

The AI operating system: A new way for CEOs to run the enterprise

An AI operating system can become an enterprise's "second brain," but only if CEOs rethink how management itself evolves.

This article is a collaborative effort by Asutosh Padhi, Brian Elliott, Daniel Swan, Kim Baroudy, and Vik Sohoni, representing views from QuantumBlack, AI by McKinsey.



At a glance

- **AI adoption is outpacing enterprise value:** McKinsey's 2026 State of AI survey shows that 40 percent of respondents at large companies report scaling AI agents, yet only 6 percent say AI generates at least 5 percent of EBIT.
- **An AI operating system (AI OS) can unlock an incumbent's hidden advantage:** It puts the company's accumulated data, expertise, and institutional knowledge to work so leaders from the C-suite to the front lines can make better decisions and take faster action.
- **The value of the AI OS is significant:** Early AI OS deployments have accelerated planning cycles by up to ten times, achieved first value in about 12 weeks, and reduced time to ROI as the system extends to new domains by 50 percent.
- **The AI OS reduces traditional management constraints, requiring shifts in how the business is run:** As AI reduces long-standing constraints, CEOs need to move from calendar-driven reviews and information flows toward continuous intervention, more distributed decisions, and greater emphasis on human judgment.

The future of management in the AI age is starting to look pretty different. Consider this scenario:

Tuesday, 8:00 a.m.: The CEO of a global insurer returns to the office after attending an industry meeting where she learned that regulators in several markets plan to cap margins on certain retail products. Her first move: Ask the company's AI OS to assess the company's exposure. Within minutes, the system has read the draft rules, matched them against the firm's product catalog, and returned an answer: About 140 products in seven countries are subject to the new regulations, and roughly a third are at risk of breaching. It also flagged an issue it cannot resolve—each regulator defines margin differently. The CEO routes the analysis to the country business unit leaders to validate by end of day.

9:15 a.m.: The German business unit head opens the AI OS, which has codified the draft rules—scope, margin formula, threshold, timing—and reconciled the regulators' definition of margin with how the business measures it across pricing, finance, and commission systems. He runs the margin model against live product data, challenges a cost allocation assumption the regulator is unlikely to accept, and reruns the model with a revised assumption: Nine products are now in breach. By 4:00 p.m., all seven country analyses have been rolled up into the CEO's view, with the assumptions behind every number visible.

Wednesday, 8:00 a.m.: Before the CEO and her business unit heads meet by video, the AI OS has already modeled the options for each breaching product—reprice, redesign coverage, restructure commissions, or withdraw—and estimated the effects on volumes, distribution partners, and capital. Leaders test scenarios live and align on a plan: protect key distribution partnerships, with no withdrawals without CEO approval. The system records those decisions as guardrails.

Thursday, 8:00 a.m.: Line-of-business heads redesign products within the guardrails. AI agents generate candidate designs and prices and test each against the codified rule and product governance policies; actuaries refine the options, compliance reviews them, and line heads decide. Agents then draft partner communications for human sign-off. As each change is approved, product-level forecasts update and roll into the group financial scenario, allowing the CFO to watch the plan reshape and stress-test it—what if regulators adopt the stricter definition? After launch, the system tracks each product against the threshold and alerts line heads before it drifts toward breach. The next regulation will be faster: The rules, definitions, and models are all now part of the enterprise's second brain.

This AI OS scenario isn't a futuristic vision; companies are building it today. An AI OS is a management layer that functions like a "second brain" for the enterprise to help CEOs, business leaders, and managers make better decisions and act more effectively and quickly. It allows incumbents to harness their greatest untapped advantage: decades of proprietary data, operating experience, and domain expertise.

The capability is crucial for companies to escape the mire of high AI adoption rates with low returns. McKinsey's 2026 State of AI survey reveals that 40 percent of respondents at large companies have scaled AI agents.¹ Yet only 6 percent attribute a significant financial impact (at least 5 percent of EBIT) to AI, about the same as the prior year.² The gap suggests that the challenge is no longer simply getting people to use AI; it is changing how the enterprise itself works.

These are still early days for AI OS, but the results have been promising. Our experience shows that companies can accelerate planning cycles by up to ten times, achieve first value in an initial domain in about 12 weeks, and reduce time to ROI as the system extends to new domains by 50 percent. Other gains include 40 percent improvements in productivity, 50 percent savings, and up to 90 percent reductions in work stream cycle times.

The full power of the AI OS comes from its ability to harness AI to generate new growth, business models, and productivity. But that potential requires companies to both build an AI OS that harnesses the intelligence it already has and rethink how to run a business when the AI OS reduces traditional management constraints. This is the pathway to developing a competitive moat that grows. The sooner a company moves, the larger that moat can become.

¹ The survey defines "large companies" as organizations with annual revenues of more than \$1 billion and "scaling" as scaling AI agents in one or more functions.

² Dan Tinkoff, Lieven Van der Veken, Michael Chui, and Tara Balakrishnan, "The state of AI in 2026: On the road to ROI," McKinsey, August 25, 2026.


What is an AI operating system?

Think of an AI OS as a second brain for the enterprise. It is a continuously operating layer that maintains enterprise context (that is, the knowledge that gives data meaning), reasons across functions, recommends or executes decisions, and learns from outcomes. It connects the CEO, C-suite, functional leaders, managers, and those on the front lines in a way that’s never been possible before (Exhibit 1).

Exhibit 1

An AI operating system drives value creation from strategy to execution.

The AI orchestration layer enables information to move up, down, and across all levels of an enterprise



McKinsey & Company

People at each level see the information relevant to their responsibilities, ask questions, use judgment, make decisions, and delegate. Information moves up, down, and across the enterprise. The system also learns and expands as the enterprise uses it. As new business domains come onto the system, they can reuse and extend existing capabilities (see sidebar, “Elements of the AI OS”).

An AI OS connects the CEO, C-suite, functional leaders, managers, and those on the front lines in a way that’s never been possible before.

Elements of the AI OS

The AI operating system (AI OS) should be built on an open architecture that keeps the company in control of its core intelligence (the enterprise's "brain") while allowing the underlying technology to be swapped out as it evolves.

The existing estate:

- *Infrastructure*: Existing compute infrastructure can support the initial deployment. While software as a service can get companies started, the target-state preference is to use a company's own cloud or on-premises deployment.
- *Storage and data*: Structured and unstructured data across existing enterprise systems and platforms, including legacy systems, can stay where they are. Add data products as needed, and connect them virtually to the semantic layer.
- *Large language models (LLMs)*: Most companies source models, including open-weight LLMs, from many providers to avoid dependence on a single one (ensure the LLM models are not allowed to train on your data, or you may slowly lose your enterprise's "brain" and proprietary knowledge to the LLMs).

AI OS on top:

- *Semantic layer*: This layer provides AI with context—a consistent understanding of the enterprise (its entities, relationships, definitions, policies, and business logic). Ontologies and knowledge graphs can make critical parts of that context explicit, persistent, and reusable. Its key purpose and measure of value are its ability to extend the context window of the agents so they can operate like experienced employees, dramatically enhancing trust in their interpretation, recommendations, and actions.
- *Analytics and models*: Deterministic machine learning models and analytical tools have a significant role in creating value and extending the capabilities of agents via tool calls. Productize existing models and add new ones as needed to reimagine high-trust workflows.
- *Agent definition and orchestration*: Agents are declarative or coded and are built to be reused. Agents leverage all the layers below to serve up answers to humans in an easy-to-digest way. Different LLM models can be selected and orchestrated according to the problem they are best equipped to solve. Dynamically optimizing across models for performance and cost remains an evolving capability.
- *Experience layer*: This layer connects intelligence to action by bringing insights, decisions, and execution to people where they work through existing systems or conversational agentic interfaces.

An AI OS is not an executive dashboard, which reports on what has happened; rather, it provides a real-time view of what is happening now and what could happen, offering recommendations and enabling collaboration across the enterprise in near real time. It connects to live systems, and the interface operates as a cockpit so executives can make decisions and act quickly (table).

Table
What the enterprise’s ‘second brain’ enables

From traditional practices	to an AI operating system
Analytics explain past performance and inform decisions	Intelligence anticipates outcomes, surfaces options, and recommends actions
Periodic meetings and reviews are informed by historical and current data	Always-on sensing of performance, risks, and opportunities enables users to test and interrogate
Insights inform decisions that are executed through established human reporting structures and digital workflows	Insights are immediately translated into decisions and coordinated human and agent execution within defined guardrails; results are fed back into the system
Learning is captured through reviews, experience, and periodic model updates	Continuous learning from decisions, actions, and outcomes improves subsequent recommendations
Activity and progress are tracked	Execution, outcomes, and generated value are tracked
Insights and decisions are tailored to individual functions and levels	Shared enterprise intelligence connects decisions from the CEO through management, the front line, and back
New AI solutions are largely built around the requirements of each use case	Reusable context, models, agents, and workflows make each new solution faster and less costly to build

At one large US retailer, for example, district managers already had extensive performance data and dashboards across more than 7,000 stores. The new AI OS capability, however, synthesized all the data into a single personalized view that diagnosed likely causes of performance gaps, recommended actions, and helped managers prioritize which stores to visit.

Top strategic priorities for building the AI operating system

CEOs often face two traps. The first is assuming that AI is so powerful that it makes transformation almost effortless. The second is believing that years of foundational technology work must happen before the company can create value.

In our experience, both traps are avoidable. CEOs can start with a few high-value business problems and build the technology they need as they go. The work is complex, but they don't need a large technology investment to move ahead; most companies can work with what they have.

The key, however, is that CEOs and their top teams need to be closely involved in designing and building the system to ensure it reflects the business's strategic priorities:

- *Protect your advantages.* All decisions should be driven by a simple rule: Own the context; rent or replace the technology. Institutional knowledge and experience need to be portable across technologies so the company can protect what makes it distinctive. There can be a strong temptation to work with a particular vendor to accelerate deployment. But if that comes at the cost of losing control over core assets or being locked into a single environment, the dangers can quickly outweigh the benefits.
- *Limit operating system sprawl.* The whole point of an AI OS is to provide an integrated view of the entire business. If each function has an independent AI OS or each use case has its own tech stack, it becomes difficult to provide an integrated view. Because functions have different levels of technical maturity and will progress at different rates, leaders should accommodate one to three AI operating systems.
- *Reimagine every process with AI.* Much of the first wave of generative AI has focused on improving existing activities. True profit-and-loss (P&L) value requires companies to redesign workflows end to end. The key is to be clear about the decisions that matter and what corresponding data, insights, and AI systems are needed to make them. The highest value often comes from reimagining workflows that cross organizational boundaries.

One manufacturer, for example, reimaged sales and operations workflows to build an AI OS that automated demand forecasting and production scheduling based on demand, inventory, and machine status data. It also created AI-powered routing, planning, and pricing capabilities that reimaged workflows across boundaries. The opportunity was worth \$140 million to \$180 million a year.

- *Build so your advantage grows and compounds.* The value of an AI OS increases as it expands. Every decision and outcome the system captures sharpens its recommendations. An incumbent that starts now can build a lead that a late mover, even one renting the same models, cannot easily close.

The foundation of this compounding advantage is reusing enterprise context, agents, models, data connections, and workflows, which requires designing for reuse from the beginning. Critical data definitions and relationships should be codified and governed to ensure consistency while allowing the system to absorb new knowledge and relationships as it builds semantic understanding.³ This approach helps agents access the relevant slice of knowledge for the task at hand rather than parse the entire enterprise model, reducing unnecessary token consumption (Exhibit 2).


This effort does not require a massive data exercise at the start. Leave the data where it is and build relationships on top as teams solve business problems, and improve data at the source over time as it creates value.

Exhibit 2

The enterprise ontology is built to reuse existing components and extend the capability easily.

Illustrative business problem and enterprise ontologies (simplified)

● New ontology component unique to each business problem ● Reusable enterprise ontology component



McKinsey & Company

- *Incorporate AI economics into the operating system.* While token unit costs have declined more than 95 percent, AI spend at many companies has increased more than three times. The choice of model, the way tasks are structured, how much reasoning is consumed, and how much human support is needed all affect the often complex new economics of AI. Case in point: The cheapest intelligence does not necessarily lower costs if lower quality creates expensive mistakes.

An AI OS should therefore incorporate AI economics into its decision-making process so that business leaders have a true view into the P&L trade-offs of any decision. Those trade-offs will evolve as large language models evolve.

³ Semantic understanding is the AI operating system's ability to understand what a company's information means in a business context. It is built from key sources, including ontologies and knowledge graphs that explicitly define and connect concepts; metadata, schemas, and semantic models that describe data and its structure; business rules and policies that establish parameters and constraints; and documents and institutional knowledge that provide additional context.

- *Build on the technology you've already paid for.* Most companies already possess plenty of usable data and platforms to start making progress. Modern approaches can work with data across systems without requiring the company to first move everything into a new environment. The practical approach is to build enough of the enterprise's second brain to solve the first high-value business problems, but design it so it can expand easily.

When the old constraints fade: Rethinking how to manage a company

As the AI OS breaks down old constraints on management, management practices themselves need to change. While the AI field is still evolving quickly, CEOs and their top teams should start to think through how they and their teams work in a few key areas.

Manage by intervention, not by calendar

Most companies still manage through artificial calendaring (for example, quarterly business reviews and monthly performance meetings). Those rhythms emerged in part because gathering and analyzing information and aligning on decisions took time.

In an AI OS world, meetings for sharing information and tracking progress become less necessary because information is continuously available. CEOs should instead redesign the management calendar around decision points. That includes eliminating standing reviews and defining the handful of signals, thresholds, and exceptions that should trigger leadership intervention.

Key CEO questions:

- What types of decisions truly need in-person leadership focus?
- Which decisions can safely and reliably be delegated to AI agents?

As the AI OS breaks down old constraints on management, management practices themselves need to change.

Manage continuously, not through checkpoints

An AI OS creates an always-on business plan. As signals come in, leaders will intervene and allocate in real time. Analysis and decision cycle times collapse.

When these dynamics take hold, the business will need to sync the speed of its control functions and frontline operations with that of its AI OS. CEOs will need to identify operational chokepoints (for example, last-mile delivery and manufacturing production issues) and determine which actions can unblock them.

Key CEO questions:

- What are the most important operational chokepoints to unblock so that decisions can be implemented quickly?
- What incentives and controls are needed to ensure that important decisions and actions happen inside the AI OS (and not in organic shadow systems)?

Manage for judgment, not information flow

As AI agents take on more execution-level tasks, the human edge comes more from exercising judgment in various situations. That includes framing difficult questions, interpreting ambiguity, challenging assumptions, innovating, and applying expertise where the answer isn't obvious.

Executives will need to spend less time gathering and transmitting information and more time exercising judgment, communicating intent, coaching people, resolving exceptions, and improving performance.

Key CEO questions:

- How can the business create an always-on learning environment?
- How can the business develop and grow the set of human skills that creates the most value with an AI OS?

Manage through direct connections, not layers

A CEO or business leader can identify an issue and engage directly with someone several layers down to address it. The AI OS can give people closer to customers and operations a much richer context for their own decisions, which historically happened higher in the organization because information or expertise was scarce. Many layers in the organization are likely to become unnecessary.

Key CEO questions:

- How does a delayed, totally transparent organization maintain a business culture?
- How does the AI OS support coaching and development so that managers can expand their reporting lines from ten to 12 people today to about 20 to 40?
- How can the business create enough trust in the AI OS so that the organization doesn't gravitate back to re-creating layers of approval?

How two companies built their AI operating systems

Before building its AI operating system (AI OS), a global OEM committed to three key principles. One was to keep its enterprise knowledge portable. It aligned on open-standard formats for its architecture and ontologies so its “second brain” would not be locked into proprietary systems.

Second, it set out to reimagine the company around AI, not simply to implement AI. That commitment was most visible when it reimaged its workflows. Teams redesigned the workflows, then built digital twins, ontologies, and semantic layers to codify the relevant processes, policies, relationships, and institutional knowledge. That allowed agents to understand the business context, enabling them to execute their tasks accurately. The company tied agent performance to business outcomes as the measure for success. Without that test, it would have been impossible to know if the context had been fully captured.

Third, the company focused on reusability and high trust from the start. It built data products as inputs, codified the knowledge layer, and productized deterministic AI/machine learning (ML) models to broaden the kind of work AI agents could do with a high degree of trust.

Consider the bill-of-materials sourcing process. Traditionally, if a product requires a particular stainless steel fastener, procurement might look for the same component at a better price from another supplier. Understanding the intent of and parameters for using stainless steel screws (for example, thermal performance, corrosion resistance, strength, and weight), agents could instead explore other options with the designer (for example, galvanized steel or plastic) and alternative suppliers to meet those requirements.

Agents extract intent, inform the design with support of AI/ML models, and run parts of the request-for-information process, while analytical models calculate the cost and implications of the alternatives. Humans review those options and make

the final decisions. For the OEM, these reimaged and agentified processes led to significant business impact, including lower costs, better designs, and faster time to market.

One US-based manufacturer followed a similar approach. It grounded its AI OS effort on three main foundations. The first was to own the crown jewels—its data, models, business rules, and institutional knowledge. It did not want a black box solution embedded within a vendor platform.

Its second objective was to build for flexibility. That didn't mean starting from scratch. It used its existing technology investments but designed the system so models and other components could be swapped out or expanded as the technology evolved.

Third was standardizing how they worked. A governance “harness”—a set of integrations and configurations—would provide consistent controls and guidelines across use cases.

To begin, the company focused its efforts on two domains: pricing and planning (finance and production). Rather than relying on the proven but time-consuming manual effort to build the ontology and semantic understanding core to the AI OS, it used AI to accelerate the process.

That meant identifying key decisions and the data needed to make them, extracting rules and experience from records, and codifying business rules and institutional knowledge that had often lived in spreadsheets or in employees' heads. The company connected about 100 internal and external data sources, cleaned and linked them, and created a semantic representation of how the business makes pricing and planning decisions.

The initial phase of the AI OS was up and running in about 16 weeks. Institutional knowledge that once sat in pockets of the organization can now be made available, consistently and continuously, across the business.

Find more content like this on the
McKinsey Insights App



Scan • Download • Personalize



First actions to consider

To put the business on a strong AI OS path, CEOs should prioritize the following actions:

- *Focus the AI OS on the greatest value opportunities.* For most companies, a small number of domains disproportionately determine value. CEOs should identify and prioritize development around the three to five business domains in which AI could materially affect growth, margins, innovation, or competitive position.
- *Set the rules for the AI OS.* Decide how many systems the enterprise will allow, who governs the company's semantic understanding, and what standards teams should follow so knowledge and agents remain reusable and portable. Put in place incentives and processes so teams work within the shared AI OS rather than creating shadow versions. This self-reinforcing mechanism is essential for building compounding advantage.
- *Fund the shared AI OS (not individual use cases).* Work on the first domain should focus on building a scalable orchestration layer that includes reusable assets: an ontology with semantic understanding, codified knowledge, consistent business logic, deterministic AI/machine learning models and agents, and a composable user interface. Funding and approvals should be tied to how well teams build those assets. The goal is to create value now while ensuring each new domain builds on what has already been developed.
- *Redesign decision rights and management rhythms.* Identify which decisions AI and agents can make or execute, which require human judgment, and when issues should escalate to senior leadership. Then begin changing the management processes around those choices. This can be uncomfortable, but overly cautious action can derail progress.

AI heralds a new management era in which the constraints that were needed to manage the traditional modern corporation fade away. Incumbents ready to move into this new era already have a huge advantage through their institutional knowledge and experience. What remains is for them to commit to harnessing it through an AI OS system that can turn those assets into the enterprise's second brain for growth and productivity.

Asutosh Padhi and **Vik Sohoni** are senior partners in McKinsey's Chicago office, **Brian Elliott** is a senior partner in the Carolinas office, **Daniel Swan** is a senior partner in the Connecticut office, and **Kim Baroudy** is a senior partner in the Copenhagen office.

The authors wish to thank Chhavi Adtani, Fabian Metzeler, Sofiya Bidochko, and Zubin Ghafari for their contributions to this article.

This article was edited by Barr Seitz, an editorial director in the New York office.

Copyright © 2026 McKinsey & Company. All rights reserved.