

Innovate at the Frontier

An Executive Framework for Scaling AI

A practical operating model for boards, CEOs, and senior leaders who want to move beyond AI experiments and turn intelligence on tap into sustained business performance.

David Beauchemin

linkedin.com/in/davidbeauchemin

Foreword

The future will not be led by the organizations with the most AI. It will be led by the ones that redesign how work, judgment, and intelligence come together.

For most of business history, intelligence was scarce and expensive. It lived inside people, and it scaled only as fast as we could hire, train, and retain them. That constraint is gone. Intelligence is now abundant, affordable, and available on demand, behaving more like electricity than like talent. It will power businesses the way electricity powered the last century of growth.

I wrote this whitepaper for executives who feel the pressure of three forces at once: boards demanding AI-driven impact, capabilities evolving faster than operating models, and regulators and stakeholders asking for assurance that this is being done responsibly. I have been working with applied AI since 2017, and have spent the years since working with executive teams across industries on Cloud & AI transformation. The pattern is consistent: the leaders who pull ahead are not the ones with the most pilots. They are the ones who treat the Frontier Firm as a practice rather than a destination.

This document distills what I have learned into one framework. It draws on independent research, on industry research, on Microsoft's own public learnings, and on the operating model my colleagues and I have shared with hundreds of executive audiences. My hope is that you can read it on a flight, share it with your leadership team, and walk into Monday morning with a plan you can actually fund.

David Beauchemin

Contents

1.	Executive Summary	04
2.	The Frontier Imperative: Intelligence on Tap	05
3.	Meet the Frontier Firm	07
4.	A New Operating Model: Three Patterns Working Together	09
5.	Iterate to Quality	12
6.	Responsible by Design	13
7.	Proof Points: Frontier Firms in Action	15
8.	The Monday Morning Plan	18
9.	Lead What's Next	20
10.	Bibliography and References	21

1. Executive Summary

AI works. The unanswered question is whether your organization can turn it into performance.

The evidence is now consistent across independent sources. McKinsey estimates that generative AI alone could add between USD 2.6 trillion and USD 4.4 trillion in annual value to the global economy.¹ Microsoft's 2025 Work Trend Index finds that 82 percent of leaders consider 2025 a pivotal year to rethink strategy and operations.² Anthropic's 2026 State of AI Agents Report finds that 80 percent of organizations already report measurable economic returns from their AI agent investments, and 88 percent expect that ROI to continue or grow in 2026.³ Bain finds that more than 80 percent of generative AI use cases are meeting or exceeding expectations.⁴

The same data shows where the value comes from, and it is not from chatbots. It comes from organizations that scale AI across the enterprise, embed it in how work flows, and learn faster than their competitors. BCG's "AI at Work 2025" study, which surveyed 10,635 employees across 11 countries, is blunt about what separates the leaders: real value comes from companies that redesign entire workflows, not from those that simply deploy tools.⁵ Microsoft's research describes these organizations as Frontier Firms: human-led, agent-operated enterprises that buy intelligence like electricity, put it to work, and compound it like interest.²

This whitepaper proposes one operating framework for getting there. The framework rests on five propositions, each grounded in public research and customer evidence:

1. **Intelligence on tap is now an executive resource.** The capacity gap between what businesses demand and what humans can sustainably deliver is widening, and digital labor is the lever that closes it.²
2. **The Frontier Firm is a practice, not a place.** There is no finish line. The edge comes from continuous visibility, guardrails, and a learning flywheel that compounds improvement.
3. **Three patterns must run together.** Elevating individual productivity, redesigning the flow of work, and AI-first work design are not sequential phases. Run in parallel on the same workflow, they produce multiplier benefits. Run alone, they produce local optimization at best.
4. **Quality is engineered, not assumed.** Frontier Firms set a quality contract before they build, fix workflow debt before they automate, and iterate in production with measurement.
5. **Responsible AI is a design constraint, not a review step.** Governance, escalation paths, and inclusive design must be built in from the beginning. EY's Responsible AI Pulse Survey finds that 99 percent of organizations have already experienced AI-related financial losses, with an average loss of USD 4.4 million per impacted company.⁶

The whitepaper closes with a Monday Morning Plan: five concrete actions any executive can take this week to fund the operating model rather than fund the technology.

2. The Frontier Imperative: Intelligence on Tap

Intelligence used to be scarce and expensive, locked inside humans. It is now abundant, affordable, and scalable, more like electricity, and it will power businesses.

2.1 The shift from scarce to abundant intelligence

For decades, intelligence was one of the most valuable and most limited assets in business, bounded by human time, energy, and cost. AI changes the equation. As Microsoft's 2025 Work Trend Index puts it, intelligence is becoming an essential durable good: abundant, affordable, and available on demand.² The leadership shift is consequential. We can now access intelligence outside of humans, through agents that can reason, plan, and act. Leaders must learn to manage and measure that intelligence the way they manage and measure capital, headcount, and technology.

The macro stakes are large. McKinsey estimates that generative AI alone could add between USD 2.6 trillion and USD 4.4 trillion in annual value to the global economy, with about 75 percent of that potential clustered in customer operations, marketing and sales, software engineering, and research and development.¹ PwC's 2025 Global AI Jobs Barometer, drawing on data from nearly one billion job ads across six continents, finds that productivity growth in industries most exposed to AI has accelerated from 7 percent (2018-2022) to 27 percent (2018-2024), nearly a fourfold increase since the rise of generative AI.⁸ The economy is repricing the value of speed and judgment in real time.

2.2 The capacity gap is real and widening

Boards are asking for more output, but the workforce is signaling that it cannot deliver more without help. Microsoft's data quantifies the gap: 53 percent of leaders say productivity must increase, while 80 percent of the global workforce, both employees and leaders, say they lack the time or energy to do their work.² The same study finds employees are interrupted every two minutes during the workday, accumulating roughly 275 interruptions per day when activity outside core hours is included.²

82%

of leaders say 2025 is a pivotal
year to rethink strategy²

53%

of leaders say productivity must
increase²

80%

of the global workforce lacks
time or energy to do the work²

This is the executive opportunity in plain language. Intelligence on tap can fill the capacity gap, but only if it is connected to work in a way that improves decisions and actions. Capability without integration is theatre.

2.3 Three forces hitting executives at once

In every executive conversation I have, the same three forces show up together:

- **Pressure to perform.** Boards want measurable AI impact, not pilots.
- **Speed of change.** Capabilities are evolving faster than operating models can absorb them.
- **Risk without guardrails.** Regulators, customers, and employees want assurance that this is being done responsibly.

When change outruns the organization, the default response is patchwork. Patchwork is where most AI initiatives stall. The way out is not more pilots. It is a new operating model.

2.4 From data to decisions: judgment is the scarce asset

More data will not solve the intelligence gap. We have largely solved for language tasks: drafting, summarizing, and explaining. We have not solved for reasoning. Reasoning does not come from the model alone. It comes from how well an organization captures context, decisions, and actions, and embeds them into workflows.

Human value is the source code of enterprise value because judgment is the scarce asset. The question is whether we are capturing the judgment of our people.

2.5 The independent evidence: agents are no longer experimental

The Anthropic 2026 State of AI Agents Report, conducted with research firm Material among more than 500 technical leaders, makes the maturity curve concrete. Fifty-seven percent of organizations now deploy agents for multi-stage workflows, with 16 percent running cross-functional processes across multiple teams. By 2026, 81 percent plan to tackle more complex use cases, and 80 percent already report measurable economic returns from their AI agent investments.³ Bain's research points in the same direction: more than 80 percent of generative AI use cases now meet or exceed expectations, even as scaling remains the primary hurdle.⁴

WHY THIS MATTERS NOW

Agents have moved from experimental tools to essential infrastructure. The leaders who invest in operating models that can absorb, govern, and scale them will widen the gap over the next 24 months. The leaders who treat agents as a procurement decision will not.

3. Meet the Frontier Firm

Frontier Firms are not defined by industry or size. They are defined by execution.

3.1 The three traits

Microsoft's research identifies a clear pattern in the organizations pulling ahead. They go wide, they go deep, and they embrace agents.²

TRAIT 1

Wide: AI scaled across the entire business

Frontier Firms do not run AI as a function. They run it as a capability that touches every part of the enterprise, from customer service to finance to engineering to go-to-market.

TRAIT 2

Deep: AI applied to industry problems

Generic productivity is a starting point, not a strategy. Frontier Firms apply AI to the high-value, often messy, problems unique to their industry: claims, underwriting, drug discovery, supply chain, regulated workflows.

TRAIT 3

Agentic: embracing systems that reason, plan, and act

Agents are not chatbots. They are systems that can complete multi-step work with human guidance. Frontier Firms design human-agent teams from the start, not as an afterthought.

3.2 Frontier Firms are already outperforming

Among the 9,037 leaders who participated in Microsoft's 2025 Work Trend Index, 844 work at companies that meet the Frontier Firm bar (defined by org-wide AI deployment, advanced AI maturity, current and projected agent use, and a belief that agents are key to AI ROI). The performance gap is striking:²

- 71 percent of Frontier Firm leaders say their company is thriving, compared to 37 percent globally.
- 55 percent say they can take on more work, compared to 20 percent globally.

- 93 percent are optimistic about future work opportunities, compared to 77 percent globally.

3.3 The Frontier is a practice, not a place

There is no finish line for AI adoption. Frontier Firms operate in perpetual beta, but not as chaos. They run three things together as a discipline:

- **Visibility** into work as it actually happens.
- **Guardrails** that keep speed safe.
- **A learning flywheel** that compounds improvement.

Every process becomes a hypothesis. Every improvement becomes fuel for the next. Scaling is not a free-for-all. It requires structured experimentation, reliable metrics, and rigorous governance so that innovation stays responsible.

On our journey to the Frontier Firm... we approach this work with a growth mindset, as learners first. Like every organization, we are still at the beginning of this journey...

Kathleen Hogan, Executive Vice President, Chief Strategy and Transformation Officer, Microsoft¹⁰

4. A New Operating Model: Three Patterns Working Together

Redesigning with AI does not begin with technology. It begins with understanding how people actually work. You cannot transform what you cannot see.

For years, information work unfolded out of sight, making it difficult to measure or improve. Unlike the output of an assembly line, you cannot watch someone negotiate a contract or draft a launch strategy. (Well, you can, but you will miss the real action.) That invisibility is the first problem to solve.

Frontier Firms run three patterns in parallel, not in sequence. If you only do one, especially if you default to chatbots, you get local optimization, confusion, and limited impact. Running all three together is what produces multiplier benefits, because you are redesigning how work, judgment, and intelligence come together.¹⁰

BCG's "AI at Work 2025" study, the largest of its kind to date, makes the point with independent data. Real value comes when companies redesign entire workflows, not when they merely deploy tools, and only 13 percent of organizations report having AI agents integrated into their workflows even though 75 percent of leaders believe agents will be vital to future success.⁵ The gap between intent and execution is the executive opportunity.

PATTERN 1

Elevate Individual Productivity: start where work happens

The fast on-ramp. This pattern is especially effective for language-heavy work such as drafting, summarizing, and classifying.

HOW TO DO IT

- **Expose how work actually happens.** Use day-in-the-life or week-in-the-life mapping to surface decisions, handoffs, exceptions, and where time goes.
- **Define escalation paths and boundaries.** Use frameworks such as RACI so that accountability and trust are explicit, not implied.
- **Turn adoption into habit.** Use rituals, nudges, and peer reinforcement so that behaviour change actually sticks.

Individual productivity is necessary, but on its own it will not fix end-to-end friction. That is the job of Pattern 2.

PATTERN 2

Redesign the Flow of Work: from fragmented processes to systems

The second pattern moves the conversation from individuals to processes. It is where most of the cycle-time gains live, because most knowledge work is delayed by handoffs, not by the work itself.

HOW TO DO IT

- **Map business processes and flows.** Use Gemba walks, process mining, value-stream mapping, and task decomposition to understand how work actually runs.¹⁰
- **Redesign flows and set guardrails with escalations.** Tighten handoffs, clarify decision rights, and make the workflow executable rather than informal.
- **Define human-agent teams.** Be explicit about what agents execute, where humans judge, and how exceptions are handled.

Pattern 2 is where speed becomes safe and where most organizations should expect the most measurable ROI in the first year.

PATTERN 3

AI-First Work Design: zero-based, embedded intelligence

The most powerful pattern, and the hardest. Pattern 3 assumes intelligence is embedded by default, and questions every assumption about roles, steps, and handoffs. It is best for higher-stakes, multi-step work where you need agentic support with clear escalations and human judgment preserved.

HOW TO DO IT

- **Design with embedded intelligence.** Start from a zero base. Add people and agents back in as needed, rather than starting from the legacy process and trying to simplify down.¹⁰
- **Be safe by design.** Build governance in from the start, not bolted on later. Use sandbox governance and policy-as-code so that teams can experiment safely without putting the enterprise at risk.
- **Build a compounding learning system.** Pilot, test, codify, and reuse patterns so that the organization learns faster over time.

4.1 The discipline that ties the three patterns together

Across all three patterns, the same disciplines apply. The recipe is consistent, the ingredients change with the pattern.

Discipline	Pattern 1: Elevate	Pattern 2: Redesign	Pattern 3: AI-First
Make work visible	Persona journeys	Gemba, process mining, task decomposition	Zero-based redesign
Define risk boundaries	Escalation paths	Domain-specific guardrails	Sandbox governance
Design the AI solution set	Human with assistant	Human-agent teams	Human-led, agent-operated
Redesign roles and structure	Adoption nudges	Job redesign	Agentic teams
Continuous experimentation	Habit building	A/B testing	Innovation flywheel

Adapted from Microsoft's "Becoming a Frontier Firm" executive playbook¹⁰ and BCG's "AI at Work 2025" study.⁵

5. Iterate to Quality

Most organizations make a predictable mistake: they place a chatbot on top of a broken workflow and amplify confusion. Frontier Firms do the opposite.

Frontier Firms treat low early quality as a stress test. When quality is low, it reveals workflow debt and shadow work, the silent patches people apply with tribal knowledge to keep things moving. The discipline is to surface and fix that debt before scaling, not after.

5.1 The five-step quality discipline

Step 1	Set the Quality Contract. Define what good looks like before you build, in business terms. Specify where judgment must remain human, what acceptable error rates look like, and when to escalate.
Step 2	Illuminate the real workflow. Surface hidden and patched work, delays, and their costs. Ask the people doing the work, not just the people who documented it.
Step 3	Fix workflow debt first. Remove the ambiguity that AI amplifies. Tighten inputs, clarify handoffs, and separate the golden path from exceptions.
Step 4	Iterate in production with measurement, not in isolation. Track cycle time, rework, overrides, and escalation effectiveness. Close the loop in weeks, not quarters.
Step 5	Codify what works so learning compounds. Build reusable patterns, checkpoints, and exception handling so that the next team starts ahead of the last one.

5.2 The principle behind the discipline

Frontier Firms do not experiment until something works. They iterate with intent, until quality is proven, judgment is preserved, and learning compounds across the enterprise. That is the difference between a portfolio of pilots and an enterprise capability.

PRACTITIONER'S NOTE

The single biggest accelerator I see in customer engagements is not a model upgrade. It is a written Quality Contract that says, in business language, what good looks like before anyone writes a prompt or buys a license. Without that contract, every conversation about AI quality becomes a debate about feelings. The independent data backs this up: Bain finds that more than 80 percent of generative AI use cases meet or exceed expectations once they reach production, but warns that scaling remains the primary hurdle.⁴ The discipline below is what closes the gap between a working pilot and a scaled capability.

6. Responsible by Design

Trust is the throughput constraint on AI scale. Without it, you cannot get from pilot to production at the pace the market is moving.

6.1 The trust gap is structural, and it is widening

The most authoritative independent measure of public trust in AI today is the 2025 KPMG and University of Melbourne global study "Trust, Attitudes and Use of Artificial Intelligence," which surveyed more than 48,000 people across 47 countries. The findings should reorder boardroom priorities. Sixty-six percent of people now use AI regularly, and 83 percent expect benefits, yet only 46 percent are willing to trust AI systems, and 70 percent want stronger regulation.⁷ The gap between adoption and trust is the strategic risk most enterprises are underpricing.

66%

use AI regularly worldwide⁷

46%

willing to trust AI systems⁷

70%

want AI more strongly
regulated⁷

6.2 Governance is a design constraint, not a review step

The cost of a single visible failure (a biased model, a hallucinated output in front of a customer, a leaked document) is now measurable in shareholder value, not just risk-register colour. EY's Responsible AI Pulse Survey of senior leaders at large organizations puts hard numbers on this. Ninety-nine percent of organizations report having experienced AI-related financial losses, with the average loss reaching USD 4.4 million per impacted company. Only 14 percent of CEOs say their AI is fully compliant with current regulations. Yet 97 percent of organizations that have implemented Responsible AI principles also report positive ROI from AI investment, a correlation that should turn governance from a cost line into a value lever.⁶

The implication is straightforward. Governance, escalation paths, and audit controls must be built into the operating model from day one, not bolted on after the fact. In the three-pattern framework, this is why "define risk boundaries" sits second only to "make work visible." The order matters.

6.3 The workforce dimension: skills, wages, and dignity

Responsible AI is not only about model risk. It is also about the people whose work AI reshapes. PwC's 2025 Global AI Jobs Barometer, drawing on nearly one billion job ads across six continents, finds that workers with AI skills now command a 56 percent wage premium over their peers, and that the skills employers seek are changing 66 percent faster in jobs most exposed to AI.⁸ BCG's "AI at Work 2025" finds the parallel risk: while

72 percent of employees overall now use AI regularly, frontline use has stalled at 51 percent, creating what BCG calls a silicon ceiling that will undermine adoption if executives do not actively invest in frontline enablement.⁵ Boards that ignore reskilling and frontline enablement will pay the bill in attrition, talent acquisition, and reputational damage.

6.4 Inclusion is part of governance, not a parallel conversation

AI systems learn from the past and make decisions that affect the future. If the populations they serve are not represented in their design, training, or testing, the systems will reproduce and amplify exclusion at unprecedented scale. Independent Canadian research underscores how uneven the practice still is. The 2026 QueerTech Inclusive AI Research Report, conducted with the support of Women and Gender Equality Canada, surveyed 100 AI product developers across the country. The headline finding: 65 percent of Canadian technology companies say their products meet the needs of the general population, but only 47 percent believe the same for 2SLGBTQI+ users, a gap of nearly 20 percentage points.⁹

97%

say inclusive AI is a moderate-
to-high priority⁹

68%

report only limited
organizational support for it⁹

39%

cite insufficient resourcing as
the top barrier⁹

The same study finds the top barriers cited by developers are insufficient resourcing (39 percent), competing priorities (36 percent), and difficulty measuring ROI on inclusive practices (33 percent).⁹ All three are executive decisions, not engineering decisions. The lesson for boards is direct: if inclusion is not funded, sponsored, and measured, it will not survive contact with delivery pressure.

6.5 Human judgment as the differentiator

The Frontier Firm is human-led, agent-operated. That phrase is not decorative. It encodes a design principle: agents scale execution; humans set direction, exercise judgment, and own consequences. Microsoft's research finds that 46 percent of leaders report using agents to fully automate workflows or processes, yet even in Frontier Firms the most demanding work, the work rooted in judgment, empathy, and creative thinking, remains human.²

*Human judgment remains the differentiator. We scale intelligence
while keeping accountability explicit.*

7. Proof Points: Frontier Firms in Action

Every figure cited in this section is drawn from a publicly disclosed customer story or earnings disclosure. Source links are listed in the bibliography.

7.1 The three patterns in the wild

The case studies below are organized by the dominant pattern each illustrates: elevating individual productivity, redesigning the flow of work, and AI-first work design. Frontier Firms run all three together. These examples show what each one looks like once it has cleared the boardroom and started producing measurable value.

7.2 Pattern 1: Elevating individual productivity

MCKINSEY (LILLI): A FIRM-WIDE GENAI ASSISTANT FOR 43,000 CONSULTANTS

McKinsey built Lilli, an internal generative AI assistant trained on the firm's 100,000-plus documents and interview transcripts, and rolled it out across its global workforce. More than 70 percent of the firm's more than 40,000 employees use Lilli on a monthly basis, with active users averaging 17 prompts a week. McKinsey reports Lilli is saving up to 30 percent of users' time on knowledge-intensive tasks, freeing consultants to focus on judgment and client work.

Source: McKinsey & Company, "Lilli: McKinsey's generative AI tool" (firm publication, 2024-2025).¹¹

MICROSOFT COMMUNICATIONS: PERSONA-LED ACCELERATION

Microsoft's Communications team mapped a day-in-the-life for communications professionals to surface recurring rhythms and bottlenecks, then deployed a custom Communications Copilot for repeatable tasks (social posts, story creation, FAQs) and an AI-driven media measurement tool inside Microsoft Teams. Reported outcomes include increased impact and efficiency on core tasks, 100 percent daily usage for mandated tasks, and a reduction in vendor spend.

Source: Microsoft, "Becoming a Frontier Firm" executive playbook, 2025.¹⁰

7.3 Pattern 2: Redesigning the flow of work

KLARNA: AI ASSISTANT DOING THE WORK OF 700 FULL-TIME AGENTS

Klarna's OpenAI-powered customer service assistant handled 2.3 million conversations in its first month live, the equivalent of the work of 700 full-time agents. The company reported it was on par with human agents on customer satisfaction, more accurate in errand resolution leading to a 25 percent drop in repeat inquiries, and was completing customer queries in under 2 minutes versus 11 minutes previously. Klarna estimated the assistant would drive a USD 40 million profit improvement in 2024 and made the service available in 35-plus languages across 23 markets.

Source: Klarna, "Klarna AI assistant handles two-thirds of customer service chats in its first month" (press release, 27 February 2024).¹²

CAPITA: TRIAGE AGENT CUTS RESPONSE TIMES BY 60 PERCENT

Capita built an agent in Microsoft Copilot Studio that triages thousands of emails a day, cutting response times by 60 percent and freeing employees to focus on empathy and resolution.

Source: Microsoft Customer Stories, "Capita."¹³

DOW: MULTI-MILLION-DOLLAR SUPPLY-CHAIN SAVINGS

Dow is projected to save millions in the first year alone with a supply-chain agent that flags misapplied shipping fees, demonstrating how agents can act as continuous auditors of high-volume transactional work.

Source: Microsoft WorkLab, "AI impact at Dow."¹⁴

ESENTIRE: THREAT ANALYSIS FROM 5 HOURS TO 7 MINUTES

Cybersecurity firm eSentire compressed expert threat analysis from 5 hours to 7 minutes using Anthropic's Claude, with AI-driven analysis aligning with senior security experts 95 percent of the time.

Source: Anthropic customer story, eSentire.¹⁵

7.4 Pattern 3: AI-first work design

MODERNA: 750 CUSTOM GPTS AND A WORKFORCE THAT BUILDS ITS OWN AI

Moderna integrated ChatGPT Enterprise across the company and built more than 750 custom GPTs to support specific workflows in just two months. Its mChat assistant reached 80 percent adoption, 40 percent of weekly active users build their own GPTs, and the legal team alone reached 100 percent adoption. Moderna's stated ambition is to operate as a fully AI-augmented life-sciences company.

Source: OpenAI, "Moderna and OpenAI: Reimagining how we work, together" (joint announcement, April 2024).¹⁶

THOMSON REUTERS: DECADES OF CASE LAW ACCESSIBLE IN MINUTES

Thomson Reuters uses Claude to power CoCounsel, its AI legal platform. Lawyers who once spent hours manually searching documents can now access more than 150 years of Westlaw case law, curated by 3,000-plus subject-matter experts, in minutes.

Source: Anthropic customer story, Thomson Reuters.¹⁷

DOCTOLIB: FEATURES SHIPPED 40 PERCENT FASTER

Healthcare platform Doctolib rolled out Claude Code across its entire engineering team, replacing legacy testing infrastructure in hours instead of weeks and shipping features 40 percent faster.

Source: Anthropic customer story, Doctolib.¹⁸

L'OREAL: 99.9 PERCENT ACCURACY ON CONVERSATIONAL ANALYTICS

L'Oreal achieved 99.9 percent accuracy on conversational analytics, enabling 44,000 monthly users to query data directly instead of waiting for custom dashboards.

Source: Anthropic customer story, L'Oreal.¹⁹

7.5 What the proof points have in common

Across every example above, three things are true. The work was redesigned, not just augmented. Governance and human escalation were defined up front. And the value was measured in business terms (cycle time, cost, revenue, customer satisfaction), not in model accuracy alone. That is the operating model in action.

8. The Monday Morning Plan

If you want multiplier benefits from AI, do not start with chatbots. Start with one real workflow and redesign how people, processes, and intelligence work together.

The five actions below are the ones I recommend to every executive team starting serious AI work. None of them require a vendor selection. All of them are sequenced to be funded, scoped, and visibly underway within the first week.

Action	What it means
1. Fund the operating model	Stop treating chatbots as the strategy. Fund the visibility, guardrails, and learning capabilities that turn AI into a sustained capability. Name an executive sponsor whose performance objective is the operating model itself, not a tool rollout.
2. Choose one workflow that matters	Pick one workflow that moves a real business metric (revenue, cycle time, cost, NPS). Avoid generic productivity. Frontier Firms get traction by going narrow and deep before going wide.
3. Run all three patterns in parallel on that workflow	Elevate individual productivity, redesign the flow of work, and apply AI-first design to the same workflow at the same time. This is what produces multiplier benefits, and it is what most organizations skip.
4. Define what good looks like	Write the Quality Contract before you build. Specify acceptable outputs, judgment boundaries, escalation paths, and the metrics that will tell you whether the workflow is improving.
5. Create a cross-functional learning squad	Stand up a small, durable team: a business owner, domain experts, builders, and governance partners. Give them authority to iterate in production, and give them a forum to share what they learn across the enterprise.

8.1 What success looks like in the first 90 days

- One named workflow with a visible Quality Contract.
- Baseline metrics captured (cycle time, rework, escalation rate, cost per transaction).
- A working human-agent design, with escalation paths defined and tested.
- A monthly executive review focused on learning velocity, not just outputs.
- At least one reusable pattern codified for the next team to pick up.

A NOTE ON METRICS

Beware vanity metrics. Number of pilots, number of licenses, and number of prompts issued tell you almost nothing. The metrics that matter are cycle time, override rate, escalation effectiveness, and the rate at which patterns are codified and reused. Those are the metrics that show whether the operating model is working.

9. Lead What's Next

We are past the moment where the question is whether AI works. The real question is whether we can turn it into performance.

Organizations stall when they treat AI as a tool, or worse, as a chatbot problem. The leaders who pull ahead redesign the operating model itself: how individuals create value, how work flows across the enterprise, and how intelligence is embedded by default. Human judgment remains the differentiator. We scale intelligence while keeping accountability explicit. The frontier is not somewhere you arrive. It is a way you operate, and it compounds.

9.1 Four commitments for executives

If you take only four things from this whitepaper, take these:

- **Focus AI on performance.** Tie every initiative to a business metric a board member would recognize.
- **Multiplier benefits come from the operating model.** Not from the model, the tool, or the vendor. Fund the model.
- **Human judgment is the differentiator.** Design human-agent teams so that judgment, accountability, and consequence remain explicit.
- **The Frontier is a practice, not a destination.** Build the learning flywheel and protect it.

The future will not be led by the organizations with the most AI. It will be led by the ones that redesign how work, judgment, and intelligence come together. That is the work in front of every executive team. The good news is that the operating model is no longer a mystery. The harder, more honest question is whether your organization will fund it.

The Frontier Firm is a practice, not a place. Every process becomes a hypothesis, and every improvement becomes fuel for the next.

10. Bibliography and References

All sources cited in this whitepaper are publicly available. Customer outcomes are drawn exclusively from the organizations' own published case studies.

1. McKinsey & Company, *The economic potential of generative AI: The next productivity frontier*, June 2023. Available at: <https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier>
2. Microsoft, *2025 Work Trend Index Annual Report: The Year the Frontier Firm is Born*, April 2025. Available at: <https://www.microsoft.com/en-us/worklab/work-trend-index/2025-the-year-the-frontier-firm-is-born>
3. Anthropic and Material, *2026 State of AI Agents Report*, December 2025. Summary published at: <https://claude.com/blog/how-enterprises-are-building-ai-agents-in-2026>. Full report PDF: <https://cdn.sanity.io/files/4zrzovbb/website/cd77281ebc251e6b860543d8943ede8d06c4ef50.pdf>
4. Bain & Company, *Survey: Generative AI's Uptake Is Unprecedented Despite Roadblocks*, 2024. Available at: <https://www.bain.com/insights/survey-generative-ai-uptake-is-unprecedented-despite-roadblocks/>
5. Boston Consulting Group, *AI at Work 2025: Momentum Builds, but Gaps Remain*, 2025. Available at: <https://www.bcg.com/publications/2025/ai-at-work-momentum-builds-but-gaps-remain>
6. EY, *Responsible AI Pulse Survey 2025*, 2025. Available at: https://www.ey.com/en_gl/insights/ai/responsible-ai-pulse-survey
7. KPMG and the University of Melbourne, *Trust, Attitudes and Use of Artificial Intelligence: A Global Study 2025*, April 2025. Available at: <https://kpmg.com/xx/en/our-insights/ai-and-technology/trust-attitudes-and-use-of-ai.html>
8. PwC, *2025 Global AI Jobs Barometer*, 2025. Available at: <https://www.pwc.com/gx/en/issues/artificial-intelligence/job-barometer.html>
9. QueerTech, supported by Women and Gender Equality Canada (WAGE), *Benchmarking Canada's Tech Industry: Responsible Practices and 2SLGBTQI+ Outcomes in the Artificial Intelligence Era*, 2026 Research Report.
10. Microsoft, *Becoming a Frontier Firm: Lessons from Microsoft's AI Journey, and Three Ways to Turn Insight into Action*, executive playbook, 2025. Includes Microsoft Communications case study among others. Available via Microsoft WorkLab: <https://www.microsoft.com/en-us/worklab>
11. McKinsey & Company, *Lilli: McKinsey's generative AI tool*, firm publication and "Generative AI: Sustaining momentum" report, 2024-2025. Adoption and time-savings figures published by McKinsey.
12. Klarna, *Klarna AI assistant handles two-thirds of customer service chats in its first month*, press release, 27 February 2024. Available at: <https://www.klarna.com/international/press/klarna-ai-assistant-handles-two-thirds-of-customer-service-chats-in-its-first-month/>
13. Microsoft Customer Stories, *Capita uses Microsoft Copilot to transform service delivery for clients*. Available at: <https://www.microsoft.com/en/customers/story/>
14. Microsoft WorkLab, *AI impact at Dow: Copilot identifies millions in cost savings*. Available at: <https://www.microsoft.com/en-us/worklab/ai-impact-at-dow-copilot-identifies-millions-in-cost-savings>
15. Anthropic, *Customer story: eSentire*. Available at: <https://www.claude.com/customers/esentire>

16. OpenAI, *Moderna and OpenAI: Reimagining how we work, together*, joint announcement, 24 April 2024. Available at: <https://openai.com/index/moderna/>
17. Anthropic, *Customer story: Thomson Reuters (CoCounsel)*. Available at: <https://www.claude.com/customers/thomson-reuters>
18. Anthropic, *Customer story: Doctolib*. Available at: <https://claude.com/customers/doctolib>
19. Anthropic, *Customer story: L'Oreal*. Available at: <https://claude.com/customers/loreal>

Disclosure. The views in this whitepaper are the author's own and do not represent the position of any current or former employer. Customer outcomes are sourced exclusively from publicly available case studies, press releases, or earnings disclosures. No confidential or proprietary information of any kind is disclosed in this document.

About the author. David Beauchemin builds and leads high-performance businesses. He has been working with applied AI since 2017, and his 25-plus-year career spans market strategy, organizational leadership, and large-scale technology transformation across Microsoft and Google. He currently works at Microsoft leading Cloud & AI and works directly with C-suite executives on AI adoption, governance, and operating-model design. Connect on LinkedIn: [linkedin.com/in/davidbeauchemin](https://www.linkedin.com/in/davidbeauchemin)