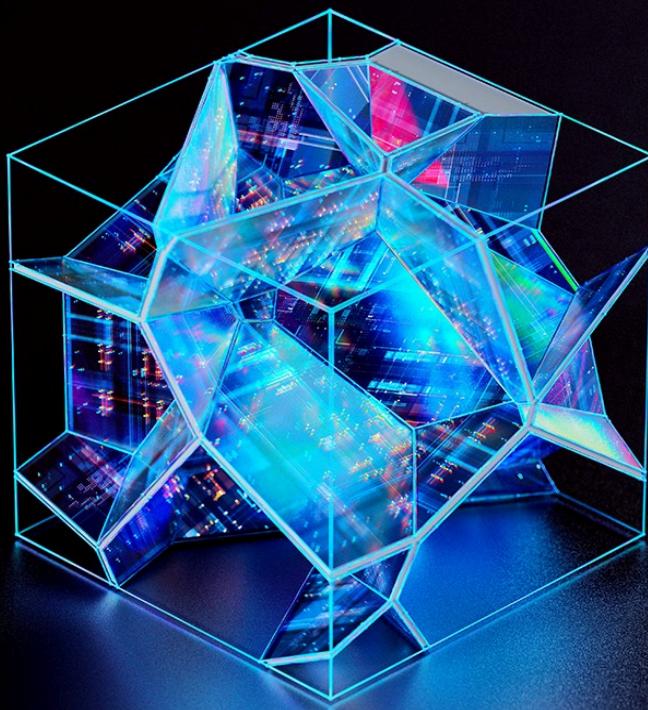


Operations Practice

When can AI make good decisions? The rise of AI corporate citizens

As agentic AI begins to influence decisions at scale, forward-looking organizations need to reimagine governance, trust, and operating models—or risk falling behind.

This article is a collaborative effort by Federico Berruti, Lari Hämäläinen, Oana Cheta, and Venky Anant, with Damian Lewandowski, representing views from McKinsey's Operations Practice.



Imagine two financial institutions. The first manages its loan origination process through a patchwork of task-level automation and predictive models. It uses historical credit scores, rigid underwriting rules, and batch processing to move applications through a series of sequential steps. While some parts of the process have been digitized, many decisions still require human intervention—whether due to exception handling, regulatory compliance checks, or risk flagging. This results in slower decision-making, inconsistent applicant experiences, and rising operational costs as application volumes climb.

Now, consider the second institution. Here, the loan origination journey is orchestrated by a network of agentic AI systems—autonomous, reasoning agents capable of executing and adapting entire workflows end to end. They don't just execute predefined workflows, as traditional AI could, or transform unstructured data into new insights or media, as gen AI can. Instead, these agents ingest real-time data across dozens of sources, from macroeconomic indicators and applicant digital behavior to regulatory changes and even sentiment analysis, all to make complex decisions (Exhibit 1). They not only assess creditworthiness but also adjust pricing, recommend optimal product bundles, and proactively flag anomalies for human review.

Exhibit 1

Although AI agents are exciting, it's important to bring the right tool for the job.

Evolution of AI capabilities

		
Traditional AI¹ Rigid but reliable implementer	Gen AI Supercharged copilot	Agentic AI Fully autonomous operator
Executes predefined workflows but doesn't adapt and is hard to scale	Understands, generates, and transforms information but lacks execution	Learns, reasons, and executes multistep workflows with minimal oversight
<ul style="list-style-type: none"> • Performs rule-based execution • Doesn't learn or reason • Requires structured inputs • Is efficient but inflexible 	<ul style="list-style-type: none"> • Processes unstructured data • Generates text-based outputs • Is context-aware but reactive • Is smart but needs human direction 	<ul style="list-style-type: none"> • Performs autonomous workflow execution • Makes decisions without predefined rules • Adapts to real-time inputs • Has fewer human touchpoints
BEST FOR: Basic repetitive and structured tasks that rarely change	BEST FOR: Human augmentation and knowledge generation	BEST FOR: End-to-end automation, process orchestration, and decision-making

¹Comprises traditional automation tools, such as natural language processing, optical-character recognition, and robotic process automation.

These agentic systems don't just execute tasks—they "think." They reason across time horizons, learn from outcomes, and collaborate with other AI agents in areas such as fraud detection, compliance, and capital allocation to continuously optimize performance. In operations, agents dynamically rebalance workloads across call centers, resolve customer inquiries with contextual and emotionally intelligent responses, and escalate only when human judgment is needed. The result: better decisions, faster cycles, and dramatically lower unit costs.

Crucially, AI agents surface insights that humans might miss. When loan performance trends diverge unexpectedly in a certain geography, an agent might detect the pattern early, analyze potential causes, and suggest mitigation strategies—before leadership is even aware of the issue. Other agents may handle regulatory reporting or stress testing autonomously, freeing up human talent for strategic innovation.

The implications are profound. Agentic AI represents not just a new technology layer but also a new operating model. And while the upside is massive, so is the risk. Without deliberate governance, transparency, and accountability, these systems could reinforce bias, obscure accountability, or trigger compliance failures.

The solution? Treat AI agents as corporate citizens. That means more than building robust tech. It means rethinking how decisions are made from an end-to-end perspective. It means developing a new understanding of which decisions AI can make. And, most important, it means creating new management (and cost) structures to ensure that both AI and human agents thrive.

In this article, we explore how organizations can navigate this complex landscape, leveraging the right tools and practices to stay ahead of the curve while adhering to evolving regulatory demands.

Agentic AI: The potential and price of entry

Agentic AI marks a sharp departure from traditional systems built on deterministic, rule-based architectures. In the past, enterprise decision-making relied on hard-coded logic and static workflows—think customer service scripts, underwriting checklists, or supply chain triggers. While useful in predictable environments, these approaches fall short when facing today's dynamic, high-volume, and context-rich realities.

Agentic systems are different. Rather than executing fixed instructions, they act more like collaborators—reasoning, adapting, and learning over time. At their core are [AI agents](#): software entities capable of perceiving environments, making autonomous decisions, and taking action to achieve defined objectives. Agentic systems typically follow one of two forms:

- *Single-agent systems* can perform end-to-end tasks independently, such as adjudicating a loan, triaging a customer complaint, or dynamically adjusting inventory levels based on incoming demand signals.

- *Multiagent systems*, by contrast, operate as decentralized networks of agents that interact and collaborate. For example, in a financial-services context, one agent may assess creditworthiness, another may model risk exposure, and a third may ensure regulatory compliance—all working together to optimize the customer journey and manage trade-offs in real time.

The potential of agentic AI lies in these systems' ability to fundamentally reshape how organizations operate. They can unlock exponential gains in speed, scale, and precision, enabling companies to reduce decision latency, eliminate handoffs, and continuously improve outcomes. Imagine underwriting decisions that can be generated in seconds, compliance reporting that updates itself in real time, or customer experiences that feel human—at machine speed and cost. In short: higher productivity, better decisions, and a more adaptive enterprise.

While the possibilities are alluring, the price of entry is high, and the transformation doesn't happen overnight. Deploying agentic AI is not a plug-and-play solution—it's a long-term commitment that requires robust infrastructure, interoperable data ecosystems, and deep integration across functions. Beyond technology, it demands a full rethinking of accountability, ethics, and governance. Leaders need to invest in operating-model redesigns, build new talent archetypes, and establish trust mechanisms that enable humans and AI to collaborate safely and effectively at scale.

AI agents as corporate citizens—who need management

To fully realize the value of agentic AI, organizations should focus less on treating these systems as experimental tools and more on managing them like they manage people. In this future-ready enterprise, AI agents become corporate citizens: accountable, governed, and expected to deliver measurable value. That means rethinking how they are funded, evaluated, and integrated. Just like human employees, AI agents require the following infrastructure:

- *A full cost structure*. Companies already understand the cost of human talent—salary, benefits, bonuses, and training. AI agents deserve the same scrutiny. Leaders should factor in the total cost of ownership, including IT systems, model retraining, orchestration layers, governance tools, and compliance. And like high-performing employees, agents should be able to excel across functions and roles—deployed where they deliver the most impact, not locked into silos.
- *Defined objectives*. Every agent needs a job description. Whether it's resolving claims, detecting fraud, or optimizing inventory, their tasks should align with business priorities, and their results need to be tracked, just like any team member's goals.

- *Performance management.* Humans are reviewed for quality, speed, and impact. AI agents should be, too. Their performance—across efficiency, accuracy, and user satisfaction—should be measured, monitored, and improved over time, with underperforming agents being retrained or retired.
- *Governance and oversight.* Humans operate under policies and cultural norms. AI agents need the same guardrails: ethical frameworks, transparency, auditability, and fail-safes for sensitive decisions. Especially in regulated sectors, this isn't optional—it's existential.
- *Cross-functional enablement.* Great employees don't do just one thing—they collaborate, adapt, and grow. Agents should, too. The best-performing AI systems are those designed for interoperability, enabling them to support multiple domains, learn across use cases, and scale throughout the enterprise.

By holding AI agents to the same standards used for people, including cost, accountability, and adaptability, organizations elevate them from tactical tools to strategic workforce assets. They don't just do work; they become part of the way work gets done.

Rethinking decision-making with ‘smart ops’

To unlock the full potential of AI in service operations, organizations need to do more than deploy technology. They need to rearchitect how decisions are made and how work is done—by building a “smart ops” structure where humans and AI agents operate in coordinated, complementary roles.

Where digital agents shine

It starts with using the right tool for the right task. Early adopters of agentic AI illustrate the value of evaluating opportunities at the end-to-end journey level and workflow level, not in isolation (Exhibit 2).

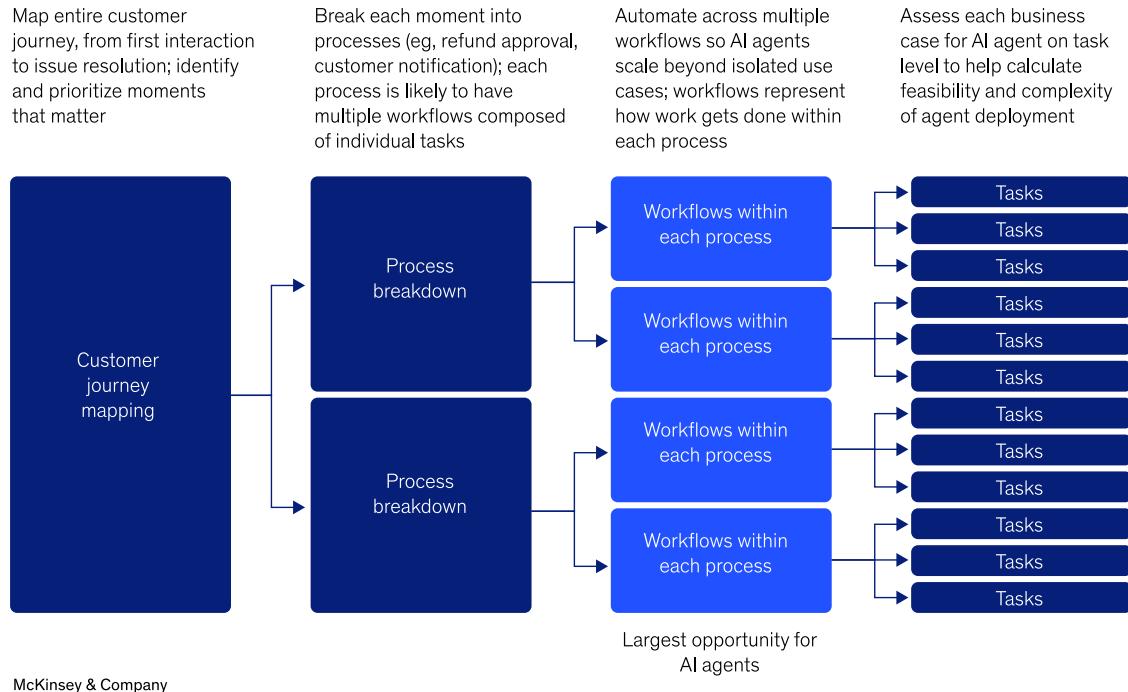
These point to a simple, actionable framework made up of the following AI agent types:

- *Task-level agents* precisely follow simple instructions to execute defined, repeatable tasks from end to end, such as processing a refund or rescheduling an appointment.
- *Autonomous problem-solver agents* perform multiple workflow steps that involve basic judgment but within defined boundaries, such as verifying subscriber eligibility, submitting a claim, or sending a follow-up.
- *Model orchestrator agents* act like digital process managers, partnering with human agents and coordinating across tools, systems, and other AI agents to surface insights, recommend actions, or summarize data in real time.

Exhibit 2

Agentic AI can achieve more than plug-and-play processes: It can transform entire experiences.

Reimagining end-to-end customer journey experience



McKinsey & Company

- *Domain-specific agents* are tailored to critical business functions—such as customer service, sales, or finance—and optimized for specific outcomes.

This modular, role-based approach allows organizations to deploy agents with precision, ensuring they are aligned to business value, operational need, and user context. But deploying agents is only half the story. To build a truly intelligent service operation, human roles should evolve in parallel.

Where humans lead

As agents take on high-frequency or transactional work, employees shift into roles that require more oversight, ethics, and judgment, including:

- *Custodians* who ensure the integrity of data, model performance, and customer outcomes.
- *Judgment holders* who handle ambiguous or high-stakes decisions where context, nuance, and trust are essential.
- *Approvers and auditors* who review exceptions, manage escalations, and reinforce compliance boundaries.

This shift demands a workforce design mindset, not just a tech implementation plan. Each digital worker—like each human worker—needs a clearly defined role and objective, a measurable impact on business performance, governance and oversight, and opportunities to evolve and learn.

It also means recognizing a simple truth: Not every digital worker will show immediate ROI—just like not every human does. What matters is the system-level performance of your human-plus-AI workforce.

That brings us to the next critical question: Once you've built the smart-ops foundation and redefined roles, which decisions should AI make, and which ones still require a human touch?

Redesigning processes: Not *what* to automate, but *which decisions*

While agentic AI offers potential across nearly any function, service operations remain the sharpest proving ground. These environments are rich with high-volume, repetitive tasks and data trapped in silos, making them ideal for intelligent automation. But the question is no longer *what* companies can automate. It's *which decisions* they should automate—and where human judgment still matters.

That's where a decision-making framework based on risk and complexity becomes essential (Exhibit 3). Rather than chasing automation for automation's sake, organizations should classify decisions based on their inherent risk and the degree of judgment required. Low-risk, low-complexity decisions, such as verifying account details or checking claim status, are prime for full automation. High-risk, high-judgment scenarios, such as fraud resolution or complex policy exceptions, may still require human oversight, supported by AI copilots.

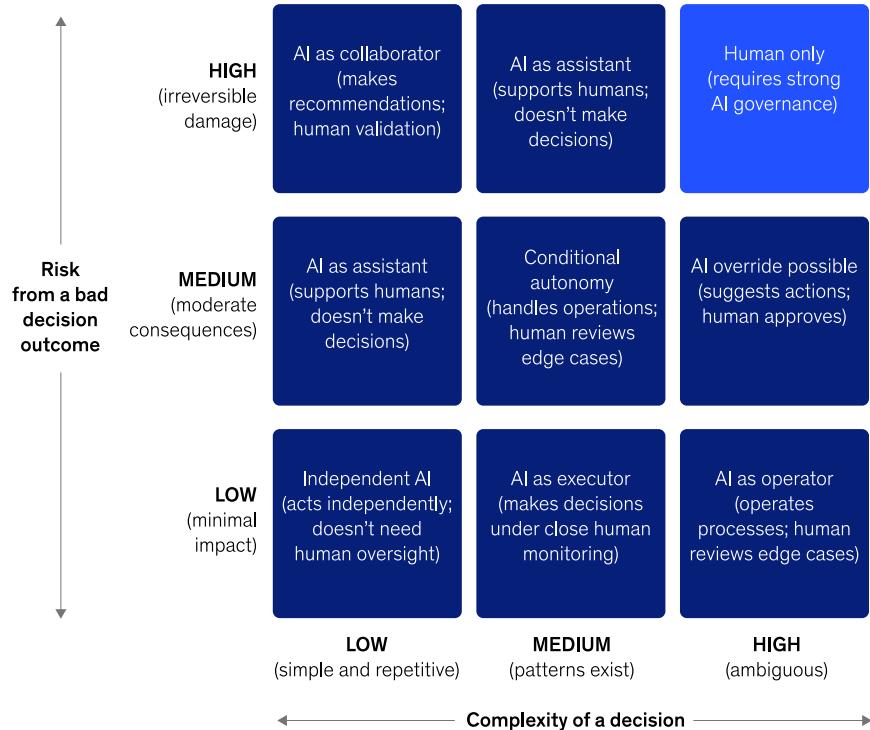
Agentic AI can already manage a wide range of frontline interactions autonomously. In healthcare, for instance, AI agents can dynamically manage appointment scheduling, predict no-show rates, and optimize clinical capacity. In utility services, they can monitor network performance, initiate preventive maintenance, and keep customers informed—all without escalation.

But the real value comes when these systems serve not just customers but also the entire organization. Every service interaction becomes a data point. AI agents can surface trending

Exhibit 3

Each decision across a process and workflow can be closely evaluated during the redesign phase for feasibility of agentic AI use.

Illustrative agentic AI decision-making framework



Implications of agentic AI decision-making on customer service leaders

Process	People	Tech
<ul style="list-style-type: none"> Assess automation suitability case by case, ensuring that value and risk analysis, not just efficiency, drive all decisions to adopt AI (not all process should be automated) Allow agents to act independently when confident; escalate complex or high-risk cases with conditional autonomy and human-in-the loop (HITAL) design Redesign workflows and processes to enable agentic decision-making (eg, implement new escalation paths, HITAL designs, and overrides) 	<ul style="list-style-type: none"> Create new roles and responsibilities to define who oversees AI, who intervenes with it, and who is accountable for AI-driven decisions Update skill matrixes and redefine job functions to address change of workforce skill requirements due to AI Enact thoughtful change management, as employee acceptance rests on achieving clarity on AI agents' role, training, and structured adoption plans 	<ul style="list-style-type: none"> Monitor agentic decisions in real time through tracking dashboards, audit logs, and observability tools Enable open, scalable architecture (eg, robust API integrations, modular systems, interoperability) for agent Ensure that infrastructure supports agentic capabilities (eg, data pipelines, compute power, security frameworks that can handle workloads at scale)

McKinsey & Company

complaints, identify breakdowns in upstream processes, and flag systemic issues before they escalate. This democratization of service data allows insights to flow seamlessly from customer touchpoints into product design, marketing, and operations, fueling faster and more connected decision-making across the enterprise.

Leading organizations can now move beyond the legacy mindset of simply reducing contact volume. In a smart-ops model, volume is not the enemy—it's the intelligence layer. If service data is being captured, interpreted, and routed effectively by AI, higher volume means more signal, faster learning, and greater value creation. What once looked like cost can now become a competitive advantage.

Of course, making this a reality requires investment. Companies need to modernize their infrastructure, clear technical debt, and establish secure, real-time data flow across business units. But those that move first will redefine how service operations contribute to growth, not just through efficiency but also through strategic insight and enterprise coordination.

Getting started

As AI agents take on more operational tasks, organizations need to rethink how work is designed, how people are supported, and how value is measured. Deploying agentic AI at scale isn't just a technical shift—it's an organizational one. To get there, companies can start by making the following moves:

- *Bridge the tech–business gap—with leadership accountability.* AI success hinges on cross-functional alignment. Embedded teams, shared KPIs, and AI product managers fluent in both business and technology ensure that initiatives are not only technically feasible but also commercially strategic. Senior leaders—starting with the [COO](#) and the [chief information officer](#)—should own outcomes, define accountability structures, and model the mindset shift AI demands.
- *Redesign roles and invest in reskilling.* As automation redefines task boundaries, [roles must shift toward exception handling, judgment-based decision-making, and customer experience](#). Companies can invest in AI literacy, data interpretation, and systems thinking to prepare talent for new, higher-value work.
- *Elevate culture and change management.* [Embracing AI requires cultural alignment.](#) Transparent communication, leader role modeling, and cross-functional ownership are critical to building trust, reducing resistance, and sustaining adoption at scale.
- *Strengthen data and architecture foundations.* AI's effectiveness depends on real-time, compliant, and connected data. [Organizations should modernize their data infrastructure](#), ensure unified governance, and build scalable systems that enable AI to operate safely and effectively across functions.

Find more content like this on the
McKinsey Insights App



Scan • Download • Personalize



To unlock the full value of AI in service operations, companies will need to shift from task automation to decision design, focusing not on what *can* be automated but on which decisions *should* be. This requires treating AI agents as corporate citizens: digital workers with defined roles, accountability, and performance metrics, embedded into the operating model. The next frontier isn't who has the most AI—it's who makes the smartest decisions about how AI and humans work together.

Federico Berruti is a partner in McKinsey's Toronto office, **Lari Hämäläinen** is a senior partner in the Seattle office, **Oana Cheta** is a partner in the Chicago office, **Venky Anant** is a partner in the Bay Area office, and **Damian Lewandowski** is an expert in the Miami office.

The authors wish to thank Eric Buesing, Greg Phalin, Maurice Obeid, and Senthil Muthiah for their contributions to this article.

This article was edited by Christian Johnson, an executive editor in the Washington, DC, office.

Copyright © 2025 McKinsey & Company. All rights reserved.