PROMPT ENGINEERING IS DYING

CONTEXT ENGINEERING IS REPLACING IT

A study of 3.8 million messages sent on the elvex enterprise AI platform makes the case plainly: prompt engineering is dying. And it never *truly* worked at enterprise scale, regardless. Better capabilities are replacing it: multiplayer context engineering and agentic planning are delivering results that prompt skill alone never could.

After elvex rolled out multiplayer context engineering and agentic planning, three things happened simultaneously: 1) Users submitted shorter, simpler messages. 2) They pursued more complex, multi-step tasks. 3) And their satisfaction with results climbed from roughly 60% to nearly 80%.

The dominant response to shallow AI transformation has been more training (prompt engineering workshops, internal guides, AI champions), which treats the symptom rather than the cause. When the platform carries the context and the planning, prompting skill stops being the bottleneck, and AI value stops being a power user story.

“Prompt Engineering,” Or, Why Enterprise AI Transformation Stalls

The tools have been adopted, but what hasn't followed is any transformation in how people actually work. The models are capable enough; what enterprise deployments lack is the infrastructure to spread their value across a team.

“AI Adoption” Is Not Transformation; Here’s Why

Most enterprise software is required to do your job. You learn it because you have to, to receive your paycheck, submit your work, communicate with your team. AI tools, by contrast, have been "nice to haves," and that distinction matters enormously for transformation.

There are several compounding reasons why AI knowledge doesn't transfer across an organization the way traditional software knowledge does.

First, these tools apply generally to all aspects of work, hitting everyone everywhere all at once. The "what do we do with this" question has been a massive system shock, and there is too much to learn all at once with no clear starting point for most employees. People then learn how to use AI at wildly different levels of expertise, and begin inputting wildly different levels of prompts; this applies even to those on the same team pursuing similar tasks.

Second, the tools themselves have been unreliable and nondeterministic. A tool that gives reliable answers is one you integrate into your work and learn deeply. A tool that is unreliable is, at best, a hobbyist interaction.

Unreliability plus overwhelming application led most executives to grant access and tell people to "see what happens." That was the AI mandate era, which is now ending. Combined with frontier lab marketing that emphasized job displacement, it meant leaders were encouraging employees to experiment with difficult, unreliable tools that might take their jobs. Enthusiasm was unsurprisingly slight.

Why Knowledge Doesn't Compound

Even among the employees who do get good at AI, the value rarely spreads. People develop workflows, prompting habits, and mental models that work for them, but the organization has no infrastructure to capture that value, share it across teams, or build on it over time. When a power user leaves or moves teams, their AI capability leaves with them.

"Millions — if not billions — of dollars have been wasted training employees how to prompt engineer, with very little ROI to show for it," says Sachin Kamdar, CEO of elvex. "The harnesses are getting smart enough that that's now unnecessary."

Two Developments That Change How All Of This Works

Two major platform-level developments are rapidly increasing the utility of AI systems, and organizations that pursue them will see a fundamentally different transformation curve.

1. Agentic Planning and Delegation

Earlier AI systems responded to a single prompt with a single response, leaving users to break down complex tasks themselves, manage the sequence of steps, and stitch outputs together. That process took enough skill and patience that it reliably stayed in the hands of power users.

Agentic AI systems can receive a large, complex task and break it down into a series of smaller subtasks, pursuing each iteratively. Smaller, more focused tasks have lower intelligence thresholds, so each step is handled more reliably than the whole problem would be, and those steps compound into a better result. It also makes complex work accessible to more people: when the platform handles the decomposition and sequencing, users do not need to orchestrate a multi-step workflow. They describe what they need, and the system handles the rest.

The data reflects this. After elvex rolled out agentic planning capabilities, average LLM requests per conversation rose from roughly 5 to more than 8. The platform was doing more on users' behalf, breaking requests into steps, pursuing each one, and returning a more complete result.

2. Context Engineering

Context engineering is the platform's ability to supplement a user's message with additional information before it reaches the model. Instead of the user having to explain the company, the workflow, the relevant history, and the desired output format every time, the platform assembles that context automatically.

This is what makes AI work for a first-week employee the way it works for your best power user: the platform already knows the business, so the user just has to describe what they need.

Multiplayer Context Engineering: How It Works

When one employee works through a problem with an AI agent, the platform captures what was learned: the approach that worked, the context that was relevant, the output format that was useful. The next time a similar problem comes up (whether for the same employee or a different one), the platform surfaces those learnings automatically.

A power user who builds a great workflow does not just benefit themselves; their work becomes an organizational asset, and the next person facing a similar problem starts where the power user left off.

The practical effect is that you do not need every employee to become an AI expert. One person figures out the right approach, the platform captures what they built, and everyone else benefits from it. AI competency spreads across the organization not because everyone went through training, but because the platform carries the institutional knowledge forward.

That requires shared workspaces where agents, context, and workflows are available to the whole team, auto-updating project context that learns from each new interaction, and an agent library where what works for one team becomes available to every team.

What the Data Show

elvex analyzed 3.8 million messages sent across its platform, spanning the phased rollout of context engineering and agentic planning capabilities (highlighted in the charts below). The results are consistent across all three measures.

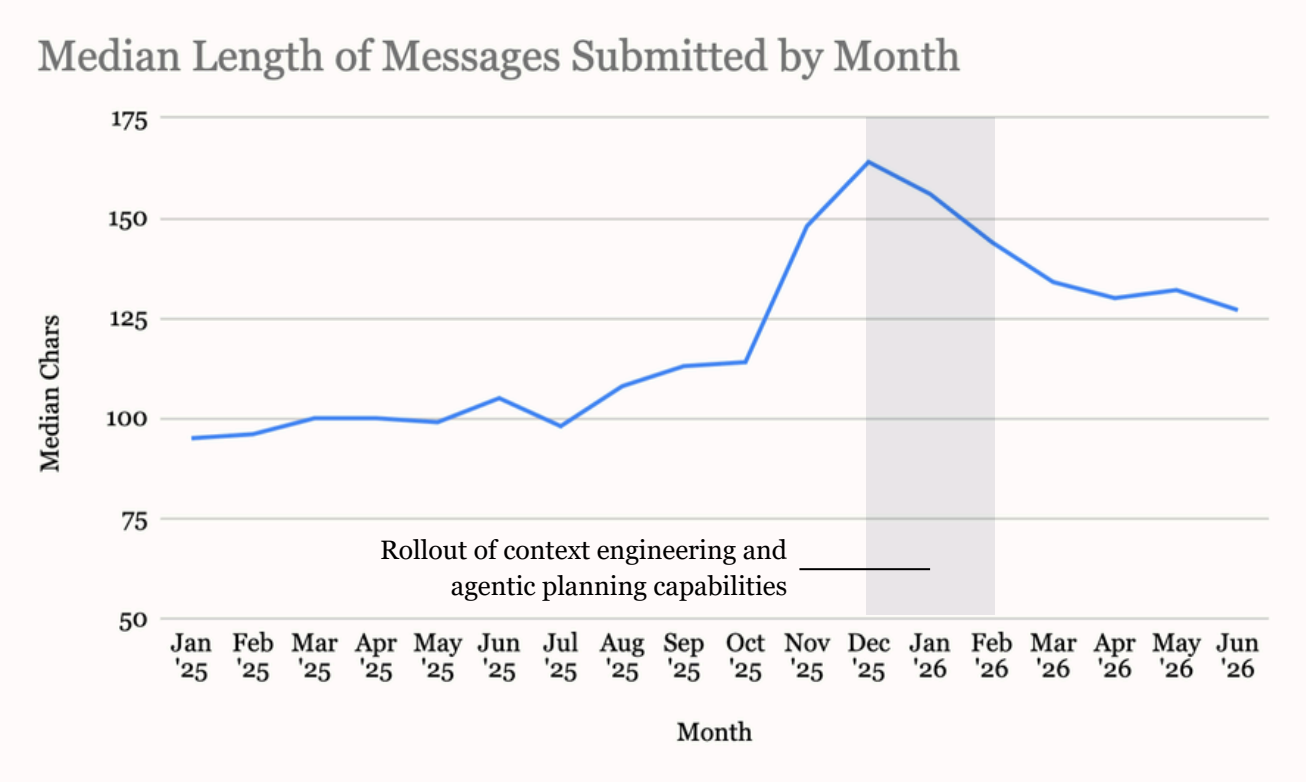

Chart 1: Median Message Length Dropped
Before the rollout, median message length was climbing, peaking at approximately 164 characters in December 2025. After the rollout, it fell steadily to approximately 127 characters by June 2026.

Users were doing less work to get good results, not less work overall. When the platform already has the context (the company background, the project history, the relevant prior work), users do not need to over-explain. They send shorter, more direct requests because the platform fills in the rest. The decline in message length is a proxy for reduced prompting burden: users stopped needing to be expert prompt engineers.

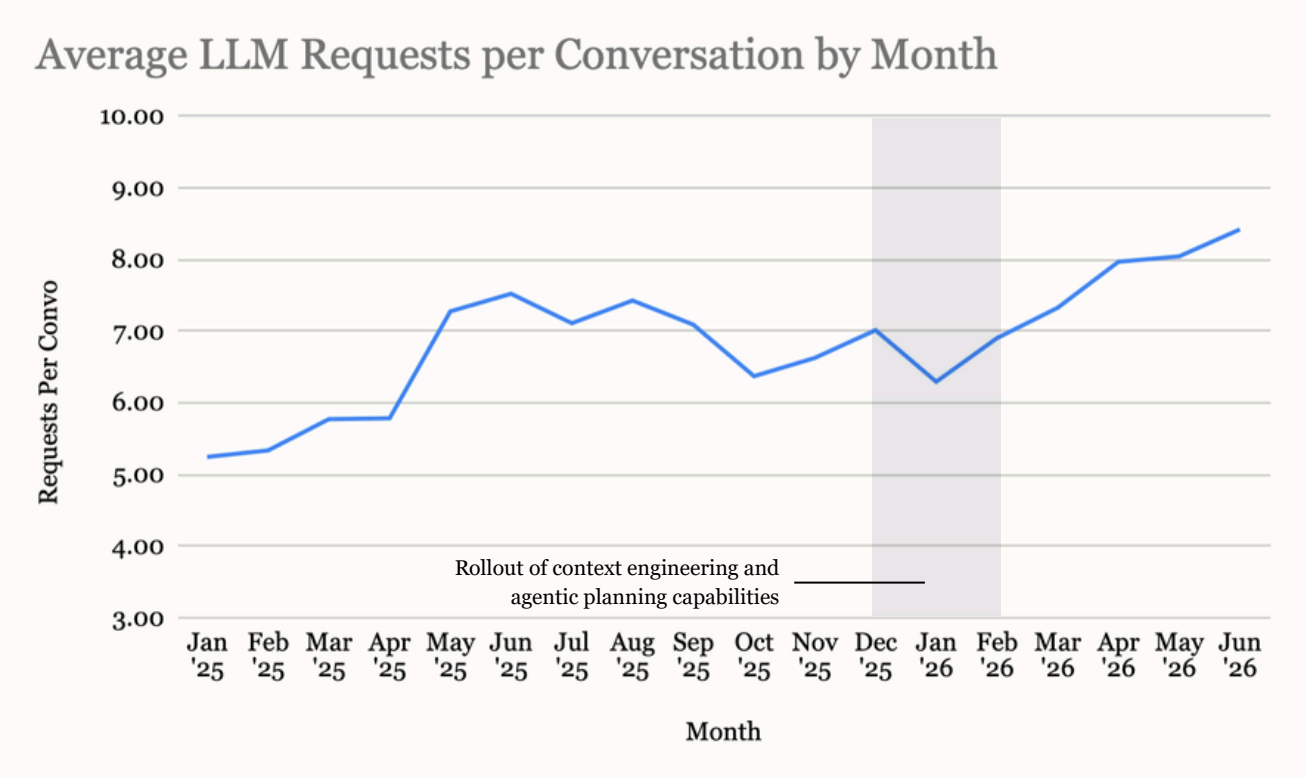

Chart 2: Requests per Conversation Increased
Average LLM requests per conversation rose from a previous high of approximately 7.2 —7.50 from May through September 2025, to 8.42 by June 2026, a 12% increase.

More requests per conversation reflects that the platform is doing more on users' behalf. Agentic systems break complex tasks into sequences of smaller requests, each handled more reliably than the whole would be. A user who asks for a competitive analysis triggers eight subtasks; the platform pursues each one and returns a synthesized result.

Rising requests per conversation, combined with falling message length, tells a clear story: users are asking for more ambitious things in fewer words, and the platform is handling the complexity.

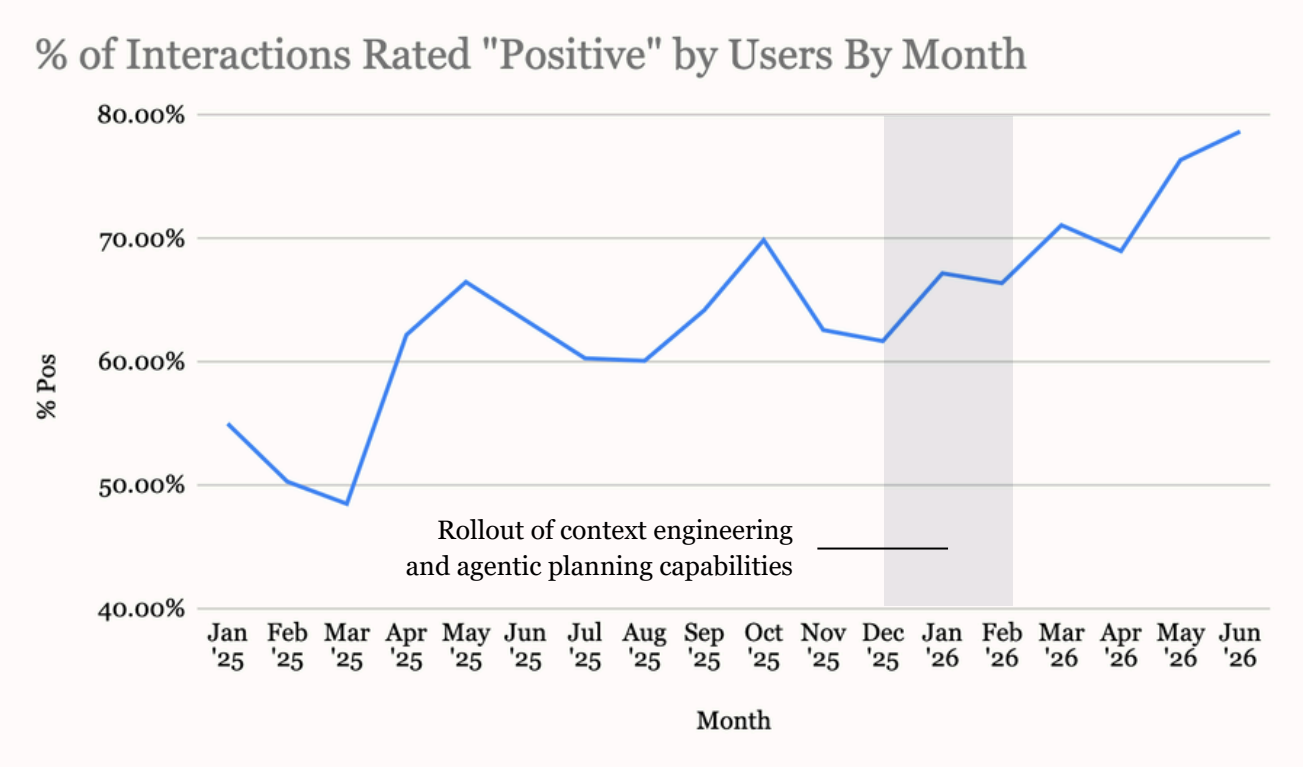

Chart 3: Satisfaction Climbed

Interactions rated "Positive" by users climbed from ranges of 62% to 70% in mid 2025 to nearly 79% by June 2026. The later numbers reflect what happens when the platform supplements the context and the planning tasks of the models. Results improve across the board, not just for power users who have invested time in learning the system, but for everyone.

Viewed Side By Side, What Does This Mean?

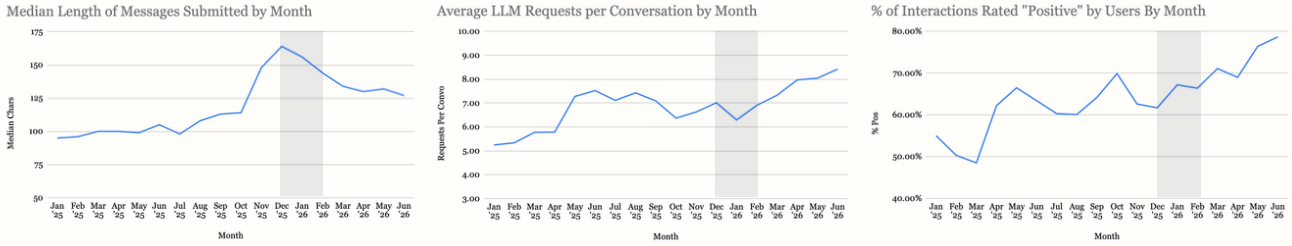

This study of 3.8 million messages sent on the AI platform elvex since rolling out context engineering and agentic planning feature sets has revealed

1. Decreasing complexity of user-submitted messages (Chart 1)
2. Increased complexity of tasks being pursued (Chart 2)
3. Increased ratings of user acceptance of results (Chart 3)

This means that "Context engineering" is replacing "prompt engineering" as more capable applications (the "harness") are able to plan multi-step workflows and also supplement user-generated messages with additional information that makes the model response more effective.

As AI Handles More of the Work Autonomously, Visibility Into What It Does Becomes the Critical Requirement

As AI systems become more reliable and more capable of handling complex work autonomously, the chatbot as the primary interaction pattern will decline. One of the main reasons chat has been the dominant interface is that it lets users inspect the interaction history and understand, at least partially, how the model arrived at its output. That was necessary when results were unpredictable.

As results become more reliable and the system handles more steps autonomously, the interaction becomes more opaque, which will drive demand for analytics, audit logs, and testing capabilities that make system behavior visible even when the interaction itself is not. Enterprise organizations will need more visibility into what their AI is doing, not less, as it becomes more capable.

Platforms built for governance from the start are better positioned for this. Governance is the foundation that makes broad transformation possible, and the organizations that treat it as such will be able to move faster as capability grows.

What to Look for in an Enterprise AI Platform

If your current AI investment is producing power user results instead of org-wide ROI, these are the platform questions worth asking:

Question	Why It Matters
Does the platform load business context automatically into every conversation?	Without this, every user either gets worse results, or re-explains the company workflow on every session
Can agents be shared and reused across teams?	Shared agents turn individual breakthroughs into org assets
Does the platform learn from interactions over time?	Static platforms stagnate; compounding platforms improve
Can non-technical users build and modify agents independently?	No-code building determines whether AI scales
Is institutional AI knowledge preserved when employees leave?	If yes, your AI investment grows over time. If no, it resets
Can you measure usage, outcomes, and ROI across the org?	As systems become more capable and more autonomous, visibility becomes more important, not less

Methodology

This analysis is based on 3.8 million messages sent on the elvex enterprise AI platform between January 2025 and June 2026.

The highlighted period in each chart represents the phased rollout of two feature sets: context engineering and agentic planning and delegation. Metrics include median message length (characters), average LLM requests per conversation, and percentage of interactions rated "Positive" by users via in-platform feedback.