



Agentic AI is revolutionizing software development

How human-machine collaboration is unlocking speed, quality, and strategic impact across the SDLC playbook

Introduction

Software development has always advanced through layers of abstraction that simplify complexity. But as hand-coding gives way to AI-driven platforms, coding agents represent the most significant innovation in software development history.

Transitioning to an agent-driven software development lifecycle (SDLC) won't happen overnight. Although KPMG has observed coding assist tools boost productivity by 10 to 30 percent, they create fragmented workflows. Developers working in separate context windows produce code that later requires stitching together, which often results in inconsistent quality. And, although this approach may suffice for new builds, it fails to truly transform the SDLC.

True modernization comes from embracing an agentic, AI-powered SDLC that eliminates drift and accelerates delivery. In this model, developers shift from writing code to engineering agents that build and maintain applications. They become engineers, last-mile problem solvers, and quality assurance specialists, ensuring agents function effectively within a single, unified context window. It's a harmonious approach, and we believe organizations that adopt this strategy can unlock cost efficiencies of 50 percent or more while also reducing time to market.

The vision is compelling, but the conflation of personal productivity tools with autonomous agents makes progress slow going. Consider, for example, the differences between classes of agents. Class 1 (RPA 2.0) agents focus on task-level automation within existing systems. Class 2 agents introduce agentic systems embedded in specific processes, such as sales portals that deliver next-best-action recommendations. And in Class 3, agents orchestrate cross-domain, transformative ecosystems that fundamentally rewire processes and roles.

As SDLC components become agent driven, the work performed by Class 3 orchestration agents becomes critical. Think autonomous driving: Without robust exception handling, even the smartest systems fail. The same applies to an AI-native SDLC.

A unified agent ecosystem transforms software development from a back-office function into a strategic engine for innovation, agility, and competitive advantage. In this paper, we'll help technology decision makers understand how coordinating intelligent agents across the SDLC can streamline workflows from initial mapping to iterative updates, and close the gap between ambition and outcome.



Agent-driven development in a modern SDLC

Imagine a world where a small team of human developers deploys intelligent agents around the clock to manage the SDLC, from planning to remediation. Other agents oversee compliance, data privacy, and business requirements. At the center are human managers guiding the agentic systems.

This shift transforms both technology and roles. Developers, instead of cranking code, evolve into supervisors of asynchronous coding agents to ensure quality. At the same time, a new role emerges: The human engineer who designs and maintains agents of varying capabilities. The result is a rebalanced workforce comprised of fewer traditional human coders but more engineers, with developers elevated into strategic positions that oversee and maintain agents.

In this refreshed model, coordinated agent teams create requirements and generate, test, and refine code at a speed and scale no individual developer can match. Specialized agents perform distinct roles: Just as organizations have different tiers of human managers, there will be hierarchical levels of agents, with some performing tasks, others managing other agents, and still others collaborating directly with human teams. This refreshed model builds on the TACO framework, which categorizes agents as taskers, automators, collaborators, and orchestrators, while also introducing clearer role and capability distinctions to meet the demands of a more defined SDLC.

The agentified SDLC manages dependencies and triggers cascading updates automatically. Even requirements-gathering changes as legacy code is analyzed to autogenerate baseline requirements, freeing experts to focus on innovation. The outcome is profound: Software development moves from a fragmented process to orchestrated automation, delivering speed, precision, and agility at peak performance.

The TACO framework

KPMG developed the TACO framework—taskers, automators, collaborators, and orchestrators—to help organizations understand and leverage the evolving landscape of agentic AI.

Taskers are single-goal executors designed to complete well-defined, repetitive tasks with minimal complexity. They follow detailed instructions provided by a human and focus on singular goals.

Automators are more advanced cross-system process integrators that manage and integrate workflows across multiple systems and applications. These agents automate predictable, end-to-end business processes with clear outcomes.

Collaborators are adaptive AI teammates that interact with human operators to achieve more complex goals that may evolve over time. These agents learn from interactions and refine their outputs based on user feedback.

Orchestrators manage and coordinate multiple agents and complex workflows across interdependent systems. Orchestrators excel in dynamic resource optimization and can manage task allocation and communication between different agents in real time.



How agentic AI transforms the SDLC

The SDLC has long relied on two dominant methodologies: Waterfall, which enforces sequential rigor and remains common in regulated environments like banking and critical healthcare systems, and Agile, which thrives on iteration and drives fast-moving sectors such as e-commerce. An agentic SDLC must combine the best of both—Waterfall's precision and compliance with Agile's adaptability and speed.

Planning AI automates feasibility analysis using historical data, market intelligence, and predictive modeling. It simulates multiple business scenarios to optimize cost, compliance, and time to market. Agents cross-reference objectives, key results, and regulations for strategic alignment, accelerating decisions and cutting executive overhead.

Requirements analysis Requirements come from legacy code, process narratives, and architecture designs while agents interview subject matter experts for further direction and insight. GenAI, with strong governance, can validate requirements against frameworks such as HIPAA and PCI DSS.

Design Agents accelerate design and create higher-quality outputs by generating architectural options optimized for cost, resilience, and compliance. They can produce multiple UX prototypes aligned with brand guidelines and user behavior analytics, and simulate load and security scenarios before coding even begins.

Development Autonomous coding agents write production-grade code, self-optimize pipelines, and predict integration conflicts. Agent guardrails enforce adherence to development frameworks and policies. Agents also manage dependencies and document their actions.

Use case: Launching a new digital payments platform

Under tight deadlines, a cross-functional team at a fintech company replaced the traditional SDLC with an agent-driven approach to boost speed and quality.

Agents handled planning by analyzing customer data, compliance rules, and past outcomes to create a prioritized roadmap, refined by product owners. They flagged conflicting requirements, proposed fixes, and generated architecture options, wireframes, and UML diagrams for designers to finalize.

In implementation, agents wrote production-ready code for core payment features, while developers validated logic and integrated APIs, cutting development time by more than half. Testing and QA were automated, catching a rare concurrency issue before launch.

Deployment and post-launch were seamless: Agents generated integration scripts, monitored performance, patched vulnerabilities, and compiled metrics and feedback for future improvements.

Quality management Agents can monitor performance, run real-time tests, and apply fixes automatically, reducing defects throughout the SDLC and accelerating release readiness.

Deployment Agents manage blue-green or canary deployments, monitor KPIs, and trigger rollbacks if anomalies arise. They even optimize timing based on traffic patterns—helping ensure safer, faster releases.


Maintenance and support Maintenance agents access and leverage a real-time knowledge base of applications, keeping them evergreen and linked to support capabilities such as real-time security reviews and compliance testing. These specialized agents are trained and qualified to automate the detection and repair of anomalies.

Retirement Automated decommissioning ensures secure data handling, compliance checks, and cost optimization—reducing risk and accelerating transitions to next-gen systems.

Overcoming challenges

As technical employees learn how to develop, train, and manage their digital coworkers, they will continue to play a vital role in reviewing and validating agent-generated code. While agents excel at speed and scale, they often struggle with complex, context-sensitive tasks that require judgment and nuance. At the same time, developers must adapt to working alongside autonomous agents, shifting from manual execution to strategic oversight.

Privacy and regulatory constraints create other challenges, especially in industries such as financial services and healthcare where PII is a concern. However, organizations should consider AI as a present-risk tradeoff as opposed to a net new risk. This is one reason why collaboration between humans and agents is especially critical. When it comes to security, agents can flag vulnerabilities, but humans own accountability.



Skills gaps and workforce readiness

Like any enterprise software, AI tools require structured onboarding, hands-on training, and ongoing support to ensure teams can use them effectively. Without this foundation, productivity gains may be delayed or diminished. Developers unfamiliar with these platforms often face a steep learning curve, which can reduce early impact.

Security and data privacy risks

AI tools often require access to sensitive data to function effectively but concerns around AI and data privacy have surged from 43 percent to 69 percent in just one quarter.¹ To mitigate these risks, organizations must implement robust safeguards such as role-based access controls, encrypted datasets, and modern governance.

Implementation costs and ROI concerns

Deploying AI across the SDLC requires upfront investment in tools, infrastructure, and training. To justify these costs, organizations should conduct ROI assessments and prioritize high-impact use cases, such as automating regression testing or streamlining manual QA processes. These areas often yield quick wins and measurable returns, helping build confidence and momentum for broader adoption.

Integrating with existing development processes

Scaling AI requires rethinking how agents fit into established methodologies such as Agile, Waterfall, and Spiral. Conducting AI-readiness assessments can help identify gaps, determine agent types, and guide phased implementation.

¹ Source: Trust, attitudes and use of artificial intelligence: A global study 2025, KPMG, 2025.

How KPMG can help

KPMG can help you transform your software development lifecycle with intelligent agents. With agentic systems, the benefits are clear—greater efficiency, reduced risk, and assured compliance. Our skills, technology, and talent will guide your AI-driven SDLC journey from vision to reality. For organizations ready to embrace the agentic era, the future of development is faster, smarter, and more dynamic.

Technology that assists with enabling an agentic, AI-powered SDLC

KPMG Blaze, powered by Ignite, is our proprietary multiagent orchestration tool that accelerates and automates the SDLC. Blaze enables context engineering, maintaining intent from the requirements phase throughout the development process, thereby reducing errors and miscommunication. Blaze also specializes in modernizing legacy systems, transforming outdated source code into modern technology stacks.

You can win with AI

- KPMG was named a “Leader” in Worldwide Data Modernization IDC MarketScape: Worldwide Data Modernization Services 2024 Vendor Assessment
- KPMG is ranked #1 for quality AI advice and implementation in the US.



Authors



Mark Shank

Principal, Advisory

Mark brings metric moving Digital Experiences to employees and customers. He specializes in building and operating Digital and Mobile capabilities within organizations. Bringing modern Digital and Mobile experiences to a workforce or customer base has presented a unique set of problems that only an integrated team of Creatives, Technologists, and Strategists can solve. Mark builds, leads and runs this capability for organizations that need many of these teams.



Chris Panneck

Principal, Advisory

Chris leads the KPMG Financial Services Technology practice and has assisted some of the largest technology organizations in the world as they navigate significant transformations of their IT and business operations. He has over 20 years of experience in technology, strategy, and operations. Chris has led multiple technology transformations over his career including developing a modern digital banking platform for a leading global bank, one of the largest infrastructure re-platforming efforts in the U.S. and a global delivery model transformation for one of the largest IT organizations in the world. He focuses on serving our Financial Services clients and assists senior technology leaders build and implement their technology agenda.

We would like to thank our contributors:

Lisa Bigelow, Whitney La Bounty, Mary Lewis, Michael Ogden, Joseph Schnering, and Rich Schulte for their contributions to this paper.

For more information, contact us:

Mark Shank

Principal, Advisory

410-949-2795

mshank@kpmg.com


Chris Panneck

Principal, Advisory


214-840-6781

cpanneck@kpmg.com

Related thought leadership:



Data governance in the age of AI



Trust, attitudes and use of Artificial Intelligence



The agentic AI advantage: Unlocking the next level of AI value

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

Please visit us:



[kpmg.com](https://www.kpmg.com)



Subscribe

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 upon such information without appropriate professional advice after a thorough examination of the particular situation.

© 2025 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.

DASD-2024-19220