

FUTURES

AI Is the New UI—But Not the Way You Think

Why agentic AI needs a different kind of interface

WRITTEN BY

Dave Wright

Chief Innovation Officer

Diana Wu David

Futures Director

CONTRIBUTORS

Grant Hulbert

Product Management Director

Sasson Jamshidi

Principal Inbound Product Manager

Eugene Chuvyrov

Principal Inbound Product Manager



AI Is the New UI—But Not the Way You Think

Why agentic AI needs a different kind of interface

Summary

- Enterprise AI investment is surging, but value is not. Most companies report no material bottom-line impact from AI.
- AI fails to deliver in part because trust is miscalibrated at every level. Workers are either withholding the autonomy that agentic AI needs to deliver or delegating blindly without accountability. Neither works.
- The prevailing solution—chatbot as UI—was a powerful start, but it remains limited by its one-size-fits-all nature and lack of visibility into the model's actual reasoning. What's needed is a new category of interface: dynamically generated, context-aware, and transparently accountable.
- Seven interface qualities separate organizations getting real value from agentic AI from those stuck in proof-of-concept: adaptable, anticipatory, instantaneous, multimodal, transparent, reliable, and delightful.
- Organizations that get this right will protect existing technology investment, unlock competitive advantage, and deliver the experience today's workers demand, leading to real enterprise value.



Early in the morning, a security alert fires. Somewhere in your enterprise, a threat is spreading across devices. Your agentic AI could handle it—pull the logs, quarantine the affected nodes, identify similar vulnerabilities across the fleet, patch them—instantly and autonomously.

This technology has arrived at an inflection point for most organizations.

The first type of company hands it over. An employee types “fix this” into a chatbot, the AI goes to work, and the threat is resolved—or so it seems. Ask that employee

what the AI actually did, and they’ll be like a passenger who bought a ticket to Paris and glanced at the map on the seatback screen without knowing if they’re headed to Charles De Gaulle, Orly, or Paris, Texas.

The second type of company uses AI for prescribed workflows only. It routes that same alert through the established process: triage opens a ticket, security reviews it, remediation gets assigned, change management approves the patch. Each function handles its prescribed slice,

each AI tool authorized only within its lane. It’s as if the employee is on a series of a dozen regional flights, each leg requiring a new boarding pass, security clearance, and gate agent sign-off. Delays compound. Unexpected events ground the whole operation for days while humans manually reroute every connection. The costs are real, but the destination never gets any closer. They may think they’re headed to Paris, but they’re stuck in Newfoundland.

Agentic AI is already generating 17% of total AI value in 2025 and is expected to reach 29% by 2028,¹ but neither of those scenarios above delivers what agentic AI actually promises. And they’re not alone in struggling to translate AI spending into measurable outcomes. McKinsey found that more than 78% of companies now use GenAI—and more than 80% of those same companies report no material





“

Agentic AI requires that business leaders evolve the way we interface with technology—dynamically, contextually, transparently—to turn every worker into a pilot who knows exactly where they’re going, why, and when to take the controls.

bottom-line impact.² The Chief Economist of Goldman Sachs says the massive investment in AI has made “basically zero” contribution to the economy in 2025.³ There are many reasons for that gap—readiness, strategy, culture, lack of vision—but one overlooked factor sits closer to the surface: the interface.


What companies are not getting right is the way a worker actually experiences, understands, and acts alongside AI in real time. The first passenger in our scenario can’t audit, course-correct, or own the outcome. The second

passenger is simply automating several inefficient steps. In both cases, companies are not empowering employees to become pilots steering their own planes, with AI as their first officer.

Many organizations think of the new UI as a hardware swap: rip out the dashboard, install a chatbot. Worse, many don’t consider the interface at all. The interface is not

just the surface between a single worker and a single AI, but an orchestration experience: the visibility layer that allows workers and leaders to see what the AI agents are doing, where value is being created, and where risk is accumulating. As today's enterprises prepare to have multiple agents operating across functions, the interface question is not "how does a worker interact with one model?" but "how does an organization keep visibility and control over an entire ecology of autonomous agents while ensuring a smooth journey for all employees?"

With AI, the user interface is not decoration; it's the accountability layer.



The seven qualities of an agentic AI interface that works

With agentic AI, the user interface must evolve beyond a screen or a chat window and serve as the full layer through which workers and leaders perceive, direct, and remain accountable for AI action. Depending on the worker's function, the visibility must extend all the way from the individual prompt to the organization-wide view of every agent in motion. Seven qualities define whether that layer enables the kind of trusted autonomy that agentic AI requires or quietly works against it.



1. Adaptable and dynamic

Static enterprise interfaces were designed for repetition: fixed menus, predetermined views, data buried in modules that employees may not know exist.

Agentic AI requires an interface that assembles itself for the task at hand. For instance, a single prompt—“How’s my project going overall?”—should be able to trigger three parallel agents simultaneously: one determines the optimal visualization format, one pulls relevant data, and a third analyzes the output and anticipates the next action. The result is a fully rendered Gantt chart that no developer built in advance. Follow-up queries produce entirely different interfaces on demand, each suited to the task and disposable once it’s complete.

This is already reaching clients. In a recent public sector engagement, an engineer fed AI-transcribed meeting notes into an agentic build tool and prompted it to produce a

mobile app prototype from scratch, delivered in roughly three days without a dedicated UI team. In the future, such custom solutions could be near-instantaneous.

Google's November 2025 findings showed that human evaluators strongly preferred dynamically generated interfaces over standard LLM outputs.⁴ Thesys's C1 engine is already rendering adaptive interfaces in real time.⁵

One constraint matters: Dynamism without discipline creates disorientation. Interfaces should shift when the task demands it and hold steady when familiarity is required.

2. Anticipatory and personalized

An interface that only responds to explicit commands keeps the worker as the bottleneck. The better model reads context and surfaces what hasn't been asked for yet. Workers should stop adjusting to the interface; the interface should adjust to the worker.

A smartwatch can now calibrate your wake time to your circadian rhythm.⁶ A personal AI birthday assistant can remind you of a friend's upcoming birthday and suggest a gift based on past purchases, before you've thought to ask.⁷ Consumer technology is priming users to expect this level of proactive personalization.

In the enterprise context, AI theoretically already knows everything about the worker: role, habits, workflow history, and past decisions. Yet most enterprise interfaces don't use any of it.

A Gartner^{®8} survey of 5,141 employees conducted from April through July of 2024 found that, "only 23% of digital workers are completely satisfied with their work applications in 2024, down from 30% in 2022." Moreover, "Digital workers satisfied with work applications are nearly three times more likely to report being much more





productive compared with those who are not satisfied.”⁹ The productivity dividend from personalization is already visible in adjacent domains on the consumer side: One European telecommunications company shifted from four customer segments to 150 using generative AI, achieving a 40% lift in response rates and a 25% reduction in deployment costs.¹⁰

The orchestration experience expands this further. An anticipatory interface doesn’t just surface what one worker needs next but also recognizes when a pattern in one part of the organization could benefit another. The interface that anticipates is also the interface that sees across silos.

3. Instantaneous

When AI takes 30 seconds to respond, workers don’t wait. They get distracted and disengage. Speed is not a feature but the psychological precondition for everything else.

When response times drop to near-zero, behavior changes fundamentally. Workers who previously avoided requesting new analyses or iterations that might take weeks

begin asking more freely, experimenting more readily, and acting more decisively. The conversational rhythm of human interaction requires a machine partner that matches it.

Humans don't pause for 30 seconds mid-conversation, so if an interface does, it has already broken the collaboration before it began. Our brain did not evolve to accept the stop-and-go rhythm of human-computer interaction. It only takes one second before users notice a computer delay and lose a sense of direct control; after 10 seconds, their attention fractures.¹¹

The standard today—several seconds for an LLM response¹² and much longer for complex queries—is not sufficient. The target should be sub-second: mid-movement on a mouse, not a loading spinner. For employees now accustomed to having real-time conversations with their cars and home speakers, only AI working at the

speed of human thought will seem like colleagues worth collaborating with.



4. Multimodal and ambient

Screen-first interfaces require workers to stop, navigate, and type. But much of work happens in motion, in warehouses, on calls, between meetings, mid-commute.

“

Gartner® estimates that “eighty percent of enterprise software and applications will be multimodal by 2030, up from less than 10% in 2024.”¹³

Workers are increasingly accustomed to pointing a phone at an image and having information surface instantly, without typing or navigating. The enterprise equivalent of that experience—frictionless, contextually triggered, modality appropriate—is where the interface is heading.

Gartner® estimates that “eighty percent of enterprise software and applications will be multimodal by 2030, up from less than 10% in 2024.”¹³ Enhancing domain-specific language models “will improve accuracy, automate operations, and drive contextual decision intelligence, enabling AI to take proactive actions across tasks.”¹³ The design goal is to let the modality match the moment—voice in the field, visual for complex analysis, touch for quick confirmation—and to make the transition between them seamless. The interface should go where the worker goes.

5. Transparently accountable

An AI that goes everywhere and acts instantly is also an AI that can cause damage everywhere, instantly—which is why transparency isn’t a nice-to-have but a load-bearing pillar. Most AI today is deployed as a text box where workers type a query, something happens, and the model returns a result. The full range of what that model did or can do stays invisible. With agentic AI, this becomes a liability.

Consider an AI assistant that autonomously orders flowers for an employee’s spouse, books a meeting for 8 a.m., and replies to an email the employee was deliberately ignoring. That agent hasn’t made anyone more productive. It’s created a crisis. Most organizations still lack clear protocols for when AI acts on their behalf. What’s needed is not more automation but a control layer: an interface that makes AI action visible and reversible before it becomes permanent. When agentic systems complete work on behalf of workers, the interface must allow those workers to understand what was done and even back out.

Interface elements—buttons, comparison views, confidence indicators, icons that invite exploration—aren’t decoration. They are the mechanism through which AI capability becomes available across the workforce and, in turn, encourages adoption and innovation. An agent that does what it says it will do, and the way it said it would do it, earns expanded autonomy over time. That credibility can only be built through visibility.



Yet while 40% of organizations cite explainability as a key AI risk, only 17% are currently doing anything about it,¹⁴ and 66% of surveyed employees rely on AI output without evaluating its accuracy.¹⁵ When workers can't see what an agent did and merely accept the results, they can't own the outcome. And in the enterprise, ownership is not optional.

The design counterweight matters equally: excessive transparency creates cognitive overload.¹⁶ The target is calibrated visibility, enough to maintain accountability, but not so much that it defeats the purpose of automation.

6. Legibly reliable

Even in tightly controlled academic settings with high-quality source material, peer-reviewed research found that leading large language models fabricated between 28% and 39% of their outputs.¹⁷ Moreover, LLMs are non-deterministic by nature, meaning identical inputs can yield different outputs, even when the model is performing correctly. That is not a minor calibration problem but a structural challenge that compounds as agents chain together: Each agent's error becomes the next agent's input, while humans' trust in the system degrades.

Trust in company-provided GenAI already fell 31% between May and July 2025; trust in agentic AI systems that act independently dropped 89% in the same period.¹⁸

McKinsey identifies trust as the foundation for AI adoption (e.g. if employees lack confidence in what a system produces, they will not use it) and points to explainability as the primary mechanism through which that trust is built.¹⁹ In a controlled study of workers using AI agents in organizational settings, the visual presentation of information reliably predicted whether users would trust the system.²⁰ Transparency and interface design are emerging as the primary drivers of trust formation, suggesting that what the system shows matters as much as what it can do.

The reliability of an AI system is only as useful as the interface that makes it legible—and workers who can see what the model is doing are both more likely to use it and more equipped to catch what goes wrong.

“

Seventy-three percent of office workers say AI and automation will increase their job satisfaction, but satisfaction is contingent on tools that actually work well.²



7. Delightful

Delight is the adoption lever that unlocks all six qualities above. Without it, the others go unused. Well-designed experiences generate trust, sustain attention, and drive adoption. For example, thanks to an approachable, non-threatening design, Japanese-developed seal robots achieved acceptance in British nursing homes that



mechanical alternatives could not.²¹ The enterprise dynamic is identical. When workers first encounter a dynamic AI interface that feels like a game rather than an ERP system, the reaction is immediate.

As enterprises formally tie compensation to AI usage metrics,²² the quality of the interface determines whether that mandate feels like empowerment or surveillance. Seventy-three percent of office workers say AI and automation will increase their job satisfaction, but satisfaction is contingent on tools that actually work well.²³ When 60% of office workers report frustration with their tech tools and 55% say negative experiences with workplace technology impact their morale,²⁴ the discipline of look, feel, and flow of the interface may determine whether workers embrace dynamically generated interfaces or dismiss them as disorienting or even dystopian.

Delightful means responsive, coherent, and worthy of a worker's attention. That is the bar consumer technology has set—one that enterprise AI must now clear.



What's at stake

As AI agents multiply across the enterprise, the orchestration layer—the control tower view of all AI activity across the organization—becomes the surface on which organizational risk is either managed or missed. Visibility into what AI is doing, and the ability to intervene before failure cascades, are not optional. Organizations that treat the interface as an afterthought are taking a risk they can't afford. For the organizations getting it right, the interface is also where the upside lives. Four areas show the clearest promises.

1. Long-term competitive advantage

BCG's September 2025 research on future-built companies is unambiguous: Organizations that build for the future today achieve 1.7 times the revenue growth, 3.6 times the total shareholder return, and 1.6 times the EBIT margin of their laggard counterparts.²⁵

The divergence is accelerating. Gartner® predicts that "Up to 40% of enterprise applications will include integrated task-specific AI agents by the end of 2026, up from less than 5% today."²⁶ The organizations that pull

ahead will be the ones whose workers can actually use, trust, and course-correct agents. The interface unlocks that capability. The organizations that aren't defining the interface model will be playing catch-up later.

2. Return on technology already owned

Organizations that have invested in digitizing workflows onto platforms now hold clean, longitudinal, cross-functional records—an accumulated asset most have barely begun to leverage. As AI agents become capable of reasoning across that entire history, the interface becomes the mechanism that makes it accessible.

Dynamic interfaces surface what the platform already knows: the relevant case, the right precedent, the workflow that solved this problem before. Workers don't have to know where to look; the intelligence comes to them. That is not a new capability so much as the full realization of one already built.

3. Talent development and retention

Workers who use AI as a coach—not a crutch—build skills, rather than dependency. The interface design can make that distinction. The productivity and satisfaction benefits of well-made AI tools are thoroughly documented above in this brief.

Organizations that design for that experience will have workers who compound in value, whose ideas become reality faster, and whose measure of a productive day shifts from tasks completed to value created. The companies that do not will have workers going through the motions of AI adoption without any of the benefits. And in a talent market where the best people have options, the quality of the tools is increasingly a decision factor.



4. Security and governance

Gartner[®] predicts that “by 2028, 25% of enterprise breaches will be traced back to AI agent abuse, from both external and malicious internal actors.”²⁷ The exposure is not hypothetical. As agents gain autonomy, the attack surface expands proportionally.

But the governance problem runs deeper than security. Workers need interfaces that make AI reasoning visible, actions traceable, and decisions reversible. Without that visibility, accountability becomes theatrical, with workers signing off on outputs they cannot interrogate, and organizations unable to identify, trace, or reverse a bad agentic action before it compounds.


Only one in five companies currently has a mature governance model for autonomous AI agents.²⁸ The organizations closing that gap are building the audit trail, the explainability layer, and the human checkpoint into the experience itself and onto the interface, not as an afterthought bolted on by compliance.

Turning your workers into pilots

Let’s return to the security alert from this morning. In the enterprise that gets this right, an interface assembles the moment the alert fires: a dynamic view of the threat’s spread

across devices, the agents already containing it, and two decisions flagged for human authorization before the AI proceeds. The employee in charge sees what’s happening, understands what the AI is recommending and why, approves or redirects, and ensures the mission continues. Nothing is invisible. Nothing is irreversible. The whole exchange takes seconds.





That is the enterprise that works thanks to the right interface for agentic AI. The same threat is resolved, not by a passenger who handed off a chat prompt and hoped for the best, not by a crew stuck in transit waiting for manual clearance at every step, but by a pilot who authorized the mission, stayed informed through a dynamic interface, and intervened precisely where human judgment was required over autopilot.

Meanwhile, the business leader has a panoramic view from the control tower: which agents are active across the organization, what they've decided, where decisions still

“

The interface is where trust is built or broken, where accountability is maintained or lost, and where the productivity dividend either reaches the worker or disappears.

require human authorization, and where value created in one part of the business can be applied in another.

Readiness, strategy, culture, and data quality are all real factors shaping whether agentic AI delivers. No interface redesign substitutes for these factors. But the interface is the mechanism through which every other lever either works or doesn't. A perfectly designed agentic workflow running behind a black box is ungovernable. A capable AI model surfaced through a static, unintuitive interface is abandoned. The interface is where trust is built or broken, where accountability is maintained or lost, and where the productivity dividend either reaches the worker or disappears.

It is also an underinvested area. Organizations that have spent years on model selection, data infrastructure, and workflow redesign are often surfacing those investments through forms designed for a pre-AI world. That mismatch is fixable, and fixing it does not require starting over. It requires recognizing that the interface is as strategic as the model beneath it.

Your pilots are ready. The cockpit needs to catch up.



Endnotes

- ¹ Jessica Apotheker, et al., "The widening AI value gap: build for the future," Boston Consulting Group, September 2025, <https://media-publications.bcg.com/The-Widening-AI-Value-Gap-Sept-2025.pdf>.
- ² Alexander Sukharevsky, et al., "Seizing the agentic AI advantage," McKinsey & Company, June 2025, <https://www.mckinsey.com/capabilities/quantumblack/our-insights/seizing-the-agentic-ai-advantage>.
- ³ Atlantic Council, "AI, supply chains, and trade resets: the global economy in 2026," January 9 2026, 46:10, https://www.youtube.com/watch?v=zZHN0-ZNe_4.
- ⁴ Yaniv Leviathan, Dani Valevski, Vishnu Natchu, and Yossi Matias, "Generative UI: A rich, custom, visual interactive user experience for any prompt," Google Research Blog, November 18, 2025, <https://research.google/blog/generative-ui-a-rich-custom-visual-interactive-user-experience-for-any-prompt/>.
- ⁵ Thesys, "C1 Generative UI Engine," product documentation, <https://www.thesys.dev>.
- ⁶ "Sleep Coach," Garmin, Accessed March 18, 2025, <https://www.garmin.com/en-US/garmin-technology/health-science/sleep-coach/>.
- ⁷ Team Ohai, "Birthday reminder tools: remember birthdays and boost productivity," Ohai, August 11, 2025, <https://www.ohai.ai/blog/birthday-reminder/>.
- ⁸ GARTNER is a registered trademark and service mark of Gartner, Inc. and/or its affiliates in the U.S. and internationally and is used herein with permission. All rights reserved.
- ⁹ Gartner Press Release, "Gartner predicts over 20 percent of workplace apps will use AI-driven personalization algorithms for adaptive worker experiences by 2028," Gartner, March 12, 2025, <https://www.gartner.com/en/newsroom/press-releases/2025-03-12-gartner-predicts-over-20-percent-of-workplace-apps-will-use-ai-driven-personalization-algorithms-for-adaptive-worker-experiences-by-2028>.
- ¹⁰ Lisa Harkness, et al. "How generative AI can boost consumer marketing," McKinsey & Company, December 5, 2023, <https://www.mckinsey.com/capabilities/growth-marketing-and-sales/our-insights/how-generative-ai-can-boost-consumer-marketing>.

Endnotes

- ¹¹ Robert B. Miller, "Response time in man-computer conversational transactions," Proceedings of the AFIPS Fall Joint Computer Conference 33 (1968): 267–277, <https://yusufarlan.net/sites/yusufarlan.net/files/upload/content/Miller1968.pdf>.
- ¹² Focal, "AI search latency metrics: monitoring & optimization guide," blog post, September 9, 2024, <https://www.getfocal.co/post/ai-search-latency-metrics-monitoring-and-optimization-guide>.
- ¹³ Gartner Press Release, "Gartner predicts 80% of enterprise software and applications will be multimodal by 2030, up from less than 10% in 2024," Gartner, July 2, 2025, <https://www.gartner.com/en/newsroom/press-releases/2025-07-02-gartner-predicts-80-percent-of-enterprise-software-and-applications-will-be-multimodal-by-2030-up-from-less-than-10-in-2024>.
- ¹⁴ Carlo Giovine, Roger Roberts, Mara Pometti, and Medha Bankhwal. "Building AI trust: the key role of explainability," McKinsey & Company, November 26, 2024, <https://www.mckinsey.com/capabilities/quantumblack/our-insights/building-ai-trust-the-key-role-of-explainability>.
- ¹⁵ Nicole Gillespie, et al., Trust attitudes and use of artificial intelligence: a global study 2025, (Melbourne: The University of Melbourne and KPMG, 2025), 75, <https://assets.kpmg.com/content/dam/kpmgsites/xx/pdf/2025/05/trust-attitudes-and-use-of-ai-global-report.pdf>.
- ¹⁶ Acuta Margondai and Mustapha Mouloua, "The transparency paradox in explainable AI: a theory of autonomy depletion through cognitive load," University of Central Florida, January 20, 2026, <https://arxiv.org/pdf/2601.13973>.
- ¹⁷ Mikaël Chelli, et al. "Hallucination rates and reference accuracy of ChatGPT and Bard for systematic reviews: Comparative Analysis," Journal of Medical Internet Research 26 (May 22, 2024), e53164, <https://doi.org/10.2196/53164>.
- ¹⁸ Deloitte TrustID survey, reported in "Workers don't trust AI. Here's how companies can change that." Harvard Business Review, November 7, 2025, <https://hbr.org/2025/11/workers-dont-trust-ai-heres-how-companies-can-change-that>.
- ¹⁹ Carlo Giovine, et al. "Building AI trust: the key role of explainability," McKinsey & Company, November 26, 2024, <https://www.mckinsey.com/capabilities/quantumblack/our-insights/building-ai-trust-the-key-role-of-explainability>.

Endnotes

²⁰ Iris Glassberg, Yael Brender Ilan, and Moti Zwilling, "The key role of design and transparency in enhancing Trust in AI-powered digital agents," *Journal of Innovation & Knowledge* 10, no. 5 (2025), <https://doi.org/10.1016/j.jik.2025.100770>.

²¹ CNBC International Live, "This robotic therapy seal is revolutionizing elderly care," YouTube video, 3:45. July 31, 2023, https://www.youtube.com/watch?v=_mj_-8yFxxA.

²² Jyoti Mann, "Meta is about to start grading workers on their AI skills," *Business Insider*, November 14, 2025, <https://www.businessinsider.com/meta-ai-employee-performance-review-overhaul-2025-11>.

²³ Eliza Sarasohn, "Despite AI enthusiasm, workforce index reveals workers aren't yet unlocking its benefits," *Workforce Lab at Slack*, October 16, 2024, <https://slack.com/blog/news/the-workforce-index-june-2024>.

²⁴ Ivanti, 2024 Digital Employee Experience Report: A CIO Call to Action (South Jordan, UT: Ivanti, September 2024), <https://www.ivanti.com/resources/research-reports/employee-experience-digital-transformation>.

²⁵ Jessica Apotheker, et al. "The widening AI value gap: build for the future," Boston Consulting Group, September 2025, <https://media-publications.bcg.com/The-Widening-AI-Value-Gap-Sept-2025.pdf>.

²⁶ Gartner Press Release, "Gartner predicts 40 percent of enterprise apps will feature task-specific AI agents by 2026, up from less than 5 percent in 2025," Gartner, August 26, 2025, <https://www.gartner.com/en/newsroom/press-releases/2025-08-26-gartner-predicts-40-percent-of-enterprise-apps-will-feature-task-specific-ai-agents-by-2026-up-from-less-than-5-percent-in-2025>.

²⁷ Gartner Press Release, "Gartner unveils top predictions for IT organizations and users in 2025 and beyond," Gartner, October 22, 2024, <https://www.gartner.com/en/newsroom/press-releases/2024-10-22-gartner-unveils-top-predictions-for-it-organizations-and-users-in-2025-and-beyond>.

²⁸ Deloitte, *State of AI in the Enterprise* (New York: Deloitte Insights, 2026), <https://www.deloitte.com/content/dam/assets-zone2/cz-sk/cs/docs/services/consulting/ai/state-of-ai-report.pdf>.

The Innovation Office

The Innovation Office is the team behind ServiceNow Futures. Our forward-looking research, thought leadership, and platform builds serve as the blueprint for AI-first business reinvention. We help our colleagues and customers see further so they can act sooner.