SERGIO CASTAGNA CONSULTING

AN EXECUTIVE GUIDE

Transforming the Enterprise for AI

The CEO's Guide to Turning Intelligence into Enterprise Performance

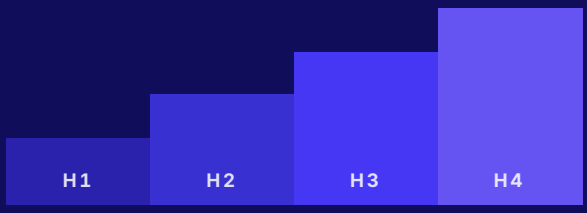

Augment. Delegate. Redesign. Reinvent.

August 2026

Contents

A Letter to the CEO

Artificial intelligence is already inside your company. Employees are using it to draft, analyse, search, summarise and solve problems. Technology teams are testing models and agents. Functions are assembling use cases. Vendors are presenting platforms, roadmaps and promises of productivity.

Yet activity is not transformation.

The strategic question is no longer whether your organisation will use AI. It is whether it will convert increasingly capable intelligence into better customer outcomes, stronger economics and a more competitive operating model — before fragmented initiatives create cost, dependency and unmanaged risk.

This is not another technology implementation. AI can support people, but it can also interpret objectives, use tools, coordinate activities and execute parts of a workflow. As these capabilities improve, AI will affect more than individual productivity. It will change how work is allocated, how decisions are made, where controls are embedded and how functions interact.

Previous transformation waves changed organisations significantly. Industrialisation transferred repeatable physical work to machines. ERP systems standardised information and processes. Shared services consolidated work and moved it to more efficient locations. Digitalisation improved access, speed and connectivity. But these transformations generally retained a fundamental assumption: people defined the work and technology executed predefined instructions.

AI begins to challenge that assumption.

The next operating model will not simply use AI. It will be partly operated by AI. That creates an opportunity to redesign the enterprise around faster decisions, fewer handoffs and a greater capacity to manage complexity. It also creates a new leadership obligation. The company must determine what AI is allowed to optimise, which actions it may take, when a person must intervene and who remains accountable for the result.

The CEO therefore needs more than an AI strategy. The CEO needs a transformation path: a way to move from employee adoption to agentic workflows, from workflow redesign to a new operating model, and from today's capabilities to a company able to absorb rapidly increasing intelligence.

This guide provides that path. It does not assume that one technology stack is right for every enterprise. It integrates business, finance and technology because none of those perspectives can lead this transformation alone. Its objective is practical: to help the CEO make better decisions, sequence investment, control autonomy and turn AI potential into measurable enterprise performance.

Executive Brief

Five conclusions for the CEO

1 **AI adoption and enterprise transformation are not the same thing.**
Access to AI can improve individual productivity quickly. Enterprise value requires more: priority workflows must be redesigned, data and architecture must support them, roles and controls must change, and performance must be measured at the level of the customer and the P&L.

2 **The employee is the entry point; the workflow is the first meaningful unit of transformation.**
Employee adoption builds literacy, reveals friction and generates evidence. Task-level assistance creates useful but often fragmented gains. The larger opportunity appears when the company redesigns an end-to-end flow of work across people, agents, systems and functions. This is where waiting time, duplicated controls, rework and organisational handoffs can be removed.

3 **AI introduces a progressive transfer of agency.**
The decisive distinction is not between one model and another. It is between AI that proposes and AI that acts. As systems move from assistance to execution and orchestration, the enterprise must introduce stronger permissions, evaluation, monitoring and human accountability.

4 **Transformation requires the integration of Business, Finance and Technology.**
Business determines where value and customer relevance exist. Finance converts ambition into economic discipline and measurable performance. Technology makes the solution secure, scalable and adaptable. Pairing any two creates a blind spot. The three must be integrated under one transformation mandate.

5 **The CEO needs an integrated enterprise perspective.**
Every internal function and external provider brings legitimate expertise, but also incentives, constraints and biases. A transformation integrator can challenge assumptions, arbitrate trade-offs and maintain coherence from strategy to execution through declared affiliations, explicit decision criteria and direct accountability to the CEO.

The governing principle is simple: if the objective is incomplete, the optimisation can be efficient and wrong.

Defining the objective, constraints and acceptable trade-offs is therefore a leadership responsibility, not a technical configuration task.

The transformation in one line

Augment people. Delegate work. Redesign the enterprise. Reinvent the company.

These are the Four Horizons of Enterprise Intelligence. They describe a progression, not four isolated programmes. The company can operate across several horizons at once, but it should not confuse activity in the first horizon with readiness for the third or fourth.

The CEO's immediate priorities

- Establish a CEO-led mandate with explicit decision rights.
- Diagnose value, workflows, data, architecture, risk and organisational readiness together.
- Select a small portfolio of outcomes that matter to customers and enterprise economics.
- Build an Enterprise AI Harness before granting material autonomy.
- Redesign priority workflows across functions rather than automating isolated tasks.
- Measure adoption, operational performance, financial value and risk separately.
- Create an operating model capable of absorbing rapidly increasing intelligence without betting on a specific AGI timeline.

PART I

The Case for Transformation

Why AI adoption and enterprise transformation are not the same thing, why this transformation is different from those that preceded it, and how AI changes the nature of organisational work.

CHAPTER 1

The AI Adoption Paradox

AI adoption is expanding rapidly, yet enterprise-level value remains uneven.

Recent research illustrates the gap. McKinsey reported that 88% of survey respondents were using AI regularly in at least one business function, while nearly two-thirds had not begun scaling AI across the enterprise. Only 39% attributed any enterprise-level EBIT impact to AI, and most of that group reported an impact below 5%.[1] Deloitte similarly found broadening access and rising expectations, but reported that only 34% of organisations were genuinely reimagining the business.[2]

The precise numbers vary by survey, population and provider. The pattern is more important than any single statistic: adoption is broad, production scale is narrower, and material enterprise value is concentrated. The statistics do not, by themselves, establish why.

At least four explanations may coexist:

Investment lag
Value may follow adoption with the delay typical of general-purpose technologies and large organisational programmes.

Attribution
Value may be captured but difficult to isolate credibly at EBIT level.

Capability limits
Models may not yet be reliable, economical or controllable enough for some high-value workflows.

Operating-model friction
Tools may be added without changing end-to-end work, ownership, controls, incentives or value-capture mechanisms.

1 McKinsey & Company, *The State of AI in 2025: Agents, Innovation, and Transformation*, 5 November 2025.

2 Deloitte, *The State of AI in the Enterprise 2026: The Untapped Edge*, 2026.

The CEO should diagnose which constraint dominates rather than assume that every disappointing result has the same cause. Waiting may be rational where capability is genuinely inadequate. Better measurement may resolve an attribution problem. Where the technology already performs but value remains blocked by fragmented ownership, duplicated work or unchanged decision rights, operating-model redesign becomes the relevant intervention.

Operating-model friction is visible when companies distribute tools without redefining work, collect use cases without choosing enterprise value pools, automate a step while leaving surrounding delays intact, or report time saved without determining whether revenue, margin, working capital, customer satisfaction or risk has improved.

The result can be an adoption paradox: the company appears busy, modern and innovative while its economics and customer proposition remain substantially unchanged.

The solution is not to stop experimentation. Experimentation creates learning, literacy and evidence. The solution is to connect experimentation to an explicit transformation architecture:

- a value agenda owned by the business;
- a portfolio of priority workflows;
- an architecture and data foundation proportionate to that portfolio;
- an Enterprise AI Harness controlling autonomy;
- an operating-model redesign;
- and a scorecard connecting adoption to enterprise performance.

The CEO should therefore ask a harder question than, “How many AI initiatives do we have?”

> THE BETTER QUESTION IS
>
> **Which customer outcomes and economic results can we attribute to AI — and is the principal constraint time, measurement, technological capability or the operating model?**

CHAPTER 2

Why This Transformation Is Different

Every major technology wave has changed the enterprise. The difference with AI should not be overstated, but it should not be underestimated either.

Industrialisation changed production economics by transferring physical execution to machines. ERP systems created common data structures, increased standardisation and made integrated planning possible. Shared services consolidated repeatable work and changed organisational boundaries. Cloud and digital platforms increased connectivity, scalability and speed. These transformations altered structures, roles and performance.

AI introduces an additional capability: increasingly, technology can interpret an objective rather than merely follow a fixed instruction.

An AI system may search, analyse, generate alternatives, select tools, plan steps, execute authorised actions, monitor results and adapt its next move. An agent may coordinate other specialised agents. A workflow can therefore become dynamic rather than fully predetermined.

This does not mean current AI is infallible, generally autonomous or equivalent to human judgement. It remains probabilistic, dependent on context and capable of producing confident errors. It does mean that the boundary between a tool and an actor is becoming operationally important.

In a traditional system, the company specifies the process and the software applies the rules. In an agentic system, the company may specify an objective, constraints and permissions while the system determines part of the path.

That changes five dimensions of management:

Work allocation
Some activities move from people to AI; others become a collaboration between them. Other activities move from software to AI.

Decision design
The organisation must determine which decisions AI may recommend, execute or escalate.

Control design
Controls must operate continuously inside the workflow, not only through periodic review.

Organisational boundaries
Agents can connect information and actions across departments, reducing the need for sequential handoffs.

Accountability
A machine may act, but the enterprise remains responsible for what it does.

> THE STRATEGIC IMPLICATION IS CLEAR
>
> **Every previous transformation changed how work was organised and executed. AI adds the possibility that technology can influence the path by which the work itself is performed.**

This is why the AI agenda cannot sit exclusively with IT, innovation or individual business functions. It affects the design of the enterprise.

CHAPTER 3

From Sequential Workflows to Multidimensional Systems

Most enterprises were designed to make complexity manageable for people.

They divide work into functions, convert uncertainty into procedures and move information through a sequence of approvals. Sales identifies an opportunity. Operations assesses feasibility. Finance evaluates economics. Risk checks exposure. Legal reviews terms. Management resolves the remaining trade-offs.

This structure is understandable. Human attention is limited, information is distributed and accountability must be clear. But the sequence creates cost: waiting, translation, duplicated analysis, local optimisation and loss of context between functions.

CHAPTER 3

From Sequential Workflows to Multidimensional Systems

AI can evaluate more variables and interactions at the same time. It can connect customer history, inventory, capacity, pricing, contract terms, service levels, credit exposure and strategic priorities in one decision environment. It can generate and compare scenarios continuously. It can identify patterns and connections that would be difficult for a person or department to see quickly.

Chess and Go demonstrated that machines can navigate enormous decision spaces under defined rules. Advances in protein-structure prediction demonstrated the value of AI in highly complex scientific problems. An enterprise is different: objectives are contested, data is incomplete, incentives conflict and consequences affect people. Nevertheless, the capability to handle multidimensional complexity has direct organisational implications.

TRADITIONAL WORKFLOW MANAGEMENT ASKS	ENTERPRISE INTELLIGENCE ASKS
“What is the next step?”	**“Given the objective, constraints, current state and possible consequences, what combination of actions creates the best outcome?”**

Figure 3.1 — The change in the governing question. The shift from a sequential view of work to a multidimensional one.

This shift can reduce functional handoffs, but it cannot eliminate governance. AI optimises what has been specified and represented in its context. It does not automatically understand the company’s true intention, ethical obligations or unstated trade-offs.

If the objective is incomplete, the optimisation can be efficient and wrong.

This is the central management problem of enterprise AI. Business understands the customer and competitive outcome. Finance makes the economic mechanism and trade-offs explicit. Technology translates the objective into data, architecture, permissions and observable system behaviour. Risk, legal, people and operational leaders define additional constraints and consequences. No function sees the complete objective alone.

The leadership task therefore moves upstream. The enterprise must define:

- the outcome to optimise;
- the constraints that cannot be violated;

- the trade-offs that require human judgement;
- the information the system may use;
- the actions it may take;
- and the evidence required to increase autonomy.

The future operating model is not simply faster. It is more connected, more adaptive and potentially less sequential.

THE DESIGN CONSEQUENCE

Traditional operating models were designed to simplify complexity for humans. AI makes more of that complexity operable.

CHAPTER 4

The Four Horizons of Enterprise Intelligence

PROPRIETARY FRAMEWORK

The AI journey is often presented as a collection of use cases or a technology maturity curve. The CEO needs a different view: how the unit of transformation expands and how responsibility changes with it.

The Four Horizons of Enterprise Intelligence describe that progression.

CHAPTER 4

The transformation in one line:
Augment people.
Delegate work.
Redesign the enterprise. Reinvent the company.

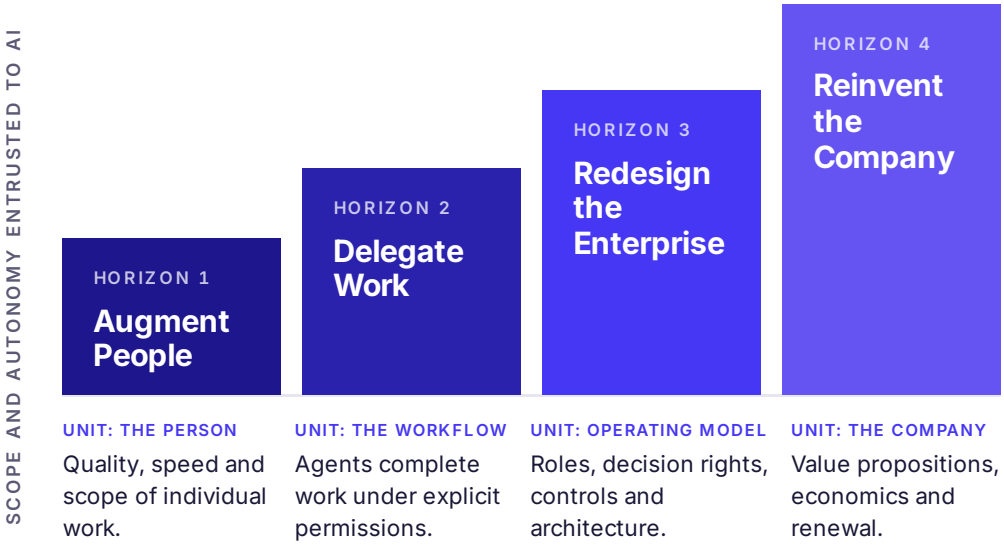

Figure 4.1 — The Four Horizons of Enterprise Intelligence. A proprietary framework of Sergio Castagna Consulting. The horizons describe a progression, not four isolated programmes; they are cumulative and may run in parallel.

Horizon 1 — Augment People

The unit of transformation is the person. AI helps employees research, draft, analyse, create, code and communicate. The objective is to improve the quality, speed and scope of individual work while building literacy and trust.

Horizon 2 — Delegate Work

The unit of transformation is the workflow. Agents complete defined sequences of work, use tools and coordinate activities under explicit permissions and escalation rules. The objective is not merely to save time but to remove waiting, rework and fragmentation across an end-to-end outcome.

Horizon 3 — Redesign the Enterprise

The unit of transformation is the operating model. Roles, organisational boundaries, decision rights, controls, performance measures and technology architecture are redesigned around people, AI and automation. The objective is structural performance advantage.

Horizon 4 — Reinvent the Company

The unit of transformation is the company itself. Leadership asks what the organisation should become if intelligence is abundant, scalable and increasingly autonomous. The objective is strategic renewal: new value propositions, different economics and an operating model capable of absorbing capabilities that do not yet exist.

The horizons are cumulative. A company does not complete employee adoption and then permanently move on. People continue to use AI while agents take on workflows and the operating model evolves. Different parts of the organisation will mature at different speeds.

But the distinction prevents a dangerous illusion. A company can distribute thousands of licences and still remain entirely within Horizon 1. It can deploy several agents and still preserve an unchanged operating model. Transformation begins when leadership connects the horizons and makes deliberate choices about where to advance.

The CEO should manage two movements at the same time:

- horizontal expansion, spreading proven capabilities across the organisation; and
- vertical progression, increasing the scope of work and level of autonomy entrusted to AI.

The further the company progresses vertically, the stronger its Enterprise AI Harness must become.

PART II

Building the AI-Ready Enterprise

The Four Horizons, the Enterprise AI Harness, controlled autonomy, value creation, workflow prioritisation, architecture, people, and the integrated Business–Finance–Technology leadership model.

CHAPTER 5

Horizon 1 — Augment People

The employee is the entry point for adoption. The workflow is the first meaningful unit of enterprise transformation.

It creates immediate access to intelligence, lowers the cost of experimentation and helps the organisation discover opportunities that a central programme would never identify alone. It also reveals practical barriers: inaccessible knowledge, poor data, unclear policies, fragmented applications and uneven management support.

OpenAI's enterprise guidance emphasises culture, ownership, quality and governance as conditions for scale, while Google and Microsoft place employee enablement at the beginning of their transformation models.[3,4,5] The common lesson is that adoption requires more than licences. People need permission, examples, skills, secure access and a clear reason to change how they work.

A credible Horizon 1 programme includes:

Safe access
Approved tools, clear data rules and role-appropriate permissions.

Leadership modelling
Executives use AI visibly and discuss both its value and limitations.

Work-based learning
Training is connected to real responsibilities, not generic demonstrations.

Use-case discovery
Employees identify recurring friction, knowledge gaps and decision bottlenecks.

Communities of practice
Strong users help peers, improve prompts and share methods.

Measurement
The company distinguishes activity from meaningful improvement.

3 OpenAI, *How Enterprises Are Scaling AI*, 11 May 2026; and *The State of Enterprise AI 2025*, December 2025.

4 Google Cloud, *Scale Your Team's Impact with Gemini Enterprise*, 2026; and *Gemini Enterprise Guidebook*, March 2026.

5 Microsoft, *2025: The Year the Frontier Firm Is Born*, 23 April 2025; and *Introduction to the Agentic AI Adoption Maturity Model*, 2026.

The CEO should resist two extremes. The first is uncontrolled experimentation, where confidential information moves into unapproved systems and teams create invisible dependencies. The second is restrictive governance that removes the opportunity to learn.

> **The objective is bounded exploration: wide enough to build capability, controlled enough to protect the enterprise.**

Metrics at this horizon should include active use, repeat use, quality improvement, cycle-time reduction, employee experience and the number of validated opportunities for workflow redesign. Self-reported time saved may be informative, but it should not be treated automatically as financial value. Time becomes value only when capacity is redeployed, service improves, cost is avoided or output grows.

Horizon 1 should not create a long detour before workflow redesign. Employee adoption and the redesign of priority workflows can progress in parallel: one builds literacy and reveals friction; the other converts that learning into structural value. The transition to Horizon 2 begins when the company stops asking only how AI can help a person perform a task and starts asking how an entire outcome should be delivered.

CHAPTER 6

Horizon 2 — Delegate Work to AI Agents

An AI agent is valuable not because it can hold a conversation, but because it can pursue an objective through a sequence of authorised actions.

An agent may retrieve information, reason over context, use enterprise tools, update systems, request approval and coordinate with other agents. That makes it suitable for workflows involving variation and judgement that conventional automation cannot handle easily.

CHAPTER 6

Horizon 2 — Delegate Work to AI Agents

The opportunity is significant, but current maturity should be described accurately. McKinsey reported that 62% of respondents were at least experimenting with agents, while only 23% were scaling an agentic system somewhere in the enterprise; no individual function exceeded 10% reporting scaled agent use.[1] Deloitte found that intended adoption was advancing faster than mature oversight, with only one in five surveyed companies reporting a mature governance model for autonomous agents.[2]

The lesson is not to wait. It is to scale autonomy according to evidence.

Good candidate workflows share several characteristics:

- the outcome can be defined clearly;
- the work occurs frequently enough to matter;
- data and tools can be accessed reliably;
- success and failure can be evaluated;
- exceptions can be recognised and escalated;
- actions can be contained or reversed;
- and a human owner accepts accountability.

A customer-service workflow illustrates the progression. At first, AI summarises the customer's history and drafts a response. Next, an agent classifies the request, checks policy, proposes a remedy and asks for approval. Later, it may execute low-risk remedies within defined commercial limits, update the customer record and escalate only exceptions. Eventually, specialised agents may coordinate service, billing, logistics, risk and retention actions around a single customer outcome.

The transformation is not the chatbot. It is the redesign of the service system.

The company should avoid creating an ungoverned population of agents owned by separate functions. Each agent introduces identity, access, data, model, evaluation, monitoring, cost and lifecycle requirements. A portfolio architecture is necessary from the beginning.

1 McKinsey & Company, *The State of AI in 2025: Agents, Innovation, and Transformation*, 5 November 2025.

2 Deloitte, *The State of AI in the Enterprise 2026: The Untapped Edge*, 2026.

AT HORIZON 2, THE CEO SHOULD ASK

Which complete outcomes are we prepared to delegate, under what limits, and what evidence would justify the next level of autonomy?

CHAPTER 7

Horizon 3 — Redesign the Operating Model

The largest value opportunity is unlikely to come from making every existing task slightly faster. It comes from removing work, handoffs and structures that no longer need to exist in their current form.

An operating model describes how the company turns strategy into results: how capabilities are organised, how work flows, how decisions are made, how performance is measured and how technology and data support execution. AI affects every component.

Research increasingly associates workflow redesign with higher reported value. McKinsey found that AI high performers were nearly three times as likely as others to report fundamentally redesigning workflows, and identified redesign as one of the factors most associated with meaningful business impact.[1] This does not establish causality: redesign may contribute to performance, result from early success or reflect stronger management capability. It does establish that workflow design is a serious management variable rather than a secondary implementation detail. The World Economic Forum similarly places end-to-end operating-model redesign and human accountability among the principles for scaled transformation.[6]

Redesign should begin with the customer or enterprise outcome, not the current organisation chart.

1 McKinsey & Company, *The State of AI in 2025: Agents, Innovation, and Transformation*, 5 November 2025.

6 World Economic Forum, *Organizational Transformation in the Age of AI*, 16 March 2026.

For each priority workflow, leadership should examine:

- which activities create value and which exist because of historical constraints;
- where information is translated or re-entered;
- which approvals are genuine controls and which are accumulated habit;
- where decisions wait for a person who lacks complete context;
- what AI can recommend, execute or orchestrate;
- where human judgement is indispensable;
- and how the economic result will be measured.

This analysis may lead to fewer sequential handoffs, broader roles, smaller coordination layers, new centres of expertise or more dynamic teams organised around outcomes. Some controls may become continuous and embedded rather than periodic and manual. Management work may shift from distributing information and checking activity towards setting objectives, developing people, resolving ambiguity and governing exceptions.

The answer will not be the same for every company. Regulation, customer trust, physical operations, data quality and strategic differentiation all affect the appropriate design.

THE PRINCIPLE IS UNIVERSAL

Do not automate the operating model you inherited before deciding whether it is the operating model you need.

CHAPTER 8

Horizon 4 — Build the AGI-Ready Enterprise

No responsible strategy should depend on a precise prediction of when artificial general intelligence will arrive or how it will be defined.

CHAPTER 8

Horizon 4 — Build the AGI-Ready Enterprise

The strategic issue is not when to place a single AGI bet. It is how the economics of the enterprise change if useful intelligence becomes progressively more abundant. The company should strengthen assets and capabilities that remain valuable across multiple technology timelines.

This changes strategic planning. Instead of extrapolating today's tools, the CEO can use a capability scenario:

> THE CAPABILITY SCENARIO
>
> **What company would we build if high-quality intelligence were abundant, scalable and increasingly autonomous?**

The question exposes assumptions embedded in the current model:

- Which activities exist because expertise is scarce?
- Which layers exist because coordination is expensive?
- Which products are valuable because analysis is difficult?
- Which customer frustrations are accepted because personalisation does not scale?
- Which decisions are slow because information is fragmented?
- Which advantages remain defensible when knowledge becomes easier to access?

If intelligence becomes more abundant, other assets may become more strategically scarce: trust, brand, customer access, proprietary data, physical infrastructure, licences, capital, distribution, relationships, judgement and social legitimacy. These complementary assets deserve attention now because their value does not depend on a precise AGI date.

An AGI-ready company does not automate everything. It develops adaptability. Its data and architecture are modular. Its workflows have explicit objectives and controls. Its agent identities and permissions can be managed. Its people know how to supervise, challenge and collaborate with AI. Its governance can respond to capability changes without rebuilding from zero. Its commercial strategy understands which forms of value remain differentiated.

The board should distinguish three types of decision:

No-regret moves
Modular architecture, explicit data rights, stronger customer relationships, workforce adaptability and governed agent identities.

Option-building investments
Bounded experiments and partnerships that preserve access to emerging capabilities without locking the company into one forecast.

Contingent bets
Larger commitments released only when defined capability, economic and risk thresholds are met.

It should test these decisions against at least three scenarios: continued incremental improvement, rapid expansion of agentic capability and a more discontinuous increase in general intelligence.

The purpose is not prediction. It is preparedness.

CHAPTER 9

The Enterprise AI Harness

PROPRIETARY FRAMEWORK

The more autonomy the enterprise gives to AI, the stronger — and more embedded — its harness must become.

Governance defines policies. Security protects assets. The Enterprise AI Harness makes objectives, permissions, controls, evidence and accountability operational inside the workflow.

DEFINITION

The Enterprise AI Harness is the system of objectives, permissions, controls, evidence and human accountability that allows AI to act with increasing autonomy without exposing the company, its customers or its stakeholders to uncontrolled risk.

It has seven components.

CHAPTER 9

The Enterprise AI Harness

1	**Objectives and constraints**	The system must know what outcome it is pursuing and which boundaries it cannot cross. Revenue, service, margin, fairness, compliance and reputation may conflict. The objective function is a leadership decision, not a technical configuration detail.
2	**Identity, access and permissions**	Every agent needs a traceable identity. Permissions should specify which data, systems, tools, transaction types and financial limits it may access. Access should be no broader than necessary.
3	**Data and architecture controls**	The company determines which data may be used, where processing occurs, which models are approved and how confidential information, intellectual property and sovereignty requirements are protected.
4	**Action and autonomy controls**	The harness distinguishes recommendation, approval-based action, bounded autonomous action and full workflow orchestration. Sensitive situations must trigger escalation.
5	**Evaluation and evidence**	The company defines what good performance means before deployment. It evaluates accuracy, robustness, bias, customer consequences and economic performance over time, not only in a one-off test.
6	**Monitoring and intervention**	Actions must be observable. The company needs anomaly detection, incident response, the ability to suspend an agent and a safe fallback when systems fail.
7	**Human accountability**	Every material workflow has a human owner. AI may execute a decision, but the company retains responsibility for the outcome.

Figure 9.1 — The Enterprise AI Harness. A proprietary framework of Sergio Castagna Consulting. The seven components operate together; the harness is weakest at its thinnest layer.

WHERE THE HARNESS BELONGS

The harness is not an additional control layer placed around AI after implementation. It is part of the operating model through which AI is allowed to act.

CHAPTER 10

The Controlled Autonomy Ladder

PROPRIETARY FRAMEWORK

Autonomy should be earned, not declared.

The Controlled Autonomy Ladder gives management a practical way to increase agency according to evidence and risk.

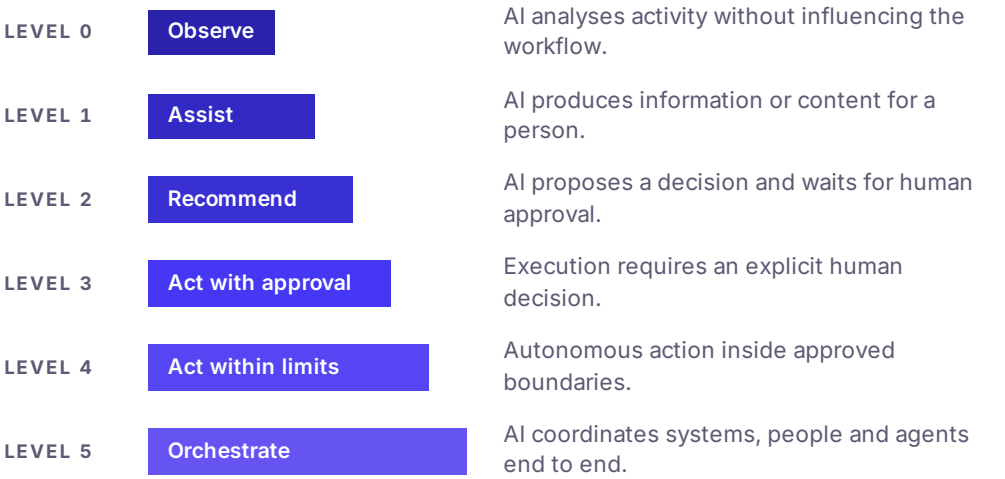

Figure 10.1 — The Controlled Autonomy Ladder. A proprietary framework of Sergio Castagna Consulting. Level descriptions are abbreviated in the figure; the full definitions follow.

Level 0 — Observe

AI analyses activity without influencing the workflow. The organisation establishes a baseline and tests whether the system understands context accurately.

Observe is also the proving mechanism between later levels. Before granting the permissions associated with level N+1, the company can run the system in shadow mode: produce the proposed decisions or actions without allowing them to affect the real workflow, then compare them with human decisions, actual outcomes and predefined thresholds. Simulation, historical replay and adversarial scenarios should supplement shadow evidence where events are rare or conditions change quickly.

Level 1 — Assist

AI produces information or content for a person. The employee remains responsible for interpretation and action.

Level 2 — Recommend

AI proposes a decision or next step, explains the relevant evidence and waits for human approval.

Level 3 — Act with approval

AI prepares or initiates an action, but execution requires an explicit human decision.

Level 4 — Act within limits

AI executes defined actions autonomously when conditions fall inside approved boundaries. It escalates exceptions or higher-risk cases.

Level 5 — Orchestrate

AI coordinates multiple systems, people or agents across an end-to-end workflow, while remaining subject to objectives, permissions, monitoring and intervention.

Progression should depend on four tests:

Performance
Does the system meet the required quality and reliability thresholds?

Containment
Are actions limited, observable and reversible where necessary?

Accountability
Is there a named owner with authority to intervene?

Value evidence
Does shadow-mode or limited-live evidence indicate that higher autonomy is likely to improve the customer and economic outcome?

Passing a gate should initially release bounded authority, not permanent autonomy. The company should define transaction limits, populations, duration, monitoring, escalation and rollback before live exposure. Evidence from controlled operation then confirms, adjusts or reverses the decision.

The appropriate level is contextual. A system may orchestrate routine internal scheduling while remaining at recommendation level for a credit, employment, medical or legal decision. Higher capability does not automatically justify higher autonomy.

The ladder also improves investment discipline. It prevents the company from building full autonomous execution before establishing whether assistance or recommendation already captures most of the value. Conversely, it prevents excessive controls from keeping a proven system permanently in a low-value pilot.

The CEO's objective is neither maximum autonomy nor minimum risk. It is the right autonomy for the outcome, supported by evidence and a harness proportionate to the consequences.

CHAPTER 11

Where Enterprise Value Comes From

AI value should be defined before technology selection.

The business case can draw from six enterprise value pools:

Growth
Higher conversion, faster product development, broader market access, improved pricing, new services and greater commercial capacity.

Customer value
Faster response, greater consistency, personalisation, improved resolution, proactive service and reduced customer effort.

Productivity and cost
Less manual work, lower rework, fewer handoffs, reduced external spend, increased capacity and better asset utilisation.

Decision quality
Faster analysis, better forecasts, broader scenario consideration, improved resource allocation and earlier detection of emerging issues.

Resilience and risk
Continuous monitoring, stronger compliance, quicker response, improved documentation and reduced operational fragility.

Strategic renewal
New business models, different sources of differentiation, redesigned organisational boundaries and the ability to compete at a structurally different cost or service level.

Figure 11.1 — The six enterprise value pools. A portfolio drawn from only one pool produces a narrower transformation than one designed across several.

A portfolio dominated by cost reduction will produce a narrower transformation than one designed around the customer and competitive agenda. The CEO should explicitly test whether AI can create new revenue, increase customer lifetime value, improve speed to market or make an uneconomic proposition viable — not only whether it can remove labour from existing work.

CHAPTER 11

Where Enterprise Value Comes From

The common mistake is to convert all reported time savings directly into cost savings. If an employee saves two hours but workload, output and cost remain unchanged, the company has created potential capacity, not captured value.

Every claimed productivity benefit should be classified:

Removed
Work or expenditure has genuinely disappeared from the cost base.

Redeployed
Released capacity has been assigned to growth, service, risk reduction or another measurable outcome.

Absorbed
Time has been saved locally, but no accountable economic or customer result has yet been captured.

Only removed and demonstrably redeployed capacity should enter the committed value case. Absorbed gains remain a hypothesis until management changes workload, output, hiring, service or structure.

Value capture requires a management decision. Capacity may be redeployed to growth, used to improve service, translated into avoided hiring or reflected in a redesigned structure. The financial mechanism must be explicit.

Every initiative should therefore connect four levels of measurement:

1. **AI performance.** Quality, reliability, latency and cost.
2. **Workflow performance.** Cycle time, throughput, rework and exceptions.
3. **Business outcome.** Revenue, margin, satisfaction, risk or working capital.
4. **Enterprise value.** Sustained P&L, cash-flow or strategic impact after implementation and operating costs.

Finance is essential not because it should restrict experimentation, but because it translates operational improvement into a credible value equation.

CHAPTER 12

How to Prioritise Workflows

The company does not need more use cases. It needs better choices.

A long list of ideas creates the appearance of momentum while distributing resources too thinly. The CEO should establish a portfolio that balances near-term proof, strategic differentiation and foundational learning.

Each candidate workflow should be assessed against six dimensions.

Enterprise value
How material is the customer, financial, strategic or risk outcome?

Transformability
Can the end-to-end workflow be redesigned, or will organisational boundaries limit the initiative to one isolated task?

Feasibility
Are the necessary data, systems, integrations and capabilities accessible within a reasonable period?

Evaluability
Can the company define success, test quality and detect unacceptable failure?

Risk and reversibility
What could go wrong, how quickly could harm occur and can actions be contained or reversed?

Organisational readiness
Is there an accountable owner, management commitment and willingness to change roles, controls and incentives?

The strongest first portfolio usually contains three types of initiative:

- **Lighthouse workflows** with visible customer or financial value;
- **Foundation workflows** that build reusable data, integration and harness capabilities;
- **Learning workflows** that test a strategically important capability in a bounded environment.

Prioritisation should penalise three patterns: a use case chosen mainly because a vendor can demonstrate it; a project with no committed process owner; and an initiative whose benefit depends on organisational change that has not been authorised.

The decision is not simply “build or buy”. It includes whether to configure, integrate, partner, redesign or stop. The company should preserve optionality where models and platforms are evolving quickly, particularly for capabilities that do not create unique differentiation.

The portfolio should be reviewed frequently. AI capability, cost and risk change faster than traditional annual planning cycles. A workflow that was infeasible six months ago may become attractive; a project that once appeared differentiated may become a commodity.

CHAPTER 13

Architecture, Data, Security and Sovereignty

The architecture decision is a business decision because it determines speed, control, economics and future freedom of action.

The enterprise needs enough standardisation to scale and enough modularity to adapt. A rigid commitment to one model or provider can create dependency. An uncontrolled multi-model environment can create duplication, risk and cost. The appropriate design is usually a governed portfolio with clear selection criteria.

The architecture should address:

- approved models and model-routing rules;
- enterprise data access and retrieval;
- agent identity and permissions;
- integration with systems of record;
- evaluation and observability;
- logging and auditability;
- security and privacy controls;
- cost monitoring;
- model and vendor substitution;
- resilience, fallback and incident response.

CHAPTER 13

Architecture, Data, Security and Sovereignty

Data quality remains important, but the company should avoid waiting for a perfect enterprise-wide data foundation before acting. Priority workflows can define which data must be improved first. This creates a value-led data roadmap rather than an open-ended cleansing programme.

Sovereignty deserves explicit CEO and board attention. Deloitte defines sovereign AI broadly around deploying AI under relevant laws, infrastructure and data requirements, with strategic independence as part of the issue.[2] For the enterprise, this means understanding where data and models operate, which jurisdictions apply, how intellectual property is protected and whether the company can change providers without losing critical capability.

Security must extend beyond the model. Agents interact with applications, credentials, tools and people. The attack surface includes manipulated instructions, excessive permissions, compromised data, unsafe tool use and unintended disclosure. Controls must therefore follow the action path from input to decision to execution.

The objective is not technical purity. It is an architecture that supports the target operating model and its risk appetite.

CHAPTER 14

People, Roles and the Human Advantage

AI transformation is not a workforce-reduction programme disguised as innovation. Nor is it credible to promise that every role will remain unchanged.

2 Deloitte, *The State of AI in the Enterprise 2026: The Untapped Edge*, 2026.

CHAPTER 14

People, Roles and the Human Advantage

Work will be redistributed. Some tasks will disappear, some will expand and new responsibilities will emerge. Microsoft describes the rise of the "agent boss": employees who build, delegate to and supervise agents.[5] McKinsey frames the future as a partnership between people, agents and robots and argues that workflow redesign is necessary to capture the potential.[7]

The enterprise should redesign roles around the future workflow, not add AI tasks to existing job descriptions indefinitely.

Human contribution becomes more valuable in areas such as:

- setting objectives and constraints;
- understanding customer and organisational context;
- exercising judgement in ambiguous situations;
- challenging outputs and assumptions;
- managing exceptions and consequences;
- building trust and relationships;
- and accepting accountability.

Managers face a particularly important transition. If AI distributes information, coordinates routine work and monitors execution, management value shifts towards direction, coaching, resource allocation, conflict resolution and governance. Some coordination layers may reduce; other roles may broaden.

The change programme should include role mapping, capability assessment, work-based learning, incentive redesign and transparent workforce planning. Employees need to know not only how to use AI, but how their work, authority and development may change.

Human oversight must also be designed realistically. A person who approves hundreds of AI recommendations without time, context or authority is not meaningful oversight. The workflow should define what the person must review, what evidence is presented and what action they can take.

5 Microsoft, *2025: The Year the Frontier Firm Is Born*, 23 April 2025; and *Introduction to the Agentic AI Adoption Maturity Model*, 2026.

7 McKinsey Global Institute, *Agents, Robots, and Us: Skill Partnerships in the Age of AI*, 25 November 2025.

The objective is not to keep a human in every loop. It is to place human judgement where it creates accountability, trust and value.

CHAPTER 15

The Leadership Model for Enterprise AI Transformation

PROPRIETARY FRAMEWORK

AI transformation requires a combination that most organisations distribute across separate leadership roles: Business, Finance and Technology.

Each pairing is valuable but incomplete.

Business	Finance	Technology
Determines where value and customer relevance exist.	Converts ambition into economic discipline and measurable performance.	Makes the solution secure, scalable and adaptable.
Finance + Technology Can create control, architecture and investment discipline, but may lack sufficient proximity to customers, markets and operational reality.	**Business + Finance** Can define priorities and economic value, but may underestimate data, architecture, security and delivery constraints.	**Business + Technology** Can accelerate innovation and experience design, but may lack the financial discipline required to convert experimentation into sustained enterprise value.

Figure 15.1 — Business–Finance–Technology integration. A proprietary framework of Sergio Castagna Consulting. Each pairing is valuable but leaves a blind spot; the three must be integrated under one transformation mandate.

Only the integration of all three can answer the complete set of transformation questions:

- Where should AI be applied?
- Which customer or strategic outcome should it improve?
- How will performance and value be measured?

- Which technology and architecture should be used?
- How will customers, data and the company be protected?
- How will the transformation be governed and executed?
- What new operating model is required to sustain the result?

This does not require one person to replace the chief operating officer, CFO, CIO, data, HR or risk leaders. It requires an integrator with sufficient understanding and authority to connect their perspectives, expose contradictions and maintain coherence from strategy to performance.

Governance remains essential, but committees do not automatically create an enterprise perspective. Business units, finance, technology, data, HR, risk and vendors all enter the agenda through legitimate responsibilities and different incentives. The resulting decisions can default to functional optimisation, platform preference, short-term savings, excessive caution or pilot proliferation.

The CEO therefore needs a leadership mechanism capable of challenging:

- a business case that treats theoretical hours as captured savings;
- an architecture selected before the target workflow is defined;
- a vendor proposal that creates dependency without strategic benefit;
- a risk position that prevents bounded learning;
- and a portfolio with no path to operating-model change.

The Enterprise Transformation Integrator performs three functions: integrating business strategy, financial value, architecture, risk, people and operating-model design; challenging assumptions and hidden dependencies; and moving the enterprise from diagnosis to decisions, redesigned workflows and measurable performance.

The capability may be internal or external, and an external integrator is not necessary in every company. Internal leadership may be sufficient when one executive has the mandate, breadth and credibility to connect Business, Finance and Technology and when the executive team can resolve cross-functional trade-offs quickly.

External integration becomes relevant when no internal role owns the end-to-end transformation logic; functional decisions use incompatible measures of value and risk; architecture or vendor choices precede definition of the target workflow; material conflicts remain unresolved; or the CEO needs a time-bound challenge during diagnosis and mobilisation.

The same discipline applied to other participants must apply to the external integrator. The mandate should specify scope, decision rights, evidence standards, declared affiliations, commercial interests, review points, stopping criteria and the conditions for transferring ownership to management. Recommendations to pause, narrow or stop should be legitimate outcomes of the work.

This role does not guarantee success or replace management accountability. It improves the quality and coherence of decisions and reduces avoidable execution failure.

PART III

The CEO Transformation Method

A practical transformation path: the Nine-Stage Transformation Roadmap, governance and decision rights, the CEO Transformation Scorecard, the readiness assessment and the first 100 days.

CHAPTER 16

The Nine-Stage Transformation Roadmap

PROPRIETARY FRAMEWORK

The transformation should be ambitious in direction and disciplined in sequence.

STAGE 1
Establish the CEO mandate
Define the strategic ambition, decision rights, governance, risk principles and accountability for the transformation as a whole.

STAGE 2
Diagnose the enterprise
Assess strategy, performance, customer journeys, workflows, organisation, data, architecture, risk and capabilities together. Identify structural constraints and existing initiatives.

STAGE 3
Define the value agenda
Select the customer, growth, cost, decision and resilience outcomes that matter. Convert them into a small portfolio of priority workflows.

STAGE 4
Build the foundations
Establish proportionate data, integration, model, security, evaluation and governance capabilities. Avoid building foundations without a defined value portfolio.

STAGE 5
Activate the workforce
Provide safe access, leadership modelling, work-based learning and mechanisms for discovering opportunities.

STAGE 6
Redesign priority workflows
Design the future flow of work across people, agents and systems. Remove unnecessary steps before automating them.

STAGE 7
Rebuild the operating model
Change roles, decision rights, structures, controls, incentives and performance measures to sustain the redesigned workflows.

STAGE 8
Scale with control
Expand proven patterns, industrialise reusable components and increase autonomy according to evidence. Monitor value and risk continuously.

STAGE 9
Design the AGI-ready company
Use capability scenarios to test strategy, scarce assets, business models, workforce and architecture. Build adaptability without relying on a precise forecast.

Figure 16.1 — The Nine-Stage Transformation Roadmap. A proprietary framework of Sergio Castagna Consulting. The stages are connected but not perfectly linear.

The stages are connected but not perfectly linear. Workforce activation can begin during diagnosis. Foundations can be built iteratively around priority workflows. AGI scenarios can influence today's architecture. What matters is that every activity belongs to a coherent transformation logic.

At each stage, the CEO should apply five gates:

Value
What measurable enterprise outcome will this create?

Evidence
What supports the assumptions?

Ownership
Who is accountable?

Risk
What could harm customers, the company or stakeholders?

Exit
Can we stop, reverse or change provider without unacceptable disruption?

CHAPTER 17

Governance and Decision Rights

AI governance should enable responsible speed, not separate risk from transformation.

OpenAI's enterprise guidance observes that early participation by security, legal, compliance and IT can accelerate later deployment by reducing reversals.[3] Databricks reports a strong association between the use of governance and evaluation tools and the rate at which AI projects reach production.[8] These findings reinforce a practical principle: controls designed early can be production enablers.

The governance model should have four levels.

Board and CEO
Set ambition, risk appetite and accountability. Review material exposures, capital allocation, workforce consequences and strategic scenarios.

Enterprise Transformation Integrator
Maintain coherence across Business, Finance and Technology. Prepare decisions, challenge evidence and escalate conflicts.

Domain and workflow owners
Own customer and operational outcomes, process redesign, adoption and performance. They remain accountable when AI is involved.

Enabling and control functions
Technology, data, security, finance, HR, legal, risk and compliance provide standards, expertise and embedded controls.

3 OpenAI, *How Enterprises Are Scaling AI*, 11 May 2026; and *The State of Enterprise AI 2025*, December 2025.

8 Databricks, *The State of AI Agents 2026*, 2026.

CHAPTER 17

Governance and Decision Rights

Decision rights must be explicit. Who approves a model? Who grants an agent permission to transact? Who can increase its autonomy? Who can suspend it? Who accepts residual risk? Who confirms that financial value has been captured?

Governance should be tiered according to consequence. Low-risk internal assistance should not face the same process as autonomous customer, financial or regulated decisions. The classification should consider data sensitivity, action scope, reversibility, customer impact, regulatory exposure and potential scale of harm.

Finally, governance must cover the lifecycle. Agents and models change, data drifts, workflows evolve and vendors update services. Approval at launch is not sufficient. Evaluation, monitoring and accountability continue until retirement.

CHAPTER 18

Measuring Transformation Performance: The CEO Scorecard

PROPRIETARY FRAMEWORK

What the company measures determines what the programme optimises.

If management measures licences and use cases, it will produce adoption and projects. If it measures theoretical hours saved, it will produce inflated value claims. The scorecard must connect six layers: adoption and capability; AI and agent performance; workflow performance; customer and employee outcomes; financial value; and risk and resilience.

Benefits should be baselined before redesign, tracked after implementation and adjusted for implementation costs, change effort and recurring model consumption. The business owner owns the operational outcome, the CFO validates the value mechanism, technology owns reliability and cost transparency, and risk monitors exposure. Where attribution is uncertain, the company should say so.

The CEO does not need a technical dashboard. The CEO needs a decision dashboard.

CHAPTER 18

Measuring Transformation Performance: The CEO Scorecard

A useful quarterly scorecard answers ten questions:

Strategic alignment	Are the priority initiatives connected to the company's strategic and customer agenda?
Value at stake	What validated value is targeted, committed and realised — and how much capacity is removed, redeployed or merely absorbed?
Portfolio concentration	Are resources focused on a manageable number of material workflows?
Workflow redesign	How many priority workflows have been genuinely redesigned rather than task-automated?
Autonomy exposure	Which agents can take consequential actions, and at what level of the Controlled Autonomy Ladder?
Harness coverage	Do all material agents have objectives, permissions, evaluations, monitoring, intervention and accountable owners?
Architecture resilience	Can the company substitute critical models or providers and maintain safe operation?
People readiness	Are roles, skills, incentives and management practices changing with the work?
Risk position	What material incidents, exceptions or emerging exposures require executive attention?
Decisions required	Which unresolved trade-offs need CEO or board action now?

Figure 18.1 — The CEO Transformation Scorecard. A proprietary framework of Sergio Castagna Consulting.

The scorecard should show direction and evidence, not only status colours. A project can be green on delivery and red on value. Adoption can be high while workflow change is low. Technical performance can improve while customer trust declines.

The CEO should review the scorecard with business, finance, technology, people and risk leaders together. Separate reporting reproduces the fragmentation the transformation is intended to remove.

THE STANDARD

Credibility is more valuable than an impressive but fragile number.

CHAPTER 19

CEO Readiness Assessment

Before increasing investment, the CEO should test whether the organisation is ready to convert it into performance.

The assessment is not a second scorecard. It is a go-or-redesign test built around seven conditions:

Strategic clarity
Can we name the enterprise outcomes AI must improve and how they affect competitive position?

Integrated leadership
Is the transformation CEO-led, with one person accountable for integrating Business, Finance and Technology?

Workflow ownership
Does every priority workflow have a business owner prepared to redesign roles, decisions and controls?

Technical feasibility
Can the required data and tools be accessed securely through an adaptable architecture?

Controlled autonomy
Do material agents have objectives, permissions, evaluation, monitoring, intervention and human accountability?

People readiness
Do employees and managers have the access, skills and transparency required to work with AI?

Value discipline
Are baselines defined, financial value mechanisms credible and stopping rules explicit?

READING THE RESULT

If the answers are mostly unclear, the company does not need a larger AI portfolio. It needs an integrated transformation diagnostic.

CHAPTER 20

The First 100 Days

The first 100 days should create clarity, evidence and momentum — not a permanent programme structure.

DAYS 1–20

Establish the mandate

- Confirm CEO sponsorship and board visibility.
- Appoint the Enterprise Transformation Integrator.
- Define decision rights, risk principles and reporting.
- Inventory active AI initiatives, vendors, contracts and shadow use.
- Agree the initial enterprise outcomes.

DAYS 21–45

Diagnose value and readiness

- Map priority customer journeys and enterprise workflows.
- Identify economic value pools and baseline performance.
- Assess data, architecture, security, skills and organisational constraints.
- Review vendor concentration and existing commitments.
- Classify current agents by autonomy and consequence.

DAYS 46–65

Select the portfolio

- Prioritise lighthouse, foundation and learning workflows.
- Assign accountable business owners.
- Define value hypotheses, evaluation criteria and stopping rules.
- Design the initial Enterprise AI Harness.
- Establish a workforce activation plan.

DAYS 66–85

Design and test

- Redesign selected workflows from the outcome backwards.
- Test AI performance in bounded environments.
- Define human roles, escalation and fallback.
- Confirm architecture and partner choices.
- Validate financial mechanisms and implementation costs.

DAYS 86–100

Commit and mobilise

- Approve the transformation roadmap and investment sequence.
- Launch the CEO Transformation Scorecard.
- Communicate what will change, why and how decisions will be made.
- Begin controlled implementation of the first workflows.
- Set the next autonomy and value gates.

Figure 20.1 — The first 100 days.

At day 100, the company should not claim transformation. It should possess something more valuable: a coherent mandate, a fact-based diagnosis, a focused portfolio, a functioning harness and evidence that the first workflows can create value safely.

CHAPTER 21

The Company You Would Build Today

Most companies are optimising organisations designed under assumptions that are beginning to disappear.

They assume expertise is scarce, coordination requires layers, work must move sequentially and technology follows explicit instructions. AI does not remove these constraints everywhere, but it weakens them sufficiently to make operating-model design a strategic question.

The greatest AI risk is not choosing the wrong model. Models can be replaced. The greater risk is continuing to operate a company designed for a world in which intelligence was scarce and work had to remain sequential.

The response is not a technology race or an indiscriminate automation programme. It is a disciplined transformation:

- augment people so the organisation learns;
- delegate bounded workflows so it develops evidence;
- redesign the operating model so value becomes structural;
- and prepare the company for increasingly capable intelligence without depending on a precise forecast.

The Enterprise AI Harness makes autonomy controlled and accountable. The integration of Business, Finance and Technology keeps customer value, economics and architecture in one decision system. Transparent transformation leadership helps the CEO challenge internal bias, external commercial incentives and the natural tendency to preserve the existing organisation.

The question is no longer simply, “How should we use AI?”

CHAPTER 21

The Company You Would Build Today

THE MORE CONSEQUENTIAL QUESTION IS

What company would we build today if we were not constrained by the operating model we inherited?

The companies that answer that question early will not merely use better tools. They will make better decisions, connect functions differently and create value through an operating model their competitors may struggle to reproduce.

ABOUT

Sergio Castagna Consulting

About Sergio Castagna Consulting

Sergio Castagna Consulting works with CEOs and leadership teams on enterprise transformation at the intersection of Business, Finance and Technology.

The approach begins with the business problem, the customer outcome and the economic result. Technology is selected and governed in service of the target operating model — not the other way around.

> Sergio Castagna Consulting has been an OpenAI Select Tier Partner since July 2026. This affiliation is disclosed because it is relevant to how clients assess technology recommendations. Formal affiliations are not symmetrical across the market: some technology providers do not offer comparable certification or partner programmes to independent consultants. The absence of an equivalent credential with another provider should therefore not be interpreted as a prior technology preference or as evidence that its solutions have been excluded from consideration.
>
> Additional technology partnerships are being developed and will be disclosed as they are formalised.
>
> Partnership status does not remove the obligation to compare models, platforms and architectures against the client's strategy, use cases, existing estate, economics, security, sovereignty and risk requirements. Any commercial interest relevant to a specific recommendation should be made explicit.
>
> In this guide, *independent* describes the discipline of judgement — transparent interests, evidence-based recommendations and accountability to the client's enterprise outcome — rather than the absence of any technology affiliation.

Independent AI Transformation Diagnostic

An independent assessment of where AI can create enterprise value, how the operating model must evolve and which decisions the CEO should make first.

The diagnostic examines:

- strategic value and competitive implications;
- customer journeys and priority workflows;
- financial value and performance baselines;
- data, architecture, security and sovereignty;

- agent autonomy and Enterprise AI Harness readiness;
- organisation, roles, skills and decision rights;
- portfolio, vendor and execution risks;
- and the first 100-day transformation agenda.

Independent Transformation Leadership

Acting as the Enterprise Transformation Integrator, Sergio Castagna Consulting connects business priorities, financial discipline and technology execution from mandate to measurable performance.

CEO Transformation Review

A periodic independent review of the transformation portfolio, architecture, providers, risks, autonomy and realised value, designed to support CEO and board decisions.

Technology providers can supply the intelligence. Sergio Castagna Consulting helps redesign the enterprise that will use it.

Website: www.sergiocastagna.com

NOTES AND SOURCES

Notes and Sources

This guide distinguishes external evidence from the proprietary analysis and frameworks of Sergio Castagna Consulting. Survey findings should be interpreted in light of their populations, definitions and provider perspectives. Company examples in the underlying reports are illustrations rather than universal proof. Superscript numbers in the text refer to the correspondingly numbered sources below.

SOURCES CITED IN THE TEXT

1 McKinsey & Company, *The State of AI in 2025: Agents, Innovation, and Transformation*, 5 November 2025. The report covers adoption, scaling, enterprise EBIT impact, agent experimentation, workflow redesign and risk mitigation.

2 Deloitte, *The State of AI in the Enterprise 2026: The Untapped Edge*, 2026. Global survey of more than 3,200 business and technology leaders across 24 countries and six industries; findings include business reimagination, readiness and governance of autonomous agents.

3 OpenAI, *How Enterprises Are Scaling AI: Practical Insights from European Enterprise Leaders*, 11 May 2026; and OpenAI, *The State of Enterprise AI 2025*, December 2025. The publications address culture, governance, workflow ownership, evaluation, hybrid work and enterprise adoption.

4 Google Cloud, *Scale Your Team's Impact with Gemini Enterprise: A Business Leader's Guide to Agentic Work*, 2026; and *Gemini Enterprise Guidebook*, March 2026. The guides describe employee productivity, agentic workflows, enterprise search, connected data, security and governance.

5 Microsoft, *2025: The Year the Frontier Firm Is Born*, Work Trend Index Annual Report, 23 April 2025; and Microsoft, *Introduction to the Agentic AI Adoption Maturity Model*, 2026. The publications cover human-agent teams, the "agent boss" concept and operating-model requirements for agentic AI.

6 World Economic Forum, *Organizational Transformation in the Age of AI: How Organizations Maximize AI's Potential*, 16 March 2026. Based on insights from more than 450 executives, the report highlights human accountability, operating-model redesign, talent, trust and experimentation.

7 McKinsey Global Institute, *Agents, Robots, and Us: Skill Partnerships in the Age of AI*, 25 November 2025. The report examines collaboration between people, agents and robots and the need to reimagine workflows.

8 Databricks, *The State of AI Agents 2026*, 2026. The report draws on aggregated platform telemetry and examines enterprise agent use cases, infrastructure, governance, evaluation and production scale.

BACKGROUND REPORTS REVIEWED

The following reports informed the preparation of this guide. They are not cited in the text and carry no reference number.

- Boston Consulting Group, *BCG AI Radar 2026: As AI Investments Surge, CEOs Take the Lead*, 15 January 2026. Survey of 2,360 executives, including 640 CEOs, addressing CEO ownership, investment and expectations for agents.
- NVIDIA, *State of AI 2026* industry survey series, March 2026. The series summarises more than 3,200 responses across financial services, retail and consumer goods, healthcare and life sciences, telecommunications and manufacturing. Respondents were drawn from audiences that had invested in or expressed interest in AI; this context is material when interpreting adoption and ROI findings.

PROPRIETARY FRAMEWORKS

Proprietary Frameworks

The following concepts and structures are presented in this guide as proprietary Sergio Castagna Consulting frameworks.

1. **The Four Horizons of Enterprise Intelligence** — Augment, Delegate, Redesign, Reinvent.
2. **The Enterprise AI Harness**
3. **The Controlled Autonomy Ladder**
4. **Business–Finance–Technology Integration**
5. **The Enterprise Transformation Integrator**
6. **The Nine-Stage Transformation Roadmap**
7. **The CEO Transformation Scorecard**