



# Architecting the Hybrid Organization

A Framework for Organizations where  
Intelligence Is No Longer Exclusively Human

---

**Tom Koulopoulos**

Lead Author | Founder, Delphi Group | Digital Futurist | Author of Fifteen Books

**Vlasios Vlastos**

Co-Author | CEO & Founder, Dalia Labs

# About This Paper

In April 2026, Delphi Group published a position paper asking what happens to leadership when the workforce is no longer only human. That paper argued that agentic AI is dismantling the fundamental assumptions of century-old leadership models and rendering traditional playbooks obsolete.

It concluded that the most successful leaders of the next decade will be those who successfully pivot from command to orchestration. It also exposed a stark reality: the urgency to adapt is absolute, as the runway for leaders to shape

this transition on their own terms is already vanishing.

This paper picks up where that one left off, addressing the vital next question: how must an organization fundamentally change to make this new way of working possible? It proposes a navigation system for the decisions that must be made before any change begins, as a direct response to the sequencing failure most organizations are making: treating structure as the goal to design toward, rather than the outcome those decisions produce.

|                                                      |    |
|------------------------------------------------------|----|
| Introduction                                         | 03 |
| The Break Point                                      | 04 |
| What We Built and Why It No Longer Fits              | 05 |
| The Architecture of the Hybrid Organization          | 07 |
| - The Shape Debate                                   | 07 |
| - The Agent Portfolio                                | 09 |
| Skills for the Hybrid Workforce                      | 13 |
| Roles in the Hybrid Organization                     | 15 |
| The Delphi Five-Step Organizational Design Framework | 16 |
| From Foundation to Form                              | 21 |
| Sources                                              | 23 |

## SECTION 01

# Introduction

In April 2026, we published a [position paper](#) [1] that asked a single question: what happens to leadership when your workforce is no longer only human? It argued that agentic AI is dismantling the four cornerstone assumptions on which a century of leadership theory was built, and with it, the playbook that shaped leaders' careers, written for an era where intelligence was exclusively human.

That paper left leaders with a hard truth: the ones who define the next decade will be those who make the transition from command to orchestration, and the window to do so on their own terms is already closing. This one starts with the inevitable question that follows. What must the organization itself become for that shift to actually work?

In answering this question, we reviewed more than fifty papers drawing from across academia, major consulting houses, and industry research. What we found was a lot of confident, competing claims about shape, about ratios, about how many layers survive, but very little agreement between them and even less evidence behind most of it.

What is true is that the traditional structure that has served organizations for decades is becoming the primary obstacle to the agentic future they are trying to build. With every passing moment, agents are becoming more capable, more autonomous, and more deeply embedded in the flow of work. Leaders are trying to adapt, develop new skills, new mental models, new ways of thinking about intelligence and accountability. But the organizations they serve still operate through structures, hierarchies, governance frameworks, and performance systems that were designed for a world where humans performed every meaningful function. The gap is between a world that is accelerating and a structure that was never designed to adapt to this degree of constant change.

To offer some clarity amid that confusion, we propose a framework for the organization of the future. This is not a blueprint for what the organization should look like, but a navigation system for the decisions that must be made before any change to the organization's structure begins. It is those decisions that will eventually shape the structure.

What is critical to understand at the outset, is that organizations do not get to opt out of this redesign, but what is required is a deliberate, values-anchored, outcome-oriented approach. This paper sets out what that redesign involves and what organizations must do to navigate the transition without losing the human qualities that make them worth building in the first place.

## SECTION 02

# The Break Point

We are living in an era of heightened ambiguity, and this is reflected in our organizations as much as it is already reflected in society as a whole. Agentic AI is one expression of it, but something about it is different. Not incrementally different but categorically different.

Every previous technology revolution augmented what humans could do, extending capacity, forcing organizations to adapt their structures and processes. None of them required organizations to reconceive what they fundamentally are. Agentic AI does just that. For the first time, we are not talking about systems that wait for instruction, but systems that set goals, perceive their environment, plan, decide, and learn with every interaction. The defining characteristic is not raw intelligence. It is autonomous initiative, and that changes the relationship between an organization and the technology it deploys.

This is not a capability that is coming. It is here, and it can absorb not just execution but the judgment and coordination organizations have always assumed required a human at the center. A single orchestrating agent can today direct a network of specialized agents beneath it, compressing into minutes what a human team took days to do. These are not hypothetical configurations. They are operational realities in organizations that have already made the shift.

*"The question is not whether agentic AI will reshape organizations. It already is. The question is whether organizations will reshape themselves deliberately and on time or discover too late that the structure they preserved was the very thing that held them back."*

The urgency we are describing is not abstract. Delphi research found that seventy-five percent of leaders consider adapting to agentic AI critical or highly urgent, yet fifty percent of their organizations are not prepared to act on it [2]. MIT Technology Review Insights confirms this finding: eighty-five percent want to be fully agentic within three years, yet seventy-six percent say their infrastructure cannot support that transition [3]. McKinsey names the same problem: eighty-six percent say their organizations are not ready, and the gap is not technological, it is structural [4,8].

BCG's research shows that what separates those capturing compound returns from those running expensive pilots that never reach the P&L is not agent sophistication [5]. MIT Technology Review argues that the real differentiator is whether organizations reimagine their operating model around agents, rather than layering them onto workflows built for an all-human workforce [3]. At Delphi, we believe it goes one step further: the differentiator is not what agents make possible, but what an organization deliberately chooses to do with them. The window for incremental adaptation has closed. What organizations need is not a technology strategy, it is a redesign of the enterprise, built for a world where human and agentic intelligence work together, and where the shape of the organization is not a legacy of the past but a conscious expression of what the organization intends to become. And this redesign does not begin with structure, not with agents, not with headcount, but with a set of foundational decisions that most organizations have not yet made.

## SECTION 03

## What We Built and Why It No Longer Fits

The organizational structures most enterprises operate today were not designed arbitrarily. They were solutions to real problems. Understanding what they solved, and why those solutions made sense, is the necessary foundation for understanding why agentic AI renders them structurally incompatible with the present moment. Two structural models have defined how organizations have been built.

### The Vertical Hierarchy

A pyramid-shaped model in which authority and direction flow downward through multiple layers of management, with information aggregating upward to give leaders the visibility they need to decide and course correct. This was the model that built the twentieth century industrial giants: General Motors, IBM, Ford, organizations that once stood several layers deep between the factory floor and the boardroom. It is still the default in large, regulated, execution-heavy organizations, such as government ministries, military command structures, and much of traditional banking, where the cost of a decision going wrong outweighs the cost of a decision going slowly.

The hierarchy existed to solve a specific problem: how does the intent of a leader at the top become action at the base, reliably, at scale, without distortion? The answer was structure, layer upon layer of management ensuring that what the leader decided eventually became execution at the base. Middle management was the organization's superglue, the connective layer holding strategy and execution together. Every layer of management existed because attention was finite and trust had to be earned at human scale.

### The Horizontal Structure

A flat structure, with fewer layers, wider spans of control, and more direct lines between leadership and execution, trading some coordination efficiency for speed and proximity to the ground. This is the model the tech industry built its reputation on: Netflix's thin management layers, Amazon's two-pizza teams, both built for markets moving too fast for a decision to survive the round trip through five layers of approval. Fewer layers meant faster decisions, and wider spans of control meant greater autonomy closer to the ground. Now strip away the shape and what remains is identical: an organization entirely dependent on human execution at every level. The model changed. The assumption did not. And it is that shared assumption that agentic AI is now dismantling.

This is not because these structures became less efficient, but because agentic AI removes the very operating limits they were designed to manage. The coordination infrastructure of the traditional organization existed because human execution was finite. When agents can execute continuously, in parallel, and across functions without those limits, the infrastructure built to manage human limitations does not become less relevant. It becomes unnecessary.

And it is here we find one of the fundamental paradoxes of agentic organizational structure: most organizations respond to this shift by layering autonomous systems onto a structure built for human execution, expecting productivity gains, and discovering instead that the agents and the structure are working against each other. MIT Technology Review calls this *the sticky tape problem*: organizations embedding agents into what remains a fundamentally human operating model, adding capabilities to a

structure that is already breaking rather than redesigning it. The result is agents constrained by the same linear, sequential, human-paced processes they were supposed to transcend.

The question in our view is not whether to change the structure. The structure will change regardless, because the forces acting on it are too powerful for inertia to hold.

*"The question is whether that change is designed deliberately, guided by a clear view of what the organization is trying to become, or whether it happens reactively, driven by the accumulated pressure of individual deployment decisions that were never connected to a coherent architectural intent."*

Deliberate redesign requires starting with a different question. Not how do we make our existing structure work better with AI, but what kind of organization do we need to be in order to create the value we exist to create, in a world where agentic AI is available to us and to every one of our competitors.




Figure 1: Ten Humans, Ten Agents, Two Architectures, Entirely Different Geometry. An illustrative example of organizational design: the same headcount and agent count, shaped into two entirely different structures.

## SECTION 04

# The Architecture of the Hybrid Organization

## The Shape Debate

The search for the right structure has produced no shortage of opinions, and remarkably little agreement. There is no consensus among the recommended structures for this transition.

PwC predicts the hourglass shape. As agents absorb execution work at the base, they also take over what the coordination middle used to do, translating direction, aggregating information, coordinating across functions. Once agents do that, the middle loses its primary rationale. What remains is a strong, expanding base, a lean middle, and a leadership tier above both, as entry-level workers become AI-fluent generalists capable of contributing at a higher level across workflows [6].

Deloitte predicts the diamond shape. The base narrows the same way, but rather than thinning, the middle expands. The middle does not disappear. It transforms demanding more judgment and accountability, not less, because what it must govern is more consequential than anything the old coordination layer managed [7].

These views are diametrically opposed. One thins the middle, the other thickens it. We think both are wrong if they are viewed as “the” correct structure. Organizational structure cannot be the one thing that holds still while everything else moves. No single shape can be prescribed in advance, because the shape each organization ends up with is not a template. It is the outcome of what that organization uniquely decides it is, what it values, and how it chooses to govern itself.

*"Our position at Delphi is that the structure is not the goal. It is an organic output. Any framework that hands a leader a precise shape and calls it the answer offers more confidence than the evidence allows."*

Organizations that begin their transformation by selecting a structure are starting in the wrong place. The architecture that emerges will be the right one only if it is the result of more fundamental decisions: about what the organization exists to achieve, about what it believes human beings are uniquely responsible for, and about how it intends to govern the autonomous systems it deploys. Two organizations in the same industry, with the same technology available to them, will arrive at different architectures because their purposes are different, their values are different, and their governance philosophies are different.

What is universal, and what we can state with confidence, is a small set of structural principles that hold regardless of the architecture any particular organization eventually discovers.

### 01 The execution base will be predominantly agent populated.

The structured, repeatable, and increasingly complex execution work that has historically occupied the base of the organizational pyramid will be largely performed by agents. This is not a future projection. It is already happening [4][8].

## 02 The judgment capability of the human layer must grow, even where headcount does not.

Regardless of whether an organization trends toward the hourglass, the diamond, or some new geometry yet to be defined, the humans who remain are not doing less. They are doing something categorically more demanding. They are responsible for interpreting ambiguity, governing agent behavior, handling exceptions that fall outside agent parameters, and maintaining accountability for outcomes produced at machine speed. That requires a level of judgment and contextual depth that the coordination middle was never required to provide and what is often underestimated [9][10].

## 03 The organizing principle shifts from authority to outcome.

For over a century, the org chart has been an authority map. It tells you who reports to whom. What the most advanced organizations are discovering is that this is no longer sufficient as the primary organizing logic. When humans and agents work together toward a shared outcome, the relevant question is not the reporting line. It is what needs to be achieved, and what combination of human judgment and agent capability produces it most effectively [11].

This shift from authority to outcome as the primary organizing principle is already taking shape in practice, and it is producing structural configurations that look unlike anything the traditional organization chart anticipated. Microsoft calls it the *work chart*: a dynamic map of human and agent capacity organized around the work to be done rather than the hierarchy of who reports to whom [12]. McKinsey describes the emergence of *agentic networks*: flat, outcome-focused configurations in which small teams of experienced humans supervise networks of specialized agents running end-to-end processes. A team of two to five people governing fifty to one hundred agents is not a smaller version of the old functional team. It is a different organizational unit with a different structural logic entirely [11]. Gartner calls these units *decision pods*: small groups of broad-skilled humans, each responsible for a specific decision domain, augmented by agents that handle the speed, scale, and data processing that no human team could match [13]. IBM frames the same idea through the lens of the *powertrain*: a concept of human-AI collaboration where people provide judgment, creativity, and oversight to keep decisions aligned with strategy and values, while agents surface insights and execute at speed and scale, both converting potential into performance by being held accountable to the same outcomes [14].

We believe these constructs point to something more fundamental. Traditional models assume that the unit of count is a human being. It is not, not anymore.

*"The addition of agents to the workforce adds a further dimension that makes prescribing any hybrid structural shape practically impossible."*

The shape of the hybrid organization reflects the configuration of both humans and agents, and once agents enter the equation, the variables multiply in ways that make any universal shape impossible to define. One organization may deploy a single orchestrating agent to manage an entire workflow that another organization breaks into ten specialized agents, each handling a discrete task. Both approaches can produce

identical outcomes. Both will produce entirely different organizational shapes. The number of agents, their autonomy levels, their specialization, and how they are configured relative to human oversight all determine what the organization looks like structurally. Two organizations with identical headcounts, identical industries, and identical strategic goals will look nothing alike structurally once their agent portfolios diverge. This is not a complexity to be resolved. It is a reality to be designed around. And it is the most powerful argument for why structural shape cannot be the starting point. The starting point must be the decisions that determine how humans and agents will work together, and the shape will follow from those decisions as a natural result.

## The Agent Portfolio

The decisions an organization makes about how many agents to deploy, at what levels of autonomy, whether to build, buy or lease them, and for how long to keep them, are among the most challenging decisions an organization faces. No existing management framework was built to guide them.

### Autonomy

Autonomy is a spectrum, not a destination. Existing frameworks describe agent capability in discrete levels, from basic assistance to full self-direction, risking the impression that organizations should climb toward maximum autonomy as a universal goal. Deloitte's Autonomy Ladder maps six levels of progression [15]. We believe the goal is not to climb the ladder. The right point on the spectrum for any given agent is determined not by what it is capable of, but by what the organization is aiming to achieve and ready to govern. The autonomy level should match the nature of the work, not the ceiling of what the technology can do. A high-autonomy agent on a simple, bounded task is not an asset, it is unnecessary risk. An agent operating at high autonomy where governance is not proportionate does not produce efficiency, it produces risk that compounds at machine speed. In our view, where that calibration is right, agents deliver transformational value; where it is wrong, they produce the accountability gaps and operational brittleness that follow whenever deployment outpaces deliberate calibration.

### Lifecycle

Agents have a lifecycle, and the most consequential dimension of that lifecycle is not deployment or decommissioning. It is what happens in between. Agents learn. They adapt. Their behavior, outputs, and decision patterns evolve through interaction with data, systems, and outcomes in ways that were never explicitly authorized at deployment. An agent running for six months in a complex enterprise environment may have drifted from its original design intent in ways that are subtle, cumulative, and consequential. The question is what has the agent become and whether it is still useful. And more than that, is it still the agent we intended? Those are different questions, and conflating them is a governance failure waiting to happen. KPMG frames agent lifecycle management as a direct parallel to human talent management, from onboarding through performance management to offboarding [16]. Our view here is different.

*"We believe agent evaluation must be a continuous integrity check rather than a periodic performance review. Agents operate at machine speed. Governance cannot operate at human pace."*

And decommissioning is not retirement. It is a governance decision triggered when an agent can no longer be verified as operating within its original design intent, or when its design intent is no longer aligned with the organization's current outcome architecture.

## Human-Agent Ratio

The right human-to-agent ratio cannot be reduced to a fixed number. It is a function, not a formula. McKinsey offers a widely cited figure, two to five people supervising fifty to one hundred agents [11], useful as a signal of scale but misleading as a design target. We believe that before any ratio can be set, an organization first has to decide what it will not delegate, regardless of how capable an agent becomes at doing it. Then comes the complexity and consequentiality of the work, which determines the autonomy level an agent should operate at in the first place. And it is that autonomy level that determines how much human judgment is actually needed, a lower-autonomy agent surfacing recommendations for approval demands a very different level of human involvement than a high-autonomy agent acting independently. And on top of this comes governance. The ratio is not decided by how many people are assigned, but by whether the people assigned actually have the time and depth of understanding to exercise real oversight. The question is not about how many agents a person can be assigned to. It is how many agents can a person genuinely govern effectively. It is all the above decisions that determine the ratio, and no organization arrives at the right number by asking the ratio question first.

In a well-designed multi-agent system, a human may interact directly with only one orchestrating agent, which itself directs a network of specialized agents beneath it. The interaction ratio at the surface is one to one. But is the accountability ratio one to one, or one to one hundred? That question does not have a comfortable answer, and organizations that do not ask it explicitly will find themselves with governance structures that look complete on paper and are inadequate in practice.




Figure 2: Interaction Ratio vs. Accountability Ratio: The interaction ratio measures how many agents a human directly touches. The accountability ratio measures how many a human is actually responsible for. The two are rarely the same number.

The accountability ratio is what determines whether human oversight is genuine or nominal. A human accountable for one orchestrating agent that governs a hundred beneath it may be exercising real governance, provided the orchestrating agent is transparent in its reasoning, the human genuinely understands how it works, knows its failure modes, and has the capability and the trained judgment to recognize when something is drifting and intervene before the consequences compound at machine speed. But if that understanding is superficial, the entire network beneath it is effectively ungoverned, regardless of how well it appears to perform on the surface.

*"Agent deployment that increases the volume of human accountability without creating the conditions for that accountability to be meaningfully discharged is not a productivity gain. It is an accountability gap."*

And accountability gaps in high-autonomy agent environments do not stay small. They scale with the system.

## Headcount or Capex

Underneath all of this sits a more basic question that organizations have not yet answered: Where does an agent actually belong in the organization? The emerging debate positions them as either human resources or capital expenditure. We believe neither is adequate. Agents are not humans, and never will be, but they are also not passive infrastructure executing predetermined logic. They perceive, decide, act, and learn, improving through use and degrading through drift simultaneously, which no capital asset does.

We believe that agents constitute a new category of workforce. One that borrows from human resource management in the areas of role definition, performance evaluation, and lifecycle governance but which also borrows from asset management in the areas of portfolio optimization and investment sizing.

## Build, Buy or Lease

And that brings us to the question: build, buy or lease? The conventional framing which suggests to build for control and differentiation and buy/lease for speed and commodity no longer holds. Most enterprise agent deployments in 2026 sit somewhere between those poles. We believe the more useful question is not build or buy/lease, it is where ownership creates advantage and where it creates drag. Commodity workflows, well-defined tasks, non-proprietary data, mature vendor solutions, are strong candidates for buying or commissioning as a service. Agents that must encode institutional knowledge, reflect organizational values, or operate where competitive differentiation depends on how the agent behaves cannot be bought off the shelf. The thing that makes them valuable is precisely what no vendor can build in advance. Those must be built, or built upon. Lease or commission is best for when the organization directs the objective and a vendor manages the infrastructure. It sits between the two, well suited to bounded tasks where speed matters more than long-term ownership. What matters is not which model an organization defaults to, but whether the choice is made consciously, based on where ownership creates genuine value.




Figure 3: The Agent Portfolio: five dimensions on which organizations must make deliberate decisions, none of which any existing management framework was built to guide: Autonomy, Lifecycle, Human-Agent Ratio, Headcount or Capex, and Build, Buy or Lease.

## SECTION 05

## Skills for the Hybrid Workforce

Every structural shift in an organization's history has demanded a parallel shift in skills. PwC, McKinsey, and Microsoft each point to a different emerging archetype. PwC points to the generalist, experienced workers moving toward broader roles, orchestrating across domains and directing agent-assisted workflows rather than executing within one [6]. McKinsey points to the T-shaped expert, a deep specialist whose value now lies in encoding that expertise into agent workflows and exercising judgment on the exceptions agents cannot resolve [10,11]. Microsoft describes the Agent Boss, the frontline worker who directs agent-assisted outcomes, evaluates what agents produce, and maintains accountability for results at a speed no human team previously managed [12].

We believe the hybrid organization needs all three. They are not competing visions of what the future worker should look like. They are complementary responses to the same structural shift, each describing a different layer of the human skills that agents cannot replace. The generalist orchestrates across the system. The specialist ensures the agents operating within each domain are well designed and well governed. And the frontline worker maintains the human accountability that no agent can carry. The doing becomes the designing. The execution becomes the governance.

Yet at the heart of this sits a paradox that every organization must keep in mind. The qualities that become most valuable as agents take over execution – most notably judgment, contextual depth, the ability to govern autonomous systems with genuine understanding – have historically been developed through early career execution.

*"The junior roles where people learned by doing, analyzing, processing, coordinating, are precisely the roles agents are absorbing first. Organizations need people with the deepest human qualities at exactly the moment they are eliminating the experiences that produced them."*

The more aggressively an organization deploys agents into entry-level and early-career workflows, the faster it erodes the foundation from which its future leaders emerge.

Beyond human skills, we believe the hybrid organization must define what qualities its agents need to function as a trustworthy workforce. Not what an agent can do, but what it must be to be deployable with confidence. Transparent in its reasoning, so humans can meaningfully oversee it. Bounded in its authority. Capable of escalating rather than resolving what falls outside its parameters. Consistent with the organization's values in ambiguous cases. Auditable, so every consequential action can be traced and challenged.

These are not technical specifications. They are organizational design requirements defined before deployment, not discovered when something goes wrong. Organizations that define them in advance govern their agent portfolios with confidence. Those that don't will instead find themselves managing consequences they never anticipated, at a speed that leaves little room for correction.

*"The measure of an agent's value is not its technical features. It is what it produces in the humans alongside it and in the business outcomes it contributes to. Organizations that evaluate agents on features alone are missing the bigger picture."*

What matters the most is what the agent contributes to the integrated performance of the hybrid workforce. That requires a measurement logic that no existing performance framework was built to provide. We call this *hybrid performance intelligence*: the capacity of an organization to measure and improve the combined output of its human and agent participants as an integrated system, rather than as two parallel workstreams running side by side. It is the missing infrastructure of the hybrid organization, and without it, organizations will struggle to know whether what they have built is working, improving, or quietly drifting from what they intended.

## SECTION 06

## Roles in the Hybrid Organization

The arrival of agentic AI is not primarily a story about new job titles. It is a story about the transformation of every existing role and the emergence of new functional needs that did not exist before.

Every role in the hybrid organization is changing in nature, not just in volume. The manager becomes an orchestrator of human and agent performance rather than a director of human effort alone, responsible for defining objectives that agents can pursue, evaluating what they produce, governing their behavior, and maintaining accountability for outcomes at machine speed. The specialist becomes a knowledge encoder, the person who translates years of domain expertise into the parameters, constraints, and judgment criteria that allow agents to operate within that domain reliably. The frontline worker becomes a director of agent-assisted outcomes, no longer performing the task but overseeing the agent that does, handling the exceptions that fall outside the agent's parameters, and maintaining the human accountability that the agent cannot carry. And leadership itself shifts from functional authority to system stewardship: responsible not only for the performance of people but for the design, governance, and continuous calibration of the human-agent system as a whole.

At the same time, new functional needs are emerging that every hybrid organization will need to address regardless of size, sector, or structural shape. The governance and oversight of agent behavior at scale. The design and continuous optimization of human-agent workflows. The measurement of integrated human and agent performance. The ethical stewardship of where and how agents operate, ensuring that the drive toward automation never overrides the values that define the organization. And the capability development of people to work effectively alongside agents, a need that is more urgent and more structurally demanding than any reskilling program the previous decade required.

How organizations staff these needs, whether through dedicated roles or distributed responsibilities, is a design choice that will vary. What will not vary is the need itself. And as technology evolves, new functional needs will emerge that cannot be fully anticipated today. The organizations that build the structural capacity to identify and respond to those needs continuously, rather than waiting for role definitions to catch up with reality, will hold an adaptive advantage that compounds over time.

## SECTION 07

## The Delphi Five-Step Organizational Design Framework

Most organizations approaching agentic AI transformation begin in the same place: experimentation and pilot deployment. They run proof of concepts, test use cases, and gradually start embedding agents into workflows before any of the important questions have been asked. As agents evolve, the questions evolve with them: what types of agents do we need, how many, and at what level of autonomy? Then they turn to people: how many do we need and what skills should they have? And almost as an afterthought, the structural question surfaces: how do we reorganize around all of this?

The Delphi Five-Step Organizational Design Framework is a direct response to that sequencing failure. It does not prescribe what the organization should look like at the end. We believe no framework can do that since the right design is specific to each organization and only becomes visible once the foundational decisions are made. What it prescribes is the foundational decisions that have to be made before any design begins.

It is designed to answer the question executive leaders are asking right now: agents are here, my organization needs to change, but where do I actually begin? Not with a transformation roadmap or a maturity model, but with five decisions that must be made before any structural or agent design begins. And the sequence of these decisions matters just as much as getting the decisions right.

### THE DELPHI FIVE-STEP ORGANIZATIONAL DESIGN FRAMEWORK



Figure 4: The Delphi Five-Step Organizational Design Framework: Five foundational decisions that must be made before any structural or agent design begins.

## Step One: Define Who We Want to Be in the Agentic Era

The first decision is not a technology decision. It is a leadership decision about organizational identity. Who do we want to be in a world where agentic AI is available to us and to every one of our competitors? What is our unique contribution when autonomous systems can perform not only tasks but create, decide, and act with agency?

This question must be answered before any structural or agent design begins, because every subsequent decision flows from it. Organizations that skip this step and begin with technology deployment will find themselves optimizing for efficiency in a direction they have never consciously chosen. The agents they deploy will make the identity decision for them, by default, through the accumulation of individual deployment choices that were never connected to a coherent view of what the organization is trying to become.

The answer to this question is not a mission statement revision. It is a genuine leadership reckoning with what agentic AI makes possible for this organization specifically, given its competitive position, its purpose, and its aspirations. What can we achieve now that we could not achieve before? What do we want to be known for in a world where the execution constraints that have always bounded ambition no longer apply in the same way? The vision that emerges from that reckoning is the reference point against which every subsequent decision is tested. Without it, the organization is not transforming. It is reacting.

## Step Two: Establish the Trust Covenant with People

The second decision concerns the people inside the organization, and it must be made with the same deliberateness as the first.

*"The Agentic AI transformation will not succeed in an organization whose workforce does not trust the intent of those leading it."*

KPMG's Global AI Pulse Survey confirms that nearly forty percent of hesitant employees cite job security worries, often alongside a lack of confidence in their AI skills, and workforce resistance is already registering as a structural challenge across industries [17].

The honest answer to the job security question is not reassuring, and organizations that attempt reassurance without honesty will suffer. Roles are changing. Some are disappearing. The pipeline of early-career experience that has built expertise for generations is being disrupted in ways whose long-term consequences are not yet fully understood. Leaders who promise that nothing will change are not building trust. They are borrowing it at a rate they will not be able to repay.

What leaders can honestly commit to, and what the Delphi Five-Step Framework proposes as the trust covenant, are three things.

1. We will invest in your capability as a continuous commitment
2. We will be transparent about what is changing and why before it changes. People who understand why their role is changing and what the organization's intent is can adapt. People who discover change has happened to them cannot.

3. We will redefine what contribution means rather than simply removing roles that agents can replicate. The human contribution in the agentic era is not less than it was before. It is different. Leadership's responsibility is to define that difference clearly and invest in the capability that makes it real.

This covenant is itself a values statement, because the willingness to make it and keep it reflects what the organization believes about the people who work within it. Organizations that make this covenant before transformation begins will find that their people become active participants in the redesign rather than resisters of it. That difference, between a workforce that is with the transformation and one that is managing around it, determines more about the outcome than any structural decision that follows.

### **Step Three: What Remains Human**

The third decision is about the human perimeter, and it must be made explicitly, formally, and before any agent is deployed. Which activities are legally, ethically, or strategically reserved for human beings? Which judgments must remain with humans not because agents cannot perform them but because the organization's values require it?

This framing is deliberate. The human perimeter is not primarily a risk management question. Risk management asks what agents cannot safely be trusted to do. The values question asks what humans must remain responsible for regardless of what agents can do. Those are different questions, and the answers are different. An organization might conclude that certain customer interactions must involve a human being not because an agent would handle them poorly, but because the organization's values around human dignity and genuine connection require it. Another might conclude that all decisions with legal accountability must have a named human owner not because agents cannot reason about legal questions but because accountability cannot be delegated to a system that cannot be held responsible in the way a person can.

Defining this perimeter before deployment is the operational and psychological unlock that makes the rest of the transformation possible. Organizations that attempt to deploy agents without drawing this line in advance discover that every deployment decision triggers a governance debate that should have been resolved upstream.

There is a further dimension that makes this step not just strategically important but culturally essential. When an organization defines what remains human, it is making a public statement about what it values. It is telling its people, its customers, and its partners that the drive toward efficiency will not override the qualities that make the organization worth dealing with in the first place. That statement, made clearly and kept visibly, is one of the most powerful trust-building acts an organization can perform in the middle of a transformation that unsettles almost everyone it touches. It will also be one of the hardest commitments to keep. As competitors move faster, as agent capability expands, and as the pressure to automate everything that can be automated intensifies, the temptation to quietly redraw the human perimeter will grow.

### **Step Four: Design Governance Architecture**

The fourth decision is the first design decision, and it is the one most consistently made in the wrong sequence. Governance architecture, the operational mechanisms of oversight, accountability, observability, and decision rights that make agent deployment safe and trustworthy at scale, is being treated across the industry as something built after deployment rather than before it. The consequences of that sequencing error are already visible. Gartner reports that only twenty-three percent of IT leaders are very confident in

their organizations' ability to manage security and governance [13]. They also project that more than forty percent of agentic AI projects will be canceled by the end of 2027 due to governance and return on investment failures [18].

Governance architecture is distinct from the human perimeter defined in step three. The human perimeter is a values and policy decision: it defines what is reserved for humans. Governance architecture is an operational design decision: it defines how that reservation is enforced, how agent behavior is made observable, how accountability is assigned when things go wrong, and how the organization maintains meaningful oversight at machine speed and at scale.

*"The critical distinction that organizations must make explicit in their governance design is between human-in-the-loop and human-on-the-loop governance. Human-in-the-loop means a person approves before the agent acts. Human-on-the-loop means the agent acts and a person monitors afterward [9]."*

This is not new terminology, it is borrowed from military doctrine on autonomous weapons and drone warfare, where the same distinction marks the shift from traditional command and control, where orders are traveling up a chain and back down, to systems making engagement decisions faster than that chain can move. These are structurally different arrangements with profoundly different accountability implications. Many organizations have quietly moved to human-on-the-loop without explicitly deciding to, because agents have been given operational authority that was never formally assigned or understood. Closing this gap requires explicit choices about which decisions require human judgment in advance, which can be monitored after the fact, and which should never be delegated to agents regardless of demonstrated capability.

The outcomes an organization can safely pursue are bounded by the governance infrastructure it is prepared to maintain. Designing outcomes without knowing what governance capacity exists is designing ambitions that cannot be safely executed. Singapore's Model AI Governance Framework, launched in January 2026 as the world's first governance framework specifically designed for agentic systems, reflects precisely this principle: organizations and the humans who oversee agents remain accountable for their actions, and risk must be assessed and bounded upfront rather than addressed after deployment [19].

## Step Five: Design the Outcomes

With identity, trust, values and governance as the foundation, the organization is ready to make the first concrete design decision: what outcomes is it actually trying to achieve with agentic AI, expressed not as technology goals but as business results?

Outcomes are the bridge between aspiration and design. They translate the answer to step one, who we want to be, into a concrete map of what the organization needs to produce differently: which decisions must be made faster, which customer experiences must change, which processes must deliver different results, which capabilities must be built that do not currently exist. It is the exercise of making aspiration tangible before any structural or agent design begins.

Without outcomes, organizations jump from vision to deployment. They ask what agents can do and then find uses for them. With the outcomes defined, the question is never what can agents do? It is what

outcomes do we need, and what combination of humans and agents produces them most effectively given our values, our governance capacity, and the trust we have built with our people?

Defining your outcomes is also the moment at which the shape of the organization begins to become visible, not as a decision but as a result. When an organization maps the outcomes it needs to achieve and asks honestly what combination of human judgment and agent capability produces each of them, the structure that is required to support that combination starts to emerge. It is not imposed. It is discovered. And the shape that is discovered will be the right shape for that organization, derived from its own purpose, values, and governance philosophy, which is the only meaningful criterion.

## SECTION 08

## From Foundation to Form

What comes next is what happens once those decisions have actually been made, and how they compound into the design itself. An important point to acknowledge is that the five decisions are weighted differently depending on the type of work an industry actually involves. Two businesses in the same sector can carry that work very differently, and a business can carry more than one of these characteristics at once, in which case more than one decision ends up carrying weight simultaneously.

Where the work itself carries physical or safety stakes, such as industrial operations, manufacturing, or physical healthcare delivery, it is Step Three (What Remains Human) that carries the most weight. When an error can hurt someone physically or trigger a serious safety failure, the answer to Step Three leaves little room for judgment and points to one clear answer. More roles get reserved for humans, both at the point where the physical work happens and at the point where it's checked and signed off. It is the nature of the work itself that pulls that decision hard in one direction, simply because the cost of getting it wrong is not acceptable. How many roles that ends up being, and where exactly they sit, is still a decision each organization has to make for itself.

When the work exposes the organization to financial, legal, or regulatory consequence, such as banking, insurance, or a hospital's compliance function, it is Step Four (Governance Architecture) that carries the most weight, leaving little room for compromise and pointing toward tighter oversight than the interaction ratio alone would suggest. The accountability ratio has to run tighter here, the same oversight capacity that lets a small team govern a two-digit or in some cases three-digit number of agents where the stakes are lower cannot stretch that far, so each team of overseers ends up governing fewer agents, and more such teams are needed to cover the same volume of work. That governance decision is what primarily determines the operating structure. How far it goes will still depend on how each organization weighs its own exposure.

In work that runs on judgment and relationship rather than volume, such as advisory practice, legal counsel, or executive consulting, the weight falls on Steps One (Define Who We Want to Be) and Three, but from a different angle. Here the value being protected is the judgment and the relationship a small number of people hold, so the answers to these two decisions are defining. Fewer roles are likely to be reserved for humans once agents absorb the research, drafting, and analysis that used to require a large body of staff to produce it. Between them, these two decisions set the human to agent ratio here, and through it, the structure that follows, though again, how far that ratio shifts will still differ by organization.

These characteristics do not sort neatly by sector, and a business can score high on more than one at once. A vaccine manufacturer, for example, carries physical and safety stakes and heavy regulatory exposure together, so Step Three and Step Four both end up carrying real weight, two decisions loading at once, not because manufacturing as a category demands it, but because this particular operation loads on two dimensions simultaneously. A different manufacturer with lower safety and regulatory exposure would feel neither pull as strongly, regardless of sharing the same industry label.

A private wealth advisory firm shows a heavier version of this same pattern. The client relationship and the fiduciary judgment push weight onto Steps One and Three, and the regulatory exposure around what counts

as advice, disclosure, and suitability pushes weight onto Step Four just as hard. That is three decisions loading at once and a firm operating in this space will feel more of the framework pulling hard simultaneously than most other businesses will.

None of this predicts the exact structure any organization lands on. What it predicts is which decisions are likely to carry disproportionate weight, given the actual characteristics of the work in each industry, and that is enough to tell a leadership team where to expect the hardest resistance before they've finished the work of deciding it.

Once these five foundations are in place, the integrated design work can begin. The agent portfolio decisions, how many agents, at what autonomy levels, built or bought or leased, and how they are governed, can be made with a clear reference point in the outcome architecture and clear boundaries in the governance design. The human configuration, who does what, in what relationship to which agents, at what level of accountability, can be designed around what the outcome architecture requires rather than what the legacy structure permits. The skills and roles that the hybrid organization needs can be built deliberately rather than discovered reactively. And the structure, the shape that so many organizations begin with, will emerge from all of these decisions compounding together, not as a choice but as a result.

*"The organizations that will define the next decade are not those that deployed the most agents or restructured the fastest. They are those that asked the right questions, made the covenant with their people before the anxiety set in, built the governance infrastructure before they needed it, and created the conditions from which the right structure could emerge rather than imposing a structure and hoping the rest would follow."*

## Sources

---

- [1] Delphi Group, “Leadership for the Agentic Age,” April 2026
- [2] Delphi Group, Leading Beyond Humans, 2026
- [3] MIT Technology Review Insights, “Rethinking organizational design in the age of agentic AI,” May 2026
- [4] McKinsey, State of Organizations, 2026
- [5] BCG, The Widening AI Value Gap: Build for the Future, 2025
- [6] PwC, No More Pyramids: Rethinking Your Workforce for the Agentic AI Era, 2026
- [7] Deloitte Insights, Reinventing Workforce Planning for an AI-Powered, Uncertain World, 2025
- [8] Microsoft, 2026 Work Trend Index Annual Report: Agents, Human Agency, and the Opportunity for Every Organization, May 2026
- [9] Deloitte, Global Human Capital Trends, 2025
- [10] McKinsey, “Rethink Management and Talent for Agentic AI,” People and Organization Blog, November 2025
- [11] McKinsey, “The Agentic Organization,” Quarterly Article, September 2025
- [12] Microsoft, Work Trend Index, April 2025
- [13] Gartner, “Gartner Says Organizations with Successful AI Initiatives Invest Up to Four Times More in Data and Analytics Foundations,” April 16, 2026
- [14] IBM Institute for Business Value, “The essential guide to agentic AI: How to compete in an autonomous future,” November 2025
- [15] Deloitte AI Institute, Agentic Enterprise 2028, 2025
- [16] KPMG, Agents of Change: New Organizational Roles in the Age of AI, December 2025
- [17] KPMG, Global AI Pulse Survey, 2026
- [18] Gartner, “Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027,” June 2025
- [19] IMDA, Model AI Governance Framework for Agentic AI, January 2026

## About the Authors

---

### Tom Koulopoulos

LEAD AUTHOR | FOUNDER DELPHI GROUP

Leading digital futurist, author of fifteen books, and founder of the Boston-based Delphi Group, a 30-year-old-think-tank advising Global 2000 companies and government agencies. He is also a columnist for Inc.com, an adjunct professor at Boston University, and has held leadership roles at Babson College Center for Business Innovation and the Dell Innovation Lab.

- ▶ LinkedIn: [www.linkedin.com/in/tkspeaks/](https://www.linkedin.com/in/tkspeaks/)
- ▶ [tk@delphigroup.com](mailto:tk@delphigroup.com)

### Vlasios Vlastos

CO-AUTHOR | FOUNDER DALIA LABS

Thought leader and founder of Dalia Labs, an agentic-first company specializing in the hybrid workforce. He holds an adjunct faculty appointment at Hult Ashridge Executive Education and has served in leadership roles with Fortune 500 companies and governments across international borders.

- ▶ LinkedIn: [www.linkedin.com/in/vlasios-vlastos/](https://www.linkedin.com/in/vlasios-vlastos/)
- ▶ [vlasios.vlastos@daliaai.com](mailto:vlasios.vlastos@daliaai.com)

