

Business Transformation in the Age of AI

Four success factors for measurable competitive advantage — a CEO perspective on leadership, human+agent teams, operating model, and governance.

Strat-Exx In Practice • Bernd Hirschle · Dr. Clemens Oertel • 9 min read

Transformations rarely fail because technologies do not work. They fail because strategy, operating model, decisions, capabilities, and responsibilities do not fit together.

The reality check:

88 % of organizations use AI regularly in at least one business function ^[1]	23 % already scale agentic AI — a further 39 % are experimenting with it ^[1]	~2/3 have not yet begun enterprise-wide scaling ^[1]	39 % report a measurable EBIT contribution ^[1]
--	--	---	--

The CEO question is therefore no longer: “How do we introduce AI?” Instead:

How do we create an *operating model* in which *people and AI agents* together generate *measurable business value*?

Business transformation thus becomes work on the entire system — on strategy, value creation, technology, capabilities, and steering. Four success factors remain decisive. Their underlying logic is timeless; in the agentic-AI era they are oriented consistently toward value, scaling, and the ability to learn.

Four levers for scalable transformation

The conditions under which ambition becomes measurable competitive advantage




Fig. 1 — Four levers for scalable transformation.

1 Smart leadership: setting direction and enabling decisions

It begins with a simple, understandable, and forward-looking target picture. It must not stop at phrases such as “We want to use AI,” “We are introducing agents,” or “We are becoming more digital.” Good leadership describes the concrete value the transformation should create: more growth, better customer experiences, higher quality, shorter cycle times, or new business models.

The central strategic question is: where do we want to improve existing processes — and where must we fundamentally redesign them? This decision takes courage. A transformation meant to avoid every interest at once loses focus and speed.

Smart leadership is visible and transparent. The CEO and top team not only explain the vision but make priority calls, resolve conflicts, protect transformation teams, and end initiatives without sufficient value contribution. They use new ways of working themselves, review progress against value and usage metrics, and do not delegate the transformation to a program or IT team.

The core team should combine experienced business experts with product, data, and technology skills as well as risk, compliance, and change expertise. External partners can accelerate capabilities; responsibility for the target picture, value, and decisions, however, stays within the company.

2 Consistent involvement: acceptance grows through co-creation

Winning the buy-in of operationally accountable stakeholders remains one of the hardest tasks. Resistance is not automatically obstruction: often it points to real risks, customer requirements, or informal processes overlooked in central planning. This knowledge must feed into the design early — without granting individual groups a permanent veto.

Buy-in does not arise from a communication campaign after planning is complete. It arises when affected employees and customers are involved in problem definition, prototypes, and rollout, help shape human+agent workflows, can follow progress, and see feedback actually lead to improvements.

Building capabilities is today part of the transformation, not an accompanying training project. **63 % of employers surveyed by the World Economic Forum** cite skills gaps as a major barrier to transformation; by 2030, **nearly 40 % of the skills needed at work** are expected to change.^[4]

This calls for role-specific learning, protected time to experiment, and clear expectations of leaders. With AI agents, new skills are added: specifying goals, delegating work sensibly, evaluating results, designing handovers, and intervening in critical situations. Employees need to understand which data and tools are used, where human control is required, and how errors or risks can be reported.^[2]

Adoption therefore belongs on the management dashboard: daily and weekly active users, repeat usage, the share of critical workflows with AI or agent support, quality of results, cycle time, and realized business value. The Microsoft Work Trend Index 2026 shows at the same time that organizational factors such as culture, leadership support, and talent practices correlate **more than twice as strongly** with reported AI value as individual factors do.^[2]

KEY POINT

Adoption shows not in the licenses purchased, but in confident daily use and better outcomes.

3

Deliberate decisions and consistent execution: from pilot to impact

Pilots deliver insights — but not yet transformation. What matters is deciding early which initiatives to scale, stop, or reorient. Quick wins are important, but must not lead to a collection of isolated solutions.

Four rules help balance speed and sustainability:

1

Start with a **relevant business problem** and name an accountable business owner as well as measurable targets.

2

Redesign workflows end-to-end instead of layering new tools onto old processes.

- 3 Create **reusable data, platform, and security building blocks** to scale further use cases faster.
- 4 **Stop initiatives without robust benefit or adoption** and reallocate resources.

Instead of a permanent “two-speed IT,” what is needed is a federated AI operating model. An AI factory provides reusable data, model, evaluation, security, and orchestration building blocks. Cross-functional AI product teams redesign end-to-end workflows and own business outcome, technology, and adoption. Human+agent teams define clear roles, handovers, quality standards, exception paths, and ultimate human accountability. This turns the center into an accelerator — not a bottleneck.^[2]^[8]




Fig. 2 — Federated AI operating model: AI factory, product teams, and human+agent teams.

Recent research shows: organizations with especially high AI value redesign workflows, embed AI into business processes, and track adoption and outcome metrics systematically. In agentic operation, agent workflows, human handovers, and quality standards must additionally become documented, repeatable, and capable of learning.^[1]^[2]

4 Excellent governance: guardrails that enable speed

Governance must not be an additional body that slows decisions. Good governance creates unambiguous decision rights, short escalation paths, and transparent criteria for funding, scaling, and stopping. It answers: who decides what, on the basis of which data, and within which guardrails?

In the age of agentic AI, the scope expands considerably. Beyond architecture, data quality, data protection, intellectual property, regulatory requirements, and cybersecurity, a controlled lifecycle for agents is needed: registered use cases and owners, approved models and tools, scoped identities and access rights, logging, evaluations, cost and action limits, defined escalations, and shutdown mechanisms. The intensity of human oversight must scale with risk, autonomy, and impact.^{[3][5][6]}

In many companies, a hybrid model proves effective: central units define principles, platforms, evaluation methods, and minimum standards; decentralized product teams own use cases, agent behavior, and business outcomes. A shared inventory creates transparency over deployed agents, data access, dependencies, and risks. This keeps guardrails consistent without removing operational ownership.

The active accountability of company leadership is measurably relevant here: in a McKinsey survey, CEO oversight of AI governance correlated with a higher self-reported contribution to results. Governance is thus not a pure compliance task, but part of the value and operating model.^[7]

Small, cross-functional teams need real autonomy within these guardrails. They should be able to make decisions quickly, own outcomes, and hand material trade-offs back to the governance body with clear recommendations.

What has changed compared to before

FROM digital channels	→	TO intelligent end-to-end processes performed jointly by people and agents.
FROM time-boxed projects	→	TO AI products, an AI factory, and reusable platforms.
FROM technology rollout	→	TO measurable daily use and effective human+agent teams.
FROM static programs	→	TO a learning operating model that masters change permanently.

Transformation as competitive advantage: learning faster than the market

Transformation is not an end state, and AI is not an end in itself. Competitive advantage arises when a company changes its value creation faster than the market — and feeds insights from every use back into better decisions, products, and processes.

For CEOs, the transformation agenda thus becomes concrete. It must show up in a few business-relevant impacts:

- faster and better decisions,
- shorter product development and time-to-market,
- more relevant customer experiences and higher sales effectiveness,
- more productive operations and lower cost-to-serve, and
- new offerings, revenue streams, and business models.

The sustainable advantage does not lie in a single model or agent. It lies in the speed with which the company redesigns work, measures quality, controls risk, and scales learning. Precisely this capability is hard to copy.^{[1][2]}

Conclusion: transformation is a leadership discipline — and a competitive system

There is still no one-size-fits-all recipe for business transformation. Industry, starting position, culture, and people differ. The four success factors, however, form a robust framework: clear leadership, genuine involvement, a scalable AI operating model, and governance that combines accountability with speed.

AI agents do not change this underlying logic — they raise the pace and reach. The difference is made by the ability to turn experiments into safe routines, daily use into measurable value, and human+agent teams into an organization capable of learning. Transformation does not end with the rollout; it becomes the enduring capability to renew competitiveness.

FOR PRACTICE

Five questions for the next transformation meeting

- 1

What concrete competitive advantage or business value do we want to create?
- 2

Which end-to-end workflow must change — and which work do people versus agents perform?
- 3

Who owns business outcome, agent oversight, and continuous improvement?
- 4

Which daily-usage, quality, and outcome metrics show whether real adoption is emerging?
- 5

Which capabilities, data, and guardrails do the teams need — and when do we scale, change, or stop?

*It is not the number of AI projects that decides. Winners connect **strategy, people, agents, and governance** into a learning operating model.*

SOURCES

A selection of current primary sources; as of 19 July 2026.

[1]

McKinsey & Company (2025): The state of AI in 2025 — Agents, innovation, and transformation. [View source →](#)

[2]

Microsoft (2026): Work Trend Index Annual Report — Agents, human agency, and the opportunity for every organization. [View source →](#)

[3]

Stanford HAI (2026): The 2026 AI Index Report. [View source →](#)

[4]

World Economic Forum (2025): Future of Jobs Report 2025. [View source →](#)

[5]

European Commission (2026): AI Act — Regulatory framework, implementation and governance. [View source →](#)

[6]

NIST (2024; upd. 2026): AI Risk Management Framework — Generative AI Profile (NIST AI 600-1). [View source →](#)

[7]

McKinsey & Company (2025): The state of AI — How organizations are rewiring to capture value. [View source →](#)

[8]

McKinsey & Company (2025): A new operating model for a new world. [View source →](#)

From ambition to *impact*.

We guide CEOs and leadership teams in translating AI agents into
a scalable operating model — as Senior Advisors for profitable
growth in B2B.

[Request a conversation →](#)