Office (AITO) to guide and execute against their priorities. The specific makeup and roles of these groups vary, but all bring together the technical, process, and people capabilities that are required to drive successful reinvention with AI, leveraging the tools that we describe in the next section. Each is also responsible for governance in their area: enlisting citizen-developers, ensuring risks are managed, and supporting the effective prioritization and scaleup of new AI solutions.

We have a small centralized team that acts as a "learning accelerator" across these groups, convening sessions for sharing learnings and best practices to help accelerate our progress overall.

The learning accelerator team also supports four key enabling areas to support progress around:

1. **People:** employee engagement and communications, upskilling, change management
2. **Methodology:** the Playbook “recipes” that optimize workflow using Continuous Improvement (CI) tools and adding new AI capability
3. **Technology Enablement:** ensuring broad access, usability, and risk management as we deploy AI tools to our employee population
4. **Measurement:** setting standards on how to measure the revenue, cost, and other impact of AI

We have organized our own transformation against the Success Framework

CEO and Top Team
Establish vision and principles

AI Transformation Office
Define, orchestrate, and support journey

Frontier Success Framework

- **Enrich** employee experiences
- **Reinvent** customer engagement
- **Reshape** business processes
- **Bend** the curve on innovation

Enablers of transformation

- People
- Methodology
- Technology Enablement
- Measurement

Three learnings that informed our approach:

Cascade goals from the top, not use cases from the bottom. It is easy to let narrow use-case statistics become the story of your transformation. Resist letting bottom-up activity become the strategic narrative.

Pair every frontier goal with a named executive owner. What is measured at the top is what moves in the business. Each goal has a senior owner accountable for the outcome, not just the activity.

Force yourselves past efficiency. On every initiative, we ask: if the only ROI here is cost out, what bigger opportunity are we missing? Sometimes the answer is "nothing—efficiency is the right play here." More often, the answer changes the design of the program.

The picture we keep in front of every initiative is a simple one: efficiency—speed, quality, cost—is nonnegotiable. Customer and employee experience, innovation, and growth—these are the ceiling.

AI drives both growth and efficiency – creating real operating leverage

New output with higher strategic value

- Revenue/ Growth
- Innovation
- Customer and employee experiences

+

Output delivered faster, with fewer errors, and more efficiently

- Speed
- Quality
- Productivity/ Cost

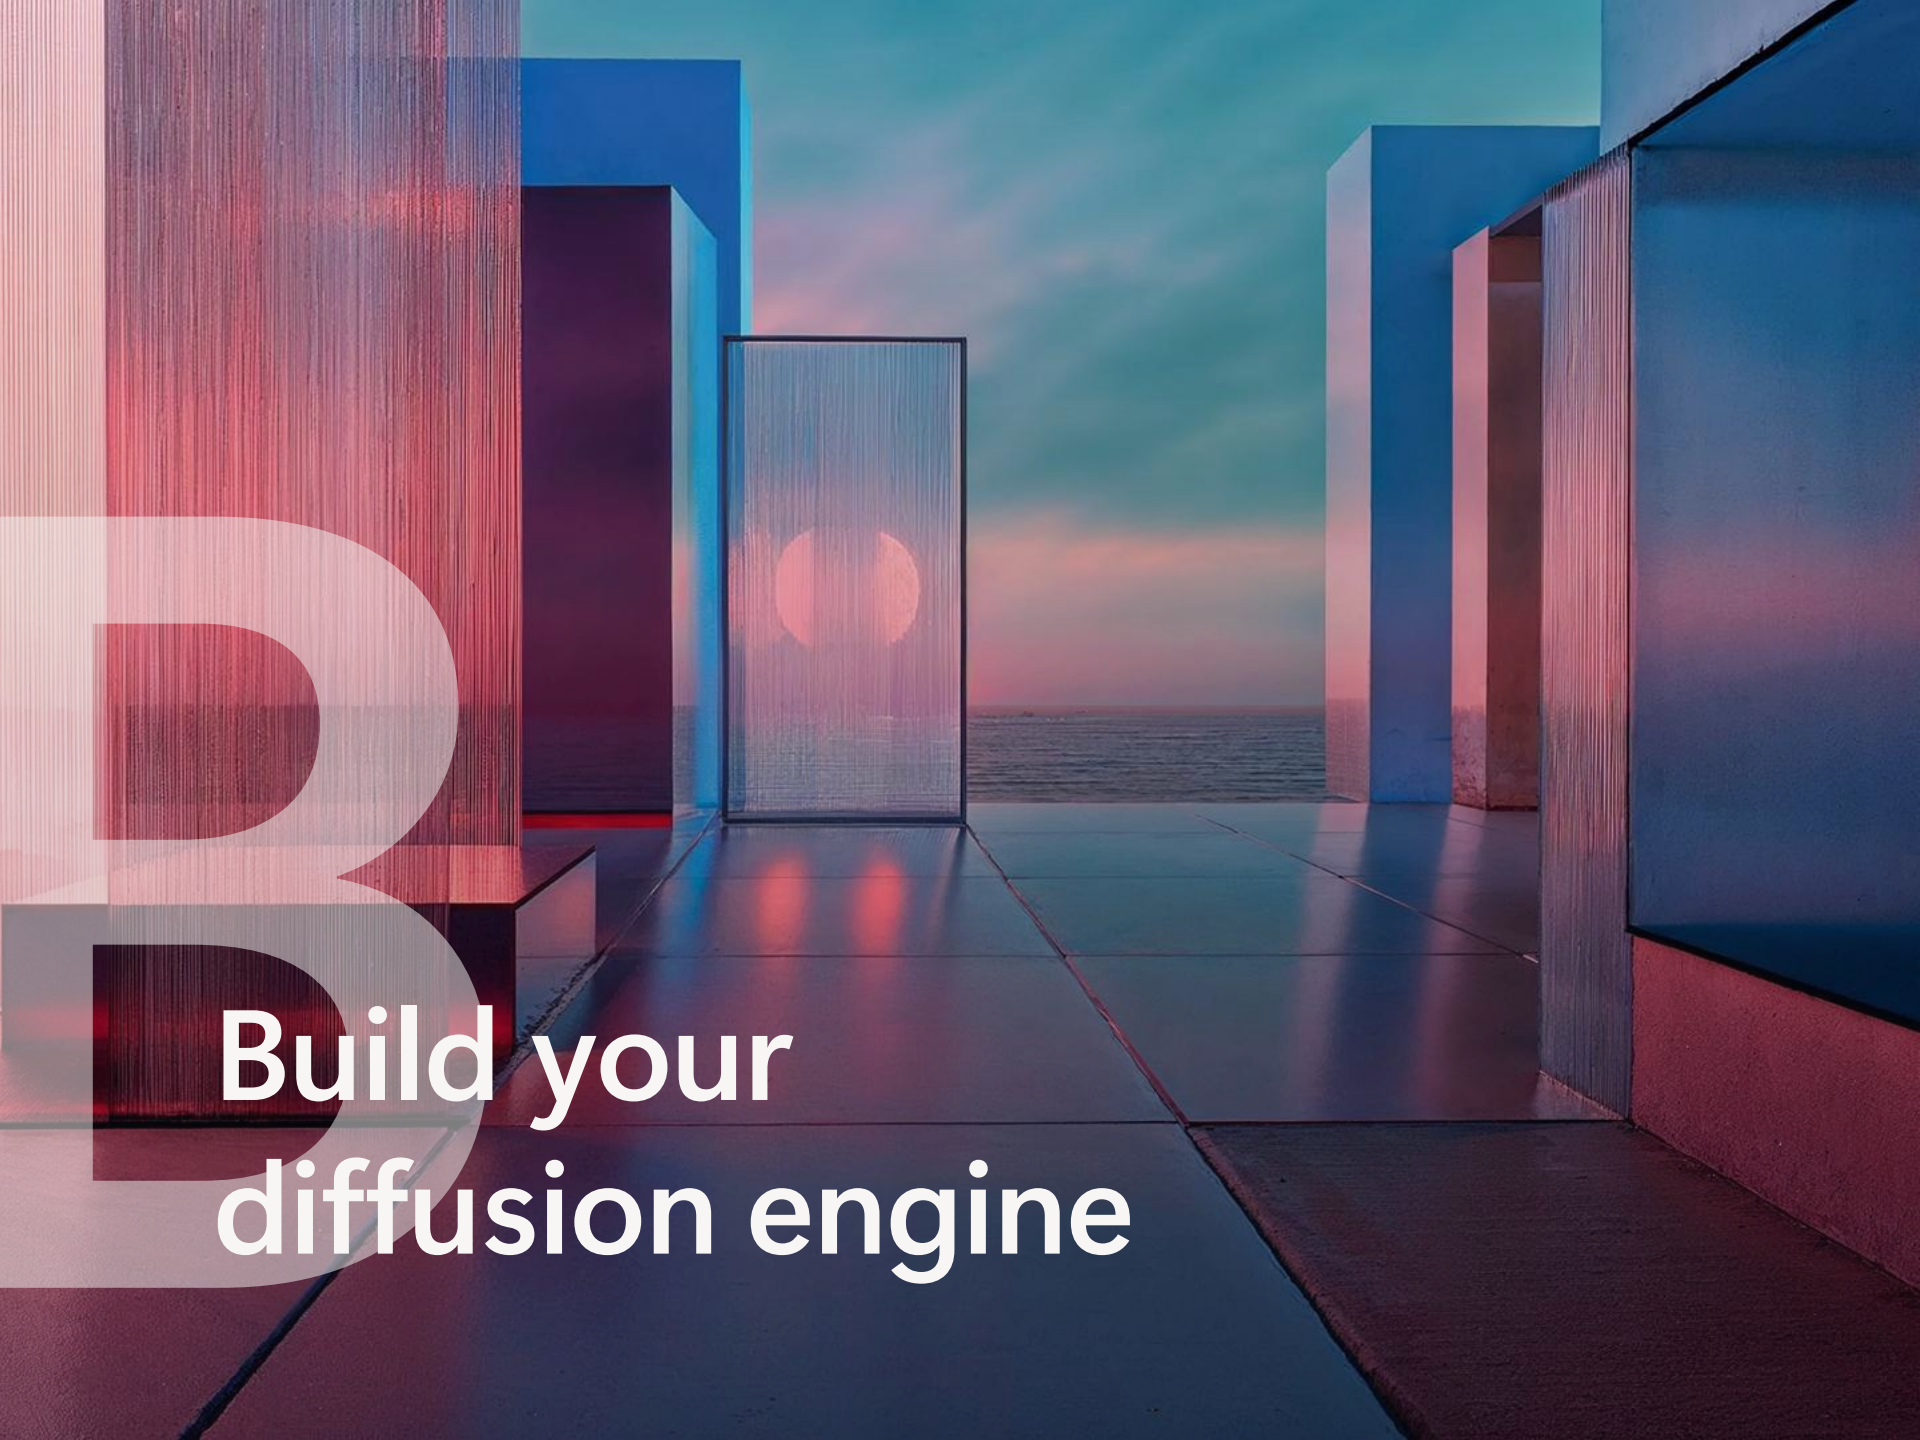
Build your
diffusion engine

Is AI adoption transforming the way your enterprise works?

Buying tools alone is not transformation. Effective system-wide diffusion is. A tool licensed and rolled out to 100,000 employees does not change how the work gets done. What changes the work is the operating system around the tool – the rituals, the redesign, the measurement, the manager behavior, the learning loops, the culture of experimentation. That is what we mean by a diffusion engine, and it is a key determinant of whether AI translates into measurable business value.

In this section, we cover:

- **Three recipes** for diffusion
- Levers for identifying strategic value through what we are calling AI's **Capability Add**
- An **approach to measuring AI's impact**, going beyond adoption statistics to ensure adoption is turning into business value

"We've learned that rolling out AI isn't the hard part—transforming the business is. That requires a methodology that creates real strategic value, clarity on what new capability AI can bring, and measuring whether it's actually delivering impact. All three have to work together."

Raghav Raghunathan,
VP, AI Transformation, Microsoft

The origins of our transformation methodology

To develop a scalable methodology, we first examined our successes and failures to date by reviewing over 100 case studies of internal AI transformation across the company, including corporate functions, our commercial organization, and our engineering teams. Perhaps our most important learning from this review was realizing that we could not treat AI diffusion as a tech product launch or an IT project. It truly is a business transformation and needs to be guided by leadership with clear business outcome goals.

We recognized the need to bring together tech expertise with process rigor and a deep understanding of people. We are bringing the rigor of Continuous Improvement methodology to this work, and we typically find significant opportunities for improvement even before we get to applying new AI prompts and agents. We think of these efforts as contributing to a "new science of knowledge work." So much of knowledge work is tacit, undocumented, nonstandard, and driven as much by organizational relationships as by process documentation. Often colleagues who are doing the work only see their part of it, optimize results locally, and don't fully understand how customer value is created.

Informed by our lived experience

100+ case studies of Microsoft AI transformation

Bringing together expertise on process, tech, and people

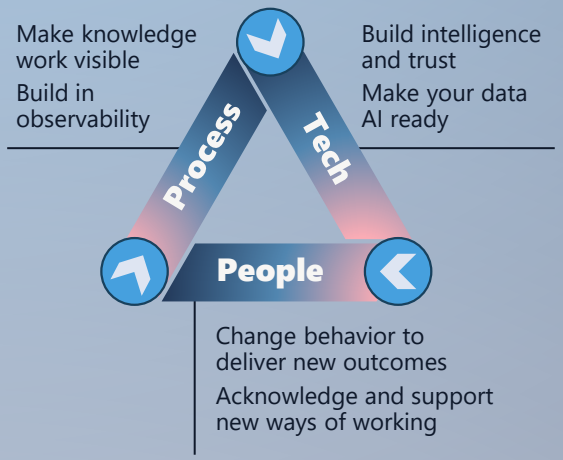

Three recipes for AI transformation

Persona Acceleration

Empower specific roles with AI tools tailored to their workflows

AI-Powered Process Redesign

Reinvent existing processes with AI to improve outcomes

AI-First Possibility

Design new work from scratch, with AI at the center of how the team operates

Persona Acceleration

Start with the role, not the tool.

In almost every organization, some people become power users quickly, while others try the tools once and fall back to old habits. The difference is rarely access. It is whether the person can see, in the context of their own work, where AI helps them do something better, faster, or with more confidence. The jagged edge—human version—means that even with the same tools within similar roles, some individuals are able to dramatically transform their day-to-day work, while others don't see enough value to learn and adopt AI in their workflow. The Persona Acceleration recipe aims to address this issue.

Persona Acceleration begins by making the work of a role visible. What does a seller, engineer, support agent, finance analyst, or manager do during the week? Where do they lose time? Where do they make judgment calls? Where do they search, summarize, coordinate, create, prepare, follow up, or reconcile? What causes them the most toil and frustrations? Once those moments are visible, AI can be mapped to the work instead of introduced as a generic tool.

What matters most in practice:

Identify the key personas and make the role visible

Build a persona- or role-level view of pain points, unmet needs, and high-value moments. This includes looking at how time was spent across internal meetings, external meetings, admin, focus blocks, and customer-facing work, then setting benchmarks on how time should be spent.

Map AI to real moments in the workflow

Do not ask people to "use AI more." Show them where AI fits in their workflow. For example, in sales, highlight activities like preparing customer materials, updating CRM, analyzing performance data, generating deal packages, summarizing meetings, and prioritizing pipeline.

Create role-based learning loops

Small-group huddles work because they are specific, social, and practical. They let people see peers applying AI to recognizable work, ask questions safely and repeat the behavior.

Engage direct managers visibly

Manager reinforcement is a key driver of adoption. Those who role-model the behavior make AI usage part of the operating rhythm rather than an optional side activity.

Instrument the change

Use huddle engagement, usage, feedback, surveys, and outcome indicators to see whether adoption is turning into habit formation and better work.

This page represents impact that we are seeing with the recipe approach, in the context of our company.

Persona Acceleration

Learn from our Sales transformation

In our Sales organization, we used this recipe to help sellers embed AI into the rhythm of their actual work and generate measurable value. The work was not simply to train sellers on Copilot. It was to identify distinct sales personas, understand where time was spent, surface the moments where AI could reduce friction or add new capability, and connect the right tools to those moments in the flow of work. The pilot combined weekly role-based huddles, multifaceted telemetry, and codified AI best practices into "Experience Cards" to turn experimentation into repeatable habits.

The results include:

3x
increase in adoption on priority use cases, as reported by survey participants

9.4%
increase in revenue per head within the pilot group

20%
faster deal closure within the pilot group

This page represents impact that we are seeing with the recipe approach, in the context of our company.

"Day in the life of" map of key personas with agents

Persona-level mapping of agents to use cases

- Utilize **Copilot** to draft discussion documents with customers
- Use **Researcher** to tailor materials and content to resonate with each client profile e.g., cloud engineer
- Leverage **Analyst Agent** to prioritize a sales pipeline instead of meeting with Commercial Engineers
- Update Microsoft Dynamics 365 with **Sales Agent** instead of holding weekly CRM hygiene touchpoints
- Leverage **Deal Agent** to automatically generate packages for T2 and T3 tranche deals
- Use **MSXi Copilot** to analyze performance data instead of meeting with a Data Support team

Time spent on activities

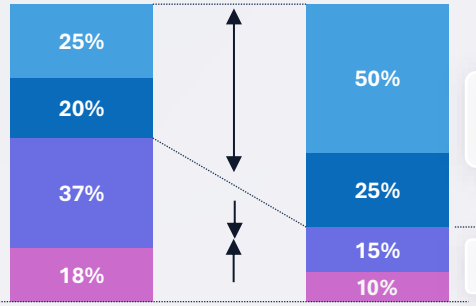

Expand seller bandwidth to allow for more customer time and be better prepared for customer conversations

Reduce internal admin time and toil

AI-Powered Process Redesign

Optimize your workflows before introducing agents.

Enterprises run on processes—some documented, many not—that evolved along with technological capabilities and customer expectations. Most current processes were not designed for the AI era, a gap that is compounded by variation across the org, drift between documented procedures and how they're executed day to day, and hidden waste. Persona Acceleration can boost individual productivity, but without the right end-to-end process, those gains often evaporate—speeding up one step frequently just creates longer queues and larger work backlogs elsewhere.

The end-to-end processes that drive function or enterprise-level outcomes often cut across roles, systems, handoffs, approvals, data sources, and decision rights. In those cases, dropping an AI agent into the existing process often just automates the dysfunction. It may make one step faster while leaving the overall flow fragmented, inconsistent, or hard to govern. The better pattern is to redesign the workflow itself. Start by asking: what is this workflow really trying to accomplish, and where does the current process make that harder than it needs to be? Then make the work visible, remove avoidable friction, clarify decision rights and handoffs, establish a shared data foundation, and introduce agents into a workflow that is ready to compound.

This is where AI transformation begins to look less like tool deployment and more like operating-model redesign. The visible change may be an agent, but the invisible work underneath is often more important: the process map, the shared source of truth, the reusable orchestration layer, the telemetry, the governance, and the cross-functional squad that owns the outcome.

This page represents impact that we are seeing with the recipe approach, in the context of our company.

What matters most in practice:

Map the workflow end to end

Use Gemba, process mining, task decomposition, user research, or structured interviews to understand how the work happens—not how the org chart and process maps says it happens.

De-waste before automating

Clarify the process before introducing agents. AI should not preserve unnecessary approvals, redundant handoffs, unclear ownership, or manual reconciliation simply because they exist today.

Build a shared data foundation

Agents reason differently when they operate from a coherent source of truth. Without that foundation, every workflow risks becoming a one-off implementation with its own context, logic, and failure modes.

Design the human-agent operating model deliberately

Some steps should remain human-led. Some can become agent-operated with human review. Some may eventually become more autonomous. The design should be calibrated to the stakes of each decision.

Use cross-functional squads with named owners

Process redesign only works when the teams that own the work, the systems, the data, and the change are operating together.

Build for reuse

Shared orchestration, observability, telemetry, and governance let the second version of the workflow become easier than the first. Without reuse, each new agent becomes another bespoke implementation.

AI-Powered Process Redesign

Learn from our Supply Chain transformation

For this recipe, the Supply Chain team pursued a "lean before agents" approach: they first mapped and simplified six end-to-end workflows and built a shared data foundation (single source of truth), then deployed 70+ purpose-built agents across planning, sourcing, fulfillment, and logistics to eliminate manual reconciliation and speed decisions.

The results were achieved not by adding tools to a broken process but by redesigning the workflow so agents could compound impact, being intentional about people roles in the new process and building in a shared agentic platform and data source as single source of truth.

The results include:

75%
reduction in cycle time across selected workflows

efficiency from reduced manual effort

measured incremental value

This page represents impact that we are seeing with the recipe approach, in the context of our company.

Map current state. Identify waste and pain points, and areas to add new capabilities

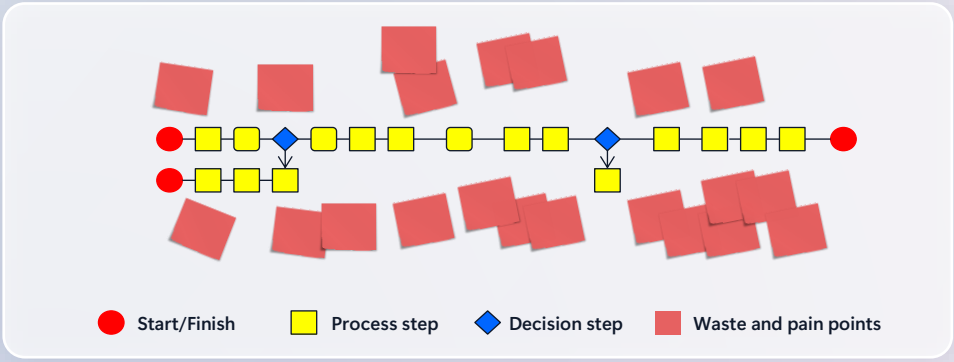

"Lean before agents" is a critical step to ensure we don't carry over the waste

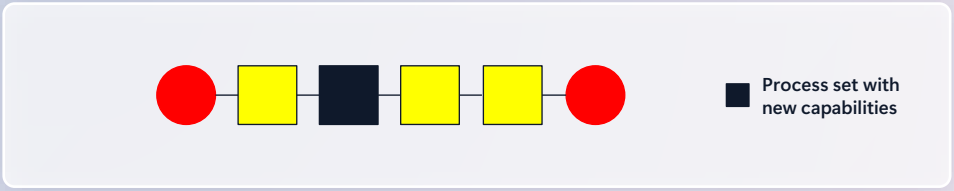

Design integrated human + agent teams that eliminate waste and bring new capabilities

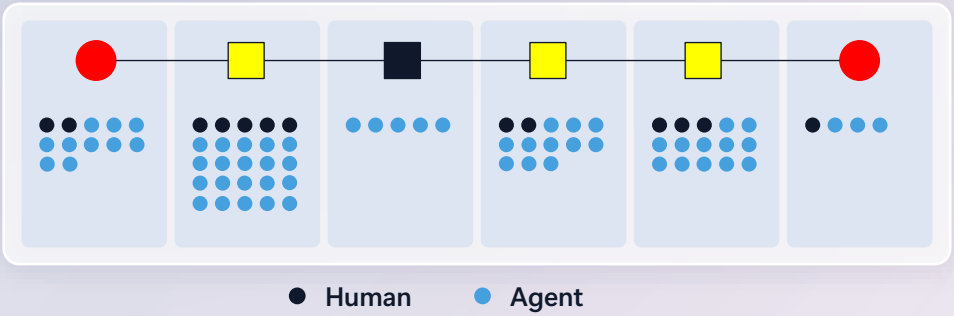

AI-First Possibility

Start from a blank page.

All enterprises need to have mechanisms to disrupt their own business model, create new sources of value, and offer differentiated experiences for customers. The risk of not doing so means that they risk being disrupted by an AI native. AI is making ideation and implementation a lot faster—especially in software—but many of these breakneck developments have happened in "greenfield" scenarios, which many established companies don't have.

Within established companies, the existing process, role design, tooling, and cadence preserve the very assumptions that AI could help break. In those cases, the right starting question is not, "How do we make today's work faster?" It is, "What would this work become if AI were assumed from day one?"

AI-First Possibility is the incubator recipe. It is not the first recipe most organizations scale broadly, but it is often where the most important breakthroughs emerge. The team is usually small, expert, and high-agency—works in a sandbox, moves quickly, writes down its standards, treats AI as the default path rather than the assistant at the edge, and most importantly, codifies what it learns so the rest of the organization can reuse it.

The prize is not only the productivity number. The prize is the reusable pattern: the spec, prompt, eval, skill, agent behavior, operating rule, or architecture pattern that can later feed back into persona acceleration and process redesign.

This page represents impact that we are seeing with the recipe approach, in the context of our company.

What matters most in practice:

Start with a small expert squad

The squad needs enough judgment and taste to know which assumptions can be safely discarded and which must be preserved. The team must also have benchmarks that the AI output will be measured against.

Zero-base the work and make AI the default route

Do not begin by asking how AI can improve the current process. Ask what the work should become if AI can reason, critique, generate, test, and repeat. If people solve issues manually, the system never improves.

Use sandbox governance

AI-first teams need room to move fast, but not outside the boundaries of safety, security, and risk. Sandboxes let teams experiment without exposing the rest of the organization prematurely.

Redesign the team operating model to be agent-friendly

Human-agent teams operate best when standards, best practices, architecture decisions, prompts, evals, skills, and workflows live where agents can learn from them—not only in meetings or tribal knowledge.

Remove velocity crushers and create an innovation flywheel

Slow builds, tribal knowledge, custom frameworks, and agents without access to logs, metrics, or feedback can crush AI-first velocity faster than they would slow traditional work.

AI-First Possibility

Learn from our Copilot Cowork transformation

Our team that built Copilot Cowork set up a sandbox with a focused objective and engaged a fully AI-first product development approach and, in the process, changed many legacy ways of working to deliver the product. This specialized nine-member squad truly "zero-based" the work by taking an "automate everything" mindset and implementing a spec-driven development approach.

The results include:

678
total commits (at 28 commits/day)

~500K
lines of code

9 member team
of meta-engineers, meta-designers, and meta-PMs, with initial product release in 28 days

Zero-base the work—don't add agents to today's process. Build the stack the agent runs on

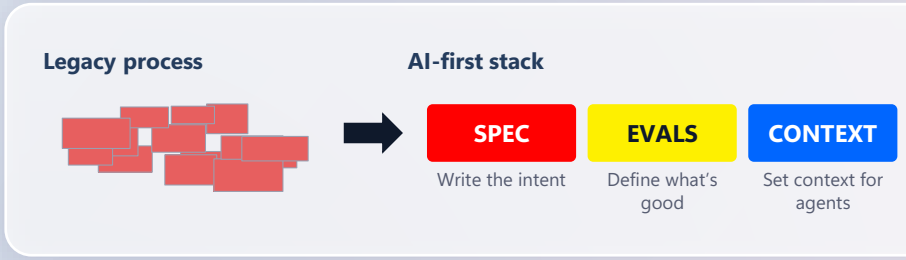

Small expert squad with agents as teammates—nine humans + agents working from a shared context and evals

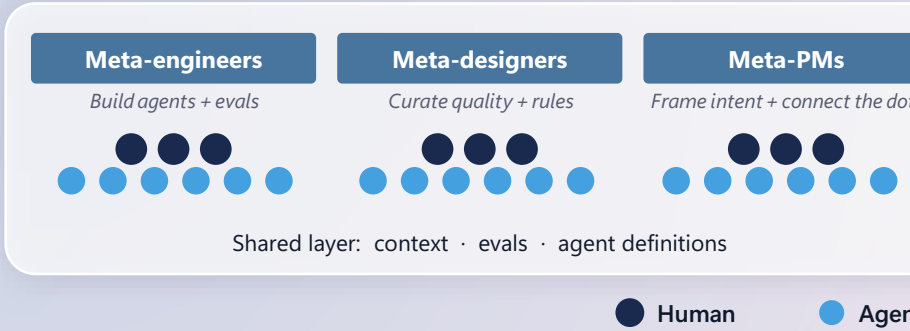

Codify the model to set up an innovation flywheel—each sandbox publishes to a shared library

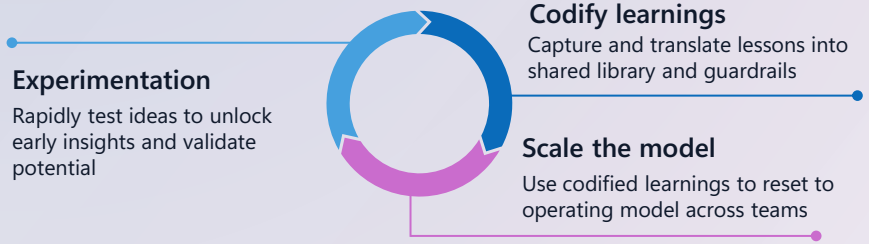

This page represents impact that we are seeing with the recipe approach, in the context of our company.

Levels of AI across all three recipes

Across the three recipes, there are three different levels in which AI can deliver transformational outcomes: 1. Human with assistant, 2. Human-agent teams, 3. Human-led, agent-operated.

Determining which level is a best fit for any given scenario is a critical part of the diffusion recipe. What to consider when making this decision:

- **Risk boundaries** that set the level of AI autonomy and delegation
- **Data and technology infrastructure** readiness for AI
- **Organizational readiness** for change

Broadly speaking, human-led, agent-operated (level 3) AI requires more autonomy and delegation, plus a higher level of data and infrastructure readiness, including access to tools. It also requires a much bigger role-change for humans.

Diffusion recipes do not have to stick to just one level of AI and can often engage multiple levels at the same time. For example, the Persona Acceleration recipe applied to software engineers can use human with assistant—level 1—for one part of their workflow (e.g., Teams Agent), and use human-led, agent operated—level 3—for another part of their workflow (e.g., agentic coding solutions for engineering).

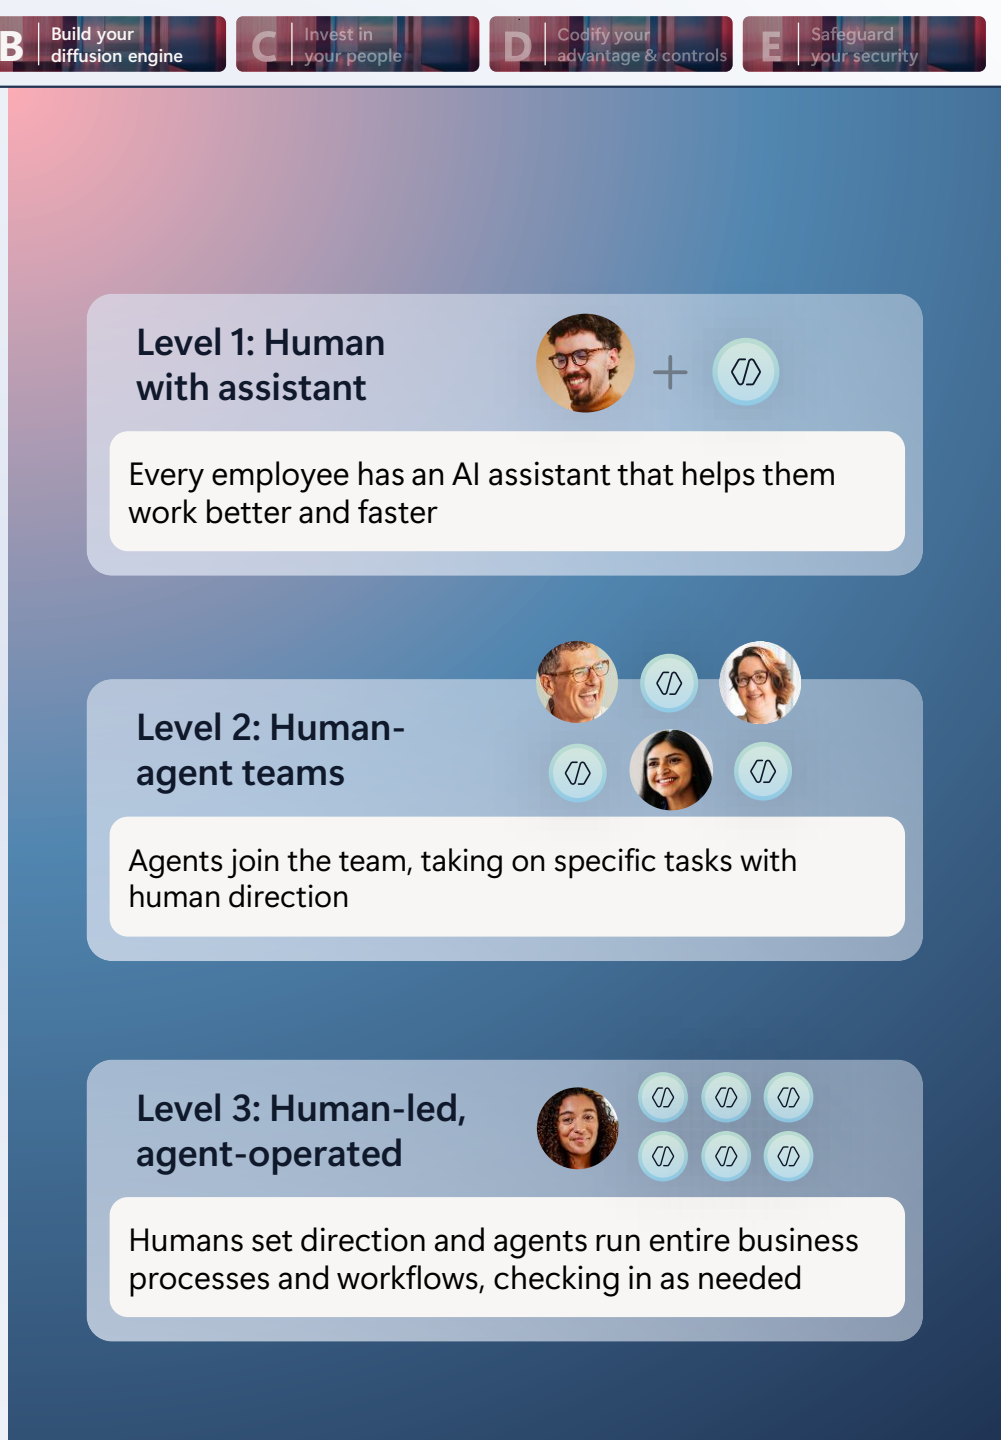

CI + AI = CA (Capability Add)

AI creates value in two ways: with traditional Continuous Improvement (CI) levers that optimize existing processes through simplification and automation, plus what we call AI's Capability Add that creates new sources of value.

AI's Capability Add comes from its unique ability to drive new output with higher strategic value—work that was not possible at the old speed, scale, or cost.

CI alone gives you a leaner version of your old company. CI + AI—the discipline to take waste out and the imagination to add capability in—compounds.

1

Traditional Continuous Improvement levers (CI):

Output delivered faster, more efficiently, and more accurately

Eliminates:

- Errors, rework, quality failures
- Unused outputs, over-documentation
- Delays, blocked work, queues
- Information handoffs, context loss
- Tool-switching, unnecessary search

2

Capability Add (CA):

New output with higher strategic value

Adds:

- Predictive and preventive quality intelligence
- Expanded solution space and structured exploration
- Real-time sensemaking and parallel cognition
- Dynamic opportunity sensing and portfolio intelligence
- Attention and cognitive-load optimization
- Precision at the right level of abstraction

Measure the pipeline from adoption to impact

You cannot measure what you cannot see. Our model for measurement separates four layers of indicators and explicitly maps the dependencies between them.

Two principles: **Start with the outcome,** identifying the leading metrics that predict the lagging ones. Then **instrument the pipeline end to end**—input through outcome. Balance automated telemetry with qualitative customer and employee pulse. The numbers will tell you what is happening, the conversations will tell you why.

How this played out in Sales

Our Sales team started with two stated goals defined by the P&L owner: A. Win the AI platform with customers. B. Improve seller experience and satisfaction (to ensure sustainability of the improvement). These objectives translated into measurable metrics that will help drive the transformation and deliver verifiable outcomes. In addition to defining the right KPIs to measure at every step, the dependencies between layers were also made explicit. This allows for data-driven troubleshooting during the transformation—if input metrics are healthy but throughput metrics aren't moving in three to six months, we know we have a habit-formation problem, not a tooling problem.

Measurement framework

Input

Activity-based metrics
Adoption, workflow efficiency, readiness. Who is using AI, where, how often. Impact timeline: 2-4 weeks

Throughput

Leading process indicators
How AI is applied, adopted, integrated. Signals of real workflow penetration. Impact timeline: 1-3 months

Output

Performance measures
Productivity and performance improvements resulting from AI use. Impact timeline: 3-6 months

Outcome

Business impact
The Frontier Firm Success Framework goals—high-level business outcomes. Impact timeline: 6+ months

Example from our Sales team

- AI Sales tool adoption
- User satisfaction (survey)
- % time with customers
- % time spent on AI platform products
- Qualified pipeline
- Win rate
- Revenue per account executive
- Customer success


Invest in
your people

How will you keep your people at the center?

This may be the most important section.

Every recipe in the diffusion engine ultimately succeeds or fails based on people. The most effective companies don't treat this as "change management" on the side—they recognize that capability building, mindset shift, and leadership alignment *are* the transformation.

What surprises many leaders is not that people matter, but that the pace and depth of change required are far greater than anything we've asked of people before. This is a moment to bring employees along deliberately—to equip them, involve them, and build confidence as the work itself changes.

In fact, human capital does not become less valuable as token capital grows. It only becomes more valuable. Humans will set ambitious goals, connect dots across domains, build relationships, and recognize patterns that matter most.

In our transformation work, we've identified four critical areas of focus to support the people-focused transitions required for success:

1. **Engage employees with honesty and optimism**—openly address what's changing, what's not, and the intent behind our leadership decisions
2. **Reimagine our people model** to unlock the full potential of human–AI collaboration
3. **Reinvent change management** into a capability that fuels continuous experimentation and adaptation
4. **Give employees agency** to shape their careers amidst uncertainty

"We found that we initially focused too much on technology alone, when the biggest effort we needed to make is with *people*."

Katy George, CVP, Workforce Transformation, Microsoft

1

Engage employees with honesty and optimism

AI-driven change creates real anxiety, and the response cannot be either denial or doom. People want honesty about what is changing, paired with genuine optimism about where it leads. It is imperative that leaders communicate amidst uncertainty, versus thinking staying silent is safer. It is not. It creates more uncertainty.

The approach we coach our leaders to take:

- Share a strategic view of your business goals and how you are driving your business with people using AI.
- Acknowledge the pace of change and uncertainty it brings.
- Assert you are shaping the future of work, not abdicating to AI (find your agency).
- Be explicit about your intent to create great jobs and career ladders.
- Create both the transparency and the opportunity for employees to transition their skillsets.
- Engage your teams in end-to-end workflow and job redesign —they are the experts! The voice of the employee matters.
- Focus more on adding AI capability into workflows than on simply taking human capacity out.
- Keep critical people-related decisions human.
- Reinforce a culture of continuous learning—what we call our Growth Mindset culture.

Messages that land well

We are helping our employees **future-proof their careers**, not their current jobs.

We acknowledge uncertainty but **our intent is to make jobs more meaningful** with higher impact (more joy/less toil).

We, as leaders, are shaping the future of work with our employees—we don't cede that future to technological inevitability.

2

Reimagine our people model

Roles are becoming more T-shaped

Vertical discipline still matters. But now horizontal integration—connecting your activity all the way to value delivery—is also a baseline expectation. Engineers redefine themselves as "full-stack product builders" rather than "master coders." Sales reps become content integrators rather than relying on specialists to bring expertise. Roles are getting richer and more fulfilling, not narrower.

Manager role-modeling is the single biggest predictor of adoption

In our pilots, manager role-modeling of AI usage is a reliable predictor of team adoption and performance. Managers who are themselves power users drive the greatest impact. Those who delegate AI use to the team do not move their organizations. We coach managers to do three things, leveraging our enduring Growth Mindset culture framework of Model-Coach-Care: role-model AI usage themselves, coach their teams to continuously experiment for exceptional performance, and acknowledge—not dismiss—employee anxiety and concerns about AI.

Higher-order human skills matter more than ever

Judgment and taste. Business acumen. The ability to scope and delegate work—including to agents. Learning agility. We are setting a higher bar sooner on these expectations, and we are adapting our learning approaches accordingly. Our Skill Advisor agent, for example, builds micro-learning prompts directly into the flow of work for our engineers. Our AI Pioneers program engages power users to help mentor colleagues.

+17 pts

lift in recognized AI value **when managers actively model** AI use.

50%

AI users identified **quality control** and …

46%

… **critical thinking** as the most important **human skills** as AI takes on more work.

Source: Work Trend Index, May 2026

2

Reimagine our people model (cont'd)

Teams are becoming more fluid and dynamic

The cadence of high-value, cross-disciplinary collaboration is rising fast. Hierarchy is decreasing. Our Forward Deployed Engineering (FDE) teams, for example, are achieving "Hypervelocity Engineering" by dramatically increasing the cadence of cross-disciplinary collaboration. Our Worldwide Learning team has redesigned its operating model so that what used to be an agile cadence for engineers now includes the entire value stream—engineers, designers, and the go-to-market function—together.

Early career roles should be redesigned, not automated away

The temptation to simply automate junior work is strong. The consequence is hollowing out the apprenticeship pipeline that produces senior talent. We are doing the opposite: compressing the development of judgment and experience by using AI embedded in work processes as a learning tool and creating multi-generational teams to foster two-way apprenticeship. Our PRAISE program—Preceptorship for AI in Software Engineering—pairs new engineers with experienced mentors for two-way learning: the senior teaches craft; the junior teaches AI fluency.

Operating models are shifting from function to value stream

Organizations are growing more dynamic as functional matrices give way to end-to-end value streams, skill-based internal talent marketplaces, and more agile resource deployment—all in service of a simple goal: creating organizations that can move at the speed of AI.

Building flexibility into the operating model is crucial for people and teams to evolve

We recognize that we cannot bring a traditional workforce planning approach to charting our course forward. Instead, we are finding ways to make our operating model more flexible to enable teams to experiment and organically evolve. In particular, we recognize the need to reexamine incentives to be sure we reward impact rather than traditional markers like tenure, and to relax team structures and span/layer guidelines to give teams the space to experiment and change.

49%

of AI work is cognitive support that elevates human effort—analysis, reasoning, and decision-making.

2x

impact on AI effectiveness from organizational interventions than focus on individual mindset and skill.

Source: Work Trend Index, May 2026

3

Reinvent change management

Change management now means creating a new culture of continuous experimentation, adaptation, and learning. The classroom-and-rollout model that worked for past technology cycles is too slow and too shallow for what AI requires. Instead, we are creating learning environments embedded in the flow of work, with peer-level practitioner coaching, and visible hands-on role-modeling from leaders.

We still rely on broad-based measures like proficiency training and storytelling, but we are supplementing these measures with new approaches that are accelerating adoption and best-practice usage, driving system-wide behavior change, and delivering on business performance goals.

The approaches that we find work best share common principles:

- Role and team-based immersive learning
- Peer practitioner-led
- Embedded in the flow of actual work
- Rooted in adult learning best practices (Learn-Do-Teach)
- A permanent part of the new rhythm of business
- Creating a culture of continuous experimentation vs. a one-time change

Train-the-trainer scaleups cement behavior change in the Teach phase

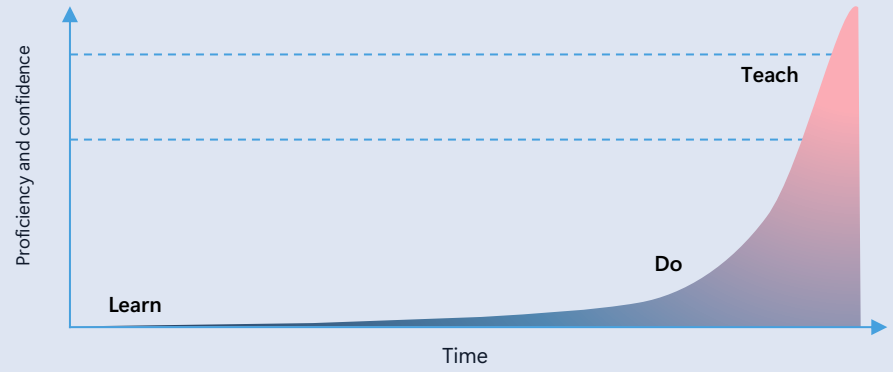

Examples from our own transformation:

A team-based approach:

Camp AIR: We are scaling a three-week immersive bootcamp across our engineering teams, led by practitioners, with time and space for repetition, struggle, and coaching. We tell participants to "leave their disciplines at the door." The cross-functional teams redesign their work together and immediately apply it to their own projects.

A role-based approach:

Sales team weekly huddles: We are making weekly huddles led by peer-level power users a permanent part of the rhythm of business. Sellers in the same role share best practices and are coached by their peers to find new ways to apply our AI sales tools to drive increased impact.

4

Give employees agency

As much as employees want to hear that their leader has agency in the era of AI, they want to feel agency too. We are focused on three areas to give employees more agency.

Create transparency regarding trends in skills and jobs

LinkedIn uses real-time data from millions of roles and hiring trends to highlight which skills are growing in demand and which jobs are emerging. It translates those insights into tools like the Skills Graph, learning recommendations, and job trend data so employees can see where the market is heading. The goal is to help people make informed decisions about how to stay relevant today and prepare for what's next.

Make it easier to find jobs internally with a talent marketplace

A talent marketplace uses AI to dynamically match employees to opportunities based on skills, capabilities, and career goals. Instead of static job ladders, it creates a fluid ecosystem where people can move across projects, teams, and roles as their skills evolve.

With AI, employees no longer need to rely on personal relationships to connect their experiences, skills, and desires to available jobs. With AI-enabled talent marketplaces, companies can create a much more dynamic and accurate "career counseling service" to help employees find new opportunities.

Provide access to learning both online and in the flow of work

Microsoft is embedding learning directly into the flow of work through tools like Viva-Learning, which offers curated content from Microsoft, LinkedIn Learning, and other providers into Microsoft Teams so employees can learn without leaving their workflow.

It also personalizes learning with AI-powered recommendations and integrates features like course sharing, manager suggestions, and calendar scheduling to make skill-building a natural part of the day.

Top 5 fastest growing skills in the US

	Skills	Most common job titles
1	AI literacy	Software Engineer, Product Manager, CEO
2	Conflict mitigation	Customer Service Rep., Admin. Assistant, Project Manager
3	Adaptability	Teacher, Admin. Assistant, Project Manager
4	Process optimization	Operations Manager, Project Manager , CEO
5	Innovative thinking	Creative Director, Teacher, Writer

Source: LinkedIn Labor Market Report—Building a Future of Work That Works, January 2026

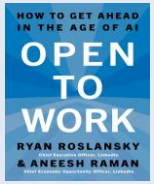

"The future belongs not to those who resist change, but to those who adapt to it—by developing uniquely human skills like creativity, curiosity, courage, compassion and communication"

From *Open to Work*, by Ryan Roslansky and Aneesh Raman


Codify your
advantage & controls

How do you accelerate and protect your secret sauce?

If we had to name one of the most undervalued moves a top team can make this year, it would be this: codify what makes your company distinct in a form AI can learn from—and protect.

Your advantage is often tacit. Every company has a "way"—a point of view on where its market is going, a body of proprietary data and IP, a sense of taste about what good looks like, and a set of risk boundaries that shape decisions. Much of that advantage lives in expert judgment, operating rhythms, and institutional memory rather than in a form AI can be measured against. When AI enters the workflow, that advantage is at risk of being averaged away. Foundation models do not know your standards, your customer philosophy, or where you would never allow a system to act autonomously. Private evals are how you prevent that. They translate your "secret sauce" into evaluable standards—concrete, testable, and repeatable—so AI can be tuned to the way you operate, not the generic average.

Codifying standards is only the first step. The real advantage comes from creating a reinforcement learning environment where your evals define what "good" looks like, your context and harness (the tools and workflows that guide execution) steer the work, and each cycle of feedback, scoring, and tuning improves the system against your standards. Next, align those codified standards to measurable business outcomes – turning what was previously tacit into explicit, accountable definition. We call this building your "hill climbing machine."

Leverage the full frontier ecosystem to build your controls. The intelligence created through your prompts, retrieval, evals, workflows, and agent decisions must stay inside your boundary. That requires architectural choices around data residency, tenant isolation, model flexibility, and IP protection so your knowledge compounds for you rather than leaking into someone else's system. When your advantage lives in your own learning loop, you can swap out any single model without losing what makes you distinct—your IP stays yours, and your standards outlast any one provider.

Private evals define your advantage. The reinforcement learning environment compounds it. Controls protect it. The companies that win will not be the ones that simply adopt AI fastest, but the ones that teach AI how they operate—and keep that intelligence inside their moat.

"If you don't define what 'good' looks like, AI defaults to something generic. Private evals are how you capture your standards so AI works the way your company does. That's how AI becomes a compounding advantage."

Charles Lamanna, EVP, Copilot, Agents, and Platform, Microsoft

Four dimensions to codifying your secret sauce

Point of view

What is your distinct thesis about where your market is going and what your customers truly value?

AI must sharpen this view, not flatten it into generic insight.

Performance

What is the proprietary data, IP, and benchmark you have that competitors don't—the workflows, relationships, telemetry, and institutional judgment customers pay you for?

AI should multiply these, not bypass them.

Taste

What does "good" look like across your most important outputs—a sales conversation, a piece of code, a customer service moment, a creative product?

Models without taste produce work that is generic and forgettable.

Guardrails and risk

Where will AI act autonomously, where will a human review, and where should not be deployed at all?

AI autonomy must be tuned to the stakes of each decision, not set globally.

The law of hill-climbing

Any enterprise whose competitive advantage cannot be measured by its own evals, improved through its own expertise, and preserved across any foundation model will be arbitraged.

The companies that will lead in the AI era will not simply be those with access to AI. They will be the ones that build systems that learn, improve, and adapt in line with how their business operates.

Your hill-climbing machine is that learning system: a reinforcement environment in which your evals define what "good" looks like, your context and orchestration guide the work, and the underlying models can evolve over time. With every cycle of feedback, scoring, and tuning, the system becomes better aligned to your standards rather than converging on generic outputs. Hill-climbing embeds a growth mindset: progress comes from iteration, where every output—good or bad—feeds the next improvement cycle.

This architecture matters because it clearly separates what you own and differentiate—your evals, context, and control layer—from what you consume externally, such as foundation models. As models continue to change, this approach ensures each new generation strengthens your company's intelligence, not just its speed.

It is also central to protecting your intellectual property. Without deliberate architectural choices, the intelligence embedded in your prompts, workflows, and decisions can drift beyond your control. A well-designed hill-climbing machine keeps your tacit knowledge, learning loops, and enterprise intelligence inside your tenant, your boundary, and your control.

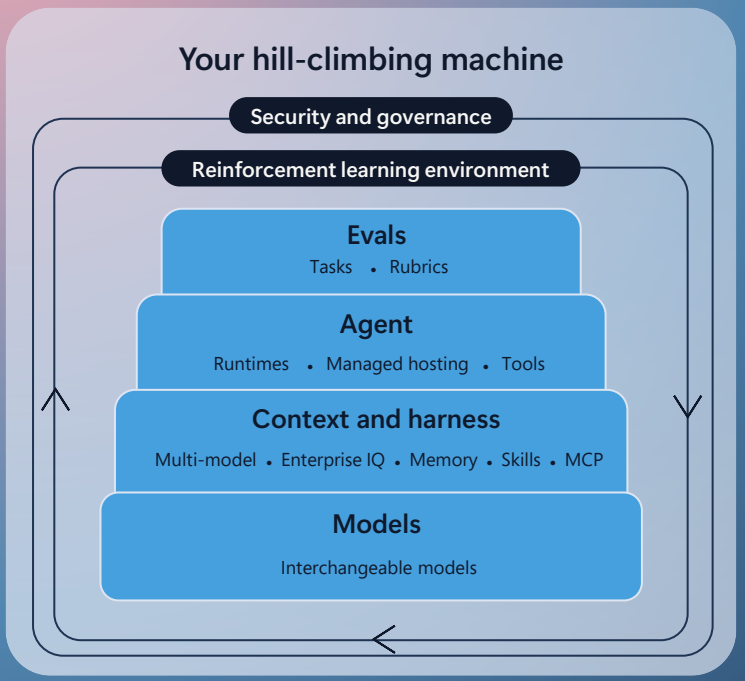

Considerations

- Your outcomes are private evals
- Context and harness are decoupled from models
- Flexibility in models
- Token efficiency controls
- Tacit knowledge in weights

In the AI era, control comes down to one decision: whether your intelligence compounds inside your boundary or leaks beyond it.

Every prompt, every retrieval, every eval, every workflow, every agent decision your company runs is intelligence. It is the fingerprint of how you think, how you sell, how you build, how you decide. The single most important architectural choice your top team makes about AI is whether that intelligence stays inside your boundary or leaks out of it.

Effective control requires ensuring all policy requirements have architectural consequences. Retrofitting data residency and tenant isolation after deployment is orders of magnitude harder than building them in from the start.

“Increasingly, customers want to be able to capture the tacit knowledge they possess inside of model weights as their core IP. This is probably the most important sovereign consideration for firms as AI diffuses more broadly across our GDP and every firm needs to protect their enterprise value.”

Satya Nadella, Chairman and CEO, Microsoft

Four factors to consider when building and safeguarding your hill-climbing machine

Hill-climbing infrastructure

Have you built the eval, scoring, and experimentation environment needed to measure outputs against your standards?

Do context, harness, models, and feedback operate in a closed loop so each cycle improves the system over time?

Can that learning system work across changing models without an architecture rebuild?

We are building the infrastructure for a hill-climbing machine: a learning system where evals define "good," context steers the work, and feedback improves performance over time.

Data residency

Are prompts, retrieval stores, outputs, and evals staying inside the geography and boundary you require?

Are regulated workloads supported by sovereign cloud options where compute and storage remain in-country?

Are AI inputs and outputs classified before they reach the model?

We are treating data residency as a design principle: AI systems should be architected so enterprise intelligence is handled within the geographic, tenant, and regulatory boundaries appropriate to the work.

Model independence

Can you change models without rebuilding the architecture around them?

Are data-residency controls applied consistently across every foundation model you use?

Are you avoiding lock-in to a single model provider as model performance changes?

We are designing for model flexibility, so organizations can benefit from model progress while maintaining architectural control and clear data-boundary expectations.

Tenant isolation

Does the same enterprise perimeter govern humans and agents?

Can AI reason over enterprise context without that context crossing the tenant boundary?

Do identity, access, compliance, and policy controls apply to agents the same way they apply to users?

We are treating the tenant as a core control plane for AI, so human and agent activity can be governed through consistent identity, access, and compliance patterns.


Safeguard your
security

Is your security ready to scale at the speed of AI?

AI is reshaping how enterprises operate, making decisions and taking action at machine speed.

AI is also reshaping the security landscape itself, accelerating vulnerability discovery, dynamically adapting attacks, and overwhelming human-centric security models.

Without rethinking their security approach, organizations will be exposed and face a stark choice: slow AI adoption to manage risk, or scale AI faster than they can control. Neither path is sustainable.

"As we embark on one of the most significant transformations in our lifetime, realizing the astonishing potential of AI will only succeed if we can secure AI solutions and make them safe."

Hayete Gallot, EVP, Security, Microsoft

Four dimensions to safeguarding security

Defend with AI

When attacks are happening at machine speed, how do we defend at machine speed?

How do we move beyond human-centric defenses that can no longer keep up?

Shift defense from reactive operations to continuous, adaptive AI protection, powered by a powerful Security Graph.

There is no success without rethinking the Security operating model.

Safely enable AI

How do we safely enable agents that are becoming independent actors with broad access and impact?

How do we evolve security systems that were originally designed for humans, not polymorphic agents?

Leverage end-to-end system to give agents identity, risk-based permissions and govern the full lifecycle, to prevent misuse.

There should not be a choice between innovation and safety.

Trust AI with your data

How do we maintain control of our data so it strengthens our AI rather than our competitors'?

How do we address the growing data debt that undermines our AI advantage?

Use AI to accelerate data classification and remove posture drift across your entire data estate.

There is no trusting AI without securing your data.

Get ready for AI

Are we prepared for human-led attacks that test response speed?

Are we prepared for agent-led attacks that exploit any remaining vulnerabilities?

Close the door before agents find it. Focus on exposure management including patching vulnerabilities, identity, network isolation...

There is no durable defense without posture management.

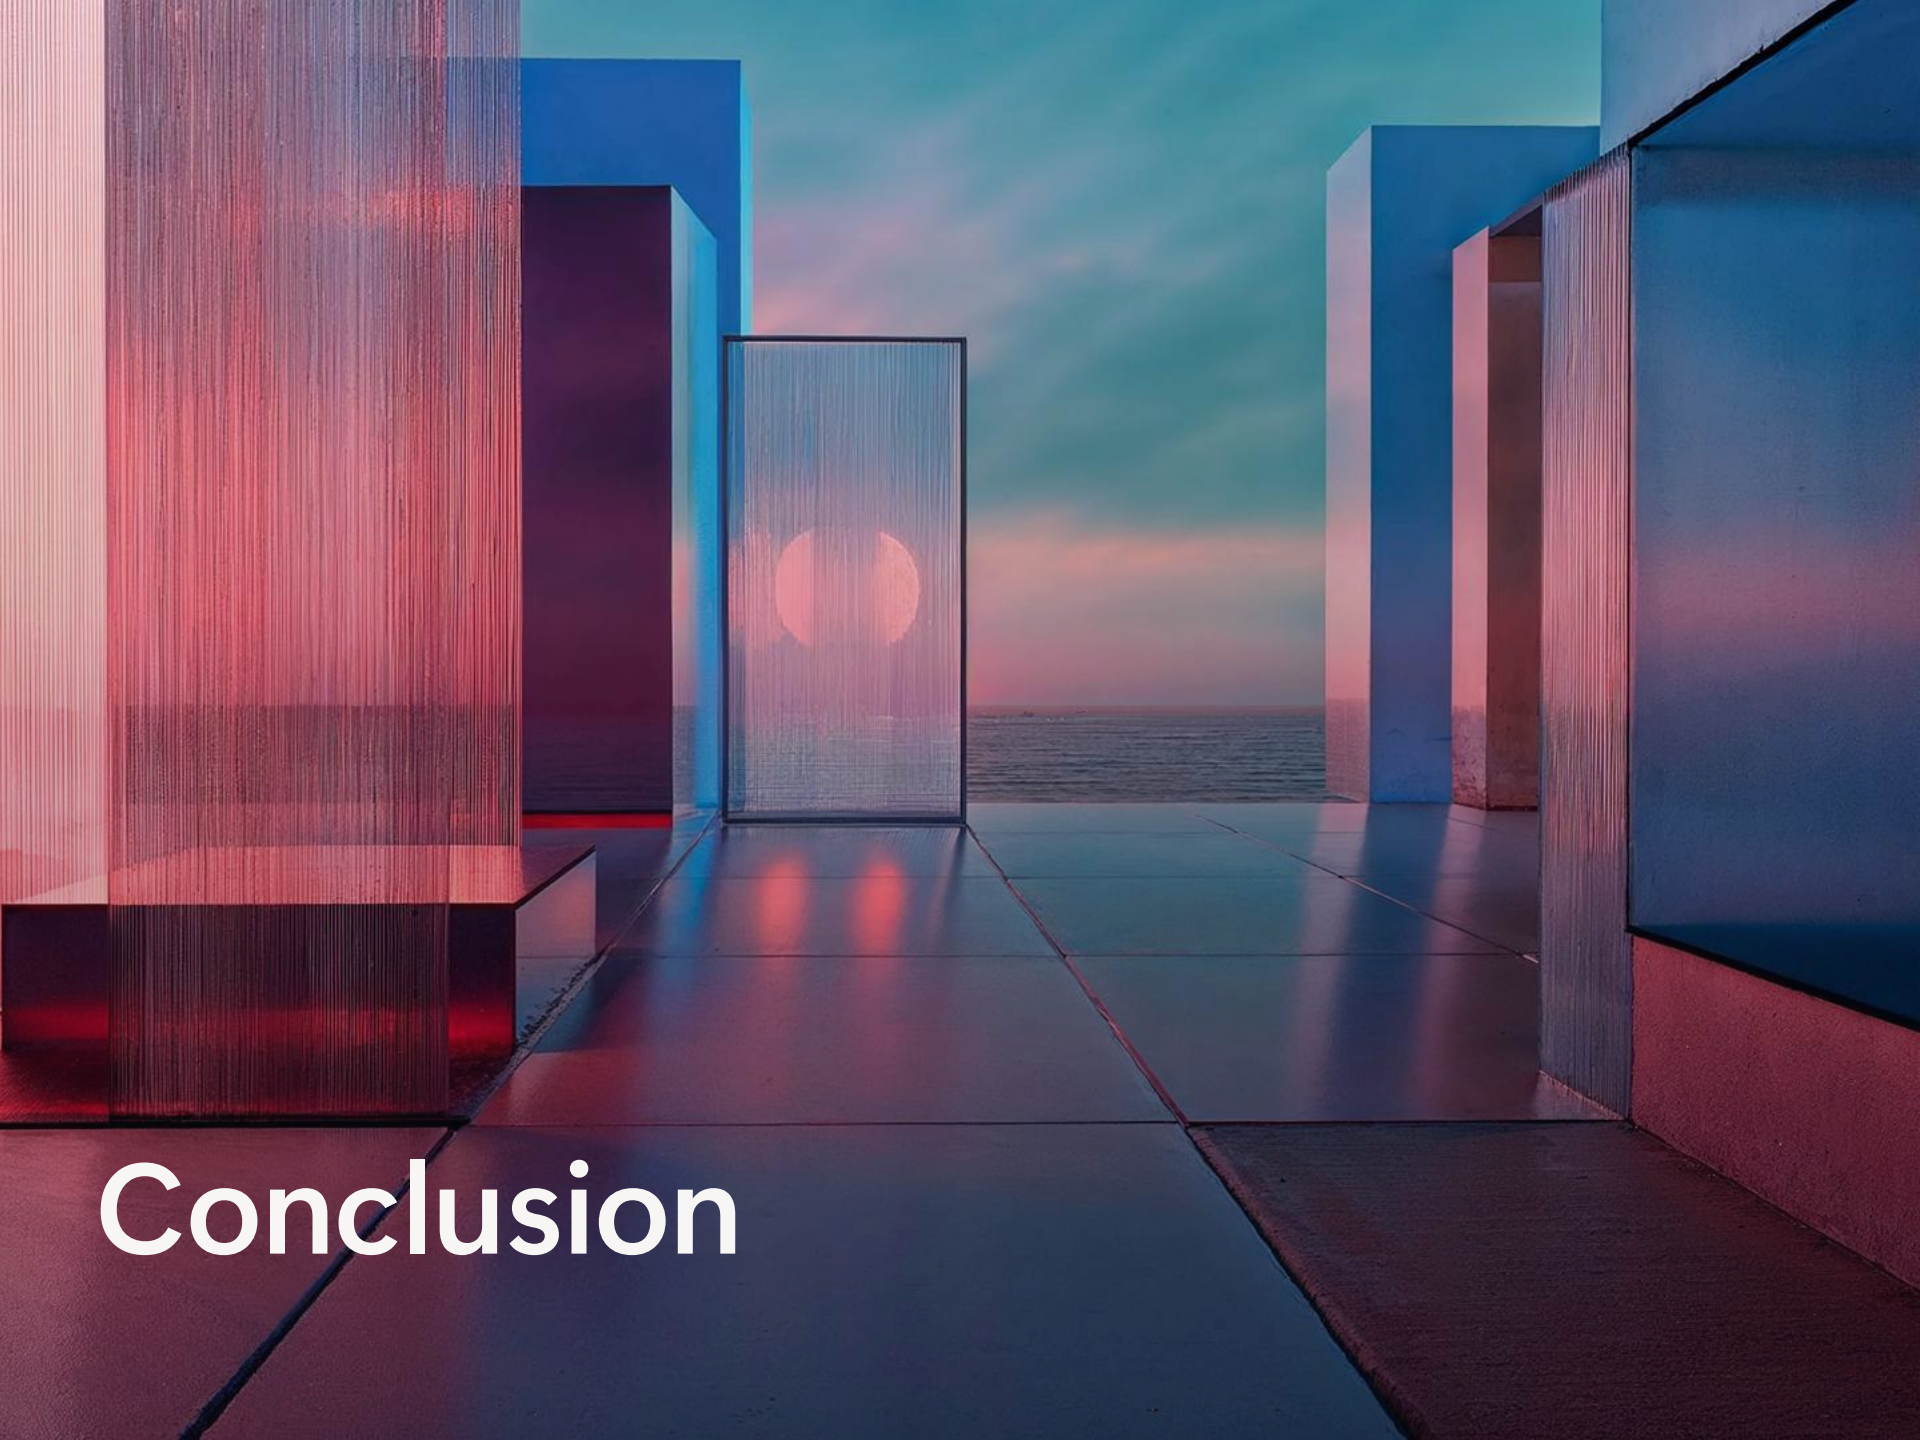

Conclusion

Conclusion

If you have made it this far into the Playbook, it's clear you are as serious about your AI transformation as we are! We hope that the ideas and insights that we are implementing ourselves have resonated with you and will help you take your impact to the next level.

As a summary, we recommend five elements to drive lasting business outcomes:

A. **Set business ambition and goals:** Anchor to enterprise value, not just productivity. Efficiency is the floor. Experience, innovation, and growth are the ceiling.

B. **Build your diffusion engine:** Tools don't change companies; people and operating systems do. Be intentional about your approach, with a focus on using AI to build the new capabilities that will enable you to meet new levels of excellence against your goals.

C. **Invest in your people:** Engage employees with honesty and optimism, reimagine your people model in a human-AI world, reinvent change management, and give employees agency to shape their careers.

D. **Codify your advantage & controls:** Codify your secret sauce before you scale. Private evals are the rubric the rest of the program is scored against. Set up the boundaries and moats that you need to protect your competitive edge.

E. **Safeguard your security:** Set up guardrails to apply to all agents and ensure AI actions are traceable, auditable, and reversible.

We are still early in this journey. We are still learning. We will get some of this wrong, and you might too. The point of this Playbook is not to claim a complete answer; it is to share the working method that has gotten us this far.

If you have suggestions and new ideas that we have not considered, we would love to hear from you! Please find us on LinkedIn and share your suggestions. We look forward to learning what you discover that we haven't, and to continue to chart the path to the frontier—together!

Kathleen Hogan
Executive Vice President,
Strategy and Transformation
LinkedIn

Katy George
Corporate Vice President,
Workforce Transformation
LinkedIn

"What we're learning inside Microsoft is that business transformation is hard, because it requires people to fundamentally rethink how they work. Reaching the frontier takes courage, intention, and sustained leadership. But if leaders are clear about their intent, bring people along, and give them real agency, people rise to the challenge—they reinvent how work gets done and unlock capabilities once thought impossible. AI then empowers every person to achieve far more than we once imagined possible."

Kathleen Hogan, EVP, Strategy and Transformation, Microsoft

For more

Frontier Firm resources

- For more Frontier Firm resources and insights, scan or click the QR code or go to: https://www.microsoft.com/worklab
- If you're interested in exploring what the Frontier Firm journey could look like for your organization's AI Transformation, we'd welcome the opportunity to convene a CXO-level discussion. Email executiveAI@microsoft.com, and a Microsoft team member will follow up to coordinate next steps.