

The Ambidextrous Autonomous Organisation

Why AI Autonomy Requires a New Architecture of Decision, Service and Responsibility

Organisations have long been required to operate two distinct logics simultaneously: exploiting today's capabilities to deliver reliable performance, whilst exploring tomorrow's opportunities to secure future relevance. Artificial intelligence — and in particular the emergence of autonomous agents capable of independent judgement — changes the terms of this structural challenge in ways that neither classical ambidexterity research nor operating-model frameworks have fully addressed.

EXECUTIVE WHITE PAPER

ALLYALLEZ GMBH · 2026

FEUSISBERG, SWITZERLAND

Executive Summary: The Argument in Five Moves

This paper develops a single, connected argument across five interpretive steps. The first two steps draw on established management research and documented Gartner analysis. The third is a structural inference drawn from those sources. The fourth and fifth represent this paper's original organisational-design contribution, informed by AllyAllez's work in mandate and responsibility architecture. Readers are asked to treat these layers as distinct: the first two are evidential; the latter three are interpretive and propositional.

1

Ambidexterity Remains the Foundational Problem

Since March (1991) and O'Reilly and Tushman (2004), management research has established that organisations must simultaneously exploit existing capabilities and explore new ones — two logics requiring different processes, metrics and cultures, reconnected only at senior leadership level.

2

Autonomous Business Is Ambidexterity Applied to Operating Models

Gartner's 2026 analysis argues that the dual path of exploitation and exploration now applies specifically to the transition toward an autonomous operating model — where decisions are increasingly contributed by AI, agents and machines acting with greater independence under defined outcomes and guardrails.

3

Autonomy Changes Decision Contribution, Not Only Task Execution

Automation replaces tasks. Autonomy repositions decision contribution. The organisational question shifts from which activities can be automated to which decisions can be made by whom or what, under which conditions and within which boundaries. This is an organisational-design question, not solely a technology question.

4

Capability Can Migrate Faster Than Mandate

An agent's technical capacity to replan, reprioritise or act does not automatically constitute legitimate authorisation to do so. Decision boundaries — concerning cost, supplier relationships, service levels, contractual interpretation and regulatory risk — require explicit mandate architecture, not merely technical guardrails.

5

Service and Responsibility Architecture Is Required

When autonomous contributions become part of service delivery, the structural questions of who defines the service, who holds the mandate, who carries liability and where human judgement must still form become acute. The autonomous organisation is, at its core, a mandate and responsibility architecture problem.

 **Source note:** Gartner material is cited as published evidence of its 2026 position. The integration with ambidexterity theory and the responsibility-architecture extension represent this paper's interpretive organisational-design argument. AllyAllez is not endorsed by Gartner, and no such endorsement is implied.

Ambidexterity: The Persistent Organisational Problem

The organisational challenge of simultaneously running and reinventing a business has generated a substantial body of research. James March, in his 1991 paper in *Organization Science*, made the foundational analytical distinction between exploitation — the refinement and use of existing knowledge, capabilities and operating routines — and exploration — the pursuit of new knowledge, options and possibilities whose value is uncertain and distant. March's contribution was to show that these two activities are not merely different in character; they are, in important respects, in competition for the same finite organisational resources, attention and commitment. An organisation that invests heavily in exploiting what it already knows crowds out the slack and tolerance for failure that exploration requires.

Charles O'Reilly and Michael Tushman developed this into a structural organisational-design argument. Writing in the *Harvard Business Review* in 2004, they described the ambidextrous organisation as one that creates structurally distinct units for exploration — with their own processes, cultures, performance metrics and, crucially, separate leadership systems — whilst maintaining tight integration at the senior executive level to prevent the exploratory unit from drifting entirely free of the core organisation's resources and direction. Their core insight was that the managerial task is not to choose between exploitation and exploration, but to prevent each from destroying the conditions the other requires.

Exploitation Logic

- Efficiency, standardisation, predictability
- Scaling known approaches and capabilities
- Errors treated as defects to be eliminated
- Short-horizon performance metrics
- Serving existing customers reliably
- Optimising cost and quality at scale

Exploration Logic

- Learning, experimentation and tolerance of uncertainty
- Searching for fundamentally new approaches
- Errors treated as information and signal
- Long-horizon option value as the measure
- Identifying emerging or future needs
- Tolerating cost without immediate return

The recurring failure mode is well documented: organisations create an innovation unit and then judge it using the metrics, decision rules and cultural expectations of the existing business. Exploration is formally permitted but practically suffocated. The organisational-design question that ambidexterity ultimately poses is not simply whether both logics exist, but which decision, resource, leadership and performance logics each mode requires — and where they must be deliberately reconnected. What Gartner's 2026 autonomous-business work suggests is that this same question now applies to the architecture of work and decision-making itself: to who decides, how decisions are formed, where machines act independently and where human judgement intervenes.

Gartner's Autonomous-Business Thesis: The 2026 Position

In a series of publications and conference presentations during 2026, Gartner articulated an increasingly concrete position on what it terms autonomous business — defined as a strategy that uses self-improving, adaptable technology to make decisions, take action and create new types of value by increasing both people autonomy and machine autonomy. This definition is notable for what it does not say: it does not define autonomous business primarily as a technology adoption question. It frames it as a question about how decisions are made, who makes them and how value is created.

Gartner reports that in a survey of 469 global CEOs and senior business leaders conducted between March and November 2025, 8 in 10 respondents expected autonomous business to become the dominant form of business by 2030. This is a surveyed expectation, not a Gartner forecast or certainty. It nevertheless signals the strength of executive conviction about the direction of travel. Gartner's own analytical position is that organisations which merely use AI to improve existing processes will not achieve competitive advantage from AI adoption alone; the larger structural opportunity lies in redesigning operating and decision models.

Autonomous-Ready Operations

A shift from sequences of automated, siloed tasks toward networks of outcome-based decisions made autonomously or augmented by AI, informed by data and human judgement.

Autonomous-Ready Intelligence

A decision stack combining data, mapped decisions, guardrails, knowledge assets and human context — enabling AI agents to support business goals within defined risk tolerance.

Autonomous-Ready Workforce

Role evolution toward more versatile functions working alongside AI, agents and robots — shifting from task execution toward outcome orchestration and oversight.

Gartner also articulates the dual-path requirement with a precision that directly echoes ambidexterity logic. Chief Supply Chain Officers, Gartner argues, need two parallel paths: exploitation — using existing AI capabilities for efficiency and near-term results; and exploration — redesigning workflows and uncovering long-term opportunities enabled by autonomy. Gartner explicitly warns that focusing only on exploitation risks optimising yesterday's operating model for tomorrow's environment.

AstraZeneca described a self-healing supply-chain framework using AI-powered decision intelligence and conversational analytics, aimed at compressing the gap between problem detection and resolution. PepsiCo announced a multi-year partnership with Gatik to deploy autonomous freight in its North American supply chain. These are cited as conference-stage examples from 2026 symposia — not as validated outcome evidence.

The Hidden Shift: Autonomy Changes Decision Contribution

The distinction between automation and autonomy is analytically important and easily conflated. Automation, as a design principle, replaces or accelerates the execution of predefined tasks. A process is mapped; technology executes the mapped steps faster, more reliably or at greater scale. The organisation retains full authorship of the decision logic; the technology is an execution instrument. The organisational question is: which activity can be automated?


Autonomy represents a categorically different relationship. An autonomous system is given a defined outcome and is permitted to — within defined boundaries — sense a situation, evaluate options, decide on an action, act on that decision, observe the effect and adapt its behaviour accordingly. The organisation delegates not only execution but a portion of the decision-forming process itself. The organisational question shifts: which decision can be taken by whom or what, under which conditions, within which boundaries?

AUTOMATION



**LINEAR SEQUENCE:
HUMAN RULES DEFINE
A to B to C.
TECHNOLOGY EXECUTES.
FULL HUMAN
AUTHORSHIP.**

AUTONOMY



**ORGANISATION
DELEGATES DECISION-
DISTRIBUTED
CONTRIBUTION.
REQUIRES EXPLICIT
BOUNDARY DESIGN.**

Gartner's concept of a decision stack addresses this directly. Rather than designing only a technology stack, Gartner argues that autonomous systems require a decision stack: mapped decision flows, data inputs, guardrails, knowledge assets and human context and judgement. The decision stack is Gartner's answer to the question of how autonomous contributions can be made tractable and governable within an operating model. It acknowledges that AI agents, to support business goals effectively, must operate within defined risk tolerance, contextual knowledge and governance structures — not merely within technical parameters.

The significance of this shift for organisational design is considerable. In an automated process, the locus of decision contribution is clear: human designers defined the logic; technology executes it. When systems become autonomous — even partially, even in bounded domains — the locus of decision contribution becomes distributed and, unless explicitly mapped, potentially opaque. Formerly implicit assumptions about who may authorise a course of action, who may incur cost, who may alter a priority or who may accept a risk must be made explicit in the architecture. They can no longer be embedded in a human decision-maker's professional judgement alone.

From Ambidexterity to Decision Architecture: The Second-Order Problem

Classical organisational ambidexterity asks how an organisation can run today's business whilst simultaneously developing tomorrow's. Both sides of that question share a broadly stable assumption: human beings, in defined roles, make decisions within recognised authority structures, and the accountability chain — however imperfectly — runs from action to role to responsibility.

Gartner's autonomous-business thesis introduces what this paper proposes to call a second-order ambidexterity problem — and it is important to be precise: this phrase is this paper's interpretive extension, not an established research term or a Gartner construct. The first order of ambidexterity concerns what the organisation does. The second order concerns how and by whom decisions are contributed: within an existing human authority structure, or within an emerging human-machine decision architecture whose boundaries have not yet been explicitly designed.

Traditional Decision Flow

- 01
- Planner assesses inventory and demand
- 02
- Dispatcher assigns carrier and route
- 03
- Logistics manager approves exceptions
- 04
- Carrier confirms and executes delivery
- 05
- Escalation follows defined authority chain
- Decision contribution is sequential, role-bound and traceable through established authority.

Autonomous Contribution Pattern

- 01
- Agent senses disruption from data and forecast
- 02
- Agent evaluates options and selects carrier
- 03
- Agent triggers rebooking and cost commitment
- 04
- Exception is detected and flagged to human
- 05
- Human judgement activated at defined threshold
- Decision contribution is distributed, partially machine-originated, and requires explicit boundary design.

The illustration above is not a claim about any specific organisation's current system. It is an analytical comparison of two decision-contribution architectures, used to make the structural point visible. The service being delivered — procure, store, transport, deliver — may remain recognisably the same in both cases. What changes is the architecture of work and decision-making: who contributes to a decision, at what point, within what authorisation and with what escalation pathway. This is the organisational discontinuity that ambidexterity theory does not fully anticipate. An organisation exploring autonomous decision contributions is exploring how its internal authority and accountability structure works — inside a live service system that cannot be suspended whilst the experiment runs.

The Gap: Capability Can Migrate Faster Than Mandate

Capability does not automatically confer mandate.

In the transition toward autonomous operations, technical capability — what a system can do — is subject to rapid, iterative expansion. AI models improve. Agents are given broader access to data. Workflow integrations deepen. The boundary of what is technically possible moves continuously, often faster than any governance cycle is designed to track. Mandate — what a system or role is legitimately authorised to decide, commit or act upon — is structurally slower. It requires deliberate organisational design, explicit allocation of decision rights, definition of service scope and considered assignment of liability.

This gap between capability and mandate is not a technology failure. It is a structural consequence of the speed differential between two fundamentally different types of change. Technical capability expands through engineering cycles, model updates and integration work. Mandate expands — when it expands at all — through organisational design decisions, governance processes, contractual revision and, in regulated sectors, regulatory review. When autonomous systems begin to contribute to decisions, the alignment between capability and mandate can no longer be assumed; it must be designed.

Freight Replanning Capability

An agent may be technically capable of identifying a disruption, evaluating alternative carriers, selecting an option and triggering a rebooking within minutes. The mandate question is separate: is the agent authorised to incur the associated cost? Alter the agreed service level? Reprioritise one customer's shipment over another's? Accept a carrier outside the approved vendor list?

Supplier Relationship Decisions

The technical capability to compare supplier options and initiate a change does not confer the mandate to alter a supplier relationship that carries contractual, compliance or strategic dimensions. Those decisions have historically been reserved — often implicitly — within specific organisational roles and approval structures.

Regulatory and Contractual Interpretation

An autonomous system operating in a domain subject to regulation, professional standards or contractual commitments may encounter situations where the technically optimal action conflicts with a legally or contractually bounded obligation. The capability to act does not resolve that conflict; only explicit mandate design can do so.

AllyAllez's public position — *When responsibility grows faster than mandate* — names the same phenomenon from a different angle. As autonomous contributions become more visible, more traceable and more attributable, the question of who carries responsibility for outcomes intensifies. If the mandate architecture has not kept pace with the capability architecture, responsibility accumulates in ambiguous spaces: potentially personal, potentially diffuse, potentially contested. The structural purpose of explicit mandate design is to prevent that accumulation.

The AllyAllez Extension: Service and Responsibility Architecture

From Decision Stack to Mandate, Service and Liability

Gartner's decision-stack concept — combining data, mapped decisions, guardrails and human context and judgement — is a significant analytical contribution. It acknowledges that autonomous systems require more than technical parameters; they require contextual knowledge, governance structure and defined conditions for human intervention. This paper does not contest that framework. It proposes an additional structural layer that Gartner's supply-chain framing does not fully address: the explicit architecture of mandate, service, liability and judgement that must bound the decision stack when autonomous contributions form part of a live service delivery.

AllyAllez defines Organisational AI as the mandate and judgement architecture of service-based organisations under AI Visibility — the condition in which AI contributions, decisions, handovers and accountability become more traceable and therefore more attributable. In this condition, the question of who authorised what, and who carries the consequences of what, becomes structurally more acute. The response AllyAllez proposes is not a compliance overlay but a structural architecture in which responsibility is carried by the design.

The service foundation for this architecture is grounded in Paul G. Huppertz's Servicialisation logic. A service is defined in full as '**a bundle of once-consumable and ephemeral beneficial effects**' (Huppertz, 2024). The operational short form is: **Service = beneficial effect at the service object**. A service type is identified through three essential identifiers: the authorised service consumer; the service object belonging to that consumer; and the service-specific beneficial effect that must be brought about at that object. Rendering a service means that the service-specific beneficial effect attributed to the called service type is deliberately brought about at a specific service object in the service transaction. Critically, **AI is a contribution layer, not the service**. Humans, AI agents, systems and external providers may all contribute to a service without thereby becoming the service itself.



Mandate

Who may decide, approve, override, stop or escalate a decision or action — whether that decision is contributed by a human, an agent or a combined human-machine process.



Service

Which service type is being rendered; for which authorised service consumer; on which service object; and which service-specific beneficial effect must be deliberately brought about at that object in the service transaction.



Liability and Responsibility

Who defines and commits to the service materially shapes where accountability lands. Legal liability remains subject to the applicable contractual and legal framework. Who must explain and carry what happens at the moment of use — not theoretically, but in the specific context of the service delivery relationship.



Judgement

Under what conditions can responsible, service-bound human judgement still form — and where must it form, regardless of what autonomous capability could in principle substitute for it.

❏ **Important:** AllyAllez's responsibility architecture is an independent organisational-design contribution that takes Gartner's documented transition as its starting context and extends the structural analysis into mandate, service and liability dimensions. The two bodies of work are complementary, not hierarchical, and neither endorses the other. AllyAllez is not endorsed by Gartner, and no such endorsement is implied.

Seven Executive Design Questions

A Diagnostic Orientation for Senior Leaders Navigating Autonomous Transitions

The following questions are offered as a diagnostic orientation for senior leaders navigating the transition toward autonomous operations. They are not a compliance checklist, a maturity assessment or a product framework. They represent the class of structural questions that must be made explicit — and answered deliberately — before autonomous contributions migrate from exploration into the operating model. The purpose of making these questions explicit is to surface the decisions that are often left implicit in autonomous-business transitions: decisions about authority, service scope, escalation design and who carries accountability when the system acts. Leaving these implicit does not make them disappear; it means they are resolved, if at all, after the fact — reactively, expensively and often personally rather than structurally.

1 Service Boundary and Reliability

Which existing service must remain reliable whilst autonomous capability is explored alongside it? What promises have been made to the service consumer — grounded in the specific beneficial effect that must be brought about at the service object — that must be honoured regardless of the operating model under exploration?

2 Decision Type and Mode

Which decisions in this domain are being optimised within the existing model — and which are being structurally redesigned? Are these two activities being held separately, with distinct metrics and oversight, or are they conflated in a single governance layer?

3 Outcome Authorisation

What outcome is the autonomous contribution permitted to optimise? What cost, risk, priority change or contractual implication may it generate — and under whose mandate does that generation sit?

4 Reserved Decision Rights

Which decisions remain reserved, reviewable, overridable or stoppable by humans — not only technically but organisationally, with explicit role assignment and a functioning override pathway?

5 Escalation and Judgement Thresholds

What event, condition or decision type triggers escalation to human judgement? Is that threshold defined in the architecture, or does it emerge ad hoc from individual interpretation at the moment of pressure?

6 Accountability at the Moment of Use

Who must explain the outcome to the service consumer at the moment of use — when something goes wrong, when a priority has been altered, when a cost has been incurred? Is that person or function structurally equipped to do so, or does accountability sit in a gap?

7 Migration Criteria

What evidence is required before an exploratory decision pattern — tested in controlled or bounded conditions — is permitted to migrate into the live operating model? Who makes that determination, under whose mandate and what structural review is required?

These questions are not asked once at the outset of an AI programme. They recur as capability expands, as agents are given access to new data, as integration deepens and as the boundary of what is technically possible continues to move. The structural architecture that answers them must therefore be a living design, not a one-time governance exercise.

Conclusion, References and Author Note

Ambidexterity describes the organisational problem. Autonomous Business describes a new manifestation of that problem. Responsibility architecture determines whether that manifestation can become a viable service system.

The autonomous era does not dissolve the organisational ambidexterity problem. It makes its consequences more structurally demanding, more visible and more consequential. Leaders must exploit AI within today's operating model — delivering near-term efficiency, reliability and value — whilst simultaneously exploring what a fundamentally different operating model looks like: one in which humans, agents and machines jointly contribute to decisions, with greater independence and within defined outcomes. This dual requirement is well-described in the management literature and clearly articulated in Gartner's 2026 autonomous-business work.

As autonomous contributions migrate from experiment toward operating model, technical capability moves faster than the organisational structures that confer mandate, define service boundaries, assign liability and preserve the conditions for responsible human judgement. The organisations best positioned to capture the value of autonomous operations will not necessarily be those with the most autonomous systems; they will be those able to integrate autonomous contributions without losing service clarity, accountability and the structural conditions under which responsible judgement can still form. The practical design priority is therefore clear: before asking what an agent can decide, make visible which service is being delivered, to which authorised service consumer, with what beneficial effect at which service object, under whose mandate, with what escalation and with who carrying the consequences at the moment of use.

Structure First. Not AI First.

References

Foundational Research

1. March, J.G. (1991). 'Exploration and Exploitation in Organizational Learning'. *Organization Science*, Vol. 2, No. 1, pp. 71–87. DOI: 10.1287/orsc.2.1.71.
2. O'Reilly, C.A. III and Tushman, M.L. (2004). 'The Ambidextrous Organization'. *Harvard Business Review*, April 2004.
3. Gibson, C. and Birkinshaw, J. (2004). 'Building Ambidexterity into an Organization'. *MIT Sloan Management Review*, July 2004.

Gartner Sources (2026)

4. Gartner Newsroom (4 May 2026). 'Gartner Highlights Three Building Blocks for Autonomous Supply Chain Future'.
5. Gartner (4 May 2026). Supply Chain Symposium/Xpo Day 1 Highlights.
6. Gartner (26 August 2026). 'How to Build an Autonomous-Ready Supply Chain'.
7. Gartner (1 May 2026). 'Build Your Supply Chain Decision Stack to Enable Agentic AI'.
8. Gartner (4 May 2026). 'How to Lead Supply Chains Into the Autonomous Era'.
9. Gartner Newsroom (18 May 2026). 'Gartner Says Chief Supply Chain Officers Must Shift to an Autonomous Business Mindset'.
10. Gartner ThinkCast (20 August 2026). 'The Autonomous Era Is Here: What Supply Chain Leaders Need to Do Next'.

Illustrative Conference and Industry Examples

11. Gartner conference session (5 May 2026). AstraZeneca: 'Self-Healing Supply Chain: AI-Driven Resilience for Smarter, Faster Decisions'.
12. PepsiCo press release (8 June 2026). PepsiCo and Gatik multi-year autonomous freight agreement.

AllyAllez Public Sources

13. AllyAllez homepage. www.allyallez.com
14. AllyAllez About page. www.allyallez.com/about
15. AllyAllez Insights. www.allyallez.com/insights

Service and Servicialisation Foundation

16. Huppertz, Paul G. (2024). 'Service-Identifizierung – Service-Typ & Wesensmerkmale'. *servicEvolution*, Leitkonzept der Servicialisierung, 24 December 2024.
17. Huppertz, Paul G. (2024). 'Service-Konzept – Service-Beitragstypen & Service-Drehbuch'. *servicEvolution*, Leitkonzept der Servicialisierung, 24 December 2024.

Author and AI-Enabled Writing Note

This paper was developed by AllyAllez GmbH with AI-enabled writing support for drafting, editorial structuring and language refinement. The conceptual framing, source selection, interpretation and final editorial responsibility remain with AllyAllez GmbH. AI-generated text is not presented as independent evidence or source material.