

How Should Management Be Redesigned for Human-AI Operations?

AI will not simply change what employees do. It will change what managers coordinate, what they supervise, how quickly decisions move, and where human judgment creates value. The executive challenge is to redesign management around the work that remains - not preserve a structure built for a different information system.

Executive Operations Insights | Saša Cibula

For most of the modern organization, management has been designed around a simple constraint: information and coordination are expensive. Supervisors collect status, managers reconcile competing priorities, directors connect functions, and senior leaders resolve exceptions that cannot be settled lower in the hierarchy. Layers exist partly because people can only process so much information, coordinate so many relationships, and make so many decisions at once.

AI changes that constraint. A planning agent can monitor thousands of signals, summarize deviations, recommend actions, coordinate routine workflows, and increasingly execute bounded decisions without waiting for a manager to gather the information first. That does not make management unnecessary. It changes the economic and organizational logic for why particular management work exists.

The risk is that companies automate the work while leaving the management system untouched. Employees receive AI tools, workflows accelerate, and more decisions become machine-assisted, yet reporting lines, spans of control, meeting cadences, approvals, incentives, and leadership roles remain largely the same. The enterprise then carries two operating models at once: a faster technical system and a management structure built for the slower one it replaced.

The executive question is therefore not whether AI will reduce management layers. That is too narrow. The better question is: **what management work still creates value once AI absorbs more information processing, coordination, monitoring, and routine decision support?** The answer should determine the structure, not the other way around.

Joint Optimization, Not Technical Substitution

Sociotechnical Systems Theory provides the strongest foundation for this problem. Developed from work at the Tavistock Institute in the 1950s, the sociotechnical tradition argues that performance depends on the joint design of the social system and the technical system. Technology should not simply determine the shape of work; people, technology, roles, coordination, autonomy, and capability should be designed as an integrated work system.

That principle is especially relevant to AI because the technology can now perform activities that were historically part of management itself: analysis, prioritization, forecasting, scheduling, monitoring, recommendation, and increasingly action. A technical design decision can therefore become an organization-design decision almost immediately. Changing what the system can do changes what people need to do around it.

Recent human-AI research has returned to sociotechnical thinking for precisely this reason. Studies of complex human-AI and human-robot systems emphasize that effectiveness cannot be reduced to a one-person, one-machine interaction. Real operations involve distributed teams, shared technology, changing conditions, escalation, trust, learning, and multiple sources of authority. The unit of design is the work system, not the tool.

SOCIOTECHNICAL SYSTEMS: JOINT OPTIMIZATION FOR HUMAN-AI OPERATIONS

Design the social and technical system together, so each reinforces the performance of the whole.




Figure 1. Sociotechnical Systems Theory applied to management design for human-AI operations.

Redesign the Work Before Redesigning the Org Chart

The easiest response to AI productivity is to jump directly to headcount and layers. If software can absorb reporting, scheduling, analysis, and routine coordination, it is tempting to conclude that managers can simply supervise more people and the hierarchy can be compressed. Sometimes that will be correct. But starting with the org chart reverses the design sequence.

Executives should first decompose the work of management. Which activities exist to collect information? Which reconcile handoffs? Which approve low-risk decisions because information historically moved slowly? Which develop people? Which manage exceptions? Which protect controls? Which create context that is not captured in the system? AI may remove substantial effort from the first three categories while increasing the importance of the last four.

That distinction matters because a larger span of control is not automatically a better span of control. If AI eliminates routine status gathering and highlights meaningful exceptions, one manager may effectively oversee a broader operation. But if the remaining work becomes more ambiguous, relationship-intensive, or risk-sensitive, the manager's cognitive burden may actually increase. The more useful concept is not simply span of control, but **span of responsible attention**: how many outcomes, exceptions, people, and machine-driven processes can a leader understand well enough to intervene when judgment is required.

The same logic applies to delayering. A layer that exists primarily to transmit information upward, translate status across functions, or approve standardized low-risk decisions may lose much of its economic purpose. A layer that develops talent, integrates competing objectives, holds cross-functional context, or manages high-consequence exceptions may become more valuable, not less. AI should expose the reason a management layer exists before leadership decides whether to remove it.

The Management System Must Move With the Work

Jay Galbraith's Star Model provides a useful supporting lens because it treats organization design as more than structure. Strategy, structure, processes, rewards, and people practices must reinforce one another. Human-AI operations make that alignment problem more demanding because the technical system can change the content and speed of work faster than the surrounding organization changes its rules.

Structure should reflect where judgment and accountability remain human. Teams may become smaller, larger, flatter, or more cross-functional depending on what AI absorbs, but reporting lines should follow the real coordination

problem rather than legacy functional boundaries. The objective is not fewer managers as an end state; it is fewer managerial interfaces that add no value and stronger leadership where integration, judgment, or development still matter.

Processes must change because AI alters decision velocity. A daily management meeting designed to collect yesterday's status becomes questionable when the system detects deviations in real time. The meeting may need to shift from information exchange toward exception review, tradeoff resolution, learning, and improvement. Escalation thresholds should also become more explicit because a human manager should not be inserted into every decision merely to preserve the old workflow.

Rewards need to follow enterprise outcomes rather than activity that AI can now perform. A manager should not be rewarded for producing more reports, approving more transactions, or maintaining a large organization. Incentives should increasingly emphasize business outcomes, quality of judgment, system improvement, talent development, responsible use of automation, and the ability to improve performance across human and digital work.

People practices may require the most significant redesign. Managers will need enough technical fluency to understand where AI is reliable, where it is uncertain, and when intervention is necessary, but their comparative advantage should not be reduced to supervising software. Coaching, sensemaking, ethical judgment, conflict resolution, change leadership, and the development of institutional capability become more important as routine coordination becomes less scarce.

A Practical Example: From Layered Coordination to Human-AI Control

Consider a regional fulfillment network in which site managers, regional managers, transportation teams, inventory planners, and finance partners spend substantial time reconciling volume, labor, inventory, service risk, and transportation capacity. The existing hierarchy evolved because no single person or system could continuously assemble the entire picture. Management became the integration layer.

Now introduce an AI-enabled control system that continuously forecasts volume, identifies capacity constraints, models labor requirements, monitors transportation conditions, and recommends or executes bounded reallocations. The technical capability may remove hours of manual analysis and many routine coordination meetings. But if every existing approval remains, every site still optimizes its own metrics, and every exception still travels through the same hierarchy, the organization captures only part of the value.

A sociotechnical redesign starts with the end-to-end outcome. Routine network balancing can be handled within explicit machine authority. Site leaders retain authority over safety, local workforce conditions, and operational exceptions. A smaller regional team focuses on cross-network tradeoffs, unusual risk, capability building, and system performance rather than assembling status. Management cadences shift from reporting what happened to examining why the system behaved as it did, which exceptions required human intervention, and what should be redesigned next.

The economic value comes from more than labor reduction. Decision latency falls. Coordination cost falls. Capacity can be used more effectively across the network. Senior leaders spend less time adjudicating routine conflicts. At the same time, the redesign may require new investment in data quality, controls, training, monitoring, and roles that maintain the human-AI work system. Those costs belong in the business case because the organization is changing the management system, not merely installing software.

The Counterargument: Management Layers Carry Hidden Value

There is a serious counterargument to aggressive redesign. Management layers often carry functions that are easy to underestimate because they are difficult to measure: tacit knowledge, informal coordination, apprenticeship, social cohesion, customer context, organizational memory, and the ability to improvise when formal systems fail. Removing a layer can eliminate visible cost while quietly weakening the organization's capacity to learn and recover.

This is particularly important for the leadership pipeline. If AI absorbs much of the analytical and coordination work historically performed by early-career professionals and frontline managers, organizations must reconsider how future leaders develop judgment. Deloitte's 2026 work on AI-driven organizational layering raises precisely this issue:

reducing coordination work can flatten structures, but it can also remove the apprenticeship experiences through which people learn to manage complexity.

The answer is not to preserve obsolete work for training purposes. It is to redesign development deliberately. Managers may need structured rotations, simulated decision environments, cross-functional assignments, exception reviews, and greater exposure to system-level tradeoffs. If AI changes the work that once developed managers, leadership development must change with it.

The Executive Design Test

Executives should resist universal claims that AI will flatten organizations, widen spans, or eliminate middle management. Those outcomes depend on what work AI actually absorbs and what work becomes more important as a result. The correct design will vary by risk, operating complexity, workforce capability, regulatory environment, and the degree of autonomy granted to the technical system.

A useful test is whether the management structure still mirrors the economics of the work. If managers spend less time gathering information but the reporting architecture is unchanged, there may be excess coordination. If AI resolves routine exceptions but approval layers remain intact, decision rights may be obsolete. If human work becomes more judgment-intensive while spans are widened indiscriminately, the organization may create a different constraint. If AI accelerates execution but learning, controls, and capability development remain slow, the technical system will outrun the management system.

The executive objective should therefore be joint optimization. Remove management work whose primary purpose was to compensate for information scarcity or fragmented workflows. Strengthen the roles that create judgment, integration, learning, accountability, and human capability. Then align structure, processes, rewards, and people practices around the new division of work.

Redesign Management Around the Value Humans Still Create

The most consequential effect of AI on management may not be that machines perform more work. It may be that organizations can finally see how much management work was created by the limitations of the previous technical system. Reporting, coordination, approvals, and information transfer consumed managerial capacity because organizations needed humans to connect fragmented processes and incomplete information.

As those constraints change, the management system should change with them. But the objective is not a thinner hierarchy for its own sake. It is an organization in which human judgment is concentrated where judgment matters, AI is used where scale and speed matter, and the two are designed as one operating system rather than parallel structures.

The executive decision is therefore not how many managers AI can replace. It is more fundamental: **what should management exist to do in an operation where information, analysis, coordination, and some decisions are no longer exclusively human work?** Once that question is answered, the appropriate structure becomes much easier to see.

Selected Sources

Oxford Research Encyclopedia of Psychology. The Sociotechnical Approach to Work Organization. 2022.

Waterson, P. E., Older Gray, M. T., and Clegg, C. W. A Sociotechnical Method for Designing Work Systems. Human Factors, 2002.

Ang, K. C. S., Sankaran, S., and Liu, D. Advancing Sociotechnical Systems Theory: New Principles for Human-Robot Team Design and Development. Applied Ergonomics, 2025.

Galbraith Management Consultants. The Star Model and Organization Design.

McKinsey & Company. The Operating Model Advantage: Why AI Winners Are Rewiring Their Organizations. 2026.

Deloitte. The Work Chart vs. Org Chart: AI-Driven Organizational Delaying and Cross-Functional Role Fusion. 2026.