gradial

AUGUST 2026

Stop Prompting.
Start Thinking With AI.

Lisa Hillmann, Gradial | Gabriel Vancea, Accenture Song

EXECUTIVE SUMMARY

Enterprise AI has produced exceptional results for a few organizations and marginal returns for most. The difference isn't access, tooling, or budget. It's how people engage with the technology.

The practitioners creating the most value aren't the ones with the best prompts. They're the ones who know what to ask, how to follow a thread, and when to push back.

That shift became practical in late 2024, when foundation models made conversational, iterative work reliably better than a single well-crafted prompt. Most enterprise training still teaches the old model: specify the inputs, craft the prompt, submit, evaluate. That discipline made sense then. Today it optimizes for answers when the real leverage is in the questions people haven't learned to ask.

One scope note: this paper is about practitioners working with agents, systems that take multi-step actions, query live data, and adapt through conversation. Not prompt engineering for developers. Most training conflates the two, and that's part of the problem.

The organizations pulling ahead work with AI like a capable teammate: iteratively, with judgment in the loop at every step. This paper explains why that works, why the old model persists, and how leaders close the gap.

95%
of organizations are realizing no return on their AI investments.

26%
of workers have received training on how to collaborate with AI, not just use it.

8%
of companies are scaling AI at the enterprise level.

01 The gap that metrics cannot explain

86% of C-suite leaders plan to increase AI spending in 2026. At the same time, 95% of organizations are seeing no returns, and only 8% have reached enterprise-scale adoption.

The usual explanations focus on data quality, change management, or governance complexity. Those are real. But they don't explain why organizations with strong foundations in all three are still struggling. And they don't explain why adoption rates plateau even after training is complete.

> *The gap isn't between organizations that get AI and those that don't. It's between the playbooks most organizations are running and the technology those playbooks were written for.*

The models have moved faster than the strategies. What's newly possible in the last six months - conversational research, iterative drafting, multi-step agent workflows that adapt as you go - wasn't reliably available before. The opportunity now is to update the approach to match the capability.

79%

of organizations face challenges adopting AI - a double-digit increase from 2025.

Writer, Enterprise AI Adoption Report 2026

80%

of CEOs say AI will force operational capability overhauls.

Gartner, April 2026

02 Why prompt engineering was right then, and is a ceiling now

To be precise: prompt engineering is still the right skill for developers writing system prompts, building model-level applications, or fine-tuning model behavior in production systems. That work is real and it matters.

What it's not suited for is practitioner-level interaction with agentic platforms - the day-to-day use of AI systems that can search, retrieve, draft, compare, and chain steps together based on an ongoing conversation. These systems don't need a perfectly engineered prompt. They need a direction and a collaborator.

The first wave of enterprise AI training treated both as the same thing. How do you write a good prompt? How do you structure your input? That focus was understandable: earlier in the model capability curve, more explicit instruction did produce more consistent results. Prompt engineering was a legitimate response to where the technology was.

But it embedded a way of thinking about AI that has outlasted its usefulness for practitioners. Prompt engineering borrows the logic of waterfall development: figure out exactly what you want before you engage. Define the inputs. Anticipate every nuance. Then submit.

In waterfall software development, that discipline made sense when iteration was expensive and handoffs went to teams who needed complete context upfront. The same logic got applied to AI, and for a while it was the right call.

> *“This is the first time we can create a real cognitive loop between people and digital systems. That is a mind-bender, because it changes how we even conceptualize work inside an enterprise.”*
>
> Satya Nadella, CEO, Microsoft

The problem now is that the models have changed. They're capable of handling ambiguity, following conversational threads, adjusting based on feedback, and sustaining context across a long exchange. The structured-input model isn't necessary anymore - and because it front-loads all the thinking before any engagement happens, it creates predictable friction:

- **Waiting to start.** “I don't know exactly what I want yet, so I'm not ready to ask.” A natural response when you believe you need a perfect input - but a barrier that wasn't necessary to begin with.
- **Treating the first response as final.** “I asked once, it wasn't quite right, so this doesn't work for this kind of task.” An understandable conclusion when the interaction model is one-and-done - but not how the technology actually works at its current capability level.
- **Staying in the shallow end.** “I'll use it for well-defined tasks and keep the complex thinking to myself.” Research, synthesis, sense-making, early-stage drafting: the tasks with the highest potential leverage stay off the table.

These patterns show up in almost every organization running AI training today. They’re not a people problem. They’re what happens when you teach teams to fill out a form instead of have a conversation. Recognizing them is the starting point.

Only **26% of workers** say they've received training on how to collaborate with AI. Not how to prompt it. How to work with it. That gap isn't a failure of effort — it reflects how recently this mode of interaction became possible.

03 The agile shift: from prompts to pods

The organizations moving fastest right now aren't necessarily smarter or better resourced. They're the ones that recognized — often in the last few months — that the interaction model had changed, and updated their approach accordingly.

The shift looks like this: from waterfall to agile in how they engage with AI. In agile development, cross-functional pods work in short cycles. They don’t wait for complete specs. They build something, show it, get feedback, and adjust. The output improves through iteration, not better upfront planning. The term “pod” here refers to this working dynamic — a human-AI pairing where the human sets direction and exercises judgment, and the agent handles execution and synthesis — not a new team structure or org design. It's a description of how the work moves, not who is on the org chart.

Working with an AI agent now follows the same logic. You bring the goal, the context, and the judgment. The agent brings speed, synthesis, and availability. Neither side needs perfect information to start. The work happens through the conversation - and the models are now good enough that this consistently produces better results than a well-engineered single prompt.

WATERFALL / PROMPT ENGINEERING	AGILE / POD
Front-load all thinking before asking	Start a conversation to figure out what you need
Judge the first output as a verdict	Treat the first output as a draft
Specify the steps the AI should take	Define the outcome; let the agent find the path
Use AI for well-defined, bounded tasks	Use AI as a research and sense-making partner
Ask once, then evaluate	Iterate until the output is shaped to your need
Measure AI by usage and output quality	Measure AI by conversation depth and iteration

Figure 1. From waterfall prompt engineering to agile pod collaboration.

This isn't about AI getting smarter. It's about how people position themselves relative to it. Which tasks they bring to it. How they respond when the first output isn't perfect.

> *The shift from prompt engineering to pods is the shift from asking AI to execute your thinking, to actually thinking with AI.*

04 What working in a pod looks like

Here's a concrete example. A team preparing for a strategic review needs to understand how competitors are talking about AI-assisted content tools.

The prompt engineering approach

The team member spends time crafting a precise query before touching the tool. Which sources to specify. Which format to request. How to scope the output. The prompt is solid. The output is decent. But it only reflects what they knew to ask for before the research began. The gap between that and what they actually needed only shows up later, usually under deadline.

The pod approach

The team member opens their marketing platform and starts a conversation: "I'm prepping for a competitive review Thursday. I need to understand how our competitors are talking about AI content tools. Where should I start?"

The agent surfaces a landscape. They follow the most interesting thread. They ask what's driving one particular shift. They ask for a concrete example to use in the meeting. They ask what a skeptic would push back on. Four talking points to close it out.

The output isn't just more complete. It's shaped by the actual conversation. It reflects what they needed, not just what they could spec in advance. And their judgment was in the loop the whole way.

> *The agent handles retrieval, synthesis, and formatting. The human brings direction, context, and the judgment to know what matters. Neither waits for the other to be ready.*

The shift is measurable. A workflow that used to take hours — spec, brief, draft, review, publish — now runs in minutes. But speed is not the real point. The output is better shaped because judgment was in the loop at every step, not applied at the end. That distinction matters: faster outputs produced without human judgment in the process are still outputs the human has to fix later. Faster outputs where judgment shaped each exchange are done.

05 Where practitioners fit in the pod

One thing gets underestimated in most AI adoption conversations: where the practitioners actually sit, and what their job is relative to everyone else.

In content operations, the pod has three layers. Upstream are the marketers, strategists, and brand leads. They own the brief. They define what good looks like for the audience, the campaign, the brand. Downstream are the engineers and infrastructure teams. They build and maintain the technical layer: the CMS integrations, the data pipelines, the environments where content actually lives.

The practitioner sits in the middle. And that position has changed more than any other in the AI transition.

In the old model, practitioners were translators. They took the brief from upstream, figured out how to execute it in the CMS, and worked with engineers when they hit technical blockers. The work was manual, sequential, and slow.

In the pod model, the practitioner becomes the agent's director. They don't need to understand how the agent retrieves content from the CMS, or which APIs are firing, or how the infrastructure is configured. That's the downstream layer's domain. What they need is the ability to give the agent a clear enough outcome, read the output, and redirect when it's off.

That changes what they need from upstream too. Instead of waiting for a fully-formed brief before they can act, practitioners can now start with a direction and let the agent help them sharpen it. “We're updating the product pages for Q3. The messaging has shifted from features to outcomes. Here's the brand guide” is enough to start. The brief gets refined through the work, not before it.

And downstream, the practitioner's relationship with engineering changes. The technical complexity of integrating AI into content workflows is real. But practitioners don't need to own it. They need to be able to articulate what they're trying to accomplish, know when something's a technical constraint vs. a configuration choice, and trust that the agent is working within whatever guardrails have been set.

> *The practitioner's job in an AI pod isn't to bridge every gap between strategy and engineering. It's to direct the agent clearly enough that the gap mostly closes on its own.*

The practitioner-as-director is a new job. Not an updated version of the old one. Retraining people while leaving their KPIs, job descriptions, and decision authority unchanged is one of the primary reasons adoption plateaus after the initial push. You can't ask someone to work fundamentally differently and then measure them the same way. The role has to change with the model.

06 What you actually need to learn (if not prompt engineering)

Daniel Pink frames the challenge in terms worth sitting with: as AI becomes an answer machine, the competitive advantage is no longer on the retrieval side. The escape hatch, in his framing, is the question side of the equation — the ability to know what to ask, how to follow a thread, and what to do when the answer reveals a better question.

Pink identifies question-asking as the most valuable human skill in the AI age, alongside iterating relentlessly and knowing how to allocate work between humans and machines. The people who win, he argues, won't be the ones with the fastest tools. They'll be the ones who know what to point those tools at.

> *Pink's argument is that as AI gets better at answers, the edge goes to people who ask better questions. The smartest move isn't to answer faster — it's to pause and question your assumptions before charging ahead.*

That reframe matters for enterprise AI adoption. Most training programs are built around the answer side: how to get better outputs, how to structure inputs, how to evaluate responses. The question side gets almost no attention. And yet that's exactly where the leverage is.

Author Warren Berger defines a great question as "ambitious but actionable" - one that reframes the problem and leads to meaningful change. That's a high bar. But it's the right bar for anyone trying to do real work with AI.

So if not prompt engineering, what? Here are the skills that actually compound.

- **Asking better questions.** Not how to phrase a request. How to know what you're actually trying to find out. What does the answer need to do? What would change if you had it? Three questions are worth internalizing as a practitioner working with AI: Why do I believe what I believe? Would I rather be right, or would I rather understand? What's the question I'm not asking that I should be? The third one is especially relevant in AI work. The unanswered version of a problem is usually obvious. The unasked version is where real movement happens.
- **Iterating relentlessly.** Pink names this explicitly. It's also what every high-performing AI workflow has in common. The first output is a starting point. The value comes from what happens next: the redirect,

the follow-up, the narrowed scope. People who treat the first response as final are leaving most of the capability untouched.

- **The redirect.** Knowing how to give useful feedback on an output that's partially right. Not “this is wrong, try again” but “the structure works, the tone is off, and the second section misses the point on pricing.” Specific and directional. Moves the work forward without starting over.
- **Knowing when you have enough.** Research that keeps expanding. Drafts refined past the point of value. A real discipline challenge in working with AI is recognizing when to stop. The agent will keep going if you ask. Developing a reliable sense of “this is good enough to move” requires conscious practice.
- **Allocating human and machine work.** Pink includes this in his framework. Some tasks are better done by a person: the judgment call, the relationship conversation, the decision where context is too layered to hand over. Most practitioners underdelegate early. Building a calibrated sense of what the agent handles well, and where your judgment still matters most, comes with exposure and reflection.
- **Output literacy.** Reading AI outputs critically is different from reading human work. Knowing where to spot-check, what to verify before it goes external, how to catch something confidently wrong rather than obviously wrong. Learnable, but it takes deliberate attention, not just usage.
- **Knowing when to hand it back.** The reverse of delegation is retrieval. Some of what the agent surfaces requires a human to decide - a judgment call, a relationship issue, a risk where the context is too layered to hand over. Knowing when to take the wheel back is just as important as knowing when to hand it over. McKinsey’s Global Managing Partner Bob Sternfels put it plainly: “We are looking for people who can think, judge, and collaborate with AI tools - maintaining ownership of the final product by making sound decisions when an AI response isn’t quite right.”

07 What leaders need to do differently

> *“2025 was about AI pilots, discovery, and experimentation. 2026 will be about delivering agentic AI ROI.”*
>
> Kris van Riper, Practice VP, Gartner

Closing the gap between AI access and AI value requires more than usage dashboards and prompt libraries. Here's where to focus.

- **Retrain for collaboration, not just prompting.** Most AI training stops at prompt quality. The skill that actually drives value is knowing how to work iteratively: how to start a conversation without full context, how to follow a productive thread, how to recognize when you have enough to act. These are teachable, but they need explicit instruction and practice.
- **Open up the ambiguous use cases.** The highest-leverage AI tasks are usually the ones where the outcome isn't yet clear: early research, sense-making across large volumes of information, first drafts in fuzzy situations. Enablement that only covers well-defined, repeatable tasks misses most of the compounding value.
- **Measure conversation quality, not just logins.** Usage rates tell you whether people open the tool. They don’t tell you whether people are working with AI in ways that make them better at their jobs.

The leading indicator to track is multi-turn conversations that improve across exchanges — that's the signal that teams are working in pod mode, not prompt mode. Most enterprise AI platforms surface conversation-level data in their analytics layer; the question is whether anyone is looking at it. Specifically: what percentage of sessions involve three or more exchanges? Are later turns in a conversation producing meaningfully different outputs than the first? These are measurable today. They're just not being measured.

- **Give teams permission to iterate.** The norm of sharing polished, finished work before it leaves your desk is directly at odds with how agile AI collaboration works. Leaders who want more value from AI need to explicitly create space for first drafts, rough cuts, and thinking out loud. Tool change and culture change have to move together.
- **Build guardrails alongside iteration culture.** Iteration culture and governance are not in tension — but they need to be designed together. For organizations in regulated industries (financial services, healthcare, legal), the question isn't whether to iterate but how to define the review gates within which iteration happens. A practical approach: distinguish between AI-assisted work that stays internal and work that goes external or into a regulated workflow, and apply different review requirements to each. That distinction lets teams move fast on research, drafting, and sense-making while keeping compliance review where it actually matters.

08 Where to start

The shift from prompt mode to pod mode doesn't require a program. It requires one person deciding to work differently and seeing what happens. Here's a starting point for each audience.

For the practitioner

Pick one task you'd normally do alone this week. Start it as a conversation instead. Don't craft the perfect input - open with what you know right now and let it move. Redirect when the output misses. Pay attention to what the third or fourth exchange gets you that the first response never would have. Do that once, with real work under real pressure, and you'll understand the difference faster than any training deck will get you there.

For the manager or team lead

Find the task in your team's workflow that takes the most time and has the least upfront clarity. The fuzzy brief. The first-draft research. The stakeholder summary nobody wants to write. Make that the pilot. Give your team explicit permission to share rough AI-assisted drafts internally - remove the polish requirement from the first pass. Measure time-to-first-draft. Not quality of final output. The iteration data will show you where the real gaps are faster than any survey will.

For the C-suite executive

Ask your team one question: what percentage of our AI interactions are single-turn versus multi-turn? If you don't know, you don't have visibility into whether your investment is producing prompt-mode or pod-mode behavior. Run a 30-day measurement sprint. Track conversation depth alongside output quality. Then use that data to reset your enablement strategy - not around access and adoption rates, but around the behaviors that actually compound. The organizations already measuring this are building a lead that gets harder to close every quarter.

About the Authors

Lisa Hillmann is VP of Transformation and Enablement at Gradial. Lisa has led digital marketing, operations and strategy teams at companies like T-Mobile and Kaiser Permanente. Her work sits at the intersection of how organizations adopt new technology and how practitioners learn to use it.

Gabriel Vancea is AI & Digital Transformation Lead at Accenture Song. Gabriel leads AI-enabled content operations and digital experience programs for Fortune 500 clients, helping organizations bridge the gap between AI investment and enterprise-scale value.

About Gradial

Gradial is a marketing execution system of work for enterprise marketing. It connects Gradial Agents to the tools marketing and digital operations teams already use - CMS platforms, ticketing systems, analytics, and design tools - enabling content operations at a speed and scale that was not previously possible. Customers include AWS, T-Mobile, Prudential, and Walmart.

References

Writer, “Enterprise AI Adoption Report 2026.”

Gartner, CEO survey on AI and operational capability, April 2026.

Gartner, Kris van Riper, Practice VP, on 2026 agentic AI ROI.