Redesigning Operating Models for Agentic AI at Scale

Point of View

ALVAREZ & MARSAL

LEADERSHIP. ACTION. RESULTS.℠

From Automation to Autonomy

Over the past couple of years, enterprises have embraced AI sporadically, from analytics-driven insights to machine learning prediction engines and, more recently, generative AI copilots. But value realization in enterprise context has been underwhelming. Despite boosting some productivity, they operated strictly within human-controlled workflows and functional silos.

Agentic AI breaks this pattern. Agentic systems interpret goals, reason through ambiguity, select tools, execute actions, learn from outcomes, and adapt dynamically. Rather than following fixed scripts, they rely on iterative planning loops.

This capability allows enterprises to shift from static, rule-based processes to dynamic, outcome-driven ones. However, traditional operating models—structured around hierarchies, functional silos, and layers of manual approval—were never designed for autonomous decision-making. This creates a widening gap between what Agentic AI makes possible and how enterprises are currently built to function.

Due to this misalignment, the challenge is not limited to IT architecture; it is organizational as much as technical. Governance, accountability, workforce models, risk frameworks, and performance management systems must be rethought as AI begins to make, or materially influence, decisions. Organizations must now ask not only “Where can AI automate?” but “Which decisions can we delegate?” and “How do we govern algorithmic authority responsibly?”

Impact on the Business Operating Model

Agentic AI reshapes the business operating model across six structural dimensions: work design, decision rights, workforce composition, digital workers, performance management, and governance.

Impact Across Six Structural Dimensions

1 Work Shifts From Static Processes to Dynamic, Outcome-Driven Execution

Traditional organizations have long optimized workflows where success is usually measured by SLA compliance, throughput, and cost efficiency. **Agentic systems optimize work around outcomes.** Agents sense signals, dynamically recompose tasks, and adapt execution in real time—without waiting for human-triggered steps. The process map no longer constrains execution; it becomes a baseline.

Organizations should define the outcomes they are optimizing for before deploying agents, not after. They should deconstruct workflows to isolate judgment-intensive tasks from those that can be autonomously executed and retire SLA-based metrics that reward compliance over outcomes.

At an American insurance company, claims processing was redesigned around a single outcome: **instant claim resolution**, rather than completion of a predefined workflow. When a customer submits a claim, AI agents operate in parallel to extract data, validate policy coverage, assess fraud risk, and benchmark against historical patterns. Based on real-time confidence thresholds, the system dynamically determines the execution path—approving and paying claims instantly, requesting additional information, or escalating to a human reviewer. There is no fixed sequence of steps; the workflow is continuously recomposed based on outcome confidence and risk context. Humans are positioned above the flow, intervening only when uncertainty exceeds defined thresholds.
The company reported more than **50% of claims resolved autonomously end-to-end**, with **settlement times collapsing from days to seconds for eligible cases.**

2 Redistribution of Decision Rights

Deploying agents is, fundamentally, the delegation of operational authority to non-human actors. Without explicit design, the boundary between what agents should do and what they actually do inevitably drifts.

Enterprises must formalize autonomy tiers with clear accountability matrices, especially in regulated sectors.

- **Tier 1 –** Advisory: AI recommends; humans decide.
- **Tier 2 –** Guard-railed autonomy: AI executes within policy thresholds; humans set rules and handle exceptions.
- **Tier 3 –** High autonomy with oversight: AI decides and acts; humans monitor and retain accountability.

Expansion across tiers must be a conscious and documented organizational decision, not a byproduct of pilots. **Business and technology leaders must co-own this boundary.**

To manage the redistribution of decision rights, a pharmaceutical company appointed a **Chief People and Digital Technology Officer,** a role created by merging the HR and technology functions. This enabled the organization to move from “workforce planning” to “work planning.” That is, instead of allocating roles by headcount alone, managers assess tasks based on whether they should be automated, augmented, or performed by a person. With more than 3,000 deployed agents, this is a decision rights model operationalized at scale: a live design discipline embedded into how the company plans and allocates work.

3 Human Roles Transform Structurally, Not Incrementally

As execution shifts to agents, human roles pivot toward design and governance. New capabilities include AI supervision, policy design, exception adjudication, and context engineering—curating the data, tools, and guardrails agents rely on to make grounded decisions. Leaders must invest in reskilling managers to supervise intelligent systems rather than manage tasks.

Enterprises should organize around **persistent, cross-functional teams aligned to customer journeys or value streams**. Each team must be accountable for the end-to-end performance of both human and AI-agent execution within its domain, integrating business ownership, technical capability, governance expertise, and economic accountability.

When a fintech company deployed AI agents to replace large portions of customer service, early gains in efficiency masked a structural flaw: while routine queries were resolved autonomously, performance degraded on complex, multi-step, and emotionally nuanced cases, exposing that the human role had been removed, not redesigned. The company subsequently shifted to operating model redesign. It reorganized customer operations into **cross-functional, domain-aligned teams,** each accountable for end-to-end performance of AI-driven customer journeys, not ticket volumes. Within these teams, human roles were redefined around exception handling, escalation, and continuous system optimization—analyzing failure patterns, refining response logic, and tuning decision thresholds.

The unit of work moved from resolving individual queries to **owning and improving the performance of an autonomous service system**, demonstrating that as AI absorbs execution, human work becomes more complex.

4 AI Agents Must Be Governed as Enterprise Actors

An agent that can query systems, execute transactions, or modify records is not a tool; it is an enterprise actor with operational authority. Every deployed agent must have a unique identity, a designated human sponsor accountable for its lifecycle, and a documented authorization scope. This is an enterprise design mandate.

Organizations need a central **AI governance function**—a senior, cross-functional authority that enforces identity standards, boundary conditions, and observability across teams, and flags when local rational decisions create enterprise-level risk no single team can see.

It should report to the CEO, COO, or CRO, independent from the teams building the agents, and bring together risk, compliance, legal, AI engineering, security, and leaders from deploying domains. Its remit should be to own the enterprise agent registry, set validation and observability standards, approve autonomy tier expansions, and lead incident response.

At a leading global bank, the scaling of AI across risk, compliance, and operations required the creation of **centralized model risk management and AI governance** functions over otherwise decentralized deployments. Rather than allowing individual teams to independently build and deploy AI systems, the firm enforced common standards for model validation, usage boundaries, auditability, and lifecycle management through centralized risk and AI governance functions.

These functions operate with institutional authority to define what AI systems are allowed to do, how they are monitored, and when they must be escalated to humans or decommissioned, ensuring consistency across business units. Every system operates within a defined scope, with clear ownership and traceability of decisions, preventing fragmentation that would otherwise emerge as hundreds of AI use-cases scale independently.

The role of this central function is not to slow deployment, but to maintain consistency across the enterprise. This ensures that locally optimized AI systems do not create systemic risk, regulatory exposure, or conflicting behaviours at the enterprise level.

5 Performance Management Must Extend To Autonomous Systems

Traditionally, performance systems evaluated only humans. They cannot explain why an agent made a decision, detect early signs of drift, or assess whether autonomous execution is generating real economic value.

With agentic AI, performance management must expand to include agents, enabling structured oversight of their lifecycle. New metrics must track autonomous decision rates, intervention frequency, outcome variance, decision latency, and risk-adjusted automation levels.

Organizations should **design observability into the system from the outset**. Define what "good" autonomous execution looks like before deployment, and instrument continuous assurance mechanisms to track decision quality and variance, intervention frequency, decision latency, and cost per decision relative to value. These metrics must sit with domain owners with economic accountability.

A global bank established a **dedicated AI monitoring function** reporting into senior risk leadership, tracking model drift, output quality, and intervention rates in real time. Performance management is treated as a **risk discipline**—and not an engineering one—embedded into operations rather than reviewed periodically.

6 Governance Shifts From Periodic Checkpoints To Embedded, Continuous Control

Traditionally, governance is either an approval model or a periodic retrospective. To be able to realize their full value, agentic systems should not be constrained by human approvals. Further, at machine speed, the risk is too high for only periodic reviews.

To effectively govern agentic systems, organizations should **embed governance upstream in execution:**

- Codify regulatory, ethical, and risk constraints before deployment
- Integrate real-time monitoring for threshold breaches and anomalies
- Define clear liability frameworks for AI-driven decisions

At a leading digital payments company, **risk and compliance are embedded directly into real-time transaction processing,** with decisions dynamically evaluated against fraud, credit, and regulatory models at the point of execution. While these are not fully agentic systems, they illustrate the operating model shift required: governance is not applied after decisions are made but codified into the decision logic itself. As organizations move toward agentic systems, where execution paths are less deterministic and decisions are more autonomous, this principle becomes more critical. Governance must be embedded upstream in how systems operate, not layered on through human checkpoints downstream.

Redesigning the Operating Model Unlocks Additional AI Value

A significant share of AI value creation depends on how an organization redesigns its operating model—reimagines processes, reorganizes its structure, enables governance, and manages workers.

Where Value Is Lost

Operating Model

- Data Foundations
- Technology & Integration
- Talent & Skills
- Transformation & Value Management

01 **Agents bolted onto legacy workflows inherit their bottlenecks,** as execution still waits on human decisions. Redesigning the process around agent autonomy can unlock up to three times the gains.

02 **Amplified risk due to governance gaps.** Without defined accountability, embedded governance, and clear human-agent boundaries, risk outpaces value and pilots never reach production. Delay compounds the cost.

03 **Functional silos prevent end-to-end optimization.** Traditional functional structures optimize departmental performance. Agents inherit those boundaries and the cross-domain value stays trapped.

Conclusion

The autonomous enterprise, human and digital agents collaborate seamlessly. Workflows self-optimize. Decisions are data-driven and near real time. Organizational hierarchies flatten as coordination becomes algorithmic.

Autonomous execution represents a structural shift in how organizations operate and demands a redesign of the enterprise operating model. Leadership roles evolve as well: the CIO becomes a Chief Autonomy Architect, the COO orchestrates human–AI ecosystems, and risk leaders oversee continuous algorithmic compliance. One of the largest Aerospace companies in the US already has a Chief AI & Autonomy Officer driving the Autonomy Capability for the firm.

Enterprises that proactively redesign around goal-based orchestration and guard railed autonomy can unlock significant gains. They cut decision latency materially, enhance resilience, and unlock significantly higher benefits.

The enterprises that win will be those that orchestrate autonomous systems with clear human purpose and accountability. Organizations that invest in governance architecture and operating model alignment today will scale AI with confidence. Those that continue to add agentic capability to unreformed operating models will compound the structural drag they are already experiencing.

The **central leadership question** is no longer whether to adopt Agentic AI, but **how to redesign the enterprise to operate responsibly and competitively with autonomous intelligence at scale.**

Disclaimer: The information and data contained in this document is of a general nature and is a triangulation of data obtained from publicly available information combined with A&M's own experience. Please contact A&M directly for further details on the specific reports used and our methodology.

AUTHORS

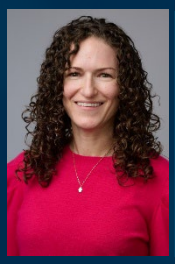

MALKA KATZIN
MANAGING DIRECTOR
Chief AI & Knowledge Officer

mkatzin@alvarezandmarsal.com

SANJEEV NARSIPUR
MANAGING DIRECTOR
Digital, AI & Technology Services

snarsipur@alvarezandmarsal.com

SONAL PRUTHI
MANAGING DIRECTOR
Performance Improvement

spruthi@alvarezandmarsal.com

SUDEEP SHARMA
SENIOR DIRECTOR
Technology Strategy & Transformation

sudeepsharma@alvarezandmarsal.com

CHAITANNYA GOEL
ASSOCIATE DIRECTOR
Business Case & Value COE

c.goel@alvarezandmarsal.com

Follow A&M on:

ABOUT ALVAREZ & MARSAL

Founded in 1983, Alvarez & Marsal is a leading global professional services firm. Renowned for its leadership, action and results, Alvarez & Marsal provides advisory, business performance improvement and turnaround management services, delivering practical solutions to address clients' unique challenges. With a world-wide network of experienced operators, world-class consultants, former regulators and industry authorities, Alvarez & Marsal helps corporates, boards, private equity firms, law firms and government agencies drive transformation, mitigate risk and unlock value at every stage of growth.

To learn more, visit: AlvarezandMarsal.com