KPMG

How AI can help reduce tech debt in M&A

2025 Technology M&A Survey

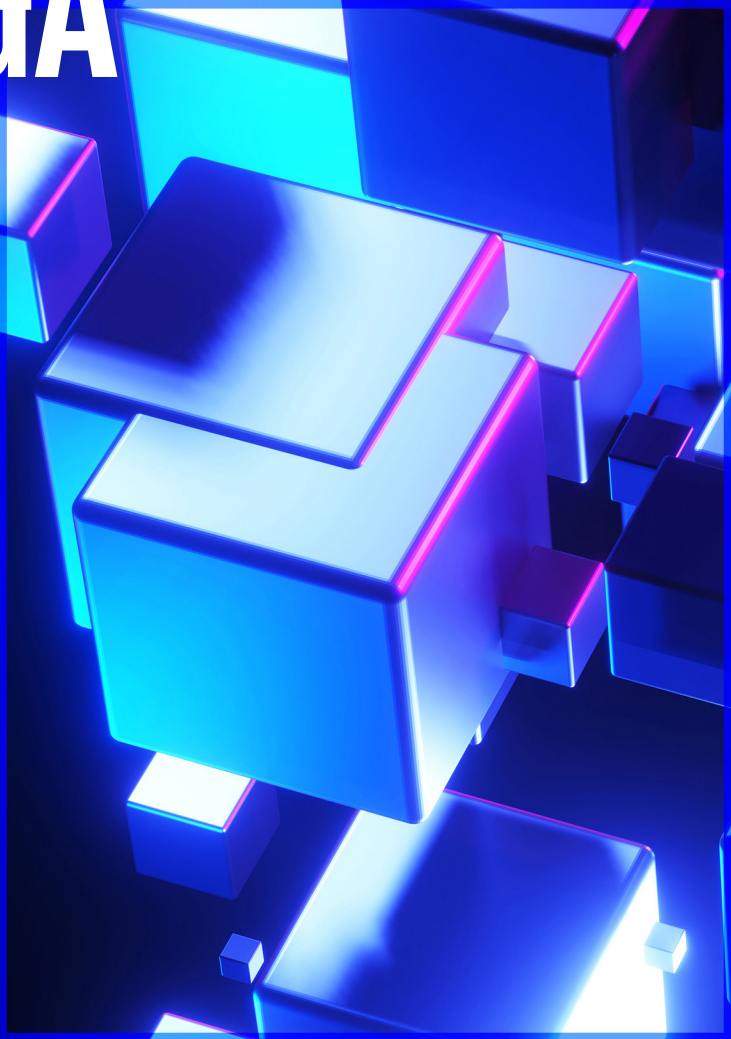

kpmg.com

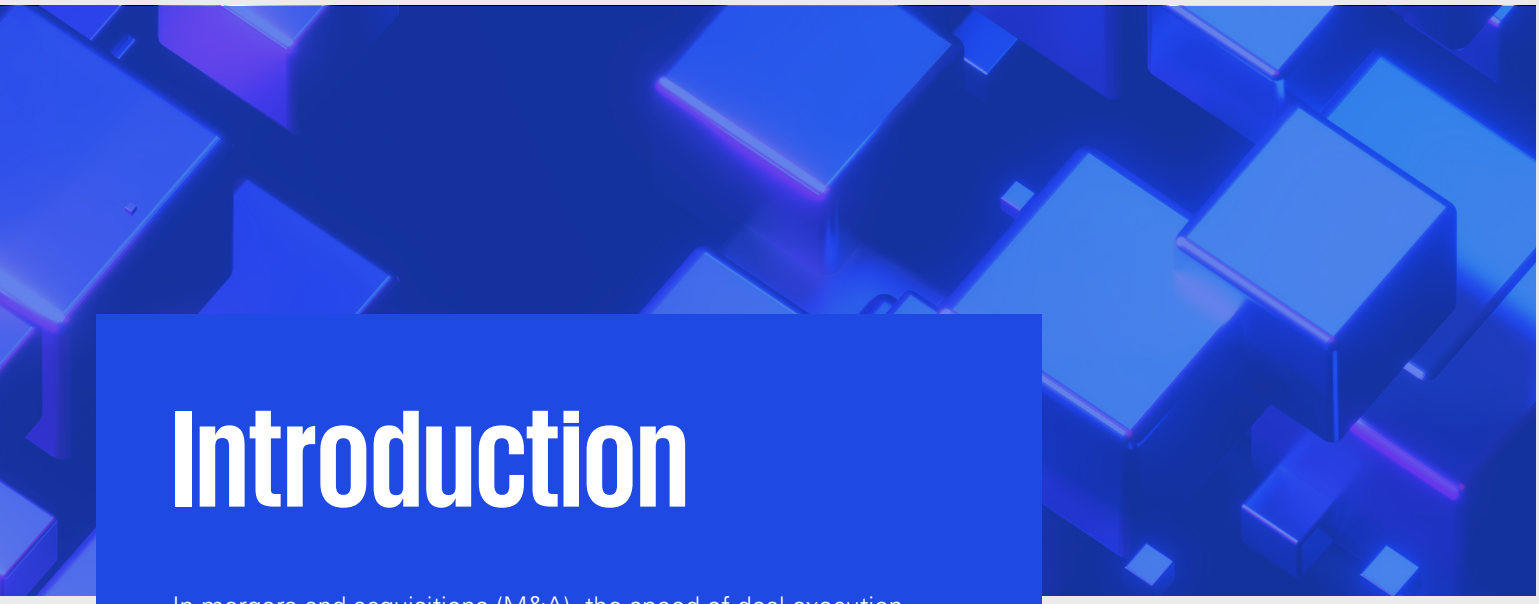

Introduction

In mergers and acquisitions (M&A), the speed of deal execution is of the essence to secure maximum value. But when it comes to avoiding technical debt or "tech debt" from the transaction, the opposite is true. The challenge of integrating and updating multiple inherited software systems may tempt dealmakers to reach for quick fixes. In the long run, however, the cost of taking shortcuts to address problems that require comprehensive, well-considered solutions will inevitably mount like financial debt. This tech debt will not only be more expensive to repay, but any delays will also impede the merged organization's ability to innovate and create lasting value.

Dealmakers and corporate strategists are increasingly becoming aware of tech debt as an M&A risk. Yet not enough players are addressing it during deal preparation, exposing themselves to unpleasant post-deal tech debt surprises. We conducted the KPMG 2025 Technology M&A Survey to better understand dealmakers' current thinking around tech debt (see Methodology) and presented the key takeaways at our annual Technology M&A Conference on November 6, 2025.

This report delves deeper into our survey findings to examine how technology and software companies, as well as investors, view the problem of tech debt in M&A—and how they are addressing it. One key development is the increasing use of artificial intelligence (AI) tools. However, while AI offers promising solutions, it is not a cure-all and brings its own set of challenges. To scale AI safely and consistently across the enterprise, organizations are beginning to explore more structured ways especially as fragmented pilot projects introduce new operational risks.

To help dealmakers navigate these complexities, we offer our perspective on leading practices for deploying AI to manage tech debt. For starters, we encourage organizations to embed AI early in the M&A process and establish strong governance, as poor implementation can actually compound tech debt.

Methodology

In September 2025, KPMG LLP surveyed 135 US-based organizations about how they are managing the problem of tech debt in M&A transactions. Respondents represented a mix of public and private companies, as well as private equity (PE) and venture capital (VC) firms and entities backed by them. About three-quarters (104) worked at technology/software enterprises, and the other one-quarter (31) came from PE/VC firms. There was limited input from independent startups.

Corporate participants were almost evenly split between mid-sized (less than $1 billion in revenue) and large companies (more than $1 billion). Among PE/VC firms, 81 percent managed more than $10 billion in assets. Respondents held senior roles, primarily in technical leadership, executive management, and investment teams.

Tech debt is rarely planned for pre-sign and often addressed post-deal close

Tech debt is a deal-shaping risk but under-addressed during the pre-signing phase. Nearly two-thirds of respondents said they discussed tech debt always or often during deal planning, but less than one-third deemed it of utmost importance to address it as part of pre-deal evaluation (Exhibit 1). This was even though most respondents believed unresolved tech debt can lead to operational disruptions (70 percent) and cyberattacks (60 percent).

Some organizations are increasingly using agent-assisted diligence approaches to surface tech-debt hotspots before signing. Agents can analyze incident history, legacy architectures, and regulatory obligations, auto-derive baseline requirements from code, and simulate modernization scenarios for cost, risk, and time-to-market. These capabilities shift M&A planning from assumptions to more evidence-driven integration roadmaps.

Exhibit 1. Early tech debt planning remains the exception

30%

of respondents indicated high importance of addressing tech debt topics during pre-deal phase

70%

of respondents cited significant disruption from tech debt—including downtime and poor user experience

60%

of respondents citing cyber risk

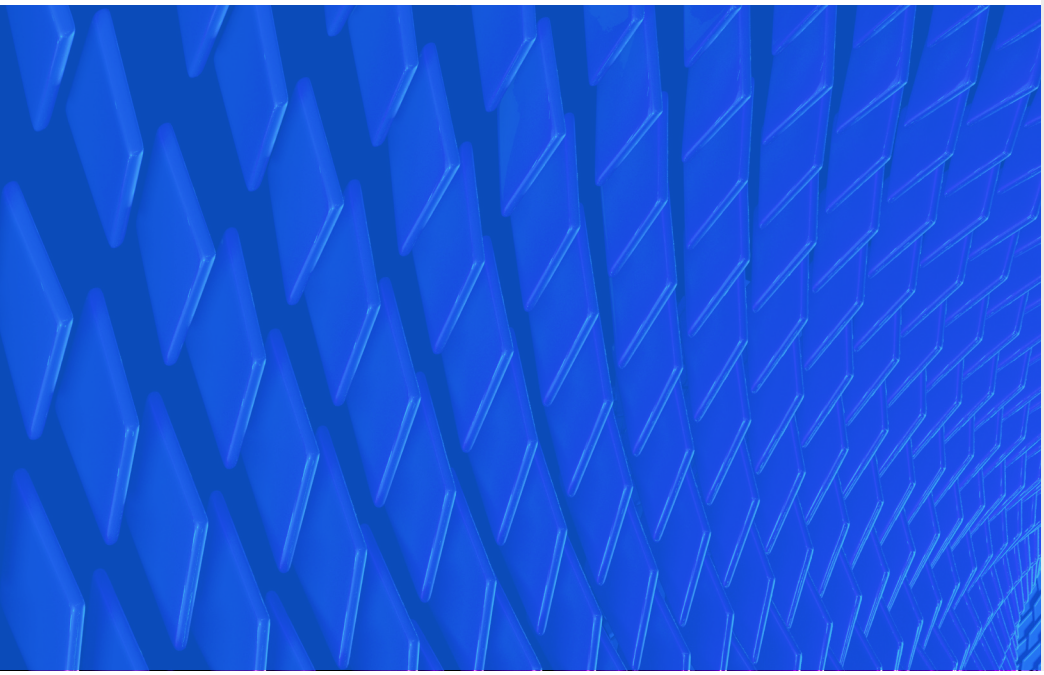

AI is tactically accelerating modernization but also amplifying fragmentation

Addressing tech debt is an expensive and time-consuming undertaking for any organization, but now they can tap the power of ever-more capable AI. Two-thirds of respondents said they are currently using or piloting AI tools to help mitigate these issues, and early results have been promising. For instance, 55 percent reported up to a 10 percent reduction in corporate IT costs after implementing AI solutions.

However, the adoption of AI to relieve tech debt remains largely tactical, not transformative. In our survey, 62 percent said they expect only incremental improvements in the next two years; just 27 percent anticipate major transformation (Exhibit 2). This is because AI deployment against tech debt is limited to specific tasks rather than for enterprise-wide applications.

Organizations looking to scale AI safely need to adopt more structured operating models to coordinate how different agents perform work. These models typically define clear agent roles (e.g., task-execution, workflow automation, human-in-the-loop judgment, and orchestration). Orchestrators provide centralized guardrails, routing, and exception handling, helping shift AI efforts from scattered point solutions to cohesive, multi-agent systems.

In addition, scaled-up deployment potentially comes with hidden risks. Organizations face cultural resistance, concerns about data governance, and integration difficulties. Worse yet, tactical AI adoption with disjointed pilot projects and hybrid in-house/external vendor tool mixes may end up creating additional tech debt, or "AI debt." Uncoordinated rollouts of AI across various departments are likely to lead to redundant tools and fragmented ownership, all of which will drive tech debt. In fact, 59 percent of respondents already see this "agent sprawl" as an emerging risk, mirroring the concerns over cloud sprawl in the early days of cloud services. During M&A integration, managing multiple AI tools will be especially difficult, and 61 percent of surveyed executives fear this sprawl will increase operational complexity.

Exhibit 2. AI adoption to relieve tech debt isn't yet transformational and can add to risk

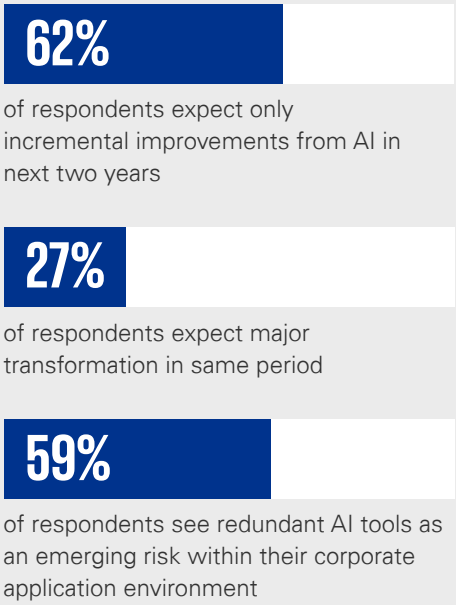

Most companies still bleed engineering capacity on legacy system issues

AI clearly can be a game changer in paying down tech debt. The most common use of AI is automated code refactoring or legacy code analysis (52 percent). But our survey found that the cost of fixing tech debt can stretch into hundreds of millions of dollars, and 55 percent of respondents said battling common problems required 20-40 percent of the engineering staff's time (Exhibit 3).

A more modern approach applies an "agentified software development life cycle (SDLC)" model, where autonomous coding and testing agents generate and validate code under policy guardrails; orchestrators coordinate workflows; and release agents manage controlled deployments. This shifts engineering effort from manual remediation to supervising agent teams, enabling modernization at scale.

Consequently, even as AI streamlines code fixes and migrations, few companies are using it to modernize architectures. For that, more advanced approaches to reducing tech debt, such as API[1] standardization, predictive maintenance, and data cleanup, would be needed.

[1] Application programming interface

Exhibit 3. Engineering staff spent too much time addressing past problems

52%

of respondents implemented or piloted automated code refactoring or legacy code analysis

55%

of respondents said up to 40 percent of engineering staff's time has been spent fixing tech debt in the past year

AI and modernization are accelerating faster than governance

AI can help organizations modernize faster, but without good governance and integration discipline, its use may exacerbate gaps in software infrastructure that contribute to tech debt accumulation in the first place. At most enterprises, AI governance is still ad-hoc and compliance-driven, lacking a disciplined framework for tech-debt prevention. In our survey, only 21 percent said they have put centralized oversight of AI in place, with the rest saying governance remains fragmented across functions (Exhibit 4).

Organizations adopting AI at scale typically need to establish role-based access rules, define context boundaries for every agent, and implement exception-handling runbooks that route edge cases to accountable humans. Effective governance also includes redundancy audits to prevent tool sprawl, change-control processes for agent skills, and logging/evidence packs to strengthen auditability and compliance.

Exhibit 4. At most companies, AI governance is too fragmented to prevent new tech debt

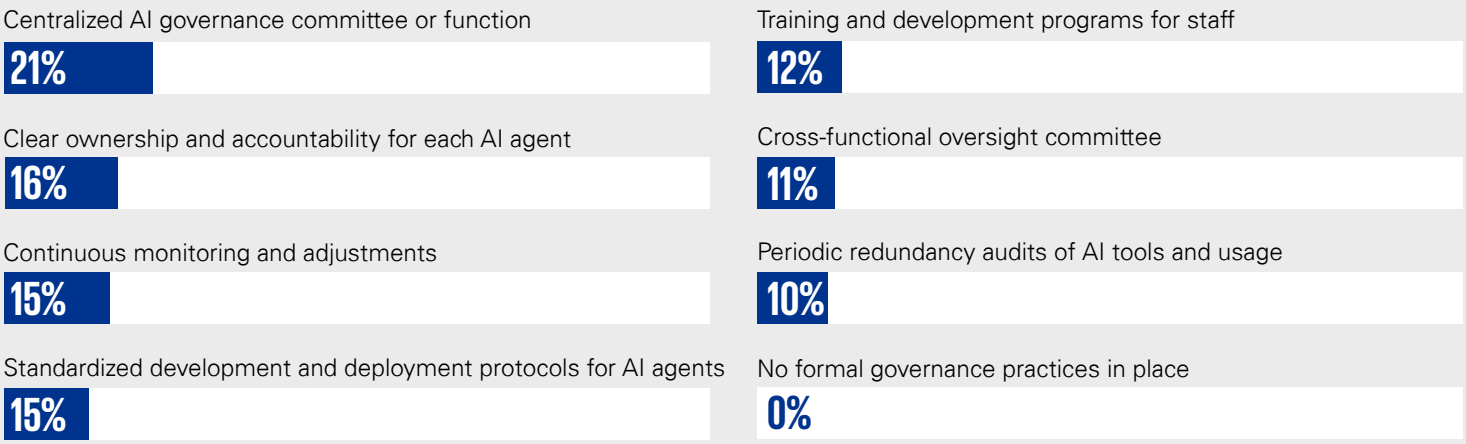

Leading practices for using AI to reduce tech debt in M&A

In M&A, tech debt is always part of the consideration but rarely a deal breaker. Scrutiny, however, is increasing, and tech debt is now a core due diligence item, influencing synergy assumptions, deal structure, and legal protections. To reduce tech debt, dealmakers in our survey indicated strong interest in using AI. But to generate cost and time savings that will maximize returns, they will need to carefully evaluate the benefits of AI deployment against the risk of adding more complexity.

Since speed is crucial in M&A, dealmakers facing strategic decisions about whether to build AI solutions in-house or buy third-party tools often opt for both. In our survey, 55 percent of respondents said they favor such a hybrid approach (Exhibit 5). Making such rushed decisions, however, is one potential move that can trigger new AI debt—overlapping agents, inconsistent standards, and unclear accountability—that will burden the combined entity down the road.

Establishing strong AI governance and clear ownership of AI-driven decisions is the best inoculation against such unintended consequences. The next competitive divide won't be between those who use AI and those who don't. It will be between those who can scale AI without compounding tech debt by incorporating governance frameworks with systematic human oversight and those who can't. An AI governance council that acts as an enterprise-wide control tower to oversee the implementation and management of AI solutions will go a long way to ensure they remain aligned with business goals. Effective governance requires centrally directed standardized processes and cross-functional coordination across product, IT, finance, and compliance departments.

Exhibit 5. Organizations are taking a hybrid approach to AI adoption

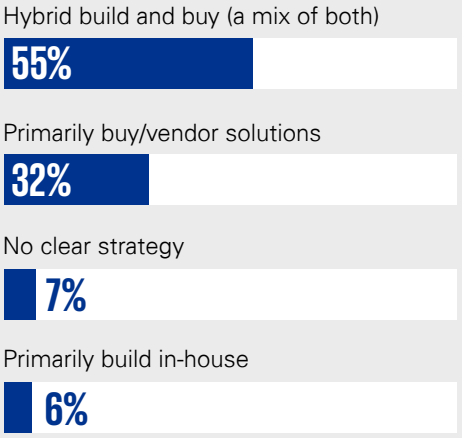

Once a robust governance framework is in place, organizations can focus on other leading practices:

- Start with targeted AI pilots to help validate performance and refine models before broader deployment.
- Embed AI early in the M&A due diligence process to improve risk detection, deal valuation, and integration planning. AI can support automated discovery, dependency mapping, and predictive risk modeling.
- Post-merger, use AI-powered controls and automated monitoring dashboards to provide extensive visibility into agent activities, usage pattern analysis, and system performance. Continuous AI monitoring is essential to effectively track integration progress and prevent the accumulation of new tech debt.

Conclusion

Addressing tech debt in M&A processes can seem daunting, but harnessing the power of AI is emerging as a promising solution. Our analysis of survey results and industry trends underscores the value of deploying AI tools early in the dealmaking process to unlock greater efficiencies and risk mitigation, while backing the initiative with an extensive AI governance framework. To tap the full potential of AI in managing tech debt, M&A players must not only act strategically but also prudently to prevent AI tools from becoming sources of additional complexity.

How KPMG can help

KPMG offers a range of services and capabilities for M&A clients looking to reduce tech debt with AI adoption. These offerings are designed to drive efficiency, mitigate risks, and unlock value throughout the M&A lifecycle.

Pre-sign technology due diligence	Assess architecture, code, dependencies, licensing, and quantify tech debt exposure to inform valuation and deal terms.
Agentic diligence sprint	Agent-assisted assessment quantifies tech-debt hotspots, auto-baselines requirements, and simulates modernization scenarios to inform valuation and deal structure.
Sign-to-close technology assessment	Define target architecture and rationalize platforms; plan and execute cloud migrations, data modernization, and automation to reduce legacy debt.
Agentified SDLC pilot	Stand up orchestrated agents—using KPMG Blaze for multi-agent coordination—to automate testing, refactoring, and release workflows with guardrails and observability. (Blaze is the KPMG GenAI modernization engine that automates the entire modernization lifecycle, accelerating transformation while dramatically cutting cost and risk.)
Post-close technology value capture	Deliver product technology and IT integration or carve-outs, transition contracts, establish AI/ governance frameworks, lead app rationalization, and drive run-cost, talent, and operating-model transformation.
Governed scale-up	Build the operating model: agent roles, exception-handling, redundancy audits, skill change-control, and dashboards to prevent AI debt and accelerate modernization across the combined entity.

Authors

For more information, contact us:

Greg Sward
Principal, Performance Transformation
415-418-9867
gsward@kpmg.com

Christopher Panneck
Principal, Advisory Offering, Technology
214-840-6781
cpanneck@kpmg.com

Rafael Barreto
Director, Transaction Strategy
415-308-5516
rbarreto1@kpmg.com

Acknowledgments:

David Hong
Manager, Performance Transformation
davidhong3@kpmg.com

Silvia Liu
Manager, Enterprise Innovation
silviazliu@kpmg.com

Megan Jaeb
Senior Associate, Transaction Strategy
mjaeb@kpmg.com

We would like to thank our contributors:
Adam Ashenfelter, Michael Gelfand, Charles Lee, and Tara Nelson

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** **Subscribe**

DASD-2025-19116