LTM

WHITEPAPER

The Agentic Edge

How Enterprises Stay +One Ahead

Bhargav Desai
AI Advisor and Transformation Expert, LTM

Michael Reid
Managing Director and Growth Executive, LTM

Table of Contents

This paper explores two arguments:

First: The challenges blocking enterprise AI adoption, and especially agentic AI, are not technology problems. They are structural and organizational and will not be solved by better models or larger budgets.

Second: A reimagined approach to Organizational Change Management (OCM), one built for continuous AI evolution, not one-time go-lives, is the most direct and durable lever available to enterprises that want to close the adoption gap. We call this operating model the +One Framework.

Executive Summary

Agentic AI systems that interpret intent, break down the goals, and take autonomous action across workflows represent a fundamentally different kind of enterprise technology. It does not simply automate a task; it changes enterprises' operations, and it is evolving at a pace that enterprise transformation programs were never designed to handle.

The failure data is clear. More than half of all AI proofs of concept are abandoned before reaching production.3 Industry-wide, AI program failure rates are estimated at 80% or higher.5 And as organizations move from generative AI tools toward autonomous AI agents, the complexity and the stakes increase further.

What is less well understood is why. The root causes are not technical. They are organizational: insufficient data readiness, absent governance frameworks, workflow designs that treat AI as an optional add-on, and — most critically, change management approaches built for stable systems and linear change. The financial consequences are real: delayed AI adoption erodes expected ROI by an estimated 15–30% per quarter of delay, while premature program restarts destroy 3–6 months of accumulated value realization with each cycle.

> AI transformation is failing because enterprises are applying static transformation models to a continuously evolving system.

This paper introduces the +One Framework — a three-layer operating model that separates what must remain stable from what must evolve — and positions Behavioral Architecture as the organizing discipline that makes adoption durable across AI capability cycles.

Part I diagnoses the challenge landscape: the structural tensions created by AI velocity, the systemic failure patterns documented across research institutions, and the specific reasons agentic AI amplifies all of them.

Part II provides the response, culminating in the +One Framework and five organizational levers that translate it into action.

Part I — The Challenge Landscape

Understanding why enterprise AI adoption is failing, and why agentic AI raises the stakes further

1. What Agentic AI Actually Means for Enterprises

Enterprises often blur a critical line: AI as a productivity tool vs. AI as an agent. While Generative AI tools augment people: humans prompt, AI drafts, humans decide, Agentic AI executes: given a goal, it plans steps, acts across systems, and adapts to outcomes, while humans supervise.

That shift changes the playbook: copilots need enablement; agents need decision-rights redesign, trust calibration, and governance for autonomous action. Treating agentic AI like a standard rollout invites resistance, failure, and costly do-overs.

2. The Velocity Mismatch: When AI Evolves Faster Than Enterprises Can Transform

A major and often overlooked challenge in enterprise AI is the widening gap between AI capability velocity and enterprise transformation capacity.

2.1 The Pace of AI Advancement

Stanford's Human-Centered AI Index Report 2025 highlights how quickly the frontier is moving, faster than most enterprise plans assume: [1]

- Model training compute doubles roughly every five months
- Training dataset sizes double roughly every eight months
- Benchmark gains are step-change, not incremental: SWE-bench rose from 4.4% to 71.7% in a single year
- Frontier models are converging fast, driving frequent capability releases and "new best" announcements

Most transformation programs still run on 18–36-month cycles. The mismatch is structural, not something teams can fix simply by "moving faster."

2.2 The AI Reset Loop — and Its True Cost

The AI Reset Loop: New capabilities mid-program force tool changes, re-scoping, and re-architecture. Each reset typically costs 3–6 months of progress.

It hits in three measurable ways:

Reset Impact	Estimated Cost	Organizational Effect
Financial loss per reset	USD 500K–USD 2M+ in sunk platform, licensing, and integration spend*	Stranded investment erodes executive confidence in AI ROI
Time lost per cycle	3–6 months of value realization per restart*	Delayed adoption compounds quarter-over-quarter ROI erosion by 15–30%*
Organizational fatigue	High — especially after 2+ resets	Change fatigue, talent attrition, and cynicism toward future AI initiatives

Legend: * LTM internal assessment

Example: A large IT services organization launched HR Agent, a self-service tool for employees. The initial launch was celebrated, but OCM was limited to typical employee announcements/outreach. While the product team pushes new features and capabilities, adoption of the agent has been very limited. This was the result of applying the old playbook to the new AI paradigm. That pattern is not unusual. It is the average outcome.

Fixed-scope, one-shot programs trigger the loop. +One breaks it by keeping what must stay stable separate from what should evolve.

3. Why AI Programs Fail: Evidence and Root Causes

AI program failure is the default, not the exception. Multiple independent research bodies report the same pattern:

- Gartner: 50%+ of GenAI projects are abandoned after proof-of-concept [3]
- RAND: Industry-wide AI project failure is 80%+[6]
- IDC: 88% of AI proofs-of-concept never reach production[7]
- Gartner: 40%+ of agentic AI programs may be canceled by 2027[4]
- McKinsey: ~2/3 of organizations are not scaling AI enterprise-wide[2]

3.1 The Real Root Causes

These outcomes are not primarily due to technological limits; they are enterprise failures. The five most cited root causes are:

#	Root Cause	What It Looks Like in Practice
1	Data Readiness	AI agents hallucinate or return unreliable outputs due to poor grounding, missing metadata, and untrusted data pipelines
2	Weak Business Case	POCs demonstrate technical feasibility but cannot articulate measurable business ROI; sponsorship erodes
3	Governance Lag	Accountability gaps, regulatory exposure, and the absence of AI risk frameworks leave programs in holding patterns
4	Workflow Integration	AI remains a standalone tool; underlying processes are not redesigned; adoption is optional rather than embedded
5	OCM Breakdown	Skills gaps, incentive misalignment, distrust of AI decisions, and adoption fatigue are the most frequent and most fatal failure modes

OCM breakdown is both the most common and the most underfunded. Enterprises spend on models, platforms, and integration, then treat change as comms after go-live. The sequence is backward.

3.2 Why Agentic AI Makes Every Failure Mode Worse

- Bad data matters more when AI acts on it, not just reports it
- Governance gaps become urgent when AI takes actions in live workflows
- Workflow integration failures are louder when agents are the execution layer
- OCM failure becomes existential when the organization must trust autonomous action

Agentic AI doesn't add new failure categories; it amplifies the ones enterprises already carry.

3.3 The Most Common Enterprise Mistakes — What NOT to Do

Research and field experience point to a small set of avoidable failure patterns. Naming them is more useful than generic transformation advice:

Mistake	Why It Fails
Restarting the entire program for every new model release	Resets adoption momentum, trust, and institutional memory, the costliest AI debt.
Over-optimizing for 'the best model'	"Best" changes quarterly; anchoring the program to a model guarantees the reset loop.
Treating AI deployment as a tool rollout	Trains features, not behavior change expected for AI
Ignoring decision rights redesign	No clear oversight boundaries = liability risk and resistance that training can't fix.
Underinvesting in adoption infrastructure	One-off change campaigns don't sustain behavior across evolving capability cycles.
Measuring success by deployment, not behavior change	Go-live metrics hide adoption failure; measure workflow embedding and decision delegation.

A Practitioner Note

Many enterprise AI governance frameworks aren't operational. They satisfy procurement and board reporting, but they don't govern what agents do.

In many organizations, the people accountable for AI risk haven't witnessed a live, end-to-end agentic workflow. That's not personal failure; it's structural lag. Behavioral Architecture is how you close it.

Part II — How to Execute OCM for the AI Journey

The +One Framework and the organizational levers that close the agentic AI adoption gap

4. Introducing the +One Framework

The name +One carries an intentional meaning. In mathematics and theology alike, there is always one more beyond the known. A limit that is approached but never fully reached. In AI, the next release is always around the corner. The enterprise is always "+one" behind; advantage goes to those built to absorb the step without restarting.

The +One Framework is a three-layer operating model: keep Stability durable, let Innovation evolve, and build Adoption infrastructure that compounds across releases, so model churn doesn't force program restarts.

	Layer	What It Contains	Design Principle
1	**Stability Layer** Foundation	Architecture, data governance, security standards, integration contracts, compliance guardrails	Does NOT change with every model release. This is the durable foundation.
2	**Innovation Layer** Evolution	AI models, agents, capabilities, orchestration pipelines, tooling, APIs	Designed to evolve with every release cycle. Updates here do not require restarting the Stability or Adoption layers.
3	**Adoption Layer** Behavior	Human behavior change, decision rights frameworks, trust calibration, champions networks, reinforcement mechanisms	Persistent and cumulative across releases. Behavioral infrastructure built in one cycle is an asset for every subsequent cycle.

+One is an operating model: update capability without rebuilding governance or re-running adoption.

4.1 How +One Prevents the Reset Loop

The Reset Loop occurs when enterprises collapse all three layers into a single "program." Every model update triggers a rebuild of Stability and Adoption, throwing away governance and behavioral capital that should compound.

+One prevents this:

Data Centre SDN Solutions	Without +One	With +One
New frontier model released	Full program re-evaluation; architecture review; adoption restart	Innovation Layer update only; Stability and Adoption Layers unchanged
New agent capability available	Scope creep, stakeholder confusion, and delayed value delivery	Modular addition to Innovation Layer; adoption infrastructure extends naturally
Regulatory change introduced	Program pause; governance rebuild from scratch	Stability Layer governance update; Innovation Layer adjusted; Adoption Layer informed
Architecture cannot support agents	Full replatform required	Stability Layer redesign; other layers preserved; reduced time to re-adopt

5. OCM as Behavioral Architecture

The usual line: "OCM manages adoption", is too small for agentic AI. Here, OCM must design how people work with systems that act.

OCM designs human interaction with autonomous systems. It is Behavioral Architecture, not a communications plan.

That changes scope, investment, and timing. Behavioral Architecture isn't downstream of deployment; it runs alongside platform engineering, and for agentic AI, it's often the harder work.

Practitioner view: Platform teams underestimate this because early artifacts are invisible, and there's no architecture diagram for trust. But the data is clear: if OCM becomes post–go-live comms, adoption losses don't recover; underuse becomes "normal."

5.1 The Three Dimensions of AI Change

AI change adds three dimensions that most OCM playbooks miss:

Dimension	What It Requires	What Fails Without It
Cognitive Transformation	Changing how people think about work, organization/ department, and their own judgment relative to AI systems	People override AI reflexively rather than selectively — undermining the value case
Behavioral Redesign	Changing how people actually interact with AI day-to-day — what they delegate, verify, and override	Adoption is superficial; workflows are not actually changed; ROI is not delivered
Trust Engineering	Systematically building confidence in AI decisions at individual, team, and organizational levels	Resistance crystallizes into policy; AI programs stall at Stage 1 indefinitely

5.2 The Psychological Tensions That Must Be Named

AI transformation triggers three predictable tensions. If you don't name them, they become resistance:

- Trust vs. Control: People must delegate to AI while staying accountable for outcomes they don't fully control
- Automation vs. Identity: When AI takes tasks, it threatens professional meaning; training won't resolve it
- Efficiency vs. Relevance: AI speedups raise hard questions about individual contribution, careers, and value

If you frame adoption as training, you'll underdeliver. The constraint is trust and control, and no deck fixes that.

5.3 What Changes in OCM Practice

Dimension	Traditional OCM Approach	Behavioral Architecture Imperative
Scope of Change	Defined system or process change	Continuously evolving AI that shifts scope mid-program
Change Target	Adopt a new tool or workflow	Rebuild how people think, decide, and supervise work
Resistance	Process friction and usability	Trust deficits, professional identity threat, and control anxiety
Training	Feature training and procedure guides	Cognitive reframing, trust calibration, and judgment development
Timeline	One-time go-live with post-launch stabilization	Continuous adoption cycles aligned to AI capability releases
Success Metric	User completion rates and system utilization	Workflow embedding, decision delegation, and outcome quality

6. Five Organizational Levers and the OCM Action Each Requires

To close the agentic AI adoption gap, you must intervene at five structural levers. Each has a specific Behavioral Architecture move that makes the change stick.

Lever	What Must Change	OCM / Behavioral Architecture Action
Process	Static SOPs give way to intent frameworks that AI decomposes and executes	Help teams define outcomes, not steps, a fundamental mindset shift from procedure to purpose
Decision Rights	Codify what AI can decide, what requires human confirmation, and what stays human	Run decision-rights workshops; make delegation boundaries visible and agreed across teams
Governance	Move from point-in-time audits to continuous runtime monitoring systems	Train oversight roles; build a governance culture alongside governance tools
Org Design	Fixed roles evolve into fluid human-agent collaboration structures	Communicate role evolution proactively; co-design new role structures with impacted teams
Workforce	From task execution to AI supervision, oversight, and outcome accountability	Address professional identity, not just skill gaps; redesign incentives to match the new accountability model

6.1 New Roles for a Human-Agent Organization

As agents take on execution, three roles tend to emerge. They won't self-organize; change teams must design, name, and enable them.

Emerging Role	Primary Accountability	OCM Implication
AI Orchestrator	Own multi-agent workflow design and end-to-end execution quality	Name it explicitly; define authority and decision boundaries
Agent Trainer	Maintain KBs, prompt libraries, and tuning pipelines	Create learning paths and career tracks; don't bolt onto existing roles
AI Risk Controller	Set guardrails, monitor runtime behavior, and ensure compliance	Stand up cross-functional oversight; bridge tech and governance

7. Decision Framework: When to Re-platform vs. Continue

A common enterprise AI mistake is restarting the program for every new capability. The +One Framework provides a decision model: default to continue and extend. Re-platform only when the architecture truly can't support the next step.

Scenario	Recommended Action	Rationale
Same-class model update	Do NOT restart	Innovation-only update.
New workflow enabled	Add-on first	Extend Innovation; replatform if forced.
New compliance rule	Fix governance	Update Stability; keep Adoption.
Orchestration blocked	Replatform Stability	Rebuild architecture; keep Adoption.
Budget pressure	Don't restart	Restart = 3–6 months of lost time.
Competitor model leap	Don't restart	Parity is innovation, not a restart.

The default answer to 'should we restart?' is no. The burden of proof lies with those advocating a restart, not those defending continuity.

8. The Financial Case for Getting OCM Right

The business case for investing in Behavioral Architecture is not aspirational. It is grounded in the documented cost of getting it wrong. The financial framing below is directional and based on industry research, but the directional signal is unambiguous

Financial Risk	Mechanism
ROI erosion from delayed adoption	Each quarter of delay erodes ROI by ~15–30%* while competitors capture value.
Stranded investment from reset loops	Each restart writes off USD 500K–2M+* in sunk spend and compounds with every cycle.
Operational risk from ungoverned agentic AI	Without guardrails, agentic AI drives regulatory, reputational, and liability exposure.
Talent and change fatigue costs	After 2+ restarts, attrition rises, and each re-engagement costs more.
Missed enterprise-scale value	Enterprises that scale AI capture up to 2x the value of those stuck in POCs.

Legend: * LTM internal assessment

The investment in Behavioral Architecture, decision rights design, trust engineering, governance culture, and continuous adoption infrastructure is not a cost center. It is the mechanism by which AI platform investment becomes recoverable and scalable.

9. Redesigning the Transformation Program Model

AI capability will keep moving; the Reset Loop is a design constraint. Treat AI transformation as a product, not a one-time program, and structure it with +One.

9.1 From Program to Product

Traditional programs optimize for go-live and stabilization, assuming a stable end state. AI has no end state.

Productized transformation replaces the go-live model with a continuous delivery model:

- Time-boxed capability releases (not phase gates)
- Rolling prioritization based on business value + model evolution
- Continuous frontline feedback for platform + adoption
- Modular +One architecture so agents update without restarts

9.2 Managing the OCM Cycle Alongside the Capability Cycle

Behavioral Architecture must ship on the same cadence as capability. Each release needs awareness, readiness, and reinforcement; one change program won't cover what comes next.

That requires persistent enablement infrastructure:

A champions network per release, near-real-time adoption dashboards, and reinforcement tied to team outcomes. Low cost, often skipped, and the skip is expensive.

The measure of a successful AI transformation is not whether employees adopted the first version. It is whether the organization can adapt continuously as AI evolves.

10. The Adoption Roadmap: Assisted to Agentic to Autonomous

AI adoption is staged, and each stage shifts autonomy, human role, and Behavioral Architecture.

Layer	AI Model	Human Role	Behavioral Architecture Imperative
Stage 1 Assisted	Copilot — AI supports and surfaces insights; humans decide	Humans decide; AI recommends	Build AI literacy and awareness; address resistance through early adopter programs and visible wins
Stage 2 Augmented	Agentic — AI executes across workflows; humans supervise	Humans supervise; AI acts	Redesign roles and decision rights; implement trust engineering; activate the champions network
Stage 3 Autonomous	Autonomous — AI operates end-to-end; humans set strategy	Humans set strategy; AI governs execution	Institutionalize runtime oversight; embed continuous adoption cycles; sustain behavioral reinforcement

Stage shifts aren't automatic. Higher autonomy requires governance, roles, trust calibration, and reinforcement before expansion.

Skipping Stage 1 foundations and jumping to Stage 2/3 reproduces Part I failures, at higher stakes.

Closing

The organizations that will Outcreate the AI-native core for seamless agentic AI adoption won't necessarily be those with the largest AI budgets or the most advanced model deployments. They are those who build a specific organizational capability: the capacity to adapt continuously as AI keeps changing.

This paper has documented the gap between AI's evolution velocity and enterprise transformation capacity; between proof-of-concept experimentation and enterprise-scale institutionalization; between technology deployment and the organizational redesign required for AI to deliver durable value.

It has also described the response: the +One Framework, which separates what must stay stable from what must evolve. Behavioral Architecture, which replaces generic change management with the deliberate design of how humans interact with autonomous systems. And a program model that treats AI transformation not as a project with an end date, but as a continuous organizational capability.

The +One name carries the organizing insight. There is always one more model release ahead. One more capability beyond the current frontier. The organizations that internalize this and build the Stability, Innovation, and Adoption layers to absorb each "+one" without restarting are the ones that turn AI velocity from a threat into a structural advantage.

> In the era of Agentic AI, the limiting factor is no longer model capability.
>
> It is the organization's ability to trust, govern, and adapt to autonomous systems at scale.

The +One Framework is the operating model that enables that capability.

About the Authors

Bhargav Desai

AI Advisor and Transformation Expert, LTM

Bhargav is an AI Advisor and transformation expert at LTM, assisting organizations in harnessing the full potential of agentic AI. He collaborates with leaders to develop organizational readiness, change strategies, and governance structures that facilitate value realization from their agentic AI investments, guiding enterprises confidently through all stages of their agentic AI journey.

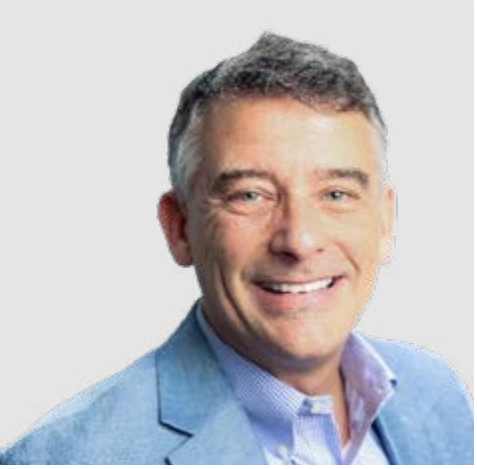

Michael Reid

Managing Director and Growth Executive, LTM

Michael Reid serves as the Managing Director and growth executive for IT services at LTM. He focuses on increasing revenue, managing strategic accounts, and developing enterprise partnerships within the communications, media, and technology sectors. Reid cultivates and expands client relationships, enhances account performance, and aligns sales efforts with delivery to deliver measurable business results. He oversees portfolios, boosts stakeholder engagement, mentors team members, and influences go-to-market strategies and transformational initiatives to maintain growth.

References

1. *Artificial Intelligence Index Report 2025*, Nestor Maslej, Loredana Fattorini, Raymond Perrault, Yolanda Gil, Vanessa Parli, Njenga Kariuki, Emily Capstick, Anka Reuel, Erik Brynjolfsson, John Etchemendy, Katrina Ligett, Terah Lyons, James Manyika, Juan Carlos Niebles, Yoav Shoham, Russell Wald, Tobi Walsh, Armin Hamrah, Lapo Santarlasci, Julia Betts Lotufo, Alexandra Rome, Andrew Shi, and Sukrut Oak, Stanford Institute for Human-Centered Artificial Intelligence, April 2025, https://aiindex.stanford.edu/report

2. *The State of AI: Global Survey 2025,* Alex Singla, Alexander Sukharevsky, Bryce Hall, Lareina Yee, Michael Chui, and Tara Balakrishnan, McKinsey & Company, November 5, 2025, https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai

3. *Why 50% of GenAI Projects Fail:* And How to Beat the Odds, Arun Chandrasekaran, Gartner, January 26, 2026, https://www.gartner.com/en/articles/genai-project-failure

4. *Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027,* Gartner, June 25, 2025, https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027

5. *Why AI Projects Fail,* James Ryseff and Anu Narayanan, RAND Corporation, April 10, 2025, https://www.rand.org/pubs/presentations/PTA2680-1.html

6. *The Root Causes of Failure for Artificial Intelligence Projects and How They Can Succeed: Avoiding the Anti-Patterns of AI,* James Ryseff, Brandon F. De Bruhl, and Sydne J. Newberry, RAND Corporation, August 13, 2024, https://www.rand.org/pubs/research_reports/RRA2680-1.html

7. *From Hype to Impact: How Agentic AI Unlocks Scalable Use Cases for Generative AI,* IDC, July 16, 2025, https://www.idc.com/resource-center/blog/from-hype-to-impact-how-agentic-ai-unlocks-scalable-use-cases-for-generative-ai/

8. *From Potential to Profit: Closing the AI Impact Gap,* Jessica Apotheker, Sylvain Duranton, Vladimir Lukic, Nicolas de Bellefonds, Sesh Iyer, Olivier Bouffault, and Romain de Laubier, Boston Consulting Group, January 15, 2025, https://www.bcg.com/publications/2025/closing-the-ai-impact-gap

9. *To Unlock the Full Value of AI, Invest in Your People, Boston Consulting Group,* November 10, 2025, https://www.bcg.com/publications/2025/to-unlock-the-full-value-of-ai-invest-in-your-people

LTM is a global technology services and consulting company and the Business Creativity partner to the world's largest and most disruptive companies. We bring human insights and intelligent systems together to help enterprises across industries rewire their business models, accelerate innovation, and drive AI-centric growth. With our integrated operations, transformation, and business AI services, we design and deliver solutions that create new productivity paradigms and new roads to value. Together with 87,000 employees across 40 countries and our global network of hyperscaler partners, LTM — A Larsen & Toubro company — owns business outcomes for over 700 clients, helping them to not simply outperform the market, but to Outcreate it.

© LTM