



Forward Deployed Engineers: Closing the gap between AI pilots and value

Why the right FDE team accelerates enterprise AI transformation

September 2026

Forward deployed engineers have become one of the most talked-about yet least-understood roles in enterprise AI. Reinvented versions of the model appear in service launches, partnership announcements, and vendor pitches, with each defining the role differently. But for all the growing attention, there is still relatively little actionable guidance on what forward deployed engineering looks like in practice, or what companies should expect from the special-forces-like teams custom-built to deliver on value.


That lack of clarity matters. For companies considering whether to engage a team of FDEs, the most important consideration is what good looks like now, and how to tell the difference between a team that will accelerate transformation and one that won't.

KPMG has observed that the most effective FDE teams share a set of common characteristics: They're multidisciplinary, embedded directly with business units, and structured to move from ambiguity to measurable impact quickly. But those traits alone don't guarantee results. How FDE teams are deployed, and against which problems, matter just as much as how they're built.

“

The strongest FDE teams are built to increase speed-to-value and are structured as adaptive pods designed for the lifecycle stage and needs of each engagement.

The next frontier: Enabling AI at scale



It's not plug-and-play, and the business cost of getting it wrong, whether through hallucination, compliance violation, or process error, can be severe.

Experimentation defined the early wave of enterprise AI investment. Organizations across industries tested genAI against dozens of use cases, funded pilots broadly, and accepted that not every bet would pay off. That approach worked as designed. It surfaced the high-potential applications, built organizational familiarity with the technology, and established the foundation for next steps.

But now, the emphasis has shifted to taking the use cases that showed early promise and deploying them as production-grade enterprise systems. Yet the gap between a successful pilot and a scaled operational system remains the central challenge: According to the KPMG AI Quarterly Pulse Survey, more than half of companies have deployed new AI systems, but only 4 percent are orchestrating AI across the enterprise.ⁱ Agentic AI systems that make recommendations, update records, trigger workflows, and execute decisions with limited human oversight are widening

the execution gap. The enterprise demands placed on agentic AI are correspondingly more rigorous: An AI agent operating autonomously within a production environment must navigate legacy system integrations, fragmented real-world data, stringent security protocols, and auditable decision trails.

Organizations recognize this. Many have moved toward more structured approaches that incorporate a federated governance framework, where business units are accountable for value realization while a central function ensures trust, compliance, and scalability. That maturity is a positive development, but it also raises the bar for delivery. Internal engineering teams, already stretched thin maintaining existing systems, may lack the specialized expertise to build and manage bespoke agentic AI solutions. Traditional delivery methods designed for predictable software development cycles struggle with the ambiguity, iteration speed, and cross-functional integration that scaling AI demands.

ⁱ KPMG AI Quarterly Pulse Survey. KPMG LLP. Available at: kpmg.com/us/en/articles/2025/ai-quarterly-pulse-survey.html

Today's FDE is built for the scaling challenge

The best version of the forward deployed engineer is purpose-built for the unique demands of the AI landscape today. While FDEs have been around for years, the FDE model that has emerged can respond directly to the specific challenges of deploying generative and agentic AI within complex enterprise environments, where federated governance and compliance requirements demand more than a plug-and-play approach.

KPMG has observed that the most effective FDE teams work in high-ambiguity environments where they can build in partnership with the business. Because they are embedded directly within functions, they can design solutions together, diagnosing complex problems and iterating potential solutions rapidly along the way. Their core mission is to bridge the chasm between a promising AI model and a fully operational, value-generating system. A key part of their methodology is to transfer capability back to internal teams, ensuring that the solutions the FDE teams create are impactful in the short term and sustainable over the long haul.

Not all FDE teams are built the same way. Some originate from the technology labs themselves, with direct access to product teams and the ability to influence the underlying model—adjusting context layers or tuning performance in ways a typical technology consumer cannot. Others come from services firms, bringing deeper business-domain expertise, value-measurement discipline, and the ability to work within an enterprise's governance and compliance frameworks. Increasingly, the most powerful arrangements combine the two: joint teams that pair platform-level access with business-value orientation.



The strength of an FDE team lies in its multidisciplinary structure, employing the right mix of specialists who can simultaneously and interactively address the expanding challenges of scaling an AI system.

Engineers, data scientists, and AI developers

This technologically expert core builds the crucial “last mile” architecture, bridging the gap between a working pilot and a scalable system complete with necessary audit trails, access controls, and compliance checks. These team members construct the robust data pipelines, API integrations, and security guardrails necessary for an AI agent to function securely and effectively within a company’s existing technology stack. They don’t just deploy a model; they build the entire ecosystem, including restructured data environments, that it needs to thrive.


Domain experts

These individuals bring business reality to what the engineers create by re-building with an AI-native approach, delivering real transformation. These insiders provide deep context that helps prevent model hallucinations and help ensure the AI agent’s logic aligns with enterprise workflows, regulatory requirements, and strategic objectives. They can also validate that the team is solving the right challenge—one that is central to the business’s operation, rather than just an interesting technical problem.

Value modelers

This role is the team’s strategic compass. Value modelers go beyond identifying potential by ruthlessly modeling and measuring ROI. They work with the organization to define success and to establish specific, concrete metrics that can be built into the system. By focusing the team on features that directly move the bottom line, they ensure that every development sprint cycle contributes to tangible, measurable business value.

AI pilot-to-value realization flow



This integrated approach allows the FDE team to **de-risk AI investments by tackling technical feasibility, business viability, and value realization** in parallel. Although FDE teams are not, on their own, the answer to the question of how to scale effective agentic AI systems, they can help **establish a path to production and ROI** when paired with the complementary capabilities—such as **AI factories, change adoption, and corporate learning**—that round out a full scaling strategy.

Setting loose FDE teams for maximum impact


An FDE team, on its own, can't solve all of the challenges of creating and scaling productive and profitable AI systems. AI factories typically are tasked with developing AI models suited to a business's requirements, creating and training data pipelines, and engineering platforms. Change adoption helps ensure that organizations understand and embrace AI agents and integrate them into daily operations. Corporate learning is needed to help employees acquire, apply, and refine new skills necessary to monitor and support AI systems.

Yet for most organizations, FDE teams can unlock the potential of advanced AI models, getting them into production quickly and realizing value. Several key questions may be part of that process.

Technological complexity: Does the project require generative AI, agentic AI, or other advanced technologies that push the boundaries of the organization's current capabilities?

Clear business problem: Has the enterprise challenge been clearly identified and articulated, even if the path to a solution is uncertain?

Commitment to collaboration: Is the company prepared to make a dedicated team of business and technological stakeholders available for rapid, collaborative prototyping and continuous feedback?



A rigorous intake process can help leaders choose which high-value challenges may call for the unique capabilities of an FDE pod.

When these conditions are met, FDE teams can unlock transformative outcomes

The challenge

01

An AI prototype works brilliantly in a controlled environment but requires complex reasoning across multiple secured, legacy databases to be truly useful.

02

A company wants to use agentic AI to automate a high-value workflow, such as vendor contract negotiation or complex financial auditing, but the business cost of a model hallucination or a process error would be catastrophic.

03

The C-suite has mandated an AI transformation, but business units are paralyzed by the sheer number of choices and lack the technological literacy to identify which workflows are primed for agentic automation.

The impact

The FDE pod builds the robust data scaffolding and security layers needed. FDEs architect the intricate connections between the new AI and the old systems, transitioning the prototype from a siloed tool into a secure, integrated enterprise asset.


Embedded FDEs build sophisticated “human-in-the-loop” (HITL) interfaces, design rigorous evaluation frameworks, and implement multi-layered guardrails. This ensures the system is not only powerful but also trustworthy enough for a production environment in which stakes are high.

An FDE team embeds into key business units for four-to-six-week discovery and prototyping sprints. The FDEs identify bottlenecks, test technological feasibility on the fly, and in multiple cases, can deliver a functioning, ROI-positive AI workflow that provides immediate value and builds momentum for broader transformation.

In situations such as these, by operating at the intersection of strategy, technology, and business operations, FDE teams are positioned to **deliver concrete results**. FDE teams from Salesforce and Palantir, for example, delivered significant improvements for their customers, including **launching agents in a week** and **cutting production time by one-third**.ⁱⁱ In fact, that acceleration is not just about moving faster; it's also about **realizing value sooner**.

ⁱⁱ Ashutosh Sharma, et al., “Forward-Deployed Engineers: Use Provider-Embedded Builders To Improve Outcomes,” Forrester, June 22, 2026, p. 6.

The strategic deployment of talent



“

The leaders who learn to deploy FDE talent against the right problems and alongside the right complementary capabilities will be best positioned to translate AI's promise into durable business advantage.

The difference between a company that merely experiments with AI and one that harnesses the ever-expanding scope and power of technology to transform its business is only partly explained by the AI models they buy. More important, in many cases, is how an organization deploys its talent. As industries move resolutely into the scaling phase of the AI revolution, strategically unleashing expert, specialized, multidisciplinary teams has become critical for success.

The FDE model provides a proven framework for navigating the scaling paradox, turning promising AI pilots into production-grade systems that deliver real, measurable business value. For business leaders, perhaps the most important challenge is to identify the high-value, high-complexity scenarios in which FDEs are most likely to have a transformative impact, and to pair them with the AI factories, change adoption programs, and learning capabilities that round out a complete scaling strategy. By reserving this model for their most essential AI missions, companies can de-risk their investments, accelerate change, and unlock the true, scalable promise of artificial intelligence.

As agentic AI matures and lab-originating FDE teams increasingly partner with services-firm specialists, the model itself will continue to evolve.

How KPMG can help

KPMG supports clients in building robust, full-stack AI architectures by providing a thorough, wide-ranging approach that spans strategy, implementation, workforce transformation, and trust. Leveraging established frameworks and advanced tools, KPMG helps organizations shape tailored AI strategies, develop actionable roadmaps, and construct sustainable AI solutions alongside the necessary data infrastructure. An AI factory approach enables rapid piloting and scaling of agentic AI agents, while streamlining integration

with emerging technologies and platforms. Throughout the journey, KPMG emphasizes ethical, secure, and compliant AI deployment by embedding trusted AI practices and risk management into every stage. A people-centric methodology helps ensure that both human and AI collaborative efforts are optimized, unlocking value, accelerating adoption, and empowering enterprises to confidently infuse AI capabilities across their operations.

Learn more: <https://www.kpmg.us/AIservices>

Authors



Matteo Colombo

Global Leader for Cloud, Data, & AI,
KPMG US

E: matteocolombo@kpmg.com

Matteo serves as the Global Leader for Digital Technologies at KPMG. In this pivotal role, he spearheads the growth, transformation, and innovation across KPMG's Digital Foundation and the global Centers of Excellence for Cloud, Data, and AI. Matteo plays an instrumental role in shaping the use of technology in the firm's collective strategy and ensuring its effective implementation in local markets. Additionally, he manages relationships with key strategic technology partners. Matteo's profound expertise encompasses both strategy and technology, with a particular emphasis on emerging trends that drive growth and innovation for clients and partners. He is a staunch advocate for the effective and ethical use of AI. With over twenty years of experience, Matteo has adeptly guided large enterprises in harnessing cutting-edge technologies to achieve significant, large-scale transformations.



Rachel Wagner-Kaiser

Director & FDE Tech Lead, AI &
Data Labs, KPMG US



E: rwagnerkaiser@kpmg.com

Rachel serves as a Director and Data Scientist in KPMG's AI & Data Labs. In this role, she leads the design, build, and deployment of agentic AI solutions that analyze risk, streamline business processes, implement automation, and identify cost savings for clients across industries. Rachel currently supports global and US forward deployed engineering efforts, overseeing FDE pod deployments to drive innovation and deliver impact at scale. Her areas of expertise include natural language processing and agentic AI solutions, with particular depth in contract management and commercial leakage analysis. With extensive experience in applied machine learning, Rachel is dedicated to helping organizations harness AI to solve complex challenges and achieve meaningful transformation.

Contributors

The authors thank Lisa Bigelow, David Bumke, Mary Lewis, John Nelson, and Rich Schulte for their contributions to this paper.

Related insights



Some or all of the services described herein may not be permissible for KPMG audit clients and their affiliates or related entities.

Learn about us:



kpmg.com

The information contained herein is of a general nature and is not intended to address the circumstances of any particular individual or entity. Although we endeavor to provide accurate and timely information, there can be no guarantee that such information is accurate as of the date it is received or that it will continue to be accurate in the future. No one should act on such information without appropriate professional advice after a thorough examination of the particular situation.

© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved. The KPMG name and logo are trademarks used under license by the independent member firms of the KPMG global organization. USCS043396-1A