AI Is Rewriting M&A's Tech and Digital Playbook

By Jacqueline Govers, Gildas Bouteiller, Claudio Di Vittorio, Daniel Friedman, and Rogier de Boer

ARTICLE MAY 19, 2026 12 MIN READ

In post-merger integrations (PMIs) and carve-outs, the performance of the technology workstream often determines whether value is realized or delayed. System integration, data reconciliation, and day 1 readiness sit at the center of execution risk. Teams must grapple with fragmented architectures, complex data environments, and interdependent systems. The stakes are high: technology typically enables more than 60% of synergies. When technology falters, consequences cascade across finance, operations, procurement, and commercial teams.

AI is emerging as a tool to address these challenges. By embedding predictive intelligence and GenAI into technology workflows, leading organizations are automating structured analysis, surfacing hidden dependencies earlier, and reducing the coordination load that impedes execution.

Transaction timelines remain largely fixed by regulatory approvals and contractual milestones. But the efforts required to meet them are not. BCG's analysis finds that AI can reduce total execution effort in the technology stream by approximately 15%, with a materially higher impact in specific tasks.

Reducing execution effort creates a strategic choice. Organizations can capture productivity gains to lower transaction costs. Alternatively, they can redeploy freed capacity to reinforce day 1 readiness, accelerate synergy capture, strengthen governance—and carry AI-enabled capabilities into steady-state operations. Either way, AI is becoming an important lever in how the technology workstream supports deal execution.

The Tech Stream Is the Natural Starting Point for AI

The technology workstream is particularly well suited to AI-enabled support. Architecture choices, integration patterns, and data-model harmonization shape not only short-term delivery efforts but also long-term cost structure and resilience. Under time pressure, these decisions are often made with incomplete visibility into cross-system dependencies. By strengthening visibility and control across the technology workstream, AI improves the reliability of integration and separation outcomes.

The nature of the work further amplifies the opportunity to integrate AI. Much of the technology workload in transactions follows consistent analytical structures: baselining applications, mapping interfaces, classifying dependencies, reviewing contracts and licenses, migrating and reconciling data, orchestrating test cycles, and consolidating business and compliance reporting. Although systems differ across deals, the underlying analytical work is often similar and data-intensive.

There is also a practical advantage to adopting AI across technology workflows. Technology teams are typically more familiar with automation, analytics, and engineering-led tools than many other functions. Embedding AI into baselining, development, testing, architecture, and program management workflows requires less change in workforce behavior and can be absorbed more quickly into live execution environments.

AI Is Relevant Across the Transaction Life Cycle

AI's impact in the technology workstream spans the full life cycle—from early blueprint and baselining work through day 1 operational transition, exits from transition service agreements (TSAs), and the long-term transformation. (See the exhibit.)

AI Has Impact Across the Tech Workstream

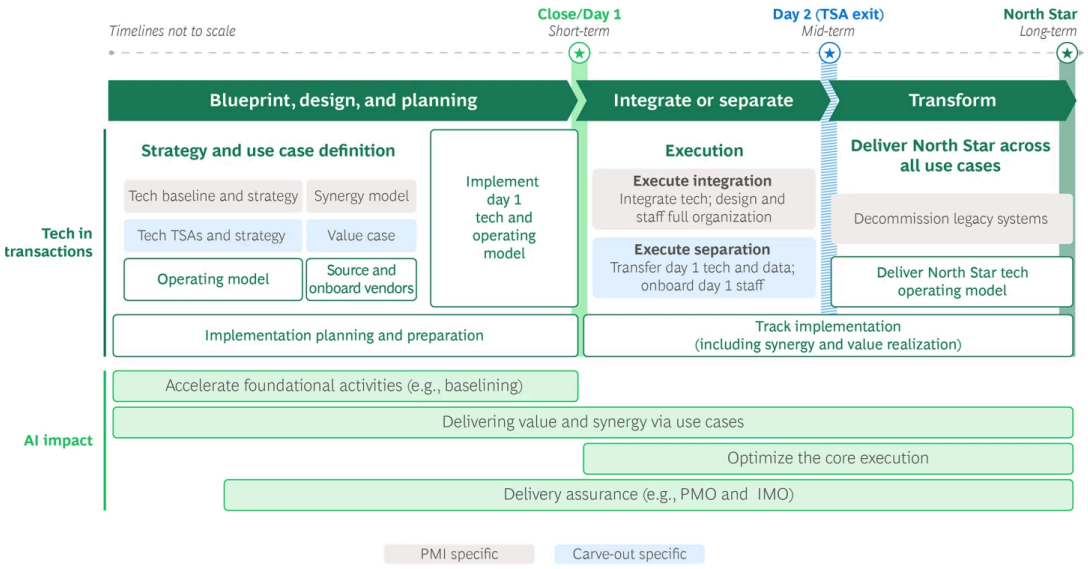

Source: BCG.
Note: IMO = integration management office; PMO = program management office; TSA = transition service agreement.

In PMIs, the focus often centers on accelerating synergy realization, rationalizing duplicated platforms, and maintaining operational continuity across merged environments. In carve-outs, the urgency shifts toward establishing independence, protecting day 1 stability, and exiting TSAs cleanly. Although the execution pressures differ, both contexts offer opportunities to reduce effort and strengthen outcomes through AI.

Delivering Impact in Four Domains

Across integrations and carve-outs, AI is being deployed in four domains where the technology workstream carries disproportionate execution effort and risk.

Accelerating Foundational Activities. Transactions often begin with an intensive discovery phase: baselining applications and reconstructing missing documentation, mapping interfaces, identifying contractual constraints, and clarifying data ownership. This work is typically fragmented across systems and stakeholders, requiring weeks of manual consolidation before

leaders have a reliable fact base. Gaps often surface only when negotiation or separation planning is already underway.

In recent carve-outs, teams have used AI-enabled agents to accelerate supplier contract baselining. Instead of manually extracting agreements scattered across contract systems and inboxes, the tools ingest contract repositories and spend data to generate a consolidated supplier inventory within days. They flag change-of-control clauses, identify contracts requiring assignment or renegotiation, and surface duplication risks before they become negotiation blockers or day 1 risks. Procurement and technology teams gain a clear view of exposure and potential synergies before negotiation positions are set—improving leverage and reducing the risk of late-stage surprises that delay execution or weaken commercial outcomes.

A similar approach is emerging in the baselining of technology landscapes. Rather than documenting the application and infrastructure footprint manually, organizations can use AI to ingest existing documentation and system extracts (such as license data and configuration management databases) and cross-reference the data with external benchmarks to generate an initial capability model of the technology landscape. Early implementations indicate that AI can automatically capture application, interface, and data inventories via direct integration with network-level information, further reducing manual effort. Highlighting overlaps and surfacing integration dependencies provides early visibility into consolidation opportunities and day 1 risks. This allows teams to move from fragmented discovery to a fact-based view of the landscape within days, rather than weeks.

> “Highlighting overlaps and surfacing integration dependencies allows teams to move from fragmented discovery to a fact-based view of the landscape within days, rather than weeks.

By compressing the ramp-up phase and reducing discovery rework, AI applications create earlier clarity and more informed decision making at the outset of the deal. Human effort shifts from manual data gathering to resolving the issues that materially affect cost, timing, and day 1 readiness.

Delivering Value and Synergies. As integration or separation plans take shape, attention shifts to how value will be captured and operationalized. In carve-outs, TSAs must reconcile legal commitments with operational realities—an effort that is often interview-heavy, iterative across functions, and prone to late-stage gaps and renegotiations.

Recent programs have deployed AI-enabled platforms to accelerate and pressure test TSA designs. By reconciling TSA clauses with system inventories and process maps, the tools generate

draft service schedules, interface documentation, and RACI (responsible, accountable, consulted, informed) matrices that are directly aligned to the underlying technology and process landscape. They also highlight missing services and unclear ownership early in negotiations—allowing deal and functional leads to resolve gaps before agreements are signed. The overall TSA design effort was reduced by roughly 50%, with drafting time for service owners cut by up to 70%.

In PMIs, teams need to identify synergies across multiple dimensions, often related to third-party agreements; for example, license reconciliation, volume increase, and rate alignment. In these scenarios, AI models can accelerate contract reviews (such as identifying change-of-control clauses) and generate precise simulations and benchmarks that illustrate how an agreement will deliver value (such as improving tiered pricing or eliminating unused licenses).

These applications not only reduce effort but also improve the quality and clarity of the value-capture plan, diminishing the risk of omitted services, misaligned responsibilities, and costly workarounds for problems that surface at day 1. Companies benefit from cleaner transitions and more predictable TSA exits, with fewer post-signing disputes and less value leakage.

Optimizing Core Execution. Build-and-test phases are typically the most effort-intensive periods of a transaction. Iteration cycles, uneven test coverage, and discovery of defects late in the process can consume disproportionate capacity as go-live approaches.

In recent integrations, teams have used AI-enabled test generators to industrialize these cycles. By integrating process flows, interface specifications, and anonymized production logs, the tools generate traceable test cases, build realistic data sets, and summarize execution results along the day 1 critical path. This can reduce manual test creation efforts by up to approximately 25%, improve traceability from requirements to defects, and surface critical failure scenarios earlier. Program leaders get visibility into threats to successful completion while there is still time to intervene.

Rather than replacing engineers' judgment, these tools shift efforts toward resolving edge cases and performance risks that are most likely to impact day 1 stability. This reduces late-stage defect triage, lowers the volume of emergency fixes, and increases confidence that critical business processes will perform as expected at go-live.

Strengthening Delivery Assurance. Even when plans are well constructed, the risk of failing to deliver often emerges from fragmented visibility across workstreams. Day 1 readiness and TSA exit depend on dozens of interdependent plans spanning IT, HR, finance, and operations. When milestones slip, cross-stream impacts are frequently identified late, requiring rework and last-minute escalation.

In complex transactions, teams have deployed AI-enabled planning agents to integrate individual workstream plans into a single dynamic task network. By recalculating critical and near-critical paths as dates shift, the tools surface hidden dependencies and provide scenario simulations and mitigation options. Program leaders gain a clear view of what is truly critical before delays cascade across workstreams, allowing governance to move from reactive tracking toward

proactive management of the dependencies that threaten timely delivery in the runup to day 1 and TSA exit.

The Impact: Reduced Effort Within Fixed Timelines

No amount of automation can meaningfully compress transaction timelines, which are often constrained by regulatory approvals, contractual milestones, or the new organization's change management needs. Rather, what AI alters is the amount of manual work required to operate within those timelines.

Transactions are labor-intensive across planning, integration or separation, and post-close stabilization. Teams reconcile data models, baseline application landscapes, draft transition documents, map dependencies, consolidate status updates, triage testing defects, and track value realization. Much of this effort is repetitive, manually stitched together, and vulnerable to rework when upstream assumptions shift.

AI targets that burden. It reduces administrative drag, minimizes manual consolidation, maintains a dynamic view of interdependencies, and surfaces risk earlier, allowing teams to deliver against fixed timelines with fewer cycles of rework. Realizing these capabilities is possible only with a usable data foundation—clear system inventories, accessible data, and defined governance—with human oversight remaining essential in how outputs are interpreted and applied.

Our top-down analysis suggests that, taken together, these applications can reduce total execution effort by approximately 15% across both PMIs and carve-outs, with reductions reaching 50% for some tasks. The distribution of impact varies by phase—often highest in building and coordination-heavy periods—but the cumulative effect is consistent: fewer manual cycles, less rework, and tighter control over risk. These results reflect early deployments; as tools mature and operating models adapt to AI-enabled ways of working, the potential impact is likely to increase.

The Strategic Choice: Bank the Savings or Reinvest to Secure Outcomes?

Reducing execution efforts creates a strategic choice. If AI removes a meaningful share of manual coordination, rework, and administrative burden from a transaction, organizations must decide how to deploy the freed-up capacity.

Bank the savings. Organizations can capture productivity gains directly as lower transaction costs. Leaner integration or separation teams, reduced reliance on external support, and simplified workflows translate into immediate budget efficiencies. In an environment where the economics of the deal are closely scrutinized, this path can protect margins and improve short-term financial outcomes.

Reinvest the capacity. Alternatively, freed-up resources can be redirected toward strengthening execution itself. Teams can devote more time to ensuring comprehensive day 1 readiness, accelerating priority synergy use cases, strengthening cyber and data controls, and improving cross-functional coordination. Rather than simply limiting costs, this approach raises quality—reducing rework, surfacing risks earlier, and increasing confidence that projected value will materialize. It also reduces the management distractions associated with integrations, allowing leadership teams to focus on running and growing the base business.

Organizations can extend this reinvestment beyond the transaction itself. The same AI-enabled capabilities used to accelerate integration or separation can be applied to redesigning how the future organization will operate—across core processes, decision making, and governance. Instead of treating execution and transformation as separate efforts, teams can carry these capabilities into steady-state operations.

> “ The same AI-enabled capabilities used to accelerate integration or separation can be applied to redesigning how the future organization will operate.

Turning AI into Execution Advantage

To scale AI deployment in the technology workstream, organizations must act deliberately in the following areas:

- **Prioritize high-impact use cases across the four domains.** Focus first on applications that materially strengthen foundational clarity, protect value realization, industrialize build-and-test cycles, or improve critical-path visibility. Each use case should be explicitly tied to day 1 readiness, TSA exit, or synergy milestones and assigned joint business and technology ownership. If an initiative cannot be linked to a tangible transaction outcome, it should not be prioritized.

- **Build the data backbone to support foundational clarity.** AI-enabled acceleration in foundational activities depends on trusted inputs. Harmonize key data sets, clarify ownership, and establish governed pipelines early in the program. Standardized definitions, validation checks, and transparent data lineage reduce the risk that automation amplifies inconsistency. Foundational clarity enables confident decisions across integration design, separation scope, and contractual remediation.

- **Redesign core execution workflows.** In the core execution domain, AI delivers impact when it reshapes how work is performed. Convert business requirements into clearly defined agent configurations, data inputs, and system access rules; specify where human review is required; and integrate outputs directly into test cycles, configuration management, release gates, and cutover planning. The objective is not to automate isolated tasks but to compress iteration cycles and reduce rework across build-and-test phases.

- **Strengthen delivery assurance through integrated governance.** AI insights must be embedded into program governance. Deploy AI agents within live delivery environments (including third-party solutions) and feed outputs (such as critical-path recalculations, readiness scoring, or value tracking) into decision forums. Formalize ownership and escalation mechanisms before scaling. At the same time, build responsible AI guardrails, including validation checkpoints, audit trails, and access controls, into the governance model. This ensures that acceleration enhances execution certainty rather than introducing unmanaged risk.

- **Make integration decisions with an AI-first future in mind**. Decisions made during the integration should support immediate execution without constraining future AI transformation. Across systems, data, processes, and operating models, favor standardization, interoperability, modular design, and governed access so the organization can scale automation and more agentic ways of working over time. Avoid quick fixes that address near-term complexity but entrench fragmentation, limit data usability, or create structural barriers to future AI adoption.

- **Carry integration capabilities into the steady state.** Do not dismantle the capabilities built to execute the transaction. Maintain the technology baseline and preserve the ways of working that enabled coordination and speed. Embedding these into the ongoing technology function allows AI-enabled practices to support not only deal execution but also day-to-day business decision making.

Organizations that embed AI into the technology workstream will execute deals with less effort and greater precision. Although AI does not eliminate complexity, it enables teams to surface risks earlier, reduce transaction costs, and redirect capacity toward higher-value activities. In high-stakes transactions, that advantage compounds.

Special thanks to Matthew Bingley, BCG alumnus, for leading the research and development of this article.

Authors

Jacqueline Govers

Managing Director & Partner
Amsterdam

Gildas Bouteiller

Managing Director & Partner
Paris

Claudio Di Vittorio

Managing Director & Partner
Miami

Daniel Friedman

Managing Director & Senior Partner; Global Leader of Transactions & Integrations
Los Angeles

Rogier de Boer

Managing Director & Partner
New York

ABOUT BOSTON CONSULTING GROUP

Boston Consulting Group partners with leaders in business and society to tackle their most important challenges and capture their greatest opportunities. BCG was the pioneer in business strategy when it was founded in 1963. Today, we work closely with clients to embrace a transformational approach aimed at benefiting all stakeholders—empowering organizations to grow, build sustainable competitive advantage, and drive positive societal impact.

Our diverse, global teams bring deep industry and functional expertise and a range of perspectives that question the status quo and spark change. BCG delivers solutions through leading-edge management consulting, technology and design, and corporate and digital ventures. We work in a uniquely collaborative model across the firm and throughout all levels of the client organization, fueled by the goal of helping our clients thrive and enabling them to make the world a better place.