


McKinsey Transformation Practice

The agentic transformation office: Redefining the economics of change

Successful transformations no longer require an army of people at the center.

by Chase Covington, David Pralong, and Tomer Slaney



The transformation office (TO) has a dirty secret. The very function charged with driving radical change across the enterprise is often one of the slowest to adopt modern tools and technologies—a notable holdout even as [agentic AI](#) is poised to transform sectors and economies. While TOs try to push an organization toward speed and agility, we often see those offices running on manually assembled slide decks and fragmented data sources.

The irony has real consequences. Resources are consumed by predictable administrative tasks. Critical insights arrive too late. Strategic narratives drift across workstreams. Leadership conversations become reactive firefighting rather than proactive steering. Straining to compensate for central capacity limits, individual workstreams begin building their own shadow tools and reports—further fragmenting data analyses and undermining consistency.

The past 12 months have seen a dramatic shift in what artificial intelligence can do. For a few leading TOs, simple chatbots have evolved into autonomous AI agents capable of executing multistep business processes from start to finish. These autonomous tools are already demonstrating game-changing value; they continuously spot performance gaps, maintain a single source of truth, and spark the right executive conversations at a rapid weekly cadence. They free up human bandwidth—a TO's ultimate constraint. Time spent on many core tasks can be reduced by 35 to 40 percent. In some cases, we've seen time savings of 70 percent or more.

As with other applications of AI, the full potential of agentic AI is evolving rapidly and, in many instances, unevenly. This article provides a first look at the opportunity for TOs. We share what we are currently seeing in transformations, what teams have already begun deploying, and the first steps toward building an early playbook.

The structural ceiling of the traditional TO

The very function responsible for driving change and improvement across the organization is often the most resistant to updating its own legacy technology and tools. That's frustrating in any context. But for agentic AI, the consequences are far greater.

The bottleneck of manual tracking

Today, a traditional TO typically relies on rigid roles, exhausting meeting rituals, and a top-heavy reporting stack. The office can require a near army of people because tracking progress manually is so slow. Consider a recent two-year global turnaround involving more than 3,000 employees across eight divisions. The transformation office required more than 15 full-time personnel, who spent more than half their time manually validating data and rerunning spreadsheet analyses to produce a mandatory 300-plus-page report every four weeks for a two-day monthly board workshop. That's common—and it commonly caps how fast the entire company can move.

The cost of data fragmentation

Because human time and attention are finite, critical operational insights emerge slowly and reach key decision-makers even more slowly—often too late to solve real-world problems. A disproportionate amount of budget is wasted on predictable administrative tasks instead of strategic problem-solving. Without tight central coordination, different workstreams start telling different versions of the same story.

Desperate to overcome central capacity limits, individual workstream teams frequently build their own makeshift tracking tools and localized reports. In one recent transformation, a fragmented landscape across eight divisions relied on more than 100 nonintegrated documents to track financial progress.

Solving for human bandwidth

Agentic AI can free two-thirds or more of a TO team to focus on higher-value work. The potential is generational. Grasping it starts with a core technological distinction: the difference between generative AI as a personal productivity *tool* and agentic AI operating as an autonomous, operational *system*.

From personal tools to autonomous systems

Today, gen AI acts like a personal assistant to help an individual complete a single task, such as drafting an email. AI agents, by contrast, do not wait for a prompt; they plan, take independent action, and close loops autonomously across multistep workflows. The TO is where this capability matters most in the company, because its automated outputs directly touch, guide, and accelerate almost every active workstream. While transformation offices have historically struggled with fragmented manual tools, McKinsey tracking across recent transformations shows that broad market adoption of core digital data layers and performance “infrastructure”—the underlying digital data layers, tracking tools, and performance architecture—is increasing (Exhibit 1).


The three automated pillars

Deploying an agentic architecture sets a faster management rhythm across the organization, shifting routine manual support to three automated pillars:

- *Performance intelligence*: Instead of waiting for periodic reviews, AI agents continuously synthesize initiative data in real time, automatically flag negative variances, and surface only the critical exceptions that senior leadership actually needs to see. For example, in one transformation, the “before” state involved tracking more than 30,000 individual improvement measures and 900,000 critical temporal data points, and required several TO team members to each spend multiple hours running manual data quality checks and conducting tedious, half-day pivot table exercises drawn from hundreds of static Excel sheets and a half-dozen IT systems. The “after” state, by contrast, is an agentic AI

Exhibit 1

AI adoption is increasing in the transformation performance ‘infostructure.’



McKinsey & Company

framework built in large part by Cursor (an AI-first code editor that understands natural language) and integrated with McKinsey's Lilli Teamspace (an agent-specific setup that provides immediate answers and insights). The new solution not only eliminates hours of manual spreadsheet manipulation and repetitive pivot table analysis but also leverages agentic AI that enables divisional teams to identify *within minutes* specific actions or omissions that had been delaying financial outcomes. The tool surfaces discrete problems early and triggers immediate, targeted escalations. This represents a fundamental shift, as advanced operating models move from tracking what already happened to flagging what will probably happen next.

- *Single source of truth:* Agentic AI enables the organization to move away from static weekly status reports and achieve a live, automatically updated picture of total value delivered, active risks outstanding, and critical decisions pending. The technology eliminates spreadsheet version-control problems and delays caused by “dueling versions.” Moreover, rather than building an expensive, centralized data lake (a process that can take more than a year), organizations can leverage AI to dynamically link separate data sets. Already, commercially proven solutions can near-seamlessly extract cross-functional insights from siloed platforms—connecting everything from Wave (McKinsey's transformation-tracking

platform) and enterprise resource planning (ERP) systems to HR infrastructure and production tools—to maintain an accurate, unified perspective. This unified view goes beyond tracking what has already happened; it allows a system to simulate what should happen next. For example, by linking IT capacity planning tools directly with project requirements in Wave, an agent can automatically sequence an entire portfolio and map out different execution tracks. If leadership needs to know how an extra funding boost in 2026 would accelerate timelines and unlock financial benefits, the agent can instantly calculate the impact using those connected systems. As these plans change, the agent handles logistical ripple effects automatically—updating executive briefings, steering committee decks, and weekly team communications without anyone having to search for the right sources or synthesize them manually.

- *Conversation facilitation:* Presentation materials for crucial weekly steering forums are near-instantaneously generated by autonomous agents rather than painstakingly assembled by analysts. For example, instead of an analyst team spending five hours manually preparing a SteerCo briefing, a team found that an AI agent could precisely analyze reams of documents and automatically generate a tailored briefing in less than ten minutes. That's not a random anecdote. We've seen TO teams extract the most vital insights from disparate reports and files to prepare meaningful, specific critical questions for leadership to ask during every part of a SteerCo review. In one successful transformation, agentic AI slashed the time spent on biweekly and steering committee reports from 12 days a month to just three—and saved about 200 additional person-hours *each month* from agentic improvements to other processes. As a result, human attention can be freed up from preparing paperwork to making considered, better-informed decisions. Ultimately, this architecture ensures that the weekly and monthly corporate cadence becomes a genuine, forward-looking management rhythm rather than an exhausting, backward-looking reporting slog.

Collectively, the three pillars create a faster and more effective management cadence. They reduce complexity, help surface insights at the pace the business needs, and free human bandwidth for strategic priorities.

Minding the guardrails

Moving from software that writes text to agents that take action introduces new operational risks. When independent AI agents connect, a single data error can quietly trigger a chain reaction of [bad decisions downstream](#), such as leaks in analyses and corrupting metrics. While leading organizations recognize these threats, [our research](#) finds that mitigation currently lags behind awareness. Moreover, at present, even the most elegantly developed agents can still have [rough edges](#). For example, agentic AI can produce too many false alarms—a drain on resources that, at least initially, can feel like two steps forward and one step back. Early-moving TOs are learning that every step an agent takes must be transparent and trackable. Yet their experience is becoming a competitive advantage. Small errors that may seem trivial in a slow-moving environment compound much more quickly amid rapid change.

A new division of labor for leaders

Operating with a leaner center does not mean the organization has a weaker core—it means that the TO can be far more intentional and achieve much greater impact. An agentic TO establishes a clear division of labor: Predictable data manipulation is automated while human capability is focused on strategy.

Elevating human judgment

There are clear boundaries to automation; human judgment remains irreplaceable for the high-stakes, nuanced work of a transformation. AI cannot program away difficult escalation calls, smooth over complex stakeholder relationships, or replace the creative nuance of organizational culture. But by handing the administrative staff work over to agents, TOs can free up a massive amount of human capacity. In a recent deployment, automating ten core insight generation and reporting processes reduced the central TO requirement from an engagement manager and nine staff members to just the manager and two core staff members. This freed up more than 25 days of capacity, allowing the remaining team to redirect their focus toward coaching, alignment, and new-idea generation (Exhibit 2).

The agentic transformation day to day


To visualize what an agentic transformation looks like on a daily basis, companies should envision long-term aspirations and account for immediate operational wins, including the following:

- *Dramatically improving the output:* The AI agent acts as an automated compliance and hygiene check, continuously scanning upcoming meeting agendas and presentation materials against the active participant lists. If it detects that the specific initiative owner responsible for sharing the vital change story has not been invited to the corresponding review meeting, it automatically identifies the omission and issues an invite. The shift is fundamental: AI moves from a compliance tool that an employee *may* use to a proactive partner that *always* pushes for better project outcomes. For instance, by analyzing current project plans against Wave's historical transformation knowledge base, agents can automatically flag underperforming workstreams and suggest specific, proven ways to accelerate delivery or boost financial impact, informed by past transformations.
- *Drafting the weekly SteerCo decks:* Instead of an analyst team spending the weekend copying and pasting charts, the AI agent automatically prepares the weekly steering committee presentations by ingesting and analyzing the prior week's presentations, data extracts from tracking databases such as Wave, and transcripts of recent alignment meetings. Once the master steering deck is compiled, the agent automatically filters and converts that data into targeted, customized presentations for individual workstream leads.
- *Automating operational reminders:* The AI agent assumes the role of project coordinator, automatically generating and dispatching every weekly reminder, follow-up message, and

Exhibit 2

Agentic AI is already freeing meaningful capacity in transformation offices.

Time spent on recurring transformation office reporting and steering activities,
number of working days per month (case example of 1 transformation, illustrative)



Note: Figures are approximate and rounded.
¹Full-time employee.
²Division.

McKinsey & Company

Find more content like this on the
McKinsey Insights App



Scan • Download • Personalize



data entry request to the broader organization without requiring human prompting.

Building the new model

While the shift to an agentic transformation should move quickly, it cannot happen overnight; it requires a disciplined approach and clear prerequisites. Before deploying autonomous tools, an organization will first need to map operational workflows, establish specific governance rules, and build agentic capabilities across the enterprise. The move also changes what TOs need from their teams. Leading organizations are rethinking [roles and responsibilities](#), repositioning staff as [workflow orchestrators](#), and sourcing and developing talent with the adaptability to manage AI systems.

The sequence of the rollout should be strategic. We observe that successful first movers begin with reporting and insight processes to automate the ingestion of data and generate basic visibility. Then, once the data layer is stable, they expand directly into true workflow automation, allowing agents to take actions across different corporate software systems. Most critically, we see that the success of deployments hinges on change management within the TO itself. Team members must understand, trust, and believe in the automated model; if they don't trust the system, they default to old, manual habits and disrupt the operational rhythm.

The TO no longer must be behind the times. Shifting to an agentic operating model removes bottlenecks that have long slowed corporate turnarounds. By automating data tracking, report making, and routine coordination, organizations free their central teams from administrative drudgery and redirect human talent where it matters most: clearing roadblocks, navigating stakeholder dynamics, and making decisive calls. The result is a leaner, faster, more powerful center—and, as we're seeing already, a genuine source of competitive advantage.

[Chase Covington](#) and [Tomer Slaney](#) are partners in McKinsey's New York office, and [David Pralong](#) is a senior partner in the Charlotte office.

Copyright © 2026 McKinsey & Company. All rights reserved.