CEIBS

NAVIGATING AI-POWERED BUSINESS TRANSFORMATION

2026 WHITE PAPER

UNLEASHING AGENTIC AI: THE STRATEGIC LEAP FROM TASK-LEVEL EFFICIENCY TO SYSTEMIC INTELLIGENCE

中欧国际工商学院 × 特赞
人工智能与商业创新研究基金
CEIBS × Tezign
Generative AI and Business Innovation Initiative

Tezign | 特赞

中国深度 全球广度

CHINA DEPTH GLOBAL BREADTH

Contents

Preface

Every technological revolution reshapes the business landscape—not just by introducing more powerful tools, but by transforming how organisations manage knowledge, delegate authority, and assign responsibility. Today, business AI stands at a critical turning point. Once confined to following instructions and aiding specific tasks, it is now being deeply embedded into workflows, involved in decision-making, and in some cases, executing autonomously. In effect, AI is evolving from a tool into an integral part of a company's operating system. As a result, the focus of corporate competition is subtly shifting: it is no longer about whether AI has been adopted, but whether it can be seamlessly integrated into management systems in a way that consistently delivers verifiable, replicable, and sustainable value.

Managers today are navigating more than a technological upgrade; they are steering a systemic transformation that reshapes strategy, organisation, and governance. Over the past two years, companies have made steady progress applying AI across functions—from content generation and customer service to R&D support and market insights. Yet success in the AI era will depend less on isolated efficiency gains and more on an organisation's ability to do three things: deploy AI agents that operate reliably within clear boundaries; weave data, domain expertise, and feedback into a coherent system for sound decision-making; and strike a new balance between authorisation, risk control, and collaboration. As AI shifts from offering answers to executing tasks, humans will increasingly focus on setting direction, defining the guardrails for AI, and taking ultimate responsibility for the outcomes.

Therefore, the competitive advantage in the AI era will be determined not by whether an organisation can access advanced models, but by how well it can integrate them. Success will hinge on the ability to move quickly to re-engineer processes, restructure for agility, and modernise governance—and to move beyond isolated experiments to establish end-to-end, closed-loop systems that convert available technology into reliable, day-to-day operations. In a world where computational power and AI models are increasingly accessible, the scarcest resources are no longer the tools themselves, but the organisational capacity to embed them, scale them, and turn them into durable competence.

AI is not developing within a single global model, but across distinct regional paradigms. The European emphasis on regulation and ethics, the Chinese focus on scale and rapid implementation, and the American drive for technological leadership are shaping a multipolar AI landscape. This intricate AI landscape makes it increasingly vital to connect and blend these diverse viewpoints—an essential capability for any organisation operating globally.

As a business school that is rooted in China, deeply engaged in Europe, and globally connected, CEIBS offers a distinctive platform where bridging regional perspectives creates a unique vantage point for exploring how business works at its core. In particular, CEIBS focuses on understanding how companies must evolve as AI shifts from a supportive tool to a proactive partner. This requires a fundamental rethinking of efficiency, the redesign of jobs, and the redefinition of the boundaries between humans and systems.

Beyond the allure of the "next big thing" in technology, the role of a business school is to equip leaders with the frameworks to translate AI concepts into actionable management solutions, scale isolated pilot projects into organisation-wide transformation, and convert short-term efficiency gains into sustained, long-term capability. This focus on the essence of business—through impactful research and dialogue—is both timely and critically important today, as technological frontiers continue to expand, and the global industrial landscape undergoes profound restructuring.

This white paper unfolds from that foundational perspective. It moves beyond treating AI as a set of tools and instead examines it within the broader context of a company's strategic goals, adoption maturity, and organisational evolution. In doing so, it provides leaders with a practical framework for assessing their current stage, clarifying their strategic direction, and mapping a credible path forward. Central to this inquiry is a critical management question that transcends technology: as companies shift from human-centric to AI-enabled efficiency, and from isolated experiments to organisation-wide collaboration, how should leaders reassess the value of human roles, redefine the boundaries of AI autonomy, and redesign governance systems accordingly? This paper not only offers practical approaches and actionable pathways, but also aims to strengthen leaders' capacity for sound judgment—and the resolve to lead through this transformation.

The true promise of AI in business lies not in simply replacing human roles, but in fundamentally reshaping, revitalising, and redefining how organisations function. Success will belong to those who can integrate algorithmic power with human judgment, balance efficiency with accountability, and translate isolated innovations into scalable, system-wide capabilities—thereby gaining a decisive edge in the next phase of competition. This white paper provides a practical framework to help leaders cut through the noise, focus on what truly matters, and navigate AI-driven transformation with clarity and confidence.

Professor Wang Hong
CEIBS President

Professor Frank Bournois
CEIBS Co-President (European)

Preface

Over the past three years, artificial intelligence (AI) has undergone a profound transformation.

In 2023, generative AI models entered the public spotlight, rapidly reshaping expectations of what software could achieve. In 2024, enterprises began embedding these models into specific business scenarios, creating the first wave of use cases. Since 2025, a more complex question has come into focus: now that more organisations are using AI, why is stable, replicable and sustainable business value still so difficult to achieve—often remaining at the level of localised efficiency gains rather than advancing towards organisation-level value creation?

This question points to a deeper shift: from expanding local capabilities to building system-level execution. For enterprises, the issue is no longer whether to use AI, but how to make it part of everyday work. As AI moves from generating content to offering insight—and ultimately to executing tasks—it is shifting from a support tool to part of the execution structure itself, shaping how work is carried out across the organisation. This is not just a technical upgrade. It requires changes to system interfaces, context structures, and organisational coordination, so that AI works within existing processes rather than sitting outside them as yet another tool. The ability to make this shift is becoming a key dividing line in enterprise AI transformation.

This shift at the enterprise level is mirrored by a broader change in industry competition. As models become more accessible and increasingly standardised, advantage no longer comes from access to better models alone. Instead, it lies in the strength of context systems, the quality of feedback loops, and the ability to integrate AI into real business processes. The real differentiator is whether an organisation can build systems that run continuously around these models—allowing agents to execute tasks reliably in real-world business environments and translate that execution into consistent, self-reinforcing business outcomes.

These changes are also reshaping organisational structure. Work is increasingly carried out by “one individual supported by a set of agents”. Teams are becoming smaller, while system capability expands. The number of roles may fall, but their variety is increasing. New roles are emerging around context system design, model behaviour calibration, and closed-loop task orchestration. As a result, enterprise demand is shifting—from buying tools to building systems, and from local efficiency gains to redesigning how work is carried out across the organisation.

As this transformation deepens, the scarce capability is no longer production speed, but judgement. The key difference lies not in how much is produced, but in how well judgements are made. Evaluation frameworks, expert demonstration data, and system design capabilities are therefore becoming essential for AI systems to operate effectively in real business contexts. Competition in AI is shifting—from technological capability to organisational capability. Ultimately, success depends not just on deploying models, but on building the structural environment in which they can operate reliably.

Against this backdrop, enterprises are adopting AI in markedly different ways. Some remain focused on incremental, localised efficiency gains. Others are redesigning business processes around agent-based systems. A smaller group is exploring operating models in which agents directly support organisation-level value creation. These differences show that AI is evolving from a set of tools into an organisational capability—one that is reshaping strategy, organisations, and governance.

In this context, isolated case studies or purely technical perspectives are no longer sufficient to explain the real changes taking place within organisations. The challenge is not whether to adopt AI, but how to understand the structural differences between stages, and how to identify sustainable directions for progress amid uncertainty.

It is in response to this challenge that Tezign has developed its approach. Drawing on years of enterprise AI deployment, and on ongoing exploration of how agents participate in organisational processes, Tezign has developed the Generative Enterprise Agent (GEA) architecture. This provides a practical framework for embedding proactive intelligence within enterprise systems. By bringing together intent recognition, task orchestration, context systems, and execution in a single operating structure, GEA enables AI to reason, decide, and act in line with business goals—while running continuously within real-world workflows.

Building on this foundation, the present research offers a broader, systematic perspective. Over the past year, Tezign has conducted cross-industry case research and ongoing field observation, in collaboration with China Europe International Business School (CEIBS) under the CEIBS x Tezign Generative AI & Business Innovation Initiative. Based on this work, the study identifies key inflection points where localised efficiency gains give way to system-level execution and organisation-level value creation. Building on the 2025 “3×3 AI Strategy Matrix”, this report introduces a framework centred on the shift from task-level human efficiency to system-level execution capability, describing how organisations change as AI becomes embedded in operational structures.

Against this analytical backdrop, the report sets out its central focus. Navigating AI-Powered Business Transformation (2026) examines how AI evolves from a support tool into part of the execution structure. It also explores how this shift reshapes the division of labour, knowledge accumulation mechanisms, and the logic of value creation. We hope this report provides organisations with a clear framework to assess their current position, identify the next stage of evolution, and understand the conditions required to move from localised efficiency gains to system-level execution and organisation-level value creation.

Fan Ling
Founder and CEO, Tezign
Professor and PhD Supervisor, Tongji University; Director, Design AI Lab, Tongji University
Chair, CEIBS x Tezign Generative AI & Business Innovation Initiative

Introduction: From Localised Efficiency Gains to System-Level Execution

If 2024 marked the moment when generative AI entered the enterprise agenda and reshaped managerial imagination, then 2025 can be seen as the year in which enterprise agents began to move into practical use. Over the past year, large models have continued to advance—particularly in reasoning, planning, tool use, and multi-step execution. As a result, expectations have shifted rapidly, from viewing AI as a content-generation tool to recognising it as an agent capable of taking on tasks. In parallel, the way organisations use AI has also changed—from treating it as a copilot for employees to exploring how it can take over parts of task chains in specific scenarios.

Beneath the momentum, however, the key question is not whether organisations have begun to use AI, but whether they are generating sustainable, replicable, and verifiable business value from it. Over the past two years, one of the most common patterns across enterprises has been an abundance of demonstrations, pilots, and proof-of-concept (PoC) initiatives—but relatively few that truly cross departmental boundaries, enter core workflows, and deliver stable returns.

Many PoCs remain confined to experimental settings, limited in scope and largely focused on individual assistance and localised efficiency gains. Some projects impress in demonstration, but once deployed in real business environments, they expose deeper structural issues—disconnected processes, unclear accountability, missing feedback loops, and limited scalability. The challenge is therefore no longer whether an organisation can build an AI application, but whether it can embed AI into business systems in a way that allows it to deliver results continuously.

This is the context in which enterprise AI enters a new phase in 2026. At one end are localised efficiency gains, where individuals use AI tools to improve performance within existing workflows. These gains are real, but remain limited in scope—largely confined to individual assistance and incremental improvements in efficiency.

At the other end is system-level execution: the ability to embed agents in business processes and enable them to take on defined tasks within established boundaries. The focus is not simply on improving individual output, but on building reliable execution capability across the organisation.

The distinction, then, is no longer whether an individual can do more with AI, but whether the organisation has built an operating model in which execution itself is redesigned. In such systems, agents participate directly in execution, processes are continuously optimised through feedback, and outcomes are captured, fed back into the system, and embedded in organisational routines. The former optimises individual actions; the latter redefines how the organisation executes and creates value.

The key dividing line in 2026 therefore lies not in model parameters, isolated features, or the emergence of new breakout tools, but in whether organisations have shifted their evaluation criteria. In the past, AI projects were often assessed based on output quality, response speed, and local cost reduction. Today, more fundamental questions matter: can an agent identify problems, mobilise resources, execute tasks, correct deviations, and deliver outcomes that meet business requirements? Can it move from building a feature to taking over a workflow segment? From supporting a role to fulfilling a role? From isolated experimentation to system-level coordination? The focus of competition is shifting—from access to models to the ability to build closed-loop business systems.

In this sense, 2025 can be seen as the first year of agents, while 2026 marks the shift from localised efficiency gains to system-level execution. The central question is no longer whether to adopt AI, but how to convert demonstration value into sustained operational value—moving beyond isolated improvements to outcomes that scale across the organisation. This requires a shift in managerial mindset: AI should no longer be treated as software procurement, an innovation experiment, or an IT enablement project, but as a system-level initiative spanning strategy, technology, and business operations. Without alignment across these layers, system-level execution cannot produce organisation-wide value.

At the same time, the rise of agents is reshaping both data logic and competitive logic. Traditionally, enterprises invested heavily in data governance before moving on to model training and deployment. This approach suited a small number of large enterprises with substantial resources and mature legacy data governance. However, for organisations with fragmented, weak, or highly unstructured data, a new opportunity is emerging: using agents to structure data while running the business. As tasks are executed, data is generated, structured, and fed back into systems, forming a real-time loop of business action, data feedback, and model optimisation. In this model, AI is no longer simply layered on top of existing data assets, but becomes an engine for generating new data, accumulating knowledge, and sustaining continuous improvement.

This also explains why, by 2026, the truly scarce capability is no longer the model itself, but the judgement framework an organisation builds around it. As models become more standardised and widely accessible, competitive advantage depends on whether an organisation can integrate proprietary data, industry know-how, and feedback mechanisms into a system that evolves continuously. Agent proactivity does not come from the model alone, but from the interaction between context, constraints, process interfaces, and feedback loops. Without these elements, even powerful agents remain surface-level tools; with them, agents can become genuine digital employees embedded within enterprise operations.

Looking back at enterprise practice in 2025, a clearer pattern has emerged: AI commercialisation is shifting from solving isolated problems to system redesign. Many organisations have achieved cost savings and efficiency gains in areas such as marketing, customer service, content, insights, retail, supply chains, and quality control. Some have also demonstrated growth potential in specific scenarios. Yet large-scale deployment remains difficult—precisely because of structural misalignment between legacy processes and new capabilities. Embedding AI into existing workflows typically yields only incremental improvements. To unlock its full potential, organisations must redesign roles, responsibilities, and processes across both operations and organisational structures.

This white paper therefore argues that the core task for enterprise AI strategy in 2026 is no longer to deploy more tools, but to move beyond isolated efficiency gains and build system-level execution that delivers sustained, organisation-wide value. The goal is not incremental improvements in productivity, but a structural shift in how value is created, how organisations coordinate, and where competitive advantage lies. The gap between leading and following organisations will increasingly be defined by three factors: who can establish trusted, closed-loop agent systems earlier, who can integrate agents into formal organisational structures faster, and who can more effectively translate AI capabilities into sustained drivers of growth, innovation, and organisational transformation.

From this perspective, system-level execution is not merely a technical matter, but also a business, organisational, and strategic one. It requires managers to rethink efficiency, redefine roles, and redesign processes. More fundamentally, it calls for a shift in mindset: the value of AI lies not in how closely it resembles human intelligence, but in whether it can reliably take on judgement, execution, and coordination tasks within clearly defined boundaries—tasks that previously required sustained human involvement. Organisations that move first will be better positioned in the next phase of competition—evolving from users of AI into organisations reshaped by AI.

1

PARADIGM SHIFT— DEFINING “MORE PROACTIVE AI”

By the spring of 2026, AI was no longer seen simply as a new technology. Business perceptions of AI had shifted from early enthusiasm to sober reassessment, and ultimately to a deeper redefinition of its role in enterprise operations. If 2023 was defined by the breakout moment of large language models (LLMs), and 2024 by widespread experimentation, then 2025 became a period of uncertainty and strategic disorientation. Despite significant investment—much of it driven by a form of technological romanticism—many organisations found that AI initiatives had yet to deliver the expected returns. In particular, many pilot projects had failed to bridge the gap between localised efficiency gains and organisation-wide value creation—that is, the ability to use AI to execute reliably across real workflows and generate measurable business outcomes at scale.

A 2025 McKinsey study highlights this gap. While 88% of enterprises reported deploying AI in at least one business function, only around 5–7% had successfully scaled its use across processes and generated meaningful contributions to earnings before interest and taxes (EBIT)[1]. In effect, most organisations had powerful technological tools in hand, but had yet to find a way to turn them into engines of growth.

Yet it was precisely during this period of uncertainty and constraint that a new paradigm began to emerge. Discussions from a series of closed-door seminars under the CEIBS x Tezign Generative AI & Business Innovation Initiative indicated that leading organisations—from global consumer goods companies to technology firms such as Tezign—were recalibrating their AI strategies. AI was no longer seen merely as a process enhancer or a “copilot” waiting for instructions. Instead, it was increasingly viewed as a “digital employee” capable, within defined boundaries, of proactively taking on tasks and exercising a degree of independent judgement.

In this context, 2026 marks a deeper shift towards what we describe as system-level execution: the operating capability required to turn AI from a source of localised efficiency gains into a driver of organisation-wide value creation. As a result, the way enterprises assess AI is evolving. The focus is moving beyond enabling individuals to work faster towards building systems that can execute tasks reliably within real-world workflows. At its core, system-level execution refers to the ability of agents embedded in complex processes to operate independently, accurately, and within defined boundaries, generating value through closed-loop interactions. The move from reactive to proactive AI therefore signals a new paradigm in which AI becomes an active participant in how work is carried out.

Capability Evolution: From Isolated Tasks to Workflows and Integrated Systems

1.1

Across successive waves of technological change, the transformative impact of general-purpose technologies rarely emerges at the point of invention. It typically occurs when they cross a critical threshold—from passive tools to proactive systems capable of initiating and executing tasks.

To understand the significance of proactive AI, it is necessary to examine the structural limitations of what we describe as “reactive AI”—that is, the prompt–response model that has dominated enterprise applications over the past three years. Since the global rise of ChatGPT in early 2023, most organisations have followed a simple linear logic: human instruction → AI response. While highly effective for discrete tasks, this model becomes increasingly constrained in complex business environments.

Reflecting on the past decade of technological evolution, Fan Ling, Founder and CEO of Tezign, argues that enterprise AI adoption is moving from isolated applications to integrated workflows and, ultimately, system-level deployment. This shift reflects not only advances in technical capability, but also a restructuring of organisational models and business logic.

1.1.1 The Limits of Task-Level AI: Efficiency Ceilings and the Patchwork Trap

At this stage, AI is deployed primarily as a standalone tool, with its core value proposition centred on supporting isolated tasks.

A typical example is the use of generative AI (AIGC) for specific tasks. Marketing teams use Midjourney to generate poster backgrounds, R&D teams rely on Copilot to produce code snippets, and administrative teams use large language models (LLMs) to draft meeting minutes. Between 2023 and 2024, such applications saw rapid growth. According to Gartner, global spending on AI software surged during this period, driven largely by tools that delivered visible impact and improvements in individual productivity.[2]

However, as Professor Wang Qi, Head of the Department of Marketing at CEIBS, noted in a seminar, this model largely reflects a “patchwork” mindset: enterprises insert AI tools into specific steps in the hope of reducing costs and improving efficiency, without changing existing workflows or organisational structures.

While this approach may initially benefit from lower marginal costs—for example, reducing content production expenses by up to 90%—it quickly reaches the limits of localised efficiency gains and may give rise to new challenges:

- System fragmentation: Employees must switch frequently between ERP/CRM systems and multiple AI tools. The cognitive burden of constant context switching often offsets the efficiency gains from AI-generated output.

● Disrupted data flows: Outputs generated by standalone tools—such as images or copy—are often stored in unstructured formats across personal devices or third-party platforms. As a result, they cannot be integrated into core enterprise data assets, interrupting the data chain.

● Value dilution: As competitors gain access to similar tools, content supply increases rapidly while marginal value declines. Enterprises are drawn into an escalating cycle: where they once produced one asset a day, they now produce one hundred, yet conversion rates fail to improve and may even decline due to information overload.

As a result, while task-level AI can address localised efficiency issues, it cannot unlock system-level growth. As McKinsey has observed, although AI adoption is widespread, only a small proportion of enterprises have succeeded in translating it into EBIT growth.[3]

1.1.2 Workflow-Level AI: Connecting Linear Workflows Through Agents

To move beyond the limits of task-level AI, leading enterprises began extending AI capabilities from isolated tasks to connected workflows in 2025—building what can be described as linear workflows. At the centre of this shift were agents, with 2025 widely seen as the first year of enterprise agent adoption.

If a copilot requires a human operator to remain in control, an agent is closer to an autonomous vehicle: it can plan a route and move towards a destination set by the user. At this stage, AI begins to demonstrate basic reasoning and planning capabilities.

Agents no longer function merely as content generators; they act as connectors between strategy and execution. They can interpret relatively open-ended business intent, break it down into structured sequences of sub-tasks, and invoke the appropriate tools to complete them. Two capabilities are particularly important:

● Connecting strategy and execution: In e-commerce operations, for example, a “new product launch agent” can link the full workflow—from market trend analysis and the identification of key selling points to product page copywriting and cross-platform distribution.

● Retaining contextual continuity: Unlike standalone tools, agents can maintain context across extended workflows. Outputs from one step—for example, refined selling points—flow directly into the next, such as copywriting. This reduces the information loss and misalignment common in manual processes.

Although linear workflows can significantly improve process automation, they typically remain confined to a single function or department. For example, a marketing agent may struggle to access supply chain data directly to refine promotional strategies, while a customer service agent may find it difficult to feed insights into product development. In this sense, these workflows operate in parallel—they have yet to be integrated into a broader, interconnected system.

1.1.3 System-Level AI: The Leap to Enterprise-Wide Coordination

In 2026, AI is entering a new stage—system-level coordination—characterised by deep integration with core enterprise IT systems, organisational structures, and external ecosystems.

At this stage, AI is no longer a bolt-on tool or an isolated workflow component. Instead, it becomes embedded in the enterprise's core operating fabric—functioning less as a peripheral tool and more as its central nervous system. Tezign describes this evolution as an "AI Full Stack" capability, in which AI spans the organisation from infrastructure-level data processing through to strategic decision-making. In practical terms, this enables:

- Cross-system interoperability: Agents can operate across traditionally siloed enterprise systems such as ERP, CRM, SCM, and HRM, accessing data and functionality as needed. Gartner predicts that by 2026, 40% of enterprise applications will incorporate agents, marking a shift from function-driven to intent-driven architectures.[4]

- Multi-agent collaboration: Agents can assume distinct roles—such as expert, executor, reviewer, and coordinator—and work together to complete complex tasks. For example, a product development agent may coordinate with market research to identify user needs, with the supply chain to assess material costs, and with legal teams to evaluate patent risks, ultimately producing a feasibility report for human decision-makers.[5]

- Organisational and technological alignment: Enterprises begin to restructure around AI-driven logic. Traditional function-based hierarchies give way to more flexible, goal-oriented networks in which humans and agents collaborate more dynamically.

At its core, the shift from task-level AI to system-level AI reflects a move from localised efficiency gains to system-level execution that generates organisation-wide value. It also requires decision-makers to move beyond viewing AI as an IT project and instead treat it as a central driver of business transformation.

1.2 Defining the Core Shift: From Reactive to Proactive AI

If AI prior to 2025 was characterised primarily by reactivity—responding to human prompts with generated outputs—then by 2026 its defining feature has become proactivity. This does not imply the kind of self-awareness depicted in science fiction. Rather, within clearly defined business boundaries, AI systems are increasingly able to generate insights, initiate actions, and coordinate across tasks without constant human prompting.

1.2.1 Redefining AI: Four Key Components of More Proactive Systems

“More proactive AI” refers to enterprise-grade agent systems that go beyond simple prompt–response interactions. Guided by defined business objectives and operating within governance rules and authorisation boundaries, these systems can continuously sense changes in their environment, generate recommendations, drive task execution, and improve performance through feedback loops.

This does not mean that AI operates with full autonomy or outside organisational control. Rather, it marks a shift from one-off tool use to continuously operating task systems. Its core value lies not in replacing human labour, but in forming a closed-loop mechanism anchored in business objectives, linking insight, decision-making and execution, collaboration, and continuous improvement.

Proactive insight

Traditionally, business intelligence (BI) has operated in a reactive, query-driven manner: users ask questions, and the system returns answers. By contrast, AI systems with proactive insight capabilities function more like a 24/7 enterprise radar, continuously monitoring market dynamics, customer behaviour, competitive activity, and internal operational signals.

These systems draw on the reasoning capabilities of large language models to identify risks before they become visible and detect early signals of emerging opportunities. Tezign’s AI-native research shows that AI is no longer limited to summarising existing data; it can also simulate consumer behaviour and preferences to generate forward-looking insights. This marks a shift from retrospective analysis to early detection and real-time alerting.

Proactive decision-making and execution

Building on these insights, AI is increasingly able to generate solutions, prioritise tasks, and drive execution within established guardrails. However, “proactive decision-making” does not imply fully autonomous decisions detached from human oversight. Rather, it refers to a form of bounded autonomy.

In scenarios where objectives are clearly defined, rules are explicit, and risks are manageable, agents can independently break down tasks, invoke the necessary tools, and carry them through to execution. By contrast, at critical junctures involving high risk, significant uncertainty, or strategic trade-offs, humans remain in the loop and typically retain final decision authority. In such cases, AI generates recommendations, orchestrates execution, and proposes options.

For example, in a supply chain context, AI can generate procurement recommendations based on price fluctuations and inventory forecasts, assess the cost implications of different options, and automatically place standardised orders within pre-authorised limits. However, decisions involving major procurement, brand risk, or cross-regional resource allocation remain firmly with management.

Proactive collaboration

The value of proactive AI extends beyond a single agent's ability to complete tasks more efficiently. More fundamentally, it lies in its ability to connect functions, systems, and agents—breaking down the operational silos that traditionally fragment business activities.

For example, when a sales agent detects rising demand for a product, it can adjust front-end marketing strategies and trigger coordinated responses across the supply chain, production, human resources, and finance. Similarly, when a legal compliance agent identifies regulatory risks, it can flag constraints for marketing, product development, and customer service.

AI is no longer a bolt-on tool attached to a single function. Instead, it is becoming a collaborative network embedded in enterprise operations.

Continuous self-improvement

A genuinely proactive AI system does not stop after a single decision or execution cycle. Its defining characteristic is the ability to continuously improve through feedback.

Each action generates data—such as click-through rates, conversion rates, return rates, customer satisfaction scores, anomaly alerts, or manual corrections. This data feeds back into the system as new context, helping refine prompts, adjust decision rules, improve tool usage, and strengthen coordination across tasks.

Through this cycle of action, feedback, correction, and renewed action, AI develops an increasingly effective closed loop. It does not simply replicate existing capabilities; rather, it continues to improve within the constraints of the enterprise judgement framework.

For this reason, the real competitive advantage of proactive AI lies not in how much content it can generate in a single burst, but in the strength of the feedback mechanisms that enable sustained improvement within the constraints of the enterprise judgement framework.

1.2.2 Strategic Pathways: Governing Legacy Data vs Generating Data Through Execution

As enterprises move towards proactive AI, they face fundamentally different strategic choices. At a closed-door session in Shenzhen, Lu Yi, Assistant Professor of Marketing at CEIBS, proposed a classification framework that divides most organisations into "Type 3" and "Type 4". This distinction marks a key dividing line between traditional and emerging approaches to data governance.

Type 1 enterprises—namely, foundational model providers such as OpenAI and Baidu—and Type 2 enterprises, which develop vertical, industry-specific models, represent a relatively small group of technology suppliers. By contrast, most organisations on the application side fall into either the Type 3 or Type 4 category. The core difference between these two groups lies in the depth of their accumulated data assets and the way they apply AI.

Type 3 Enterprises: Legacy Data Governance (Traditional Paradigm)

This category includes large industry leaders, including major international consumer goods companies. Over many years—often more than a decade—these enterprises have invested heavily in comprehensive data governance, building up substantial digital assets in the process.
Their AI strategies are built on high-quality, static historical datasets, often referred to as legacy data. They follow a traditional machine-learning logic: governance first, application later—starting with data cleansing, followed by model training or fine-tuning, and then deployment. With strong data foundations, these enterprises are well positioned to develop highly accurate decision models.
However, this approach also has clear limitations:

- **High cost**: Data governance involves significant sunk costs. Our research indicates that leading enterprises often invest billions of renminbi annually in proprietary IT systems and data infrastructure, while staffing costs for data governance teams typically run into the tens of millions. For most mid-sized enterprises, this is prohibitively expensive and difficult to replicate.

- **Data latency**: Models built on historical data are inherently backward-looking—relying on past patterns to predict future outcomes. In fast-changing markets, even data from six months ago can quickly lose predictive value.

- **Constrained innovation**: Highly structured data governance systems can limit agility. Efforts to maintain data quality and compliance often come at the expense of speed in business execution.

Type 4 Enterprises: Data Generation Through Execution (Emerging Paradigm)

This category consists of enterprises with relatively weak data foundations, or whose data is highly unstructured and dispersed across fragmented workflows—for example, customer service call recordings, sales communication logs, non-standard documentation, and meeting minutes. These organisations represent the majority of the Chinese market.

For such enterprises, it is neither necessary nor affordable to complete comprehensive data governance before deploying AI. Instead, they embed agents directly into operational workflows, where data is structured in the course of execution—for example, during content generation or customer interaction. In this model, agents continuously generate structured, high-value data in real time.

This represents a new learn-through-execution paradigm—one in which systems improve continuously as they operate—and provides a key pathway for Type 4 enterprises to catch up and compete. Its advantages include:

1. Bypassing data governance to access business insight directly: Leveraging the semantic capabilities of large language models, enterprises can process unstructured data directly without going through lengthy data cleansing. This allows them to generate structured insights more quickly.

2. Real-time feedback and closed-loop learning: A real-time data loop emerges—linking business activity, data feedback, and model improvement. Each action generates new data, which feeds back into the system to refine decision rules and prompt design.

3. Rapid iteration and market responsiveness: Models are no longer trained primarily on historical data from months ago, but on feedback generated as recently as yesterday—or even within the past hour. This enables faster adaptation to changing market conditions.

In this paradigm, the value of data is no longer determined by volume alone, but by the speed at which it circulates and the quality of the feedback it generates. Professor Lu Yi emphasises that if AI is used only to produce content—for example, generating thousands of pieces of copy each day—without feeding operational data back into the system, it will not deliver meaningful efficiency gains. System-level execution creates value when business feedback—such as click-through rates, conversion rates, dwell time, and negative reviews—is continuously fed back into the system to refine agent strategies in real time.

Accordingly, enterprises shift their focus from how much data they possess to how much feedback-driven data they can generate. For Type 4 enterprises, this may be the most viable pathway for survival and growth in the era of large language models.

1.3 Strategic Foundation: Building an Enterprise Judgement Framework for Defensible Advantage in the AI Era

As AI becomes more actively embedded in business operations, the main risk for enterprises shifts from inefficiency to misdirected execution. If not properly guided and constrained, highly proactive agents can produce serious errors at scale—for example, generating non-compliant content or sending incorrect automated responses to customers.

For this reason, building an enterprise judgement framework to guide and constrain proactive AI is becoming a foundational element of AI strategy in 2026.

1.3.1 Strategic Orientation: Context as the Foundation, Feedback as the Engine

Large language models may provide general-purpose capabilities, but whether enterprise agents can be truly proactive and reliable depends less on the model itself and more on the presence of a continuously operating enterprise judgement framework. As Fan Ling, founder of Tezign, argues, the essence of proactive agents lies not in the strength of any single model, but in the stable coupling between the context system and the feedback loop.

To translate this into an operational architecture, Tezign has proposed the Generative Enterprise Agent (GEA) framework. Designed to operate within real business processes, GEA is neither a single model nor a standalone tool. Rather, it is a system-level architecture built around business intent, using orchestration and reasoning capabilities to coordinate multiple models and agents in continuous interaction with the enterprise context.

Viewed through its four-layer architecture, the enterprise judgement framework is not a simple assembly of technologies, but a complete mechanism spanning goal setting, task orchestration, proactive execution, contextual knowledge accumulation, and continuous evolution.

The first layer is Intent. Users no longer issue step-by-step instructions; instead, they define business objectives directly—for example, identifying growth opportunities in Southeast Asia next quarter or refining product positioning based on recent user feedback. This marks a shift from command–execution to goal-driven problem solving. The Intent layer not only initiates the system, but also sets direction, boundaries, and priorities for subsequent reasoning and resource allocation.

The second layer is Orchestration. This serves as the system's reasoning core, powered by a Creative Reasoning Model. Based on business objectives, the system breaks tasks into structured sub-tasks, explores multiple execution paths, and determines which require insight generation, content creation, or access to external systems and proprietary knowledge. It then dynamically assigns tasks to the most appropriate combination of models, tools, and agents. What enterprises need is not simply a model that can answer questions, but a system that can organise capabilities and coordinate action around business objectives.

The third layer is Agent Skills, encompassing proactive agents, MCPs and APIs. This layer forms the system's execution foundation. Rather than waiting for instructions, proactive agents continuously monitor changes in the operating environment, detect anomalies, and initiate responses. Agent Skills modularise capabilities such as insight generation, content creation, analysis, and distribution into callable functions. MCPs and APIs connect the system to CRM, ERP, content platforms, and external data sources, enabling agents to operate within real business processes and generate outcomes that can be tracked, evaluated, and reviewed through a feedback loop.

The fourth layer is the Context System. This forms the true foundation of the enterprise judgement framework and serves as GEA's unified memory. It captures not only explicit documentation, but also tacit knowledge such as brand tone, creative assets, project history, product information, user models, communication practices, and historical decision logic. It is not simply a knowledge base or vector database, but the enterprise's single source of truth, enabling all agents to reason and act based on a shared enterprise context. Without this layer, AI produces answers that may be technically correct but contextually misaligned. With it, AI can determine what counts as the right course of action within that particular organisation.

Within the GEA architecture, multi-model capabilities are not treated as a standalone layer. Instead, they function as underlying infrastructure, dynamically invoked by the Orchestration layer as needed. System performance depends not on the models themselves, but on the clarity of intent, the effectiveness of orchestration, and the quality of the context system.

These four layers operate through an end-to-end feedback loop. While the context system determines whether AI understands the enterprise, the feedback loop determines whether it can improve over time. Business outcomes, user behaviour, conversion data, satisfaction scores, and human review decisions continuously feed back into both the context system and the orchestration layer, refining decision logic, task pathways, and execution strategies.

In this sense, the enterprise judgement framework is far more than a knowledge base attached to a model. It is an integrated system in which Intent, Orchestration, Agent Skills, and the Context System operate within a continuous feedback loop. This allows AI not only to execute tasks, but to understand why they should be performed in a particular way—and whether they are appropriate for the organisation.

1.3.2 Reshaping Value: Three New Dimensions of Scarcity in 2026

As the cost of accessing general-purpose LLMs falls rapidly—at a pace reminiscent of Moore's Law—and these models become as accessible as basic utilities, the advantage of early technological adoption will quickly erode. As Guo Wei, Associate Professor of Strategy and Entrepreneurship at CEIBS, argues, business leaders must return to a fundamental question: what remains constant amid rapid change? It is this that forms the basis for building defensible competitive advantage.

In the 2026 competitive landscape, general-purpose models are no longer scarce. What will continue to differentiate enterprises is the distinctive combination they build across three dimensions:
Proprietary data assets: This does not mean publicly available online data, but high-quality internal operating data accumulated over time—such as historical business records, failed R&D experiments, and distinctive customer interaction histories. Crucially, this data must be transformed into dynamic

institutional memory, providing a single source of truth for agents when making decisions under uncertainty.

Deep industry insight: Exploratory knowledge work is inherently nonlinear and uncertain, with critical decisions often relying on the tacit expertise of domain specialists. The ability to systematically extract this know-how from experienced employees and codify it into agent-accessible capabilities will become a key source of competitive advantage—and a significant barrier to entry.

High-frequency feedback loops: The most advanced AI systems are shaped not only in controlled environments, but through continuous exposure to highly competitive, real-world market conditions. Enterprises must therefore embed AI deeply within business operations, capture feedback without loss, and rapidly adjust model strategies through short iteration cycles.

When these three dimensions are integrated into an enterprise judgement framework, AI is no longer a replaceable external tool. Instead, it becomes a core component of the organisation's competitive capability.

1.3.3 Implementing AI Through the Strategy–Technology–Business Triangle

Bringing these dimensions of scarcity together into an enterprise judgement framework that supports agent-driven operations cannot be achieved by any single function alone. Based on extensive field research, Professor Wang Qi proposes a "golden triangle" framework—the strategy–technology–business triangle. Efforts to achieve system-level execution and value creation through single-point interventions—such as purchasing expensive AI systems or issuing ambitious strategy statements—are likely to prove ineffective. Meaningful impact depends on aligning three interdependent layers:

Strategy Layer: Strategic Direction and Organisational Alignment

Led by top decision-makers such as founders and CEOs, this layer goes beyond setting an "AI-first" vision. It involves reallocating resources, breaking down organisational silos, and fostering a culture that accepts measured experimentation and failure.

- **Core responsibilities**: Define the strategic intent and boundaries of AI deployment. Clarify which elements must not change, and must in fact be strengthened, such as core values and brand ethos. Establish the core principles underpinning the enterprise judgement framework.

- **Key actions**: Determine whether AI is used primarily for cost reduction, revenue growth, or new business models. Define the scope of AI autonomy. Set ethical guidelines for AI use.

Technology Layer: Capability Enablement and Architectural Redesign

Led by the Chief Digital Officer (CDO) or Chief Technology Officer (CTO), this layer focuses on building a flexible, secure, and interoperable AI platform. Beyond core engineering challenges such as model orchestration and resource allocation, it requires infrastructure capable of managing complex context systems so that AI can reliably access enterprise knowledge.

- **Core responsibilities**: Build the infrastructure that supports the enterprise judgement framework, including the selection of foundation models, agent orchestration platforms, machine learning operations (MLOps) systems, and connected data pipelines.

• **Key actions**: Enable real-time data capture and processing. Ensure system stability and security. Provide low-code or no-code tools that allow business teams to participate in configuring agents.

Business Layer: Real-World Deployment and Accountability

This is the most critical—and often most overlooked—layer. Effective deployment must be owned by front-line business leaders operating under real market pressure. AI only delivers value when it is embedded in concrete business challenges—such as supply chain inventory optimisation or front-line sales performance—and held accountable for measurable outcomes such as revenue and profit growth. Without this, it risks remaining a technology showcase rather than a source of real productivity.

• **Core responsibilities**: Provide real business scenarios and continuous feedback. Front-line teams must participate in training and refining agents, converting tacit knowledge into explicit prompts and rules.

• **Key actions**: Define concrete application scenarios. Feed real business data into AI through day-to-day operations. Review and correct outputs on an ongoing basis. Only when business experts are deeply involved in training, reviewing, and refining agents—and when outputs are continuously corrected through human-in-the-loop processes—can AI avoid becoming a showcase deployed “for its own sake”.

Professor Wang Qi emphasises that AI initiatives **create real business value only when they are actively used in operations**. If driven solely by the technology function, without deep engagement from business teams, AI risks becoming a showcase—impressive in form, but limited in impact.

1.4 Redesigning the AI Maturity Journey: Towards L3 Proactive Systems

As proactive AI becomes a defining force in reshaping the competitive landscape, traditional digital maturity models—based on metrics such as cloud migration rates, the proportion of structured data, or model size—are no longer sufficient. In the era of AI agents, the key measures of enterprise evolution shift to the scope of AI autonomy, the level of responsibility it assumes for outcomes, and the way human–machine collaboration is structured within the organisation.

Based on extensive CEIBS research into more than one hundred pioneering enterprises worldwide, this white paper redefines the AI maturity journey as a four-stage progression from L1 to L4. This framework draws on both the logic of autonomous driving classifications and insights into organisational evolution. As enterprises move from one level to the next, each transition reflects not just incremental gains in model performance, but a deeper shift in operating logic and management philosophy.

Table 1: Evolution of Enterprise AI Maturity (2026 Edition)

Maturity level	Core characteristics	Typical application models	Human–machine interaction model	Core business value and capability indicators	Current status
L1 Assistive	Humans remain primary actors; AI functions as a standalone tool supporting specific tasks.	Copilots, conversational chatbots, single-task generation tools	**Prompt–response model**: humans issue explicit instructions; AI generates a single output for a defined task.	**Task-level efficiency gains**: reduces repetitive work, alleviates local capacity constraints, and improves individual productivity.	**Widespread adoption**: by 2024–2025, most enterprises had reached this baseline.
L2 Collaborative	AI is embedded in workflows and collaborates with humans, linking tasks into linear business processes.	Linear workflow agents (e.g. product launch assistants, compliance agents, quality inspection agents)	**Human-in-the-loop**: AI executes multi-step processes; human approval is required at key decision points.	**Workflow-level automation**: reduces context-switching costs and enables seamless information flow across extended processes.	**Deepening exploration**: leading enterprises are actively tackling implementation challenges and expanding adoption at this stage.
L3 Proactive	AI senses changes in the business environment and generates recommendations, which are reviewed by humans. Supported by an enterprise judgement framework, it operates with contextual awareness.	Multi-agent systems capable of perception, planning, and cross-system coordination	**AI-initiated, human-authorised**: AI proactively generates insights and recommended actions; humans review and approve.	**System-level execution and growth**: enables cross-system coordination, breaks down silos, and ties AI directly to KPIs and business outcomes. Frees humans to focus on strategic and creative work.	**Key strategic target**: in 2026, leading enterprises should aim for broad adoption at this level.
L4 Autonomous	AI operates with a high degree of autonomy within defined boundaries, taking on most routine decisions and coordinating across functions.	AI-driven organisations (e.g. DAO-like structures or connected enterprise intelligence systems)	**Human-on-the-loop**: humans define goals and guardrails, while AI operates end-to-end and humans intervene only in exceptional cases.	**Business model transformation**: enables fully closed-loop operations and continuous self-optimisation based on environmental feedback.	**Vision stage**: currently limited to a small number of high-frequency, low-risk scenarios (e.g. algorithmic trading).

In 2026, enterprises should not overreach by pursuing an L4 Autonomous model in which operations proceed without human oversight. The complexity of real-world business dynamics, together with the ethical risks involved, means that humans will retain ultimate decision authority for the foreseeable future. At the same time, organisations should not remain at the L1 Assistive stage. The marginal efficiency gains delivered by standalone tools are being rapidly commoditised.

The most urgent and practical priority is to move towards L3 Proactive—building agent networks that are both proactive and governed by an enterprise judgement framework. This shift enables organisations to move from task-level efficiency gains to system-level execution and value creation.

This is not an incremental upgrade, but a shift in operating logic centred on proactivity. Enterprises must first make the cognitive shift from reactive response to proactive operation. They must then choose a strategic pathway suited to their circumstances, with Type 4 enterprises particularly well positioned to adopt a learn-through-execution approach. Finally, they must establish a robust judgement framework based on the strategy–technology–business triangle, providing a stable foundation for progress along the AI maturity journey.

The ultimate objective is not superficial cost reduction or short-term efficiency gains, but a deeper transformation of how the business operates and creates value.

Endnotes

[1] McKinsey 2025 AI Report : 88% of Companies Are Failing at AI | by Analyst Uttam | AI & Analytics Diaries, Nov 28, 2025， https://medium.com/ai-analytics-diaries/mckinsey-2025-ai-report-88-of-companies-are-failing-at-ai-053f1ac746d3

[2] Manfred Bremmer. Gartner: IT spending will exceed $6 trillion in 2026 - CIO, Feb 3, 2026， https://www.cio.com/article/4126847/gartner-it-spending-to-exceed-6-trillion-by-2026.html

[3] MJ Smith. McKinsey's State of AI 2025: What Separates High Performers from the Rest, Jan 7, 2026， https://www.colabsoftware.com/post/mckinseys-state-of-ai-2025-what-separates-high-performers-from-the-rest

[4] Gartner Predicts 40% of Enterprise Apps Will Feature Task-Specific AI Agents by 2026, Up from Less Than 5% in 2025, Aug 26, 2025， https://www.gartner.com/en/newsroom/press-releases/2025-08-26-gartner-predicts-40-percent-of-enterprise-apps-will-feature-task-specific-ai-agents-by-2026-up-from-less-than-5-percent-in-2025

[5] Gartner Identifies the Top Strategic Technology Trends for 2026, Oct 20, 2025， https://www.gartner.com/en/newsroom/press-releases/2025-10-20-gartner-identifies-the-top-strategic-technology-trends-for-2026

[6] Intelligent Agents in AI Really Can Work Alone. Here's How. - Gartner, https://www.gartner.com/en/articles/intelligent-agent-in-ai

2

DEEPENING THE BLUEPRINT: THE 3×3 STRATEGIC MATRIX FOR THE PROACTIVE AI ERA

In the 2025 report Navigating AI-Powered Business Transformation, we introduced the “3×3 Strategic Matrix” as a framework for assessing enterprise AI strategies. The matrix examined AI strategy along two axes. The horizontal axis captured the breadth of AI’s role in the business, from cost reduction and efficiency improvement to growth and business model innovation. The vertical axis measured the depth of AI adoption, from proof of concept (PoC) to scaling and organisational transformation.

The 2025 matrix reflected a copilot-based view of AI. Enterprises were mainly asking how AI could help people complete specific tasks more efficiently. In other words, the focus was still on task-level human efficiency and localised efficiency gains.

By 2026, more proactive AI capabilities were changing this logic. As agents became more capable of taking on defined tasks, the question shifted from how AI could assist individuals to how it could participate in system-level execution and generate organisation-wide value. To establish a shared basis for this shift, the following table compares the core definitions and underlying assumptions of the 2025 and 2026 matrices.

Table 2. Core Logic Comparison of the 2025 and 2026 Matrices: From Localised Efficiency Gains to System-Level Execution

Strategic Dimension	Matrix Position	Localised Efficiency Gains	System-Level Execution	Core Transition Marker
Breadth of AI Strategy (Horizontal Axis)	Cost Reduction and Efficiency Improvement	**Closed-Loop Assistance**: AI acts as an enabler within existing workflows, helping employees work faster and at lower cost, while humans remain the ultimate control point in the process.	**Closed-Loop Delegation**: AI independently takes on and completes defined end-to-end business workflows, becoming the primary executor in the process.	The focus shifts from “hours saved” to the proportion of complete workflows handled by AI.
	Driving Growth	**Passive Response**: Organisations respond to customer needs, using AI to improve marketing conversion rates or response speed, such as through intelligent customer service.	**Proactive Engagement**: AI proactively anticipates trends, identifies demand, and initiates commercial interactions, with performance directly tied to growth KPIs such as sales volume.	AI moves from back-end support to a front-line role in driving sales and customer traffic.
	Business Model Innovation	**Product Feature Add-on**: AI is incorporated into existing standardised software or hardware products as an additional feature or selling point.	**Outcome-Based Delivery**: The revenue model changes fundamentally, shifting from one-off product sales to the continuous delivery of AI-enabled capabilities and outcomes.	Revenue models evolve from hardware sales or one-time licensing towards subscriptions, service-based delivery, and shared ROI arrangements.
Depth of AI Adoption (Vertical Axis)	Proof of Concept (PoC)	**Quality Validation**: Outputs generated by large language models (LLMs) are tested in sandbox environments to assess whether they are accurate, polished, and fit for demonstration.	**Agency Validation**: Organisations evaluate whether agents can independently plan, reason, and act within defined boundaries to achieve closed-loop execution and establish a trust contract.	Evaluation criteria shift from output quality to operational stability without human intervention.
	Scaling	**Tool Proliferation**: Enterprises purchase and deploy software licences at scale, giving every employee access to a Copilot.	**Digital Headcount**: Organisations formally delegate authority to AI by assigning defined “digital employee” roles and integrating AI as a stable component of operational structures.	AI is assigned a defined role, reporting line, and resource-access authority within the organisational structure.
	Organisational Transformation	**Behaviour Management**: Processes are redesigned to standardise how employees use AI tools more effectively.	**Outcome-Oriented Management**: As organisations move towards self-directed models, dedicated AI-native functions are created and management shifts from monitoring behaviour to managing outcomes.	Traditional business–IT silos break down, giving rise to an integrated human–AI operating model.

2.1 Breadth of AI Strategy (Horizontal Axis): From Tool-Based Assistance to AI-Enabled Value Delivery

The horizontal axis asks how broadly AI is expected to contribute to the business. In the proactive AI era, this is no longer just a question of where AI tools are deployed. It is a question of how far enterprises are prepared to let AI systems take on responsibility for value delivery.

2.1.1 Cost Reduction and Efficiency Improvement: From Helping Humans Work Faster to Autonomous Closed-Loop Execution

In the reactive AI phase, cost reduction and efficiency improvement usually followed a simple human-in-the-loop model: AI produced a draft, and people reviewed, revised, and approved the final output. This made work faster, but only up to a point. The process was still limited by human judgement and approval cycles.

By 2026, leading organisations were moving beyond this model. In clearly defined scenarios, agents were beginning to take over entire closed-loop workflows that had previously depended on human intervention.

Shiji Kaiyuan, a customised printing provider serving small and micro enterprises, illustrates this shift. The company operated a long-tail e-commerce platform for customised print products and handled more than 50 million non-standard enquiries each year, covering everything from brochure sizes and business card materials to layout requirements. Its large customer service and design teams had long made the business highly labour-intensive. In early AI trials, the system merely supported customer service agents by suggesting responses. Staff still had to intervene frequently, and the impact on costs was limited.

At the end of 2024, Shiji Kaiyuan rebuilt its AI customer service system to take advantage of stronger large-model reasoning capabilities. This time, AI was no longer used as a “typing assistant”. It was given the authority to handle customer conversations independently.

On public e-commerce platforms such as Pinduoduo, AI digital employees took over the overnight shift from midnight to 8 a.m. without human supervision. When a small business owner asked late at night about custom coffee cup sleeves, the AI could interpret the vague request, draw on the company’s knowledge base to recommend a suitable design solution, and guide the customer towards a transaction. According to the company’s data, the conversion rate for automated overnight customer service reached 45–50%, almost matching the performance of daytime shifts staffed by human experts.

In the proactive AI era, the key measure of cost reduction and efficiency improvement is therefore beginning to shift: from how many hours employees save to how many standardised business workflows can run as fully automated closed loops without human intervention. Once AI can independently process orders, it is no longer merely a cost-optimisation tool. It becomes a direct engine for protecting margins.

2.1.2 Driving Growth: From Passive Response to Proactive Prediction and Engagement

In the growth dimension, enterprises are moving beyond using AI as a passive dashboard for querying data. Instead, AI is beginning to act as an active growth aide: predicting market shifts and identifying sources of incremental demand. Its performance is now tied directly to measurable growth KPIs.

Swire Coca-Cola, a beverage giant managing millions of offline retail outlets across China, illustrates this shift. Traditionally, product distribution relied heavily on the personal experience of front-line sales representatives. But the tacit knowledge of experienced salespeople was difficult to replicate at scale. Nor could human judgement precisely calculate the best assortment across hundreds of SKUs, different store catchments, and varying outlet types. To address this growth bottleneck, Swire Coca-Cola introduced an AI-based Recommended Order system.

The agent integrates multiple data inputs, including store attributes, real-time weather, historical sell-through performance, and competitor promotions. Each morning, it proactively generates optimal SKU ordering recommendations for individual stores and delivers them to front-line sales representatives. During the National Day holiday period, for example, the AI challenged the sales team's habitual assumptions and unexpectedly recommended that a university campus store increase its stock of mini-can Coca-Cola bundled with snacks. This precise prediction translated directly into an uplift in retail sales.

Here, AI was no longer a passive statistical tool. It had become an active growth engine.

2.1.3 Business Model Innovation: From Selling Products or Projects to Delivering AI-Enabled Capabilities and Outcomes

The highest level of strategic breadth is business model innovation. At this level, enterprises move beyond selling one-off products or projects and begin delivering dynamic outcomes or knowledge-based capabilities that AI can provide on a continuous basis. By combining deep industry know-how with AI, organisations can move beyond their original competitive boundaries and create new sources of growth built around services and capabilities.

Ipsos China, part of the global market research and consulting organisation Ipsos, provides a representative example. Historically, the company generated revenue mainly by billing consultant time and delivering one-off, bespoke offline research reports. This model was inherently constrained by the size and capacity of its human research teams. Faced with the democratisation of knowledge brought about by generative AI, Ipsos chose to transform rather than defend its existing model.

Ipsos China used decades of accumulated qualitative interview data to build its proprietary AI Community, a digital twin community of consumers. Today, the company not only sells customised research projects, but also offers brands such as Procter & Gamble and Yili Group access to "digital consumers" that are online 24 hours a day and capable of multi-turn conversations. These capabilities are delivered through SaaS subscriptions or private deployments.

This shift—from selling labour-intensive projects to monetising standardised AI-enabled assets—shows how traditional service providers can use AI to reshape their business models.

2.2 Depth of AI Adoption (Vertical Axis): From Isolated Use Case Validation to Human–AI Symbiosis

If the horizontal axis defines where organisations seek value, the vertical axis measures how deeply AI is embedded in the business. In 2025, depth of adoption was often assessed through indicators such as technology coverage, tool adoption rates, and the amount of computing power procured.

By 2026, however, the basis for evaluation was shifting. The key question was no longer simply how widely AI tools had been deployed, but how much trust and authority organisations were prepared to grant to AI systems. Depth of adoption was also reflected in how far organisational structures had been redesigned to accommodate agents.

2.2.1 PoC: Validating Agency and Establishing a Trust Contract

During the earlier experimentation phase, proof-of-concept (PoC) projects typically focused on technical feasibility: whether large models hallucinated, or whether generated code could actually run. In the era of proactive AI, however, the purpose of PoC shifts fundamentally. It is no longer simply a test of whether the technology works, but a test of agency: whether AI can safely and independently complete closed-loop tasks within real business workflows shaped by noise, uncertainty, and compliance requirements. In this sense, PoC becomes the starting point for establishing a "trust contract" between the enterprise and AI.

In high-stakes domains, this trust-building process is especially cautious. Merck's human resources team, for example, explored the use of AI for flight-risk prediction—identifying employees at risk of leaving. In this context, predictive accuracy alone was not enough. Because the project involved sensitive employee data, Merck first introduced a Digital Ethics Check tool customised by its internal digital ethics team.

The project team then set strict boundaries around fairness, data privacy, and algorithmic transparency. AI-generated predictions were defined as advisory only; they could not be used as the sole basis for dismissal or other major personnel decisions. Management approved the pilot only after the system had demonstrated not just predictive accuracy, but also compliance with these boundaries and a highly explainable decision process.

In the agent era, the purpose of PoC is therefore to build confidence through rigorous validation, so that business teams can begin entrusting selected decision rights to this new "digital colleague".

2.2.2 Scaling: Establishing "Digital Headcount" and System-Level Delegation

With the trust contract in place, AI adoption moves into the scaling phase. At this stage, organisations no longer treat AI as an experimental plug-in managed by an innovation team. Instead, they formalise AI's role by giving it "digital headcount" status and granting it system-level authority within core business processes.

SFC, a cross-border e-commerce exporter built around large-scale product listing and distribution, illustrates the power of this kind of delegation. In cross-border e-commerce, waves of overseas intellectual property (IP) claims and temporary restraining order (TRO) lawsuits—often filed aggressively or opportunistically—pose a persistent threat to sellers. A single infringement can lead to the immediate freezing of store operations and funds, creating an existential risk for the business. With tens of thousands of new SKUs listed each day, manual compliance checks are both impractical and insufficient.

To address this challenge, SFC developed ERiC, an AI-powered compliance platform combining computer vision with large-model reasoning. Rather than asking employees to check products manually in the system, SFC embedded ERiC directly into the mandatory product-listing workflow as a gatekeeper. In this formal role, the platform automatically scanned large volumes of products each day, checked them against tens of millions of global design patents and extensive trademark databases, and flagged infringement risks before listings went live.

The impact was substantial. By delegating mission-critical compliance review to AI at system level, SFC reduced compensation payments related to TRO litigation by more than 80% within two years.

Atypica is an AI-native, multi-agent market research system that introduces multiple "digital employee" roles into enterprise user research and market insight generation. Traditionally, companies had to invest heavily in long-cycle market research projects, often relying on consulting firms to conduct the work. Atypica was built on a different premise: market insight is a form of exploratory knowledge work. Its starting points are often ambiguous, its process is non-linear, and its outcomes are difficult to predict. As a result, it cannot be standardised through conventional SOPs, and traditional retrieval-augmented generation (RAG) systems or simple multi-agent solutions are often insufficient.

To address this complexity, Atypica was built on the Generative Enterprise Agent (GEA) architecture, allowing the system to respond flexibly to complex research needs. It operates through two core agents:

1. Reasoning Agent: Integrates contextual information and performs continuous reasoning and judgement, providing directional guidance for the overall research process.

2. Execution Agent: Carries out specific tasks under the guidance of the Reasoning Agent, including social media observation, data collection, and report generation.

For example, in a commercial study on young consumers' coffee preferences, the system might begin with the system analysing thousands of posts on RedNote (Xiaohongshu) and other social platforms. Atypica then builds virtual consumer personas, identifies potential contradictions in the data, and dynamically initiates follow-up interviews. It ultimately synthesises these findings into an insight report that provides targeted support for brand decision-making.

Through this dual-agent design, Atypica not only improves research efficiency, but also adapts and improves continuously in response to different enterprise needs. This helps companies better understand consumer behaviour and make more data-driven decisions.

When companies measure the risk mitigation and business value generated by AI systems with the same rigour they apply to human employee output—and even place the return on investment of “digital employees” under HR oversight—large-scale delegation is no longer a concept. It has become part of the organisational practice.

2.2.3 Organisational Restructuring: Towards a Self-Directed Organisation—From Managing Behaviour to Managing Outcomes

The deepest form of AI adoption is organisational restructuring. Once multi-agent systems become deeply embedded in the business and begin to take over closed-loop workflows, traditional hierarchies and functional silos become a major constraint on efficiency. At this point, management must shift from supervising how people perform specific tasks to managing the business outcomes the system is expected to deliver. To make AI work effectively, leading enterprises are beginning to create new organisational forms.

Mengniu Group offers a clear example of bold strategic transformation. As the company pushed AI deeper into its core operations, it faced a familiar problem in large-scale digital transformation: technology teams often lacked direct business ownership, while business units lacked technical expertise.

Rather than relying on the conventional model of loose “technology enablement”, Mengniu redeployed key talent from its digital and other back-office functions into its largest core unit: the room-temperature dairy business. These employees were embedded directly in the business unit, both physically and organisationally, leading to the creation of a new department dedicated to the room-temperature dairy operating model.

The team’s performance model changed fundamentally. It was no longer assessed against traditional IT-style vanity metrics such as system launch rates or lines of code. Instead, it became directly accountable for core business KPIs, including revenue and profit in the room-temperature dairy business.

The team was no longer a software provider working from the sidelines. Equipped with AI technologies and data-driven insights, it became a specialised task force responsible for redesigning the business’s end-to-end workflows. This organisational design brought technology and business into much closer structural alignment.

A similar form of organisational restructuring took place at **Shinyway Education**, a leading international education services group in China. To accelerate the rollout of its AI exam-preparation product ChillPrep, the company broke down traditional divisional boundaries and created a lightweight virtual organisation made up of business specialists, technical experts, and external advisers.

The team also changed how it was managed. Instead of relying on conventional KPI scoring, it adopted an OKR system that translated broad strategic goals into concrete AI training outcomes and revenue targets. This agile, goal-oriented structure allowed AI capabilities to be translated quickly into a commercial offering. The project team was eventually spun out of the original business unit and established as an independent subsidiary built around the new business logic.

At the highest level of AI adoption, transformation is not a matter of adding a layer of technology to the old organisation. It requires the organisation to build a new operating structure around the proactive capabilities of AI. When departmental boundaries, reporting lines, performance metrics, and authority-allocation mechanisms are redesigned so that agents can complete business workflows more effectively, technology and business become structurally integrated. Only then can an enterprise move towards a genuinely human–AI symbiotic, self-directed organisation.

The arrival of proactive AI has therefore rewritten the logic of business transformation. Across the redesigned 3×3 Strategic Matrix—whether horizontally across cost reduction and efficiency improvement, growth, and business model innovation, or vertically through PoC, scaling, and organisational restructuring—the core shift is the same: enterprises are beginning to grant AI system-level authority.

Once agents move beyond the role of assistive tools and become “digital employees”—with measurable commercial KPIs, formal headcount status and the ability to reshape traditional hierarchies into self-directed organisations—technology and business are no longer separate domains. In this new phase, competitive advantage will not come from layering more tools onto existing structures. It will come from embedding proactive AI deeply into core workflows and organisational architecture.

3

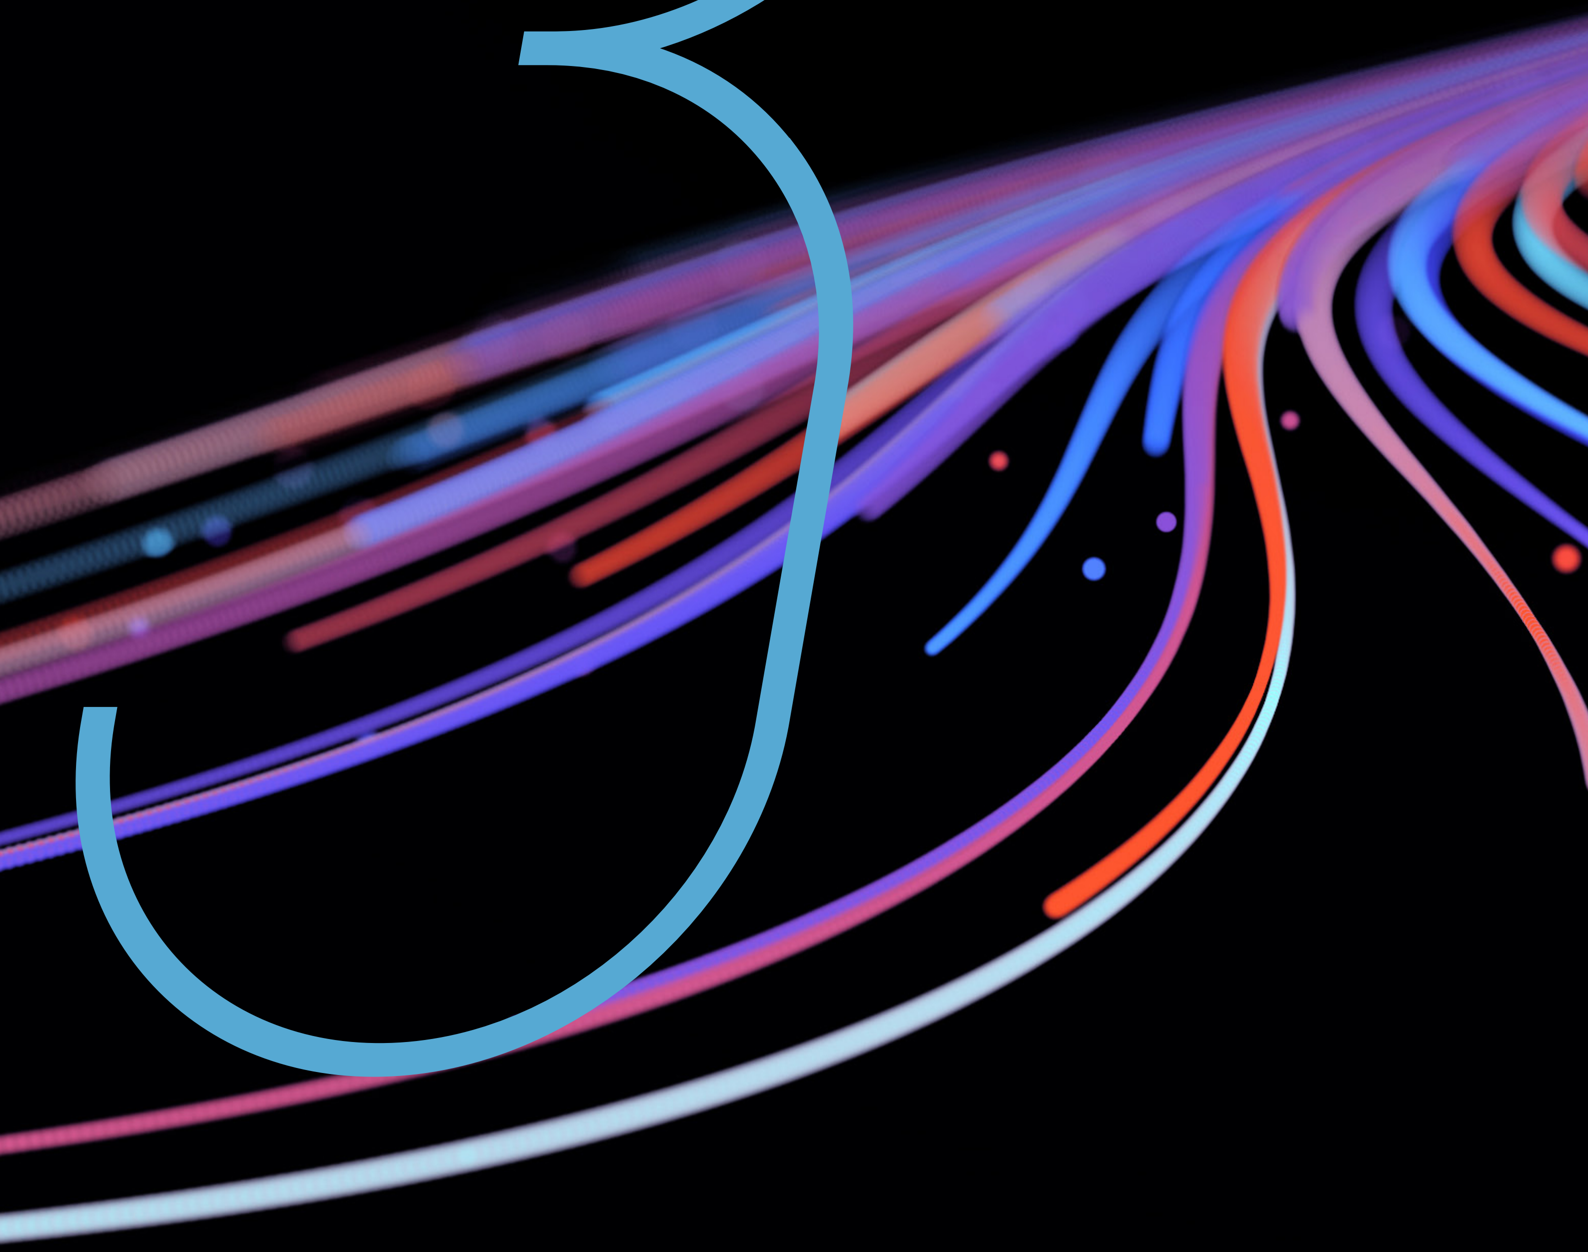

RECONFIGURING SYSTEMS: AGENT-DRIVEN VALUE LOOPS AND ORGANISATIONAL RECONFIGURATION

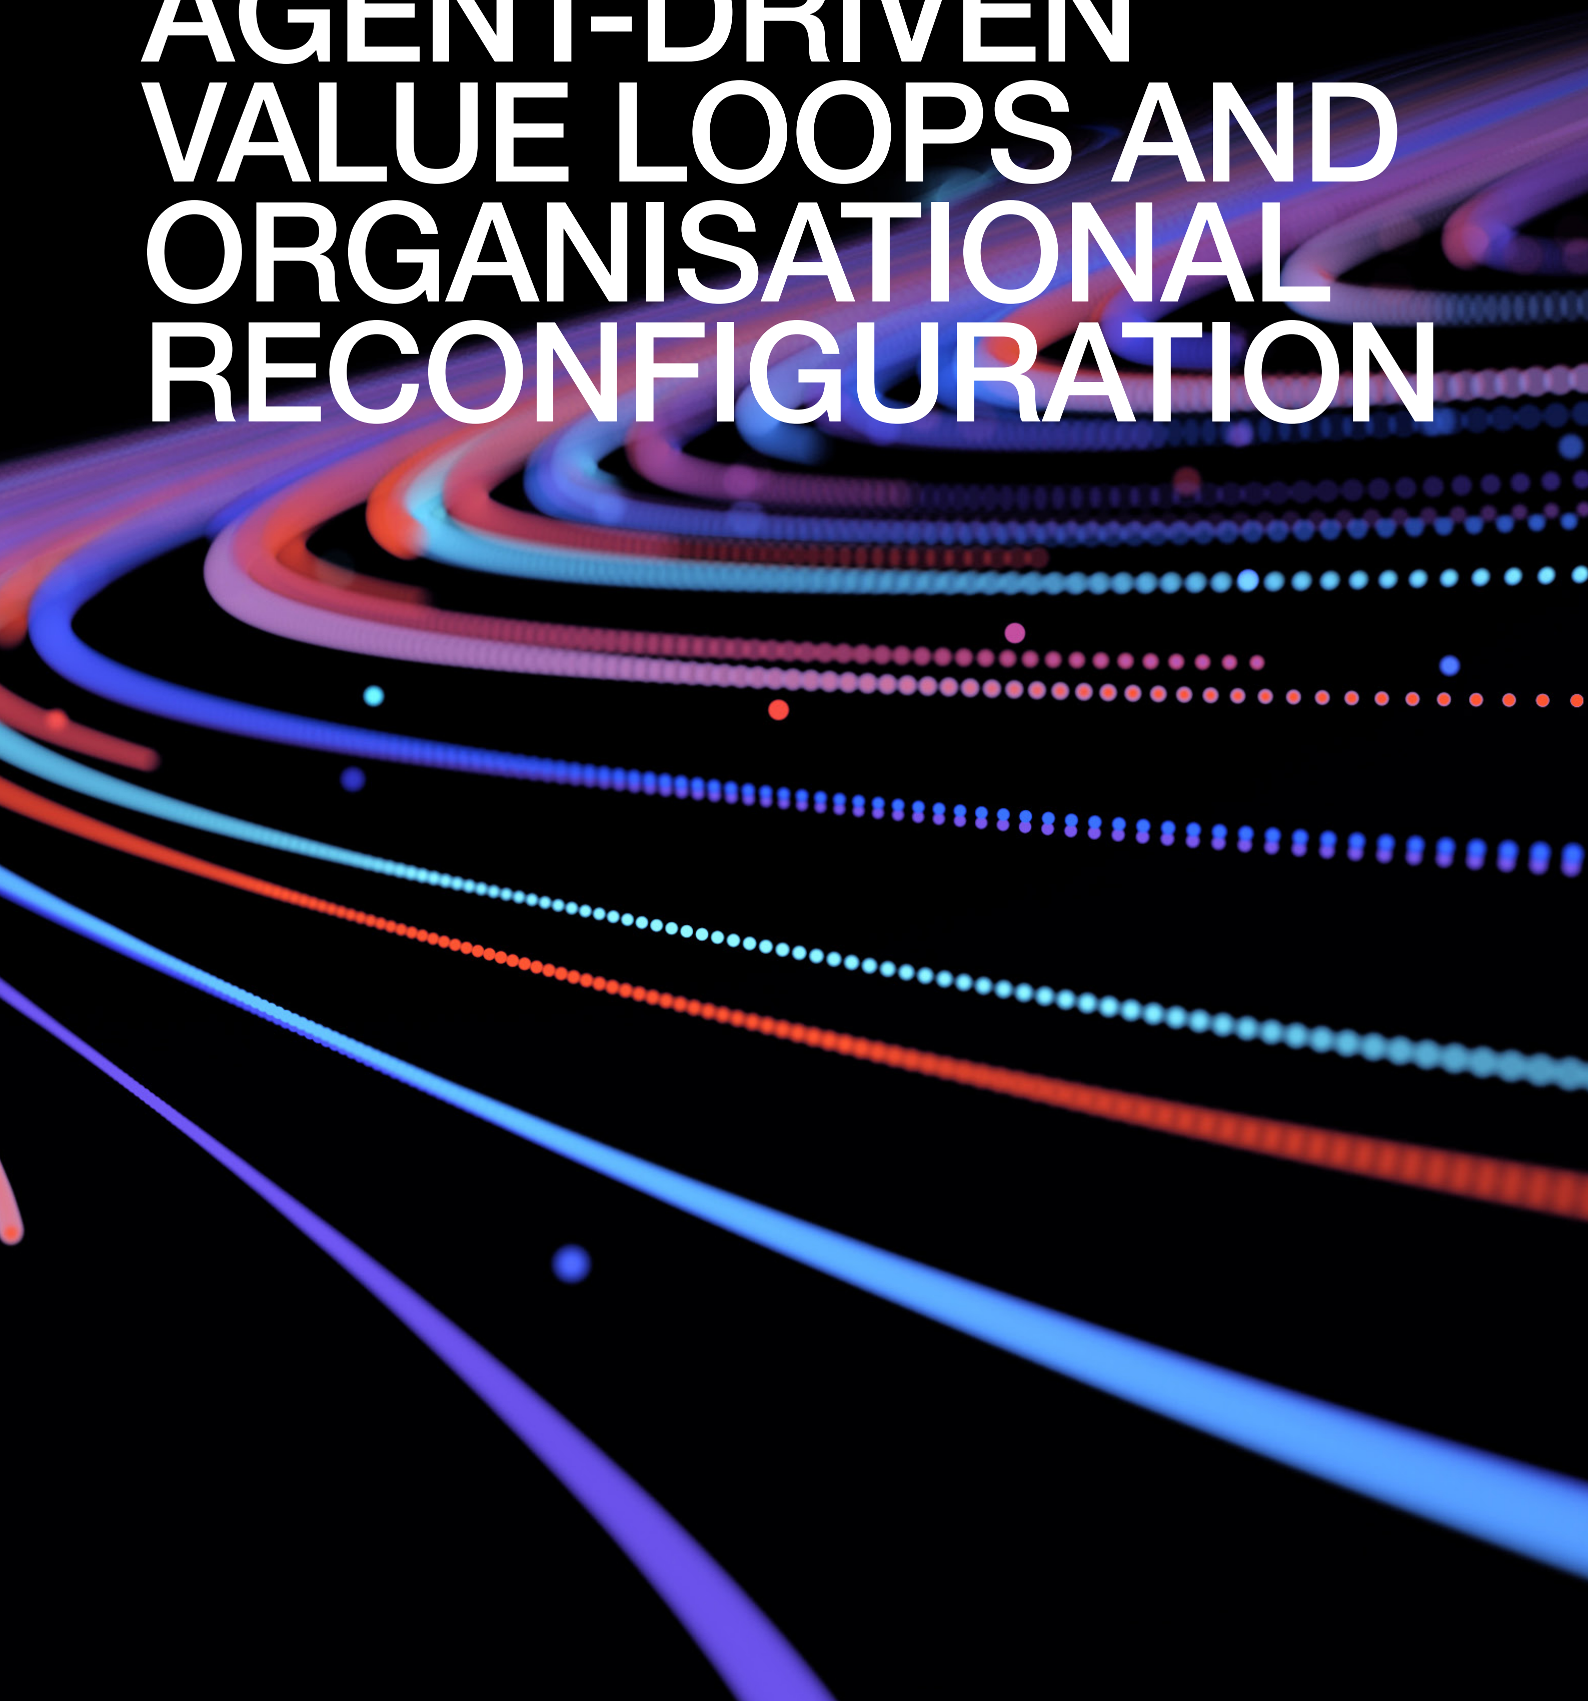

Part I: Agile Breakthroughs and the Growth Engine—Reducing Trial-and-Error Costs and Driving Revenue Growth

Case 1: **Mars: From the Hero-Product Model to Digitally Enabled Probabilistic Innovation**

Case 2: **Semir: Embedding AI and Building a Second Growth Curve**

CASE 1

Mars: From the Hero-Product Model to Digitally Enabled Probabilistic Innovation

I. Market Shifts and the Challenge to the Hero-Product Model

Mars, Incorporated is a global, family-owned business spanning snacking, pet care and food and nutritional products. In 2025, its global annual revenue exceeded US$65 billion. For large fast-moving consumer goods (FMCG) companies such as Mars, competitive advantage has long depended on scale. The core model is to develop “hero products” with broad mass-market appeal—such as Dove, M&M’s, Snickers and Extra—and distribute them widely through centralised distribution networks and mass media, reducing unit production costs and maximising economies of scale.

The product innovation process behind this model is typically linear and highly structured. Consumer insights teams gather market feedback through focus groups, in-depth interviews and questionnaires. The findings are then tested against multiple data sources to determine where product development should focus. Only then do product ideas move through concept approval, formulation development, packaging design and preparation for mass production. In a traditional market environment, developing a single new product can take close to a year.

In recent years, however, China’s retail channels and consumer behaviour have changed significantly. Research has shown that the growth of e-commerce and the fragmentation of offline retail channels have reshaped competition in the FMCG sector, while consumers have come to expect more differentiated products and faster product iteration.[1] Different retail formats also have distinct product requirements. Warehouse retailers such as Sam’s Club tend to favour high-quality, large-pack customised products suited to their fee-paying members. Fast-growing bulk snack chains, by contrast, need products that fit the price points and everyday buying habits of lower-tier markets.

These shifts have weakened the traditional one-size-fits-all product strategy. Social media has also accelerated the turnover of consumer trends, raising the risk that a product with a long R&D cycle may miss its market window by the time it launches. To respond to this faster-moving market, Mars Snacking’s digital team worked with the R&D and marketing departments to rethink the company’s approach to product innovation.

As Elva Liu, Digital Director of Mars Snacking, saw it, fragmented demand meant product innovation could no longer be approached as an engineering problem: a search for the right answer through a single, lengthy research cycle. Instead, it had to become a process of rapid co-creation and iteration with the market and consumers. One way to improve the success rate of new products was to generate a broader range of ideas and validate them more quickly in the market. Against this backdrop, Mars began building a digital product innovation platform named Crystal Ball.

STRATEGIC SHIFT ANALYSIS

From the Hero-Product Model to Probabilistic Innovation

Core contrasts in FMCG product innovation models

Traditional Hero-Product Model	DIMENSION	Probabilistic Innovation Model
KEY OBJECTIVE		
Improve forecast certainty before committing resources to a product with broad channel appeal. Big-bet	→	Generate a broader range of ideas and run frequent, low-cost tests to improve portfolio-level success. Probabilistic
INNOVATION CYCLE		
~12 months Focus groups → data validation → concept approval → formulation development → packaging design → mass production	→	Significantly compressed AI-assisted insights → multiple concepts → small-batch trial production → feedback loop
CHANNEL STRATEGY		
Scale a one-size-fits-all product strategy through centralised distribution channels and mass media.	→	Channel-specific customisation: high-quality large packs for warehouse retailers; specific price points and everyday consumption needs for bulk snack chains.
SOURCE OF INSIGHTS		
Specialist market research teams or external agencies; focus groups, in-depth interviews and questionnaires	→	Structured e-commerce data and unstructured social media signals analysed through Crystal Ball; accessible to front-end business teams
COST OF EXPERIMENTATION		
Turning early ideas into concrete product concepts requires substantial time and resources, making failure costly.	→	Generative AI sharply lowers the marginal cost of producing insights, concepts and visual prototypes, making frequent testing more feasible.

Source: Mars Snacking's digital team; "The Consumer Sector in 2024," McKinsey, 2024.

II. Early Data Platform Exploration and Barriers to Adoption

Around 2021, Mars launched Crystal Ball 1.0. The platform marked an initial attempt to build a cross-category data insight architecture. Its logic was to integrate multiple data sources and give R&D staff a broader view of the market.

Crystal Ball aggregated structured transaction and review data from major e-commerce platforms, while also incorporating unstructured trend signals from social media. Its monitoring scope extended beyond Mars's core categories, such as confectionery and chocolate, to include data from fast-iterating categories such as beverages, snacks and health-oriented products. The aim was to draw inspiration for new flavours and ingredients from trends emerging outside Mars's traditional category boundaries.

In practice, however, the platform encountered barriers to adoption inside the organisation. Under the previous model, consumer insights work was handled mainly by specialist market research teams or external consulting firms. Crystal Ball 1.0, by contrast, required R&D and marketing staff to operate the system themselves.

Built around a conventional business intelligence (BI) logic, Crystal Ball 1.0 mainly presented complex charts and dashboards. R&D and marketing staff had to set their own query parameters and have sufficient data literacy to extract usable product concepts from the system.

As Elva later observed, this operating model placed relatively high demands on employees without a data background. It also disrupted established ways of working, creating resistance to adoption. Although data aggregation enriched the information available to the business, Crystal Ball 1.0 could not directly turn that information into business recommendations or product prototypes. It remained largely query-based, requiring business staff to set parameters and analyse the outputs themselves. As a result, its impact on the overall speed of product innovation was limited.

III. Introducing Generative AI and Upgrading the Interaction Architecture

As generative AI and large language models (LLMs) advanced, Mars upgraded the Crystal Ball platform around 2024, moving it into its 2.0 phase. The aim was to change how employees interacted with the system, lower the barrier to use and give the platform greater autonomy in data processing and content generation.

The upgraded Crystal Ball platform included three main functional modules. The first was a Chat BI module for structured data queries. Because business decisions required a high degree of data accuracy, Mars needed to reduce the risk of logical errors, or “hallucinations”, in LLM-generated outputs. Its technical team therefore worked with vendors to add validation rules that checked the LLM’s outputs against the structured database. This allowed business teams to generate market trend reports through natural-language prompts, such as “Extract the characteristics of low-sugar snacks recently popular among young women in tier-one cities”, significantly lowering the technical threshold for data analysis.

The second module, AI Innovator, focused on unstructured small-data signals. Broad e-commerce sales data is often a lagging indicator, while early consumer trends are more likely to emerge in the text and images shared on social media. Mars introduced an enterprise-level agent architecture, incorporating its internal marketing methodology into the system. The system analysed and summarised dispersed text and image content, then distilled potential directions for product concepts. To some extent, this simplified the previous process of manually reading large numbers of posts and turning them into research reports.

The third module supported visual design generation. The platform integrated image generation technology powered by large models and fine-tuned the model using each Mars brand’s visual guidelines and historical design assets. Once business teams had defined the core selling points and product concept, they could instruct the system to generate initial packaging design prototypes aligned with the brand’s visual identity.

Together, these modules brought several previously separate workflows into a single system that supported natural-language interaction—from market trend insight and product concept development to early-stage visual design.

PLATFORM EVOLUTION

The Two-Stage Evolution of Crystal Ball

From BI dashboards to a generative AI three-module architecture

1.0
BI Phase
2021

Data Aggregation and Cross-Category Insights Platform

Integrated structured e-commerce transaction data with unstructured trend signals from social media. The platform moved beyond confectionery and chocolate to include cross-category data from beverages, snacks and health-oriented products, helping R&D teams draw inspiration from broader market trends.

Cross-Category Data Integration | E-Commerce + Social Media Data | BI-Based Analysis Dashboards

Required Data-Analysis Skills | Query-Based with Limited Recommendations | Barriers to Internal Adoption

Advances in generative AI → platform architecture upgrade (2024)

2.0
AI Phase
2024

Generative AI Three-Module Platform

Introduced large language models and image generation technology, replacing complex dashboard queries with natural-language interaction. The platform brought market trend insight, product concept development and visual design generation into a single workflow that front-end business teams could use directly.

Chat BI
Natural-language queries for structured data
Business teams enter natural-language prompts, while validation rules help reduce hallucination risks before the system generates market trend reports.

AI Innovator
Analysis of unstructured social media signals
Analyses text and image content from social media, combines it with Mars's marketing methodology and distils potential directions for product concepts.

Visual Design Generation
On-brand packaging prototype generation
Fine-tunes image outputs using brand-specific visual guidelines, then generates initial on-brand packaging design drafts from product concepts.

Core capability shift: from query-based data retrieval to system-generated recommendations. Insights, product concepts and visual prototypes can be generated within a single natural-language interface, enabling front-end business teams to use the platform without a data-analysis background.

Source: Mars Snacking's digital team; compiled from interview with Elva Liu.

IV. Reshaping the Front End: Shorter Workflows and Channel Co-Creation

The upgraded platform simplified Mars's front-end business processes. Under the previous process, identifying demand and approving product concepts required cross-functional coordination and could take several weeks. After Crystal Ball was introduced, business teams could obtain analytical insights more quickly, while the system generated multiple product concepts in parallel for comparison and screening. According to Mars's internal assessment, the upfront cycle for innovation insight and concept generation was significantly shortened.

Crystal Ball also changed how Mars worked with retail channels. Traditionally, business development followed a largely one-way process: the brand prepared product proposals internally, presented them to retail buyers and then returned to internal teams for revision and refinement if feedback was received. This made the cycle relatively long.

Mars experimented with a more interactive approach in a negotiation with a bulk snack chain in China. The retailer came to the meeting with an initial request for an entry-level New Year gift box. Mars's sales and R&D teams used Crystal Ball to translate the client's requirements—including product positioning, target audience, packaging style and cost structure—into concrete design directions. The platform quickly generated several packaging prototypes for the client to review, each with a different design style and product mix. The retailer's representatives compared the options directly in the meeting and selected one on the spot. In this process, the platform served as a design and analysis aid in the early sales process. This real-time, on-site co-creation model shortened the cycle from aligning on requirements to making a commercial decision.

V. Back-End Manufacturing: Operating a Dual-Track Flexible Supply Chain

Faster front-end concept generation required Mars's manufacturing operations to become more agile. To support frequent market testing and channel-specific customisation, Mars adjusted its production operations and established a dual-track supply chain model.

For core products with stable sales and clearly defined demand, such as M&M's, Dove and Snickers, Mars continued to rely on fully automated factories in China for large-scale production. This allowed the company to keep unit production costs under control while maintaining consistent quality.

At the same time, for regional test and channel-customised concepts rapidly generated through Crystal Ball, Mars used a flexible production track to provide additional capacity where needed.

Dove's tea-flavoured chocolate, sold in a dark green tin, showed how this dual-track model worked in practice. After Crystal Ball detected interest in tea-flavoured chocolate and generated a product concept, Mars did not move directly to large-scale production. Instead, the team used a flexible production line for small-batch trial production and supplied the product exclusively to Sam's Club to gather market feedback.

After launch, Crystal Ball collected data on sell-through rates, inventory turnover and consumer feedback, then fed the data back into Crystal Ball's underlying data layer for analysis. This operating chain—AI-assisted concept generation, small-batch physical trial production, market data feedback and scaled production—created a closed-loop validation process.

OPERATING MODEL AND STRATEGIC TAKEAWAYS

Dual-Track Manufacturing and the AI-Assisted Innovation Flywheel

How digital generation and agile physical production reinforce each other

DUAL-TRACK SUPPLY CHAIN MODEL

Standard Track: Large-Scale Production		Flexible Track: Small-Batch Trial Production
• Applicable products: core products with stable sales, such as M&M's, Dove and Snickers • Production approach: large-scale production in fully automated factories in China • Key objective: keep unit production costs under control while maintaining consistent quality **Priorities:** scale efficiency · lower unit production costs · consistent quality	Dual Track	• Applicable products: regional test and channel-customised concepts rapidly generated through Crystal Ball • Production approach: flexible production capacity where needed • Key objective: validate hypotheses through lower-risk trial production and real market feedback **Case:** Dove tea-flavoured chocolate → small-batch trial production → exclusive supply to Sam's Club for market testing

AI-ASSISTED CLOSED-LOOP VALIDATION FOR PROBABILISTIC INNOVATION

Four-step Validation Process CLOSED LOOP

1. **AI-Assisted Concept Generation**
Crystal Ball generates market insights, product concepts and visual prototypes for comparison and screening.

2. **Small-Batch Trial Production**
Flexible production lines produce small batches for selected retailers, such as Sam's Club, to test market response.

3. **Market Data Feedback**
Sell-through rates, inventory turnover and consumer feedback are fed back into Crystal Ball for analysis.

4. **Scale-Up**
Validated concepts can move towards scaled production, while market data strengthens the platform's underlying data layer.

Co-creation in practice: In a negotiation with a bulk snack chain in China, Mars's sales and R&D teams used Crystal Ball to translate the client's requirements into concrete design directions. The platform quickly generated several packaging prototypes for review, each with a different design style and product mix. The retailer compared the options in the meeting and selected one on the spot, shortening the path from aligning on requirements to making a commercial decision.

KEY STRATEGIC TAKEAWAYS

01
Reduce the Cost of Experimentation
Generative AI sharply lowers the marginal cost of producing insights, concepts and visual prototypes, allowing companies to move from long-cycle bets on a single blockbuster product to frequent, low-cost probabilistic tests.

02
Match Digital Generation with Physical Responsiveness
Computing power alone does not create business value; it must be matched by a manufacturing system agile enough to turn digital concepts into market-tested products.

03
Move Decision-Making Closer to the Front End
Natural-language interaction gives front-end business teams direct access to data and generative capabilities, making Mars less dependent on centralised back-end analysis when responding to business opportunities.

Source: Mars Snacking's digital team; compiled from case interview.

VI. Organisational Adjustment and Routine AI Use

As Crystal Ball became part of day-to-day operations, Mars also adjusted its management routines, organisational roles and ways of working around the platform. During the rollout, Mars Snacking's digital team adopted a guiding principle: business teams would lead, while the digital team would support them. This meant that new features had to be shaped around the real needs of R&D, marketing and other business teams, ensuring that technology development remained anchored in specific business use cases.

As the platform became easier to use, Mars also changed its internal standard operating procedure (SOP) for product review. In the past, R&D teams would typically submit one or two mature product concepts after a long preparation period. After Crystal Ball was introduced, management adjusted its output expectations. Within the same timeframe, teams were expected to use AI tools to generate and submit multiple alternative concepts for comparison and evaluation. This procedural requirement helped embed the platform into everyday business operations.

Collaboration across teams also became easier. Previously, when brand teams briefed external advertising or packaging design agencies, text-based briefs often left room for ambiguity, leading to misunderstandings and repeated rounds of revision. With the new platform, internal teams could first generate an initial visual proposal and send it to the agency as a concrete point of reference, while specifying the details that needed to be refined. This prototype-based approach reduced communication friction in early-stage collaboration.

As a result, employees could devote more attention to formulating business hypotheses, reviewing outputs for brand compliance and coordinating with retail channel partners.

VII. Key Takeaways: Reducing the Cost of Experimentation and Synchronising Digital Generation with Physical Manufacturing

The evolution of Mars's Crystal Ball platform reveals the deeper commercial logic behind how traditional manufacturers can use AI to respond to volatile markets.

First, AI changes the economics of experimentation in business innovation. In the traditional FMCG development model, developing a new product concept and turning it into something tangible requires substantial time and resources. This meant companies had to be highly confident in their forecasts before committing resources to a project. Generative AI sharply lowers the marginal cost of producing market insights, product concepts and visual prototypes. As a result, companies can move from making long-cycle bets on a single blockbuster product to running frequent, low-cost probabilistic tests.

Second, digital generation only creates value when physical production can keep pace. Software can generate dozens of product concepts in minutes, but that speed has little commercial value if the production system still takes months to complete trial production and evaluation. Mars therefore paired its AI tools with small-batch trial production lines and adjusted the way production tasks were allocated across its supply chain. The case shows that computing power alone does not create business value; it must be matched by a manufacturing system agile enough to turn digital concepts into market-tested products.

Third, lowering the barrier to interaction pushed decision-making closer to the front line. By replacing complex dashboard queries with natural-language interaction, Crystal Ball made data and generative capabilities directly accessible to front-end business teams. Once sales and R&D staff could work directly with those capabilities, Mars's response to business opportunities became less dependent on centralised back-end analysis and more rooted in distributed, agile collaboration at the front end.

In this sense, effective technology adoption is not simply about introducing new tools. It requires companies to redesign the organisational rules, channel relationships and supply chain processes that allow those tools to create value.

Endnote

[1] McKinsey & Company, "The consumer sector in 2024: Navigating market fragmentation," McKinsey, 2024, https://www.mckinsey.com/industries/consumer-packaged-goods/our-insights/the-consumer-sector-in-2024

CASE 2

Semir: Embedding AI and Building a Second Growth Curve

Semir is a large Chinese apparel company with brands including Semir and Balabala. Founded in 1996, it launched Balabala in 2002 and was listed on the A-share market in 2011. For a mature company like Semir, AI is significant not because it adds new tools, but because it addresses long-standing structural tensions in the apparel industry.

In practice, these tensions manifest in several recurring challenges. Product development cycles are long, even as consumer preferences change rapidly, making it hard for traditional planning to keep pace. Visual shoots and content production are costly, yet e-commerce platforms demand more frequent updates and stronger conversion. Companies must also manage expanding SKU ranges alongside mounting inventory pressure—balancing product variety with per-item efficiency. Finally, large organisational scale and complex operational chains mean that, without sustained momentum and the right mechanisms, AI initiatives can easily remain stuck in isolated pilots.

In recent years, Semir has moved beyond standalone applications. Instead, it has progressively integrated AI across the product launch cycle—linking market insights, product planning, design, visual content, marketing, customer service, and operations. The aim is to transform AI from a supporting tool into a new foundation for how the business operates.

I. From Experimentation to End-to-End Embedding

Semir began exploring AI at an early stage. By the end of 2022, it was already experimenting with tools such as GPT and MidJourney, starting with small-scale, isolated trials. In 2023, the company attempted model training, but the technology was not yet mature enough for these efforts to translate into viable commercial applications.

By 2024, as generative AI tools such as Doubao, Jimeng, and Huiwa became more usable, AI initiatives moved back onto a commercial footing. After February 2025, progress accelerated significantly, entering a phase of full-scale rollout. These early experiments were gradually brought together, forming a more continuous, end-to-end workflow spanning market insights, consumer segments, supply chains, operations, and customer service.

Throughout this process, Semir developed a clear prioritisation of AI value: revenue growth first, followed by cost reduction, and then efficiency improvement. The reasoning is straightforward—cost savings and efficiency gains have natural limits, whereas growth does not. For a large apparel company, AI adds limited value if it merely reduces time spent on visual design, copywriting, or data processing. It becomes truly strategic only when applied to core business levers such as product success rates, conversion, traffic efficiency, and turnover.

Guided by this logic, Semir structured its AI efforts within a clear framework. At the top, AI supports growth and business model innovation. In the middle, it delivers operational and customer value. Across day-to-day operations, it is embedded in product development, content, customer service, marketing, supply chain, and operations. This is supported by underlying tools, data, platforms, and organisational alignment.

By 2025, according to the company's own reporting, AI had generated over RMB 100 million in incremental revenue, delivered cost savings of tens of millions, and improved efficiency by around 20%. The revenue uplift came mainly from virtual host livestreaming, intelligent ad placement, and improvements in click-through and conversion. Cost savings were concentrated in visual production and parts of its software and hardware development. These results suggest that AI at Semir has moved beyond proof of concept and is delivering measurable, attributable business outcomes.

STRATEGIC FRAMEWORK ANALYSIS

Semir's AI Rollout Framework: A Growth-First Strategy and Four-Stage Evolution

From isolated pilots to full value-chain deployment, delivering measurable results

AI VALUE PRIORITIES

1 Revenue Growth

AI is applied to **core business drivers**—product success, sales conversion, and traffic utilisation—rather than peripheral support functions.

Virtual host livestreaming generated over RMB 10 million in incremental revenue. Intelligent ad placement improved conversion rates.

With no natural ceiling, growth is the top priority.

2 Cost Reduction

AI replaces **high-cost activities** such as visual production and software and hardware R&D, reducing both fixed and variable costs.

Cost savings reached tens of millions of RMB. Visual production costs declined significantly.

Cost reduction has limits, ranking below growth but above efficiency gains.

3 Efficiency Gains

AI takes over **repetitive tasks**—image creation, copywriting, and data processing—freeing staff to focus on higher-value decisions.

Overall efficiency rose by about 20%. Process automation covered multiple scenarios.

Efficiency gains are valuable but limited on their own.

FOUR-STAGE EVOLUTION OF AI APPLICATION (2022–2025)

2022 Pilots

Isolated pilots
Experimentation with tools such as GPT and Midjourney in small-scale trials
Proof of concept

2023 Constraints

Model constraints
Model training limited by immature technology, preventing commercial deployment
Technical bottleneck

2024 Maturity

Tool maturity
Tools such as Doubao, Jimeng, and Huiwa became usable, bringing AI back into commercial use
Commercial rollout

2025 Scale

Full-scale rollout
Rapid expansion across insights, planning, visual content, marketing, customer service, and supply chain
Scaled deployment

RMB 100 million+
Incremental revenue
Virtual host livestreaming | Intelligent ad placement
Higher click-through and conversion rates

RMB 10 million+
Cost savings
Reduced visual production costs
Lower software and hardware R&D costs

~20%
Efficiency gains
Expanded process automation across multiple business scenarios

Source: Semir Group internal evaluation data; 2025 performance attribution analysis

II. Virtual Host Livestreaming: From Extended Hours to Incremental Revenue

Virtual host livestreaming was the first AI use case at Semir to form a clear, self-sustaining business model. In 2025, it generated tens of millions of yuan in incremental revenue, which the company tracks separately. Crucially, this revenue did not come from improving existing livestreams with human hosts, but from extending coverage into hours that were previously unavailable—or difficult to sustain. Because virtual hosts can operate 24/7, most of the uplift came from expanded airtime, not higher conversion within existing streams.

This changes how livestream supply is organised. In the past, supply was constrained by host schedules, time availability, and on-site arrangements. With virtual hosts, it is increasingly determined by system capacity. This allows off-peak periods—such as late-night and low-traffic hours—to be brought into regular operations. Many long-tail time slots that were previously difficult to sustain are now commercially viable. In this sense, virtual host livestreaming functions primarily as a capacity expansion tool. It extends operating hours and broadens the business's ability to capture traffic and transaction opportunities.

The second shift is in cost. Longer livestream hours no longer require extra hosts or full teams. That turns what was once a high-cost model into a low marginal-cost one, making previously unviable time slots workable. As a result, success is no longer measured simply by comparing single-session GMV with that of top human hosts. Instead, it is evaluated on whether it can, at lower cost, reliably cover long-tail traffic, fill supply gaps, and convert previously unserved demand into incremental sales.

The third shift is in interaction. Virtual host livestreaming is no longer limited to scripted narration. By 2025, visual realism—from lighting to fabric texture—was approaching that of real hosts. At the same time, stronger interactivity allows systems to respond to common user queries on coupons, links, and product details. This is critical in apparel, where conversion relies on real-time explanation of style, sizing, fabric, and pricing. The ability to generate attributable incremental revenue indicates a move beyond basic broadcasting toward transaction-level interaction.

At the same time, Semir does not treat virtual host livestreaming as a full replacement for human hosts. Even as the technology improves, human oversight remains essential at key stages, including planning, content review, validation, and risk control. Even small deviations in how products and the brand are presented can directly affect commercial outcomes. In practice, virtual host livestreaming is therefore not "unmanned broadcasting," but scaled livestreaming under human oversight.

CORE BUSINESS SCENARIO ANALYSIS

The Triple Impact of Virtual Host Livestreaming and the Forward Shift in Product Planning

How AI reshapes front-end sales supply and early-stage product decision-making in apparel retail

RMB 10m+

Projected incremental revenue in 2025 **from virtual host livestreaming**, coming primarily from previously untapped time slots rather than incremental gains within existing human-hosted streams

24/7 operational capability
Covering late-night, off-peak, and other low-traffic periods

THREE STRUCTURAL SHIFTS IN VIRTUAL HOST LIVESTREAMING

1 Supply Model Shift

Past
Supply constrained by host schedules, time slots, and venues; off-peak periods difficult to sustain

Now
Supply driven by system uptime, incorporating off-peak, late-night, and long-tail hours

Role: Scalability tool expanding operating hours and traffic capture

2 Cost Structure Shift

Past
Extending hours required additional hosts and support teams, resulting in high incremental costs

Now
Extended hours operate at low marginal cost, making long-tail periods commercially viable

Role: Incrementality engine capturing long-tail demand at low cost

3 Interaction Upgrade

Past
Limited to scripted, repetitive broadcasting with little real-time interaction

Now
Improved realism supports responses to frequent queries on coupons, product links, sizing, and more

Role: Transactional capability enabling real-time, commerce-driven interaction

PRODUCT PLANNING WORKFLOW: FROM COMMIT-FIRST TO VALIDATE-THEN-SCALE

Traditional Workflow	→ Post-AI →	AI-Enabled Workflow
• Relied on limited platform data, reviews, and competitor inputs; insights were slow and fragmented		• Aggregates multi-platform reviews, VOC, and competitor data to identify opportunity categories and bestseller features
• Market testing only after physical samples and formal photo shoots		• AI-generated product visuals, virtual try-ons, and product detail pages enable testing before physical production
• Peak-season product development cycles of up to 240 days		• In some scenarios, testing cycles reduced to around 15 days
• Commit first: scale production, then validate through market sales		• Validate first: market feedback informs scaling, improving product–market fit

240 days
Traditional peak-season product development cycle

↓ Cut by ~94%

~15 days
Average AI-assisted testing cycle in key scenarios

Source: Semir Group internal assessment; 2025 virtual host livestreaming attribution analysis

III. From Market Insight to Product Planning: Moving Decisions Upstream

Unlike virtual host livestreaming, which operates at the sales front end, Semir's use of AI in product planning and style testing reworks how upstream decisions are made. A core challenge in the apparel industry is that market trends shift quickly, while internal information gathering and processing struggle to keep pace. In the past, product planning relied on a limited set of platforms, a small number of reviews, and a narrow competitor set, resulting in decisions that were both slow and incomplete.

With AI, market insight is no longer confined to single platforms or small samples. Instead, Semir can draw simultaneously on multiple platforms and channels, analysing a wide range of inputs—from reviews and usage contexts to user pain points and competitor performance. These inputs are categorised, compared, and used to generate recommendations on opportunity categories, bestseller characteristics, and specific product decisions—from core product attributes such as fabric, colour, and functionality to sizing and go-to-market approach.

More importantly, these capabilities extend beyond automated reporting into the core of product planning. They are applied end to end—from insight generation and opportunity identification through to design, virtual imagery, copywriting, and product page development. These outputs then feed directly into launches and front-end sales. In other words, AI no longer supports isolated tasks; it underpins a continuous workflow—from insight to planning, visualisation, and front-end testing.

The most critical shift in this process is the earlier timing of product testing. In the past, a peak-season product could take up to 240 days from development to launch; now, in some cases, the average testing cycle has been reduced to around 15 days. The key change is not simply faster image generation, but the ability to produce product visuals, virtual model imagery, product pages, and front-end display content before physical samples and formal shoots are ready. This allows the company to test market response earlier and make production decisions based on actual demand.

This shifts product decision-making from “commit first, validate later” to “validate first, scale later.” At the same time, this does not lead to uncontrolled SKU expansion. Instead, AI improves single-item efficiency: by enabling broader sampling, faster feedback, and lower trial-and-error costs, it increases alignment between each product and consumer demand while reducing ineffective SKUs. This improvement is already reflected in internal feedback.

IV. Visual, Copy, and Process Automation: Turning Tasks into System Capabilities

Visual content has long been a high-cost component of apparel retail. Semir has applied AI extensively to virtual model imagery, outfit swapping, image compositing, short video production, and product page content creation. Work that once relied on location shoots, travel, set design, and post-production can now be completed using AI. On platforms such as Tmall, the company uses 15-second dynamic videos to showcase products, with AI-generated content increasingly used in categories such as technical outerwear. These changes have directly reduced production costs while significantly increasing the frequency of content updates.

Copywriting has also been reorganised. In the past, producing full product page copy was time-consuming; today, Semir has embedded its knowledge base into AI agents. Staff can upload basic materials and parameters, after which the system quickly generates the required copy and supports export. Most product page copy is now produced in this way. At the same time, advances in image generation tools have lowered the barrier to use. Tasks that previously required specialised design software and technical expertise can now be carried out by a broader range of staff through AI agents.

Behind this, process automation has also progressed in parallel. A growing share of routine, rule-based operational tasks—such as pricing, inventory control, and performance monitoring—are now handled through automated workflows. In this context, AI acts as a continuously operating digital assistant: it can automatically gather data, perform initial processing and analysis, and deliver results directly to relevant teams through internal platforms. In turn, AI not only reduces manual effort, but also reshapes how tasks are triggered and how workflows are executed.

V. Organising for Scale: Turning Tools into Business Actions

Semir's AI adoption has moved from isolated pilots to scaled application across multiple business scenarios. This shift has been driven not only by improvements in tool availability, but also by strong organisational push. Over the past three years, the company has incorporated AI into its formal transformation agenda, emphasising that it is not a standalone system built from scratch, but a means of supporting existing strategy and execution. Organisational coordination therefore becomes critical: for a company of this scale, effective implementation depends on both top-down and bottom-up alignment.

In practice, Semir has not confined AI to the technology team, but has steadily extended it into the business through platforms, training, consulting, and ongoing implementation support, including customised development and AI agents. Functions such as merchandising, visual content, customer service, livestreaming, and operations are each developing their own practical applications. AI is no longer just an additional tool used by a small number of roles, but is becoming part of day-to-day business activity.

At the same time, the role of people within the system has not diminished. Whether in product planning, design, visual validation, and the review of customer service scripts and livestream content, final responsibility remains with human operators. AI expands sample size, shortens cycles, takes on repetitive tasks, and amplifies organisational capacity; judgement, validation, risk control, and accountability continue to rest with people. This division of roles is a key reason why Semir has been able to scale AI across multiple business scenarios relatively quickly.

VI. From Internal Use to External Offering: The Launch of SENCHUANG QIRUI

Another key development in the Semir case is the shift from internal use to external deployment of AI capabilities. In March 2025, Semir established SENCHUANG QIRUI as a standalone subsidiary, initially incubated as an internal venture. Its services focus on AI-driven automation across e-commerce operations, visual content, product planning, and customer-facing functions, delivered through training, consulting, ongoing implementation support, customised workflow development, and AI agent development.

The establishment of this subsidiary shows that Semir is moving beyond treating AI as an internal support function. The methods, processes, and tools developed across product planning, market insight, virtual host livestreaming, visual generation, customer service, and automated operations are now being consolidated into solutions for external clients.

By 2025, SENCHUANG QIRUI had served more than 30 clients, most with annual revenues between RMB 100 million and 1 billion. Around 70% are in the apparel sector, with others spanning home appliances, furniture, mother and baby products and dairy. Internally, the company also uses market-based pricing, treating Semir itself as a client.

This marks a shift in how these capabilities are positioned. What was previously embedded in internal processes, management practices, and operational roles is now being made explicit, modular, and productised. Through a dedicated entity, these capabilities are turned into training, consulting, implementation support, and AI-driven tools and agents that can be sold, replicated, and delivered.

Semir is therefore no longer only an apparel retailer, but is also becoming a provider of AI-enabled business capabilities.

STRATEGIC EVOLUTION ANALYSIS

From Internal Application to Capability Spillover: How AI Creates a Second Growth Curve

How Semir turns internal AI deployment into scalable, market-facing capabilities

THREE-STAGE PATH OF AI CAPABILITY SPILLOVER

1 Internal Capability Accumulation

Retail operations as a live testing ground

- Digital human livestreams generate trackable incremental revenue
- AI-enabled product planning compresses testing cycles
- Visual generation and workflow automation scale across operations
- Results consolidate into repeatable methods

Internal capability base

2 Capability Externalisation and Productisation

From experience to systemised offerings

- Operational know-how is codified into deliverable solutions
- Workflows and agents are modularised and standardised
- Training, consulting, and implementation form a service system
- A dedicated subsidiary is incubated (March 2025)

Market-ready capabilities

3 Empowering New Business Externally

SENCHUANG QIRUI scales beyond the core business

- Serves 30+ external clients with deep engagement
- Apparel dominates (~70%), with expansion into home appliances, maternal and infant products, and dairy
- Internal units are charged on a market-based basis
- The firm evolves from retailer to AI capability provider

Second growth curve

SENCHUANG QIRUI: THE COMMERCIAL VEHICLE FOR CAPABILITY SPILLOVER (FOUNDED MARCH 2025)

SENCHUANG QIRUI
Independent subsidiary of Semir Group
AI-enabled e-commerce solutions provider

30+
Deep-engagement clients

RMB 100M–1BN
Typical client size

~70%
Apparel client share

Service Portfolio
- End-to-end AI enablement for e-commerce operations
- Visual content generation at scale
- Data-driven product planning and insights
- Customer service QA and process optimisation

Service Delivery Model
- Capability-building through onsite training
- Diagnostic consulting and solution design
- End-to-end implementation support
- Custom workflow and agent development

Industries Served
- Apparel (~70%, including children's wear and casual wear)
- Home appliances & furniture
- Maternal and infant products, and dairy
- Other consumer goods categories

Business Model Innovation
- Market-based pricing, including for Semir's internal units
- Semir is treated as an external client
- Capabilities designed to be replicable, marketable, and deliverable
- Operates independently as an internal startup

STRATEGIC INSIGHT: AI AS THE ENGINE OF SEMIR'S SECOND GROWTH CURVE

The Evolution of AI—from Efficiency Tool to Growth Engine STRATEGIC INSIGHT

Tool → Growth Driver → Infrastructure → Generative Mechanism

AI as an efficiency tool
Automates repetitive tasks such as content and visual production

AI as a growth driver
Integrates into core operations, directly shaping conversion rates and product success

AI as business infrastructure
Connects the full value chain—from insight and planning to content, marketing, and service

AI as a second-curve generator
Internal capabilities spill over into independent new business models

Source: Semir Group; SENCHUANG QIRUI company data; 2025 business data compilation

VII. Key Insight: AI as a Second Growth Curve

The key takeaway from the Semir case is not how widely AI has been applied across the business, but the fact that these capabilities have not remained confined to internal use. Virtual host livestreaming, AI-enabled product planning, visual generation, and process automation first produced measurable results within the company's core retail operations. Building on this, Semir went on to establish an independent subsidiary to offer these capabilities externally, turning what were originally internal tools into replicable, marketable products and services.

The role of AI at Semir has accordingly expanded from optimising retail operations to enabling new business forms. It is no longer just a tool for efficiency or growth, but a mechanism for generating a second growth curve. The Semir case shows how a company can use its own operations as a testing ground, consolidate proven AI capabilities and extend them into new business models.

Part II: Deep System Integration and Central Orchestration—Establishing the Methodology and Advancing Organisational Evolution

Case 3: **Anker Innovations: From AI Tools to an AI Operating System**

CASE 3

Anker Innovations: From AI Tools to an AI Operating System

Anker Innovations (hereinafter “Anker”) is a global consumer electronics company with a highly international operating model. Founded in 2011, the company has built a multi-category business spanning smart power, home automation, and intelligent audio-visual products, serving markets worldwide. In 2024, Anker reported total revenue of RMB 24.71 billion (US$3.6 billion), up 41.14% year on year, and net profit attributable to shareholders of RMB 2.114 billion (US$310 million), up 30.93%.

For a business operating across categories, regions, and platforms, incremental efficiency gains matter—but they are not decisive. What ultimately determines competitiveness is whether core functions can be integrated into a coordinated system—namely R&D, marketing, customer engagement, supply chain, and organisational operations.

It is in this context that Anker’s investment in AI should be understood. Rather than rolling out standalone efficiency tools, the company is embedding AI into its operational backbone, driving a shift from a human-driven model to an AI-driven system.

Once AI moves beyond answering questions and begins to operate within defined boundaries—as a closed loop that continuously senses, decides, executes, and improves—it ceases to be a tool and becomes part of the operating system itself.

Anker’s experience reflects this transition. AI is no longer confined to assisting employees with discrete tasks; it is increasingly embedded in live business workflows, taking on a growing share of execution responsibilities previously carried out by human labour.

I. A Three-Phase Evolution: From Capability Exploration to Organisational Integration

Effective AI transformation does not come from rolling out AI everywhere at once, but from adopting it in stages as capabilities mature. At Anker, this meant introducing AI in clearly defined phases rather than as a single rollout.

The company moved through three successive stages. In 2023, it focused on exploring AI capabilities and achieving targeted breakthroughs. In 2024, it began embedding AI more deeply into business workflows. By 2025, it had shifted towards workforce and organisational transformation, laying the foundation for a more AI-native operating model.

This progression reflects a simple view: not all AI capabilities mature at the same pace. Rather than rolling AI out broadly, companies should focus on a small number of high-maturity, high-ROI use cases, deepen adoption there, then expand from that base.

The first phase focused on making AI visible across the organisation. Anker prioritised three relatively mature domains: marketing content generation, customer service, and software engineering—where early results could build organisational trust.

In marketing, the company launched a centralised AIGC content platform in 2023, bringing together internal and external capabilities to automate large-scale production of digital avatars, short videos, and other marketing assets. Efficiency in key scenarios improved by more than 60%.

In customer service, AI began handling portions of global email and live chat interactions.

In R&D, the company focused from the outset on a hard operational metric: the proportion of AI-generated code ultimately committed to the codebase. Rather than evaluating AI based on how it performs in demonstrations, Anker measured its impact through production outcomes.

The underlying logic was straightforward: start with a small number of high-impact use cases so the organisation could first see AI in action, then build trust, and only then expand adoption.

The second phase focused on integrating AI directly into business workflows and enterprise systems.

If 2023 was about providing tools for employees to use, then 2024 was about embedding AI into the flow of work itself—so that employees would use AI as part of their daily work, often without even realising it.

Anker described this shift as moving from “AI-assisted” to “AI-driven” operations. To support this transition, the company advanced along two fronts.

First, it developed an enterprise-wide AI agent platform, which at its peak in 2024 supported more than 600 active AI agents across multiple business functions and stages of the value chain.

Second, it re-engineered its AIGC platform by embedding core organisational knowledge—such as brand guidelines, user profiles, design standards, competitive information, channel characteristics, and style templates. As a result, the platform not only generated content faster, but also produced outputs that were more accurate, more consistent, and better aligned with downstream distribution.

The third phase shifted focus to people and organisational design.

By 2025, Anker’s view was clear: as AI adoption deepens, the key constraint is no longer model capability, but whether people, processes, systems, and organisational structures are ready for it.

The company therefore focused on building AI-friendly and AI-native conditions, driving coordinated changes across four areas: employee capabilities, organisational structure, IT infrastructure, and knowledge management.

At this stage, the transformation extends beyond tools to the enterprise operating system itself.

STRATEGIC EVOLUTION ANALYSIS

Anker's Three-Stage AI Transformation: From Capability Exploration to System Integration

Focus on high-maturity, high-ROI use cases first, then scale

A THREE-PHASE EVOLUTION PATH (2023–2025)

2023 · Phase One
Capability Exploration
Targeted Breakthroughs
Make AI visible across the organisation

Strategic Priority: Building Organisational Trust

- Started with three high-maturity, high-ROI domains
- **Marketing:** Launched an AIGC content platform, improving efficiency in key scenarios by over 60%
- **Customer Service:** Introduced AI into email and live chat workflows, with around 35% of requests fully automated
- **R&D:** Tracked AI code adoption, with around 35% of code generated by AI and accepted into the codebase
- Started small—letting the organisation see AI in action and build trust

Model: Make AI tools accessible for employee use

2024 · Phase Two
Business Integration
System Embedding
From AI-assisted to AI-driven operations

Strategic Priority: Embedding into Business Systems and Processes

- Embedded AI directly into business systems as part of day-to-day work
- Built an enterprise-wide AI agent platform, at one point supporting **600+ active AI agents**
- Rebuilt the AIGC platform with embedded brand standards, user profiles, design guidelines, and other core knowledge
- Customer service automation reached 45–50%
- Developed an intelligent advertising system, with around 20% of product campaigns fully managed end to end

Model: Embed AI into processes to enable autonomous system execution

2025 · Phase Three
Workforce Transformation
AI-Native Organisation
Build AI-ready and AI-native conditions

Strategic Priority: Full Alignment of Organisation and Talent

- The key constraint shifted from model capability to organisational readiness—across people, processes, systems, and structures
- Launched the **NEW Talent Initiative** to shift employees from AI users to AI builders
- In selected scenarios, customer service automation reached **80–85%**, supported by AI voice calling
- In selected new projects, **95%+ of code** was generated by AI
- Drove AI-enabled transformation of business products alongside full alignment of IT infrastructure

Model: AI as a stable part of the operating system

CORE EVOLUTION LOGIC

Evaluation Criteria
Hours saved by employees

→ Upgraded Evaluation Criteria
Number of workflows AI can run independently

→ Ultimate Goal
From using AI to being reshaped by AI

600+	80–85%	95%+	3.5 million+
Number of active AI agents on the enterprise platform	Customer service automation rate in selected scenarios (2025)	AI code adoption rate in selected new projects	Cumulative marketing assets generated by the AIGC platform

Source: Internal assessments, 2024 Annual Report and interviews (Anker Innovations)

II. From Task Efficiency to Closed-Loop Takeover: Three Key Use Cases

The key measure of AI value is no longer how much time it saves employees, but how much of a workflow it can take over independently. Anker's use cases illustrate this shift clearly.

1. Customer Service: From Assisted Responses to Workflow Takeover

Among Anker's use cases, customer service is one of the clearest examples of deep AI integration into core workflows. With around 95% of its business serving global markets, its support operations span multiple time zones and are high in volume and standardised, delivered primarily through email, live chat, and phone.

In 2023, Anker introduced AI agents into email and live chat workflows, enabling around 35% of customer requests to be handled autonomously. By 2024, this had risen to 45–50%. By 2025, with deeper integration across knowledge bases, internal systems, and agent capabilities, AI takeover reached 80–85% in selected scenarios. Combined with AI voice calling, this enabled a 24/7 global service model.

Routine service tasks—such as refunds, returns, case handling, and routine enquiries—are increasingly executed directly by the system. Human agents are correspondingly freed to focus on more complex and higher-value interactions.

This progression reflects a clear shift in focus. Traditional customer service automation was designed to assist human agents and accelerate response times. Anker's approach, by contrast, aims to systematise the handling of repetitive and standardised scenarios, allowing AI to directly take over parts of the service workflow.

In other words, AI is no longer positioned alongside human agents to offer suggestions. It is increasingly taking on execution responsibility, becoming the primary actor within parts of the service process.

As a result, the evaluation criteria have also shifted: the focus is no longer on how much time employees save, but on how many standardised workflows can be fully executed end to end without human intervention.

2. Marketing: From Content Generation to Workflow Automation

Marketing was one of the earliest domains in which Anker validated the commercial value of AI. In 2023, the focus was on improving content production efficiency. Through an AIGC content platform that integrates external capabilities, the company enabled the rapid generation of images, videos, digital avatars, and other marketing assets, significantly improving productivity across marketing and design teams.

By 2024, the focus had shifted from content generation to advertising execution and strategy optimisation. Anker developed an "intelligent advertising" system designed to create an end-to-end loop—from strategy design and value proposition analysis through to asset generation, campaign optimisation, performance feedback, and post-campaign review.

According to internal interviews, after sustained A/B testing, the system has outperformed human execution on key metrics such as win rate and ROI. Today, around 20% of product campaigns are fully managed end to end by the system, while in roughly 80% of cases, AI generates the initial strategy, which is then reviewed and refined by human teams.

More importantly, Anker does not view marketing AI simply as a way to generate more content. Instead, it has built a system that embeds enterprise knowledge and brand standards directly into the workflow. Beyond model capabilities, the platform incorporates brand guidelines, user profiles, design standards, platform-specific requirements, and template systems.

As a result, what AI produces is not generic content, but production-ready assets that align with internal standards and can flow directly into the company's global asset management and distribution system.

Today, more than 95% of design, R&D, and marketing materials are generated through this platform. To date, it has produced over 2.2 million marketing assets, covering around 90% of all marketing content production.

What this reflects is not simply faster content production, but a restructuring of the marketing workflow itself under AI.

3. R&D: From Code Assistance to Workflow Orchestration

In R&D, Anker's experience shows what AI looks like once it is embedded in core operations. Since 2023, the company has tracked a key internal metric: AI code adoption—the share of AI-generated code that is ultimately accepted and committed to the repository. This figure reached around 35% early on. Today, in some new projects, more than 95%—and in certain cases nearly all—code is generated autonomously by AI.

More importantly, AI's role has expanded well beyond code generation. It now extends across the full development lifecycle—from product requirements and feature design through to project management, testing, and operations.

In other words, AI is no longer involved in isolated steps. It operates across multiple stages of the workflow, forming a coordinated network of agents.

The internal evolution follows a clear path: from single-task agents, to role-level agents, to process-level coordination across multiple agents and onwards towards higher levels of business autonomy.

The real step change lies not in better outputs at the task level, but in AI's ability to operate within defined boundaries—perceiving context, planning actions, coordinating resources, and reliably driving workflows forward.

IN-DEPTH ANALYSIS OF CORE SCENARIOS

AI Takeover Across Three Core Scenarios: From Assistance to Autonomy

AI value is measured not by time saved, but by how much of the workflow it can take over

Customer Service

24/7 global service

From **assisting human agents to respond faster to AI takeover of standardised service workflows**—refunds, returns, case handling, and routine enquiries are handled by the system, freeing up human agents for more complex, high-value work

End-to-end takeover

CURRENT AUTOMATION RATE
80–85%
2025 (selected scenarios)

2023: AI automation rate ~35%

Handled email and live chat, helping build initial trust

2024: AI automation rate 45—50%

Integration with knowledge bases and systems expanded scope

2025: AI automation rate (selected scenarios) 80—85%

AI voice calling enabled 24/7 global service coverage

Marketing

Global multi-category content and campaigns

From **faster content production** to **an automated end-to-end workflow covering strategy, creative generation, campaign optimisation, and performance review**—embedding brand standards and user profiles to produce assets ready for direct use in the asset management system

Process intelligence

ASSET COVERAGE RATE
90%
Marketing production coverage

- Over 2.2 million assets generated, covering 90% of total production
- Around 20% of product advertising fully automated end to end
- Around 80% of strategies generated by AI, then reviewed by humans
- A/B testing shows higher win rates and ROI than human-managed campaigns
- 95%+ of design, R&D and marketing assets generated on the AIGC platform
- Built-in brand guidelines, competitive intelligence, and channel characteristics ensure more consistent output and smoother distribution

R&D

End-to-end software development

From **code assistance** to **an end-to-end development workflow covering requirements, development, project management, testing, and operations**—AI operates continuously across stages, forming a linked network of agents

End-to-end workflow

CODE ADOPTION RATE
95%+
Current stage (selected new projects)

2023: AI code adoption rate (AI-generated code committed to the repository) ~35%

Tracked the share of AI-generated code committed to the repository, validating value through production outcomes

Current stage: AI code adoption rate (selected new projects) 95%+, approaching 100%

AI extends beyond coding to the full development lifecycle, including requirements, project management, testing, and operations

Sources: Anker Innovations internal assessment; interview materials; 2024 annual report

III. Building AI Advantage: Turning Tacit Knowledge into Usable Assets

The essence of AI capability lies not in stronger models, but in whether an enterprise can turn its proprietary data, business rules, and feedback into a decision system that AI can use.

The real value of Anker's approach lies not in the number of use cases or tools, but in its early recognition that enterprise AI advantage comes not from the model itself, but from a combination of context, constraints, rules, and feedback.

Whether embedding brand standards, user profiles, and design guidelines into its AIGC platform, or integrating knowledge bases, workflows, and feedback loops across customer service and advertising systems, Anker is doing the same thing: turning tacit knowledge—once scattered across employee experience, business processes, and organisational routines—into structured, reusable data assets.

As differences between models continue to narrow, competitive advantage no longer comes from access to a particular model, but from the ability to build a system that continuously integrates proprietary data, domain know-how, and feedback.

Without this, even powerful agents remain surface-level tools. With it, they begin to function as true digital workers embedded in business operations.

In this sense, Anker is not simply using AI to improve efficiency. It is systematically converting enterprise experience into machine-executable processes and integrated workflows.

SYSTEM COMPETITIVENESS ANALYSIS

Turning Tacit Knowledge into Assets and Organisational Transformation: The Real Moat in the Deep End of AI

As performance gaps among foundation models continue to narrow, enterprise competitive advantage is shifting from access to models to the accumulation of proprietary knowledge and organisational redesign.

TURNING TACIT KNOWLEDGE INTO ASSETS: CONVERTING ENTERPRISE EXPERIENCE INTO MACHINE-EXECUTABLE, PROCESS-BASED, AND COORDINATED WORKFLOWS

1 Sources of tacit knowledge

Dispersed across employee experience, business processes, and organisational routines, making it difficult for AI to directly access and use

- Employee experience and decision criteria
- Brand guidelines and design standards
- Target audience profiles and competitive intelligence
- Process interfaces and feedback rules
- Channel characteristics and style template systems

2 Systematic knowledge codification

Making tacit knowledge explicit and structured, and embedding it into AI platforms and agent systems

- AIGC platform with built-in brand guidelines and user profiles
- Customer service system integrating knowledge bases and process interfaces
- Advertising system integrating value proposition analysis and outcome feedback
- R&D agents linking the requirements – development – testing workflow
- 600+ agents across business domains and value chains

3 AI-accessible data assets

Building a continuously evolving proprietary decision system where outputs align with enterprise standards and can be directly deployed into business workflows

- More precise generation aligned with brand tone and commercial judgement
- More stable outputs with reduced need for manual review and iteration
- Smoother distribution through direct integration into the global asset management system
- Continuous evolution through closed-loop feedback
- Agents evolving from "smart tools" to "digital employees"

THREE DIMENSIONS OF ORGANISATIONAL RESTRUCTURING: SYSTEMIC TRANSFORMATION IN THE DEEP END OF AI

01
Hybrid employee capabilities

Before
Highly specialised by function, focusing on single-role tasks
Now
Moving towards cross-role collaboration enabled by AI

NEW Talent Initiative: from tool users to AI capability builders

02
Smaller and more agile organisational units

Before
Siloed functions and hierarchical handoffs, with repetitive work handled manually
Now
Systems take over repetitive work, freeing humans for higher-order judgement and creativity

Reduce unnecessary hierarchical handoffs and adapt structures to AI-enabled collaboration

03
Shift in management focus

Before
“Activity-based” management: focusing on whether each step is performed manually
Now
“Outcome-based” management: focusing on targets, clear boundaries, and controllable risk

As AI takes over more processes, managers focus on outcomes and risk

SHIFT IN THE SOURCES OF ENTERPRISE AI MOATS

Before	Reduced differentiation in model capabilities	Now
• Foundation model integration (yes/no) • Number of AI tools deployed • Efficiency gains in single-task scenarios • Demo performance and proof of concept		• **Proprietary data** and industry know-how turned into structured, reusable knowledge • **Business rules and operational context** embedded into AI-accessible decision systems • **Closed-loop feedback mechanisms** driving continuous system evolution • **Organisation, process, and talent** systematically restructured to align with AI-enabled collaboration

Sources: Anker Innovations internal assessment; compiled interview data

IV. Organisational Transformation: How AI Changes People and Organisations

As AI transformation moves into deeper stages, the key differentiator is no longer tools or platforms, but whether the organisation can redesign its core organisational elements around AI—roles, responsibilities, and capabilities.

In interviews, Anker repeatedly emphasised that at this stage, the challenge is not building more agents, but upgrading employee capabilities and reshaping how work is organised and coordinated. The focus shifts from helping individuals use AI tools more effectively to ensuring that the organisation itself can operate alongside AI systems.

To support this shift, Anker has outlined several priorities, including making business processes AI-driven, advancing organisation-wide AI transformation, aligning infrastructure with AI requirements, and exploring AI-native product models. It has also launched the NEW Talent Initiative (an internal upskilling programme) to help employees transition from users of AI tools to builders of AI capabilities.

This organisational shift has three key implications.

First, employee roles are becoming more hybrid. Highly specialised, function-based work is giving way to more flexible, cross-role collaboration supported by AI.

Second, organisational units are becoming smaller and more agile. As systems take over repetitive tasks, human effort shifts towards higher-level judgement and creative work. This requires reducing unnecessary layers and breaking down functional silos.

Third, management is shifting from overseeing activities to managing outcomes. As more work is executed by AI, the question is no longer whether each step is performed by humans, but whether results are delivered, boundaries are clear and risks are controlled.

The Anker case is best understood not as a story of extensive AI deployment, but as a broader transformation pathway. It begins with a small set of high-maturity use cases, where value is clear and outcomes are measurable. Capabilities are then consolidated through a unified platform, bringing fragmented efforts onto a shared foundation. Over time, AI becomes embedded in end-to-end workflows, ultimately driving a reconfiguration of the operating model, processes, talent, and knowledge systems.

For companies still experimenting with tools or running fragmented pilots, this pathway offers a clear and practical reference.

V. Key Takeaways: How AI Is Reshaping Enterprise Competition

In the AI era, the key question is no longer which companies are using AI, but which are integrating it early as a stable part of the operating model. Anker's experience shows that AI's most important impact is not greater technical intelligence, but greater operational ownership.

This is reshaping how value is measured. The focus is no longer on how much time employees save, but on how many standardised workflows AI can run independently.

It is also changing the logic of growth. AI is no longer just a back-end support tool; it is increasingly becoming a front-end engine for insight, content, campaign execution, and optimisation.

Most importantly, it is redefining organisational logic. As AI moves deeper into decision-making and execution, companies must redesign roles, workflows, responsibilities, and how work is organised.

For global hardware companies, this shift is especially significant. The business already operates under four structural pressures: short product cycles, volatile demand, complex cross-regional coordination, and extended service chains—areas where isolated digital tools deliver only marginal gains.

Real transformation occurs only when AI is embedded in core workflows across R&D, marketing, customer service, and operations, enabling companies to move from using AI to being fundamentally reshaped by it.

What Anker is driving is a broader evolution: AI moves from a peripheral tool to the operational backbone, from augmenting human work to owning parts of the workflow, and from local efficiency gains to system-level performance.

For Chinese enterprises today, the value of this case lies not only in the speed of adoption, but in the direction it sets for redesigning the enterprise operating model.

Part III: From Digital Rehearsal to Physical Reinvention—Multi-Agent Coordination and the Reconfiguration of Asset-Heavy Operations

Case 4: **Bosch Power Tools: Building a Virtual Customer Ecosystem and Reshaping the Global R&D Pipeline Through Multi-Agent Simulation**

Case 5: **Midea: From Lights-Out Factories to Agent Factories**

CASE 4

Bosch Power Tools: Building a Virtual Customer Ecosystem and Reshaping the Global R&D Pipeline Through Multi-Agent Simulation

I. The Complex Decision Chain in Cross-Border B2B Business and the Physical Constraints of Traditional Insight Generation

Bosch Power Tools is a core division of Robert Bosch GmbH, selling professional power tools, garden tools, and measurement instruments worldwide. According to Bosch's official financial disclosures, the division generated sales of approximately 5.1 billion euros in 2024 and employed around 18,700 people. About 90% of its sales came from markets outside Germany.

As a highly globalised hardware manufacturer, Bosch Power Tools operates in around 150 countries and regions. Its product portfolio covers most specialist trades recognised in UN standard classifications, including carpentry, plumbing and electrical work.

Within Bosch's broader industrial portfolio, Power Tools is not the group's largest business. Yet it has long served as a testing ground for innovation: large enough to support the validation of new methods, but organisationally flexible enough to pilot approaches that can later be extended across the wider group.

Bosch's Shanghai-based global product R&D and digital hub is particularly representative. Deeply rooted in China, it connects product development, manufacturing, logistics and marketing. For a mature and complex industrial system of this kind, the value of AI does not lie in simply adding tools to isolated process steps. It lies in reshaping how products are developed, how institutional knowledge is captured and reused, and how global teams collaborate.

Bosch Power Tools operates in highly complex B2B purchasing environments, where decisions are rarely made by a single end user. Instead, they are typically shaped by multiple stakeholders with different priorities, each influencing the final decision.

Consider a construction contractor procuring a batch of professional-grade rotary hammers. The workers using the tools care about weight, grip comfort, and vibration control. The finance manager approving the budget focuses on total lifecycle cost and volume discounts. The compliance or safety manager overseeing the site, meanwhile, needs assurance that the equipment meets relevant dust-control and noise standards.

In traditional R&D and marketing processes, understanding how a new product concept will be received across global market segments often requires cross-border field research conducted by external market research agencies. Such research is expensive and difficult to execute at scale, which means companies can usually draw only limited samples from a small number of core markets.

More fundamentally, linear surveys and standalone focus groups are poorly suited to capturing the dynamic internal trade-offs among the different roles involved in a B2B purchasing decision. When front-end market signals arrive late or are distorted by sample bias, the risks build up later in the process: tooling, mass production, and global distribution all involve substantial sunk costs.

As product iteration accelerates across global markets, this slow and expensive path of physical validation has increasingly become a structural bottleneck, constraining the agility of multinational manufacturers.

II. Shifting the System Mindset: Moving Beyond the Industrial Fixation on "100% Precision"

With the rise of generative AI and large language models, Bosch Power Tools began exploring how these technologies could be integrated more deeply across the value chain. In the early stages of adoption, however, the organisation had to rethink some deeply held assumptions.

As a German company with deep roots in industrial manufacturing, Bosch has long cultivated a management culture built on precision, zero-defect standards, and rigorously codified SOPs (standard operating procedures). Large language models, however, work differently. They are probabilistic by nature, meaning their outputs may vary and cannot be judged the same way as traditional IT software or precision machinery. During the early rollout, expecting AI to generate a flawless, ready-to-use business output from a simple prompt reflected a standard that AI, by its very nature, could not meet.

If an AI-generated report contained factual errors or gaps in reasoning, teams could quickly begin to question the system's business value and revert to traditional manual workflows.

Lv Zhichao and his digitalisation team recognised the risk in these expectations. If a company demands absolute perfection from generative AI without human intervention or deep contextual grounding, the technology will struggle to move beyond low-risk, low-value peripheral tasks such as basic translation or text polishing. Its value will remain trapped at the margins of the organisation.

The team's internal assessment was that, at that stage, AI's primary impact was not to overturn the company's underlying business model overnight. Its more immediate significance lay in reshaping the operating model: how day-to-day work was executed, how workflows were structured, and how people and systems collaborated.

To move beyond this perfectionism trap, Bosch worked to reset internal expectations around the boundary between human and machine. The management message to business teams was clear: at the current stage, AI systems were not positioned to provide flawless final answers. Instead, they should serve as auxiliary computational engines that could quickly produce a professionally informed working baseline. The AI system took on the heavy computational work: large-scale cross-border data retrieval, multidimensional scenario simulation, and the development of initial solution frameworks.

But the critical decisions still belonged to people. Experienced professionals were needed to judge complex business contexts, manage compliance risks, and refine the core business innovations. In other words, the model would perform best with humans in the loop.

Therefore, for large industrial companies with heavy asset commitments, this shift in mindset is essential. Only when teams accept a reasonable margin of error and understand what AI should and should not be expected to do can the technology move from experimental pilots into the heart of the business.

III. Building a "Virtual Customer Company": Using a Multi-Agent Network to Reconstruct B2B Decision Dynamics

With this new understanding of AI in place, Bosch focused its implementation on one of the most time-consuming and expensive parts of the process: market validation. Rather than using AI as a one-way Q&A tool to generate broad market reports from prompts, Bosch developed a multi-agent "Virtual Customer Company" system.

In consumer marketing, companies often use a single user persona to simulate customer behaviour. In Bosch's complex B2B context, however, a single persona cannot capture organisational purchasing behaviour. To reflect this reality, Bosch uses large language models to recreate an entire customer organisation in a digital sandbox, such as a construction contractor or manufacturing company. Within this virtual organisation, the system generates multiple agent roles by injecting different instructions and industry background data. Each agent is assigned a specific professional background, standpoint, and set of resource constraints: a frontline worker focused on efficiency, a cost-conscious procurement manager or a safety officer responsible for risk control.

When Bosch's R&D or marketing teams propose a new product concept, they no longer need to immediately launch expensive offline field research across global markets. Instead, they input the product's core selling points and commercial strategy into the system. The virtual roles then interact and debate internally, each responding based on its own commercial logic. A virtual finance manager might object to a higher unit purchase cost, while a virtual project manager might support the product because its IoT asset-tracking function could reduce equipment loss.

The system records how these virtual roles challenge one another, persuade one another and reach compromises, then quickly produces a diagnostic report that integrates market feedback from multiple perspectives. By adjusting parameters such as macroeconomic data, labour costs, and industry characteristics for different countries, Bosch can run simulated research across multiple global markets in parallel within a few hours. This proactive multi-agent simulation moves early-stage concept validation into a digital environment, compressing a process that once took months and significantly reducing the cost of trial and error in hardware innovation.

CORE INNOVATION MECHANISM

Virtual Customer Organisation: Using Multi-Agent Simulation to Reconstruct B2B Decision Dynamics

Cross-border market research moves from months-long, high-cost fieldwork to parallel digital simulation completed in hours, reducing the physical constraints on market feedback for industrial firms.

TRADITIONAL FIELD RESEARCH VS. VIRTUAL CUSTOMER ORGANISATION SIMULATION: A FIVE-DIMENSIONAL COMPARISON

	Traditional Cross-Border Field Research Market Research Agencies · Field Interviews · Focus Groups	**Virtual Customer Organisation and Multi-Agent Simulation** Multi-Agent Collaboration · Digital Sandbox · Parallel Simulation
RESEARCH CYCLE	Weeks to months; constrained by the practical limits of field execution **High time cost**	Simulations across multiple countries completed in parallel within hours **Compressed cycle**
MARKET COVERAGE	Usually limited to a few core markets and small sample sizes **Geographic constraints**	Markets in multiple countries can be simulated in parallel by adjusting macroeconomic parameters **Fewer geographic constraints**
STAKEHOLDER SIMULATION	Relies on single surveys or focus groups; struggles to capture multi-role dynamics **Distorted context**	Multiple agents play roles such as workers, procurement managers and safety officers, simulating real organisational decision dynamics **High-fidelity simulation**
TRIAL-AND-ERROR COSTS	High financial cost; delayed feedback creates major sunk-cost risks in later stages **Rigid high cost**	Early validation is shifted into digital simulation, reducing sunk-cost risk before physical tooling and mass production **Lower trial cost**
ITERATION FREQUENCY	Long cycles hinder agile response and make high-frequency experimentation difficult **Low-frequency, linear**	High-frequency experimentation and agile iteration become viable practices in asset-heavy manufacturing **High-frequency, parallel**

AI-Enabled Reconfiguration

VIRTUAL CUSTOMER ORGANISATION ARCHITECTURE

Virtual Customer Organisation
Simulating a construction contractor or manufacturing enterprise in a digital sandbox

Multi-Agent System

INPUT LAYER
Product Concepts and Commercial Strategy
- Core selling points of new products
- Target-market pricing strategy
- Country-level macroeconomic data
- Labour costs and industry characteristics

SIMULATION LAYER
Multi-Agent Decision Dynamics

Frontline Workers
Weight, grip comfort, vibration control

Finance Managers
Lifecycle cost, volume discounts

Safety Officers
Dust emissions, noise compliance

Procurement Managers
Unit cost, supply chain stability

Project Managers
IoT tracking, equipment loss

Compliance Officers
Industry regulations, certification requirements

OUTPUT LAYER
Multi-Perspective Market Feedback Report
- Records of challenge, persuasion and compromise among roles
- Procurement decision prediction based on integrated perspectives
- Country-specific market feedback
- Product improvement directions and risk alerts

Months → Hours
Concept validation cycles compressed from months-long fieldwork to hour-level simulation

150+ Countries
Bosch's operating markets can be simulated in parallel through parameter adjustment

Sunk Cost ↓
More validation completed digitally before physical tooling and mass production

Source: Bosch Power Tools Digitalisation Team; interview materials; Bosch Group FY2024 financial data

IV. Embedding Expert Models: Making Tacit Knowledge Explicit and Enabling Flexible Human–Machine Collaboration

Beyond external market feedback, multinational manufacturers face another challenge inside the organisation: much of their most valuable expertise is held by a small group of senior specialists. When less experienced employees are asked to prepare market plans or demand analyses, they often lack the judgement, context and analytical structure needed to do the work well. Generic AI chatbots do not close that gap on their own. Unless users already know how to frame the task and supply the right context, the output is likely to remain broad and superficial.

To address this, Bosch developed an internal Expert Model. The idea was not simply to use AI to empower external users, but to help Bosch's own senior experts extend their impact inside the organisation. These specialists understood customers and the operating logic of the industry in depth. Bosch's team therefore worked to embed their thinking frameworks, analytical models and standard operating procedures into AI agents. This enabled human experts and agents to work together more flexibly, while allowing that expertise to be applied at a greater scale.

In practice, the model changes the direction of support. When an employee needs to complete a complex professional task—such as drafting a new product launch business plan to Bosch's standards—the employee no longer has to begin by knowing how to instruct the AI. Instead, human experts first train the agent in how the task should be approached. This reverses the usual interaction: instead of the employee having to instruct the AI, the agent guides the employee through the work, asking structured questions based on the embedded expert methodology. For example: What is the product's core differentiating feature? What price ranges do competitors occupy in the target market?

Through this multi-round exchange, the agent helps the employee clarify the problem, fill gaps in the analysis and organise the necessary information step by step. Only when the inputs are sufficiently complete and logically coherent does it draw on the underlying large language model to generate a formal document that meets Bosch's internal standards.

The significance of this shift goes beyond faster document production. Bosch was effectively turning tacit expert judgement and experience into an explicit system-level asset—one that could guide employees through complex work, raise the baseline quality of output across the organisation and reduce the cost of cross-functional support.

KNOWLEDGE MANAGEMENT INNOVATION

Expert Model: Turning Tacit Knowledge into Reusable Organisational Capability

By embedding expert frameworks and SOPs into AI agents, Bosch enabled employees to complete complex tasks through agent-led questioning, improving consistency and baseline output quality across the organisation.

FROM EXPERT BOTTLENECK TO SYSTEM-BASED GUIDANCE

Traditional Knowledge Transfer Challenge Expertise concentrated among a small number of senior specialists		Expert Model Solution Turning expert know-how into reusable system capability
• Specialist knowledge and working methods remain difficult to access and reuse across the organisation	▶ Expert Model	• Embeds senior experts' thinking frameworks, analytical models and SOPs into AI agents
• Less experienced employees lack the context and structure needed for complex analytical work		• Uses agent-led questioning to guide employees step by step, helping them clarify their thinking and close information gaps
• Generic AI tools produce shallow outputs unless users know how to frame the task		• Employees no longer need advanced prompting skills to work through complex tasks in natural language
• Expert guidance is costly to provide across functions and difficult to scale		• Captures expert know-how in the system, making it reusable and reducing cross-functional support costs
• Organisational capability depends heavily on individuals, creating knowledge-loss risks when people move on		• Converts tacit knowledge into explicit digital assets that support more stable organisational capability

EXPERT MODEL WORKFLOW: NEW PRODUCT LAUNCH PROPOSAL

Four-Step Expert Model Collaboration Workflow

Human-in-the-Loop

Human experts define the framework → agents guide employees through structured questions → employees supply the inputs → the system produces a Bosch-standard document

1	2	3	4
Expert Framework Embedded Senior experts embed analytical logic, SOPs and evaluation criteria into the agent interaction framework	**Agent-Led Questioning** Based on embedded expert methodology, the agent guides the employee through structured questions	**Iterative Q&A Refinement** Through natural-language responses, employees clarify their thinking and fill information gaps	**Bosch-Standard Document Generation** Drawing on the underlying large language model, the agent generates a formal document that meets Bosch's internal standards
Expert-Led	Agent-Led	Human–Machine Collaboration	Agent-Generated

HOW THE EXPERT MODEL STRENGTHENS ORGANISATIONAL CAPABILITY

Dimension	Before the Expert Model	After the Expert Model
Knowledge Base	Held in the memory and intuition of a small number of senior specialists	Captured as system-level digital assets for scalable reuse across the organisation
Task Threshold	Complex tasks require direct expert involvement or lengthy training	Employees can work through complex tasks through agent-guided natural-language interaction
Baseline Output Quality	Output quality varies widely by individual and across functions	Baseline output quality improves, with greater consistency across teams
Knowledge Transfer Risk	Critical know-how is vulnerable to staff turnover and costly to pass on	Critical know-how is embedded in the system and retained by the organisation, not just by individuals

Source: Bosch Power Tools Digitalisation Team; interview materials

V. Reversing the Workflow: How the Enterprise Empowers AI

As agents became part of Bosch's day-to-day operations, the company began to rethink the workflows around them. Bosch's digitalisation team framed this as a reversal in management logic: the question was not only how AI could empower the enterprise, but also what the enterprise needed to do to empower AI.

This shift matters because agents do not create business insight from vague instructions alone. Their output depends on the depth and structure of the context they are given. If business teams give the system only a broad prompt, they are likely to receive broad, generic analysis in return, with little distinctive commercial insight. To make agents useful in a specialised industrial setting, Bosch had to provide them with the substance of the business: accumulated industry expertise, supply-chain constraints, safety regulations, and concrete market strategy objectives.

Bosch therefore redesigned its existing workflows by working backwards from what AI needed in order to perform. The team reviewed the full chain from product concept to market launch and broke role-level responsibilities down into finer-grained tasks. Within Bosch's specialist roles, the core value of AI lay in helping employees do what they could not previously do. By decomposing roles into micro-tasks, Bosch could distinguish between work suited to agents' probabilistic generation and multidimensional exploration, and work that still had to remain with people because it required physical validation and absolute precision.

The result was a new operating logic across Bosch's global R&D network. AI was placed earlier in the process, where it could support high-frequency, divergent exploration and generate multiple possible directions. Human judgment remained downstream, setting direction and making the final decisions. At the same time, business units had to make their own operating logic explicit, turning experience that had previously depended on individuals into parameters and rules the system could recognise. In this way, AI moved from a siloed testing tool into an integral part of the company's core operating model, facilitating a seamless transition from a purely human-driven workflow to a human-machine collaborative paradigm.

OPERATING LOGIC TRANSFORMATION

Mindset Reset and Workflow Reversal: AI Explores Upfront, Humans Decide Downstream

Bosch moved beyond the expectation that AI must be "100% precise", turning it from a peripheral support tool into a strategic driver of more agile global operations.

MINDSET RESET: MOVING BEYOND THE PERFECTIONISM TRAP

Old Mindset: Industrial Perfectionism

Expecting AI to turn a simple prompt into a 100% perfect final business solution

- Judging AI tools by the standards of precision manufacturing equipment
- Questioning the system's business value as soon as factual errors or reasoning gaps appear
- Confining AI to low-value peripheral tasks such as translation and text polishing
- Risk of reverting to traditional manual workflows

New Mindset: Fault-Tolerant Human–Machine Collaboration

Machines generate a multidimensional working baseline; humans retain final commercial judgement

- AI is positioned as an auxiliary computational engine that provides a professionally informed working baseline
- AI handles the heavy computational work of large-scale data retrieval and multidimensional scenario simulation
- Complex business judgement and compliance risk control remain within the human decision loop
- The immediate impact lies in reshaping the operating model, not overturning the business model overnight

WORKFLOW REVERSAL: AI UPFRONT EXPLORATION VS HUMAN DOWNSTREAM DECISION-MAKING

Granular Task Decomposition and a Redesigned Division of Labour Between Humans and AI

SOP Redesign

Reviewing the full chain from product concept to market launch to identify AI-suitable tasks and human-critical decisions

AI Upfront Tasks: Probabilistic Generation and Multidimensional Exploration

Handled by Agents

Global Market Data Retrieval and Synthesis
Processes macroeconomic and industry data across countries to quickly produce a baseline report

Multi-Role Scenario Simulation
Simulates internal decision dynamics in a virtual customer organisation to generate market feedback from multiple perspectives

Structured Document Drafting
Generates a Bosch-standard business plan draft based on the expert framework

Employee Guidance and Q&A
Uses agent-led questioning to help employees clarify their thinking and lower the task threshold

Human Downstream Tasks: Physical Validation and Absolute Precision

Retained by Humans

Final Judgement in Complex Business Contexts
Applies professional judgement to the AI-generated baseline and refines the core business ideas

Compliance Risk Control
Final review and decisions on safety regulations, industry certification and supply-chain compliance

Core Business Ideas and Strategic Direction
Makes core strategic decisions on product differentiation and market entry

Physical Prototype Validation and Mass Production Decisions
Makes final decisions on tooling, mass production and global rollout after digital simulation

WORKFLOW REVERSAL: AI UPFRONT EXPLORATION VS HUMAN DOWNSTREAM DECISION-MAKING

01 Resetting the Cost of Trial and Error

Shifts part of early market validation into digital simulation, making high-frequency experimentation and agile iteration economically viable in asset-intensive hardware manufacturing

Digital Simulation Before Physical Validation

02 Turning Tacit Knowledge into System-Level Assets

Transforms advanced experience that once depended on individuals into digital assets embedded in the system, enabling more stable transfer and scalable reuse of organisational capability

Knowledge as Organisational Capability

03 Moving Beyond Dependence on Certainty

Reframes AI as a tool for exploration rather than perfect answers, and turns it into a strategic driver by giving it the knowledge, context and rules it needs to perform

New Human–Machine Collaboration Model

Source: Bosch Power Tools Digitalisation Team; interview materials

VI. Core Takeaways: Resetting Trial-and-Error Costs and Linking Digital Simulation with Physical Manufacturing

Bosch Power Tools' experimentation with multi-agent AI networks in complex global operations reveals how industrial giants should rebuild their competitive core as they enter the most advanced stages of digital transformation.

First, proactive AI lowers the cost of trial and error in global hardware innovation. In the physical world, testing a B2B business hypothesis—whether by shipping engineering prototypes, organising cross-border focus groups or conducting multi-level channel interviews—has traditionally required rigid commitments of time and money. By building "Virtual Customer Corporations," Bosch has partly moved this field-based understanding of complex business environments into digital space, where large models can run simulations concurrently. Research cycles that once took months can now be compressed into hours. More importantly, Bosch has reduced the physical constraints that previously limited access to high-value market feedback, making high-frequency experimentation and agile iteration economically viable even in asset-intensive hardware manufacturing.

Second, embedding tacit knowledge into agent networks turns core enterprise capabilities into system-level digital assets. For multinational organisations, a long-standing challenge is that advanced industry experience and working methods often depend on the memory and intuition of a small number of individuals, making them difficult to transfer intact across a large organisation. Bosch's approach was to make expert knowledge part of the workflow itself. By building senior specialists' analytical logic and SOPs into agent interactions, the company enabled employees to work through complex tasks with expert guidance built into the system. Knowledge that once depended on individual memory and intuition could therefore be reused more consistently across the organisation, making Bosch's internal capabilities more stable, transferable and scalable.

Finally, the leap in system-level performance requires a management reset: companies must move beyond an industrial-age dependence on certainty. The competitive strength of industrial manufacturers has traditionally rested on extreme process precision and zero deviation. Generative AI, by contrast, works through probabilistic, divergent exploration. Bosch's practice shows that companies need to find a balance between the pursuit of maximum precision and the gains available from system-level efficiency. By loosening the insistence that tools must be "100% correct" and establishing a collaborative rule in which machines generate a multidimensional working baseline while humans retain final commercial judgement, companies can move generative AI beyond peripheral typing assistance and turn it into a strategic engine for more agile global operations. The underlying purpose remains to empower people to do what they could not do before.

CASE 5

Midea: From Lights-Out Factories to Agent Factories

Midea Group is a global technology group that grew out of home appliance manufacturing. According to its 2024 annual report and 2025 interim report, the company recorded total revenue of RMB 409.1 billion (US$57.6 billion) in 2024, including RMB 104.5 billion (US$14.7 billion) from its B2B business. As of mid-2025, it had more than 400 subsidiaries worldwide, 38 R&D centres, 63 major manufacturing bases, over 190,000 employees and operations in more than 200 countries and regions.

Midea Cloud, a Midea subsidiary established in 2016, was created to turn the group's accumulated manufacturing, management and digital practices into software and platform-based solutions. It has also participated in the development of multiple lighthouse factories across Midea.

For a company with such a large manufacturing network, complex product portfolio and deeply integrated global operations, the significance of AI does not lie simply in adding a new set of productivity tools. It lies in bringing AI into the operational backbone of the business: planning, manufacturing, quality and supply.

Before this latest phase of AI adoption, Midea had already been pursuing consistency across the organisation through its "One Midea, One System, One Standard" initiative. It had also reshaped its manufacturing logic through the T+3 model and the Midea Business System (MBS), its lean business system. The T+3 model is driven by real end-market demand. It creates an end-to-end order-based delivery process from customer order placement (T0), to materials preparation (T1), manufacturing (T2) and delivery (T3). Midea has publicly disclosed that the model can shorten full-cycle lead times to as little as 12 days.

MBS, or the Midea Business System, is a lean management system built around pull-based operations, strategy deployment, talent development and daily management. Its purpose is to support full value-chain operations through high quality, short lead times and low cost.

I. From "Let Everyone Try" to Full-Scale AI Adoption

Midea's AI rollout was not a one-off, top-down deployment. It followed a staged process: broad experimentation first, then filtering for high-value use cases, and finally integration into the core of the business.

Jin Jiang, President of Midea Cloud, recalled that around 2023, Midea adopted what it called a "let everyone try" approach. The company opened AI platform capabilities to all employees without predetermining what the right use cases should be. Employees in different roles were encouraged to identify possible applications based on their own work.

By 2024, Midea had moved beyond open-ended experimentation. The group recognised that AI's value would not be proven by models, algorithms or computing power in isolation, but by whether they could solve real business problems. It therefore began identifying high-value use cases across R&D, manufacturing, supply, sales, management and finance. By 2025, as some of these use cases began to deliver clear operational value, AI adoption expanded across the company and entered a full-scale phase.

What best illustrates this process is not the slogan, but the outcome of use-case selection. Jin disclosed that during the “let everyone try” phase, more than 15,000 agents emerged inside Midea. These included applications directly related to operations, such as translation, HR, legal affairs, manufacturing and R&D, as well as smaller scenarios connected to employees’ own work and daily lives.

Yet most of the efficiency gains and cost savings came from just over 100 high-value scenarios. Midea therefore quickly shifted its focus from general-purpose scenarios to deeper operational scenarios. The functions that truly determine enterprise efficiency and manufacturing competitiveness—R&D, manufacturing, quality, planning and supply—were also the areas where AI adoption was more difficult and slower to take hold.

Midea also developed a clear view of the different levels of AI application. The first stage is use-case-based application, where one scenario solves one problem, such as contract review, quality inspection, translation or equipment parameter optimisation. The second stage is agent collaboration, where different agents interact around the same task chain, expanding the scope of what agents can cover. The third stage is the agent factory, which Jin described as the “ultimate form of A2A”.

For example, a demand forecasting agent may generate demand trends based on sales, weather, pricing and market indices. It can then interact with a production scheduling agent to determine which factory should take an order, which production line should handle it and what production rhythm would deliver the best balance of cost and delivery performance.

STRATEGIC EVOLUTION ANALYSIS

Midea's Three-Stage AI Adoption Path: From Broad Experimentation to Full-Scale AI Adoption

Midea moved from broad experimentation to high-value use-case selection, narrowing 15,000 agents to 100+ value-creating scenarios before extending AI into core operations and agent factories.

THREE-STAGE AI ADOPTION PATH (2023–2025)

2023 · Phase 1

Broad-Based Experimentation

Open experimentation before targeted selection

Strategic Focus:
Broad Exploration

- AI platform capabilities were opened to all employees without prescribing specific use cases.
- Teams explored possible applications based on their own workflows.
- Applications spanned translation, HR, legal, manufacturing and R&D, among others.
- Other applications reflected employees' everyday work and personal needs.

15,000+
internal agents

2024 · Phase 2

High-Value Use-Case Selection

From models to business value

Strategic Focus:
High-Value Use-Case Filtering

- Shifted focus from models, algorithms and computing power to concrete business use cases.
- Identified high-value applications across R&D, manufacturing, supply, sales, management and finance.
- Shifted focus from general-purpose applications to deeper operational scenarios.
- Focused on core functions such as R&D, manufacturing, quality, planning and supply.

100+
high-value scenarios contributing the vast majority of value

2025 · Phase 3

Full-Scale AI Adoption

Integration into core operations

Strategic Focus:
System Integration

- A set of scenarios began to show clear operational value, expanding AI use across the company.
- AI moved into core operational domains such as planning, manufacturing, quality and supply.
- The agent factory model began to take shape, with A2A coordination emerging as a new manufacturing paradigm.
- A global replication mechanism began extending deployment across 60+ production bases.

80%+
average efficiency improvement in core manufacturing scenarios

USE-CASE VALUE FUNNEL: FROM BROAD EXPLORATION TO DEEP FOCUS

15,000 agents → 100+ high-value scenarios → deep integration into core operations

Value-Driven Selection

The functions that truly determine enterprise efficiency and manufacturing competitiveness are R&D, manufacturing, quality, planning and supply—areas where AI adoption is more difficult and slower to take hold.

15,000+ → **100+** → **Deep Integration**

Broad Exploration
Open to all employees · bottom-up discovery
General-purpose and employee-led applications

High-Value Scenarios
Account for the vast majority of efficiency gains and cost savings

Operational Core
R&D, manufacturing, quality, planning and supply
A2A coordination

THREE-LEVEL AI APPLICATION ARCHITECTURE: FROM SINGLE USE CASES TO THE AGENT FACTORY

01 Single Use-Case Application

Each application targets a defined task, helping teams solve specific problems and validate business value quickly.

Typical scenarios: contract review, quality inspection, translation and equipment parameter optimisation.

02 Agent Coordination

Agents begin to interact across the same task chain, expanding AI from isolated applications into connected workflow support.

Typical scenario: a demand forecasting agent works with a production scheduling agent to determine the factory, production line and rhythm that best balance cost and delivery.

03 Agent Factory

Advanced A2A coordination enables agents to sense, decide, execute, feed back and optimise within defined manufacturing boundaries.

Typical scenario: 14 agents covering 38 core production scenarios, with production scheduling response speed improved by 90%.

Sources: Midea Group 2024 Annual Report; Midea Cloud interview materials; official disclosures from Midea Group

II. After the Lights-Out Factory: Why Manufacturing Systems Still Need to “Think”

The shift from lights-out factories to agent factories marks a critical step in Midea’s AI transformation. Zhang Xiaoyi, Vice President and Chief Digital Officer of Midea Group, has publicly described the agent factory as a “new species”. Jin Jiang explained the distinction in operational terms: in a lights-out factory, each action is governed by rules that humans have set in advance, and the system executes those rules strictly. In an agent factory, more actions begin with AI sensing the environment, assessing the situation, invoking models and generating the next instruction. The former is instruction-driven; the latter begins to move towards goal-driven operation.

Midea did not move beyond the lights-out factory because automation had become unimportant. Rather, long-term operations had made the boundaries of automation increasingly clear.

The Jingzhou refrigerator factory offers a useful example. This factory, with a history of more than 30 years, faced rising product complexity as the number of production models increased from four to twelve. The mixed production of mid- and high-end products, as well as domestic and export orders, continued to make manufacturing and production scheduling more difficult.

In 2022, the factory upgraded through flexible automation, IoT and AI. The transformation increased labour productivity by 52%, shortened lead times by 25% and reduced quality defects by 64%. It also enabled real-time inspection across 199 quality control points, while one sheet-metal production line could complete a model changeover in one second with a single click. This is precisely the strength of a lighthouse factory: it pushes automation, digitalisation and flexibility to a high level.

In Midea’s view, however, the lights-out factory does not automatically solve the problem of change itself. When factories face rush orders, non-standard equipment failures, material bottlenecks or mixed-flow changeovers, systems based on hard-coded rules still depend heavily on people to diagnose abnormalities, reset logic and reschedule operations.

Jin offered a simple example. In a lights-out factory, equipment generally runs according to predefined rules. In an agent factory, an equipment agent can monitor operating parameters, assess the risk of failure and generate a maintenance work order before a breakdown occurs. Maintenance therefore shifts from passive response after a stoppage to proactive intervention before it happens. This captures Midea’s broader vision for the agent factory: not simply a factory with more automation, but one in which more parts of the operation can sense, judge and respond.

MANUFACTURING PARADIGM SHIFT ANALYSIS

Lights-Out Factories → Agent Factories: From Instruction-Driven to Goal-Driven Operations

Automation keeps processes running reliably; agent factories help systems adapt, decide and coordinate as conditions change.

LIGHTS-OUT FACTORIES VS AGENT FACTORIES: SIX CORE DIFFERENCES

Generational shift →

	Lights-Out Factories (Lighthouse Factory Stage) Automation · Digitalisation · Flexibility · Instruction-driven	Agent Factories (A2A Coordination Stage) Perception · Decision-making · Execution · Feedback · Goal-driven operations
OPERATING LOGIC	Systems execute predefined human-set rules **Instruction-driven execution**	AI senses operating conditions, assesses system status, invokes models and generates the next instruction **Goal-driven execution**
EXCEPTION HANDLING	People must diagnose anomalies, adjust logic and reschedule operations **Reactive response**	Agents detect anomalies, adjust workflows and coordinate proactive responses **Proactive intervention**
EQUIPMENT MAINTENANCE	Equipment follows fixed rules; maintenance is triggered after stoppages occur **Post-failure repair**	Equipment agents monitor operating parameters, assess failure risk and generate maintenance work orders before breakdowns occur **Predictive maintenance**
PRODUCTION SCHEDULING	Mixed-flow changeovers still require human intervention under hard-coded rules **Static rules**	Demand forecasting and scheduling agents work together to dynamically optimise factory and production-line allocation **Dynamic optimisation**
CORE STRENGTH	Executes predefined actions more steadily, faster and with fewer people **Stable execution**	Handles judgement-intensive work once reliant on human experience, including anomaly detection, quality loops and equipment fault prediction **Judgement and coordination**
CAPABILITY BOUNDARY	Struggles to adapt when conditions change; rush orders and material bottlenecks still require manual handling **Dependent on predictability**	Helps factories move from reliable automated execution to adaptive response under changing conditions **Adaptive response**

THE JINGZHOU WASHING MACHINE FACTORY: AGENT FACTORY OUTCOMES

Midea Jingzhou Washing Machine Factory · Agent Factory Outcomes
WRCA-certified as the “world’s first multi-scenario agent factory” · 14 agents · 38 core production scenarios

Publicly launched in 2025

80%+	90%	15min → 30s	81
average efficiency improvement from agents across core manufacturing scenarios	increase in production scheduling response speed	first article inspection time reduced	AMRs coordinated by a logistics agent for cross-zone operations

First Article Inspection
AI glasses use historical first-article inspection data to flag error-prone areas. The system then coordinates with R&D and quality agents to automatically retrieve design drawings and verify them against the physical components.
Efficiency: 15 minutes → 30 seconds

Fastening in Mixed-Flow Production
KUKA collaborative robots coordinate closely with the planning agent to identify product models in mixed-flow production and automatically update the screw-fastening program.
Automated mixed-flow changeover

Logistics Scheduling
81 AMRs are coordinated by a logistics agent for cross-zone transfer, dynamic route sensing and autonomous obstacle avoidance.
Autonomous coordination across the whole site

PREVIOUS VALUE MEASURE
Labour hours saved

→

CURRENT VALUE MEASURE
End-to-end tasks AI can take over
Human judgement captured and reused

Sources: Official disclosures by Midea; interview materials from Midea Cloud; WRCA certification announcements

III. The Jingzhou Washing Machine Factory: Bringing Agents onto the Manufacturing Floor

The next stage could be seen in Midea's Jingzhou washing machine factory, which was publicly launched in 2025. If the Jingzhou refrigerator factory represented the lighthouse factory stage, the washing machine factory marked the move into the agent factory stage. According to Midea's official disclosures, the factory had 14 agents covering 38 core production scenarios, coordinated through the Midea Factory Brain. It also received WRCA certification as the "world's first multi-scenario agent factory".

The factory did not use agents merely as separate point solutions. Its architecture was based on distributed multi-agent coordination: agents communicated through A2A protocols, coordinated autonomously and drew on a large-model reasoning engine built for industrial applications. This created a closed loop running from perception and decision-making to execution, feedback and continuous optimisation.

These agents were not operating as abstract software layers; they were already embedded in specific shop-floor tasks. In first article inspection, for example, AI glasses drew on market issue data and historical inspection records to flag common error-prone areas. The system then coordinated with R&D and quality agents to retrieve drawings automatically and compare them against the physical product, cutting first inspection time from 15 minutes to 30 seconds.

Other examples showed agents moving from analysis into direct shop-floor coordination. At the automated dryer rear-cover fastening workstation, KUKA collaborative robots worked with the planning agent to identify product models in mixed-flow production and update the fastening programme automatically. In the injection moulding workshop, 81 AMRs were coordinated by a logistics agent, enabling cross-zone transport, dynamic route sensing and autonomous obstacle avoidance. Midea's humanoid robot "Meiluo" and AI inspection robot "Yutu" extended this agent-based coordination into inspection and on-site execution.

Midea disclosed that across multiple core manufacturing scenarios, agents had delivered average efficiency gains of more than 80%, including a 90% increase in production scheduling response speed.

The contrast with lights-out factories is clear. A lights-out factory is designed to execute predefined actions more steadily, faster and with fewer people. An agent factory begins to take on judgement-intensive work that previously depended more heavily on human experience: anomaly identification, route adjustment, mixed-flow changeovers, closed-loop quality management and equipment fault prediction. For manufacturing systems, this means AI is no longer just generating reports or making recommendations. Within defined boundaries, it is beginning to assume more responsibility for on-site execution.

IV. From Single-Factory Breakthrough to Global Replication: The Role of Midea Cloud and Organisational Mechanisms

An advanced factory can prove that a direction is viable, but it does not automatically prove that the organisation can replicate it. One important reason Midea was able to move the agent factory from concept to implementation is that, before AI entered deep operational scenarios, the company had already standardised factory management and data governance.

Jin noted that Midea had more than 60 production bases around the world. If each factory had been left to explore on its own, the process would have been highly inefficient. During the digitalisation phase, the group had already unified factory processes, data, standards and KPIs. It could then summarise best practices and promote them across the organisation, allowing new technologies and methods to be replicated across more production bases.

Within this mechanism, Midea Cloud was more than an IT implementation provider. It served as the software and platform layer through which Midea codified its manufacturing experience. It drew on long-accumulated methodologies such as T+3, MBS, lighthouse factory development and industrial internet practices, and translated them into industrial software, solutions and templates that could be replicated across sites.

Midea has publicly disclosed that MBS now covers all its factories in China and has accelerated its overseas rollout since 2023. Midea Cloud's official website also shows that it has participated in the development of multiple Midea lighthouse factories. As a result, the agent factory does not have to remain a single-factory showcase. It has the potential to become a replicable organisational capability.

This replication logic had already begun to extend overseas. In 2025, Midea's refrigeration equipment factory in Si Racha, Chonburi Province, Thailand, was selected by the World Economic Forum as a Supply Chain Resilience Lighthouse Factory. According to the World Economic Forum, the factory deployed 72 digital and AI solutions, reducing order lead time by 43%, cutting customer complaints by 32% and increasing employee certification speed by 62%.

In talent training, Midea also introduced the "Meike" system at the factory. The system converts documents into online courses with one click, supports four languages and has shortened the time required to obtain core skills certification from eight days to three days. This shows that agents and generative AI at Midea are no longer confined to individual production lines. They are beginning to enter quality loops, training systems and supply chain resilience building across the global manufacturing network.

ORGANISATIONAL CAPABILITY AND GLOBAL REPLICATION ANALYSIS

From Single-Factory Breakthrough to Global Replication

Midea standardised factory management and data governance before scaling AI from individual use cases to A2A coordination across 60+ production bases.

MIDEA CLOUD: CODIFYING MANUFACTURING KNOW-HOW FOR REPLICATION

Midea Cloud: Turning Manufacturing Know-How into Replicable Software Assets
Connecting established manufacturing methods to industrial software, solutions and replication templates.

Software Layer for Manufacturing Know-How

METHODOLOGY INPUT

Established Manufacturing Know-How

- T+3 order-based delivery model (12-day end-to-end lead time)
- MBS lean business system (pull-based operations)
- Lighthouse factory development practices
- Industrial internet capabilities

CORE ROLE

Midea Cloud

- Established in 2016 as a manufacturing digitalisation subsidiary
- Turns manufacturing practices into software and platform-based solutions
- Supported multiple Midea lighthouse factory projects

REPLICATION OUTPUT

Replicable Organisational Capability

- Industrial software and solutions
- Standardised replication templates
- Unified processes, data, standards and KPIs
- Best practices scaled across production bases

Standardised Foundation (Prerequisite)

Standardised factory management and data governance as the foundation for AI deployment in deep operational scenarios
Unified processes | Unified data standards | Unified KPI system | Unified business language | One Midea · One System · One Standard

Capability Validation (Pilot Factories)

Single-factory breakthroughs validate the direction—Jingzhou refrigerator factory (lighthouse factory) + Jingzhou washing machine factory (agent factory)
Labour productivity +52% | Lead time -25% | Quality defects -64% | Scheduling response +90% | FAI time 15min → 30s

Global Replication (Scaling Mechanism)

Best practices replicated across 60+ production bases—full coverage in China and accelerated overseas rollout
MBS deployed across all factories in China | Overseas replication accelerated since 2023 | Si Racha Lighthouse Factory in Chonburi, Thailand | Meike training system with four-language support

GLOBAL REPLICATION: DOMESTIC BENCHMARK VS OVERSEAS EXPANSION

Domestic Benchmark: Jingzhou Refrigerator Factory (Lighthouse Factory Stage)
2022 transformation · Flexible automation + IoT + AI

52%	25%	64%
increase in labour productivity	reduction in lead time	decrease in quality defects

Real-time inspection across 199 quality control points; one-click, one-second model changeover on a sheet-metal line; production models expanded from 4 to 12, supporting mixed production of mid- and high-end models and domestic/export orders

Overseas Expansion: Si Racha Refrigeration Equipment Factory, Chonburi, Thailand
2025 · WEF Supply Chain Resilience Lighthouse · 72 digital and AI solutions

43%	32%	62%
reduction in order lead time	reduction in customer complaints	faster employee certification

The Meike system converts documents into online courses in four languages with a single click, reducing core skills certification time from eight days to three.

STEP 01
Unified Business Language
Digital and lean systems standardise processes, data structures and rules

STEP 02
Validation Through Individual AI Use Cases
High-value use cases validate ROI and replicability

STEP 03
A2A Coordination Integration
Agents enter core manufacturing processes, enabling a perception–decision–execution closed loop

STEP 04
Global Replication and Scaling
Replicable capabilities scaled across 60+ production bases on a standardised foundation

Sources: Midea Group 2024 Annual Report; WEF Lighthouse Factory certifications; Midea Cloud website; interview materials compiled by the authors.

V. Core Takeaways: How Proactive Intelligence Enters the Manufacturing Core

Based on public materials and interviews, Midea appears to be advancing along a relatively clear path. First, digitalisation and lean systems unify business language, data structures and process rules. AI then moves from individual use cases into A2A coordination, before extending towards the agent factory.

The lights-out factory answers the question of whether a system can execute reliably. The agent factory goes one step further: can the system continue to judge and coordinate execution amid change? For manufacturers, this means the measure of AI value is shifting—from how many labour hours it saves, to how many task chains it can take on and how much replicable judgement it can embed in the organisation.

Jin also stated clearly in an interview that a fuller version of the agent factory will still take time. Companies should not understand it as an overnight “unmanned myth”. Midea’s case therefore does not present a finished end-state, but a practical path for manufacturing’s evolution towards proactive intelligence.

In this model, people define objectives, boundaries and constraints. Within those boundaries, systems sense, coordinate, judge and adjust. The factory gradually moves from being able to “run automatically” to being able to “respond to change”. The significance of this step lies not only in higher production efficiency, but in a shift in the operating logic of the manufacturing system itself.

Part IV: Paradigm Disruption and Ecosystem Advancement—Rewriting Underlying Logic and Capitalising on Core Capabilities

Case 6: **Schneider Electric: Redefining Growth Through AI-Enabled, Software-Defined Hardware**

Case 7: **Matrix Design's AI Journey: From Industry Know-how to Vertical Models and an Expanding AI Ecosystem**

CASE 6

Schneider Electric: Redefining Growth Through AI-Enabled, Software-Defined Hardware

Schneider Electric is a global energy technology company with operations spanning buildings, data centres, industry and infrastructure. Its portfolio ranges from low- and medium-voltage power distribution solutions to automation control, building management and data centre infrastructure, covering much of the chain by which electricity is generated, distributed and used.

In 2025, Schneider Electric's global revenue reached €40 billion, with software and services accounting for 19% of annual revenue. System integration and services were becoming an increasingly important source of growth, generating revenue from integrated solution delivery, system upgrades, operations support and ongoing services—not just from standalone equipment sales.

For a company long known for its strength in hardware, AI and digital transformation were not simply about developing more digital products. They raised a more fundamental question: in energy, power and automation—industrial sectors long defined by physical assets and hardware systems—was software still just a control layer attached to hardware, or had it begun to define the hardware itself, shape system architecture and change the company's growth logic?

Schneider Electric's answer was already moving beyond the traditional industrial sequence of "hardware first, then intelligent hardware, then software control and analysis". Instead, the company was starting from the value it wanted to create—such as efficiency and customer experience—and using that value proposition to shape the software functions and architecture. Those software choices then determined what the hardware needed to do, leading ultimately to integrated application systems in which software and hardware were designed together. AI was embedded throughout this process as a technical capability, creating value across the chain rather than appearing as a separate layer.

I. From Hardware-First to Software-First: A Shift in Product Logic

Schneider Electric's latest transformation began with a shift in product logic. In the traditional development model for industrial companies, hardware came first: devices were connected, and software was then added for control and analysis. Schneider Electric placed greater emphasis on starting from customer value, including business development, customer experience and operational efficiency. It first designed the software application layer, then used that layer to define the underlying architecture and determine what the hardware needed to include—and what it no longer needed.

On the surface, this looked like a change in the order of R&D. At a deeper level, it reflected a shift in competitive focus from "hardware first" to "software first".

This did not mean that hardware was becoming less important. Physical devices remained essential wherever electricity had to be transmitted, converted or controlled in the field. The change was more specific: hardware was no longer automatically the centre around which the whole system was designed.

Software did not replace the physical distribution and transmission of electricity. It replaced—or increasingly absorbed—the logic that sat around those physical processes: monitoring, scheduling, control and optimisation. As embedded computing became more widely available and cloud computing moved closer to the edge, this logic no longer needed to sit in a separate middle layer. Time-sensitive tasks that required real-time response could move down to the edge, while less time-sensitive optimisation tasks, especially those requiring data from multiple management systems, could move up to the cloud.

This was the real meaning of Schneider Electric's "software-first" strategy. The company was not simply adding a software business line. It was making software the starting point for system design. For a traditional industrial company, this was a deeper shift than launching a new software product, because it changed where value was created.

As hardware became more standardised, software became harder to copy and more important as a source of differentiation. Its algorithmic logic could help defend against competition and deepen customer relationships. At the same time, Schneider Electric began to pursue a strategy of standardised hardware, software platforms and customised applications to reduce SKU complexity. In this model, software coordinated the system and delivered differentiated experiences, while standardised hardware continued to provide the essential physical execution layer and support the company's core profit streams.

By 2025, software and services accounted for 19% of Schneider Electric's revenue. This did not mean that hardware was in decline, but it did show that software and services had moved from the margins towards the centre of the business.

II. From Product Logic to System Architecture: Open Automation and Architectural Leadership

If "software first" changed how Schneider Electric defined products, its push into open automation changed the industry rules on which it had long relied.

Industrial automation has long faced a structural problem: systems are relatively closed, centred on hardware and confined to proprietary ecosystems. Programming software is often tied to hardware devices such as controllers, and products from different vendors are often incompatible. Once customers select a supplier, switching becomes difficult. For leading vendors, this closed model has traditionally been an important moat. But it also limits efficiency, flexibility and innovation at scale.

Software-defined automation points to the future direction of the industry. In industrial settings, it brings three main advantages: open ecosystems, virtualised control and data continuity. Schneider Electric was an early mover in open automation, and its EcoStruxure Automation Expert (EAE) platform was positioned explicitly around open, software-defined automation.

Schneider Electric's move into open automation was driven by customer pain points, guided by technology trends, validated by value creation and supported by an ecosystem strategy. It was not only a response to the limitations of closed industrial systems. It was also a key step in actively shaping a more open, flexible, intelligent and sustainable next generation of industrial automation.

The need for this shift was clear. Industrial scenarios are highly fragmented: the process parameters for dairy fermentation, for example, differ significantly from those for brewing. No single company can develop every type of customised software or manage every category of automation hardware and software system. Schneider Electric therefore focused on building a unified software foundation and a standardised, modular control system covering control applications and functional components. It also connected data links across distributed control nodes and provided open interfaces for upper-layer IT applications, enabling ecosystem partners and end users to use AI models for more efficient and intelligent management and control across the manufacturing process.

In this sense, open automation was a form of self-disruption. Once standards and interfaces were opened, the customer lock-in that traditional equipment vendors had built through closed systems would inevitably be weakened. Service revenues attached to those closed ecosystems could also come under pressure.

Schneider Electric nevertheless continued to pursue this path. The underlying logic was not simply strategic positioning, but a recognition of the value of open automation. This was not an altruistic act of giving up control. Its value lay in using open software architecture, interoperable connectivity and compatible industry standards to break down the closed barriers of traditional industrial automation. By removing collaboration bottlenecks between equipment and systems from different vendors, open automation could activate innovation across the ecosystem and move industrial automation towards greater efficiency, flexibility and intelligence.

STRATEGIC TRANSFORMATION ANALYSIS

From “Hardware-First” to “Software-First”: A Shift in Product Logic

Schneider Electric has moved beyond the traditional sequence of “hardware first → intelligent hardware → software control and analysis”. It now starts with the value to be created, then uses software functions and system architecture to define the role of hardware. Software is no longer a supporting layer; it is the starting point for system design. This supports a strategy of standardised hardware, software platforms and customised applications.

PRODUCT DEVELOPMENT LOGIC: FROM HARDWARE-FIRST TO SOFTWARE-FIRST

A reversal in development logic: from “hardware-driven software” to “software-defined hardware.”
On the surface, this looks like a change in development order; at a deeper level, it signals a shift in competitive focus.

Software-First Strategy

Traditional Hardware-First Path Hardware at the centre, with software added for control and analysis

First Step: **Hardware R&D** as the core driver	Second Step: **Hardware Interconnection** for device networking and data collection	Third Step: **Software Control and Analysis** as a supporting layer	Result: **Hardware Remains at the Centre;** software depends on hardware

New Software-First Model Starting point: customer value · Software defines what hardware is needed and what can be removed

First Step: **Define Customer Value** efficiency, experience and operational optimisation	Second Step: **Design Software Application Layer** interface, architecture and functions first	Third Step: **Software Defines Hardware Requirements,** determining what is needed and what can be removed	Result: **Software Becomes the Starting Point;** AI creates value across the chain

THREE SOFTWARE-FIRST EXECUTION STRATEGIES: REDUCING SKU COMPLEXITY AND RESHAPING THE VALUE HIERARCHY

Strategy 1: **Hardware Standardisation—** Reducing SKU Complexity	Strategy 2: **Software Platforms—** Building a Hard-to-Replicate Moat	Strategy 3: **Application Customisation—** Differentiated Experiences and High-Value Services
• Greater hardware standardisation lowers R&D and manufacturing costs. • Standardised hardware with essential physical execution functions continues to support core profit streams at the base layer. • Execution-layer hardware remains indispensable: electricity still has to be transmitted, converted and controlled by physical devices. **Hardware: The Foundation of Core Profits**	• Algorithmic logic is difficult to copy, creating a stronger moat. • Software deepens customer relationships through subscription- and service-led models. • EcoStruxure Automation Expert (EAE) breaks down closed-system barriers, enabling open ecosystems, virtualised control and data continuity. • Open automation weakens traditional lock-in and shifts advantage from device binding to ecosystem leadership. **Software: Moat and Ecosystem Leadership**	• Software coordinates the system and delivers differentiated experiences for different industry needs. • AI is embedded in the application layer to support sensing, diagnosis, prediction and prevention. • By 2025, software and services accounted for 19% of revenue, moving from the margins towards the centre of the business. **Application Layer: Differentiated Value and Customer Entry Point**

SHIFT IN COMPETITIVE FOCUS: FROM PRODUCT COMPETITION TO ARCHITECTURAL LEADERSHIP

Past Competition Logic **Hardware Provider: Competing on Product Performance**		Current Competition Logic **Architecture Shaper: Competing at the System Level**
• Competition focused on hardware specifications and performance. • Closed systems created customer lock-in by binding devices to supplier ecosystems. • Revenue was dominated by one-off capital expenditure hardware sales. • Software remained an accessory to hardware and struggled to demonstrate independent value. • System integration and services stayed peripheral rather than becoming core growth drivers.	▶	• Software and hardware integration defines competition. • Open automation replaces device lock-in with ecosystem leadership. • Subscription-based software and services create recurring revenue. • System integration, integrated solutions and operations support are becoming key sources of growth. • Software shapes product development sequence, ecosystem rules and high-value service integration.

Source: Company interviews and publicly available information

III. AI as an Embedded Capability Across Product Development, Customer Service and Marketing

Throughout Schneider Electric’s digital transformation, AI served as an embedded technical capability, becoming increasingly integrated into different parts of the business. What was most notable about the company’s use of AI was that it did not appear mainly as standalone “AI products”. Instead, AI was built into the full path from defining value and designing software to developing integrated software-hardware application systems, and was applied across product development, customer service, marketing and other business functions.

In product development, AI appeared mainly through embedded algorithms. Equipment was moving from simply executing commands towards intelligent hardware capable of sensing, diagnosing, predicting and preventing faults. At the same time, AI capabilities were built into software from the earliest stages of design. Through AI-native system architecture, Schneider Electric connected data flows and intelligent feedback loops across the system, allowing intelligence to be built in from the start and enabling systems to evolve and upgrade as the technology advanced.

In customer service, AI was integrated with distributed computing, open automation and data-centre energy-efficiency management, supporting optimisation, scheduling and predictive maintenance. In other words, AI was built into Schneider Electric’s customer solutions as a core capability, extending the value those solutions could deliver.

In marketing and sales, AI became part of Schneider Electric China's B2B digital marketing model. The China team had built a WeChat-centred system with two layers. Public-facing content was delivered through WeChat service accounts, video accounts and mini-programmes, allowing Schneider Electric to maintain regular contact with a broad customer base. More segmented engagement took place through WeCom, Tencent's enterprise communication platform, where the team operated private communities and service platforms tailored to different customer roles. As this model matured, performance metrics shifted from follower numbers to more meaningful measures such as monthly active users and high-engagement users. Schneider Electric China contributed nearly half of the Schneider Electric's global monthly active users, ranking first globally.

Here, AI did not replace sales. Instead, it was applied to content production, customer insight and lead identification, allowing the long and complex customer decision-making process for industrial products to be managed with greater precision.

AI AS AN EMBEDDED CAPABILITY

AI as an Embedded Capability Across Products, Services and Marketing

Schneider Electric's AI applications do not appear mainly as standalone "AI products". Instead, AI is embedded across the path from value definition and software design to integrated software-hardware systems. The less AI appears as a standalone layer, the closer it comes to the real role of AI in industrial companies.

Core Insight: AI Creates Value by Becoming Embedded
Schneider Electric did not present AI as a separate "new story". AI mattered because it was built into products, systems and customer relationships—not as a visible add-on, but as a capability that steadily reshaped how value was created.

AI ACROSS THREE KEY AREAS: PRODUCT DEVELOPMENT × CUSTOMER SERVICE × MARKETING

Product Development:
From Simple Execution to Intelligent Hardware
Embedded Algorithms · AI-Native Architecture

AI Role: Embedded Algorithms and AI-Native Architecture
- Hardware moves beyond simple execution towards sensing, diagnosis, prediction and prevention
- AI is built into software from the earliest design stage
- AI-native architecture connects data flows and feedback loops across the system
- Systems can evolve and upgrade as the technology advances
- EcoStruxure Power Operation uses prompts and Skills to support AI-native development

AI Positioning: Core Infrastructure in System Architecture

Customer Service:
Extending Solution Value
Optimisation · Scheduling · Predictive Maintenance

AI Role: Core Capability Within Solutions
- AI combines with distributed computing, open automation and data-centre energy management
- Supports optimisation, scheduling and predictive maintenance
- Embedded within solutions rather than offered as a standalone product
- Coordinates energy systems and IT loads under fluctuating AI demand

AI Positioning: Core Capability Within Customer Solutions

Marketing:
B2B Digital Marketing and Customer Insight
WeChat Ecosystem · Segmented Engagement · Monthly Active Users

AI Role: Content Production, Customer Insight and Lead Identification
- WeChat public-facing channels, such as service accounts, video accounts and mini-programmes, maintain broad customer reach
- WeCom supports segmented engagement through private communities and service platforms
- AI supports content production, customer insight and lead identification
- Metrics shifted from follower numbers to monthly active users and high-engagement users

AI Positioning: Precision Management of Industrial Buying Decisions

DIGITAL MARKETING IN CHINA: A BENCHMARK FOR PRIVATE-DOMAIN OPERATIONS

Two-Layer B2B Model Centred on WeChat:
China Contributes Nearly Half of the Group's Monthly Active Users
AI-driven content production and customer insight support more precise management of the industrial buying process

Global Leader

Public-Facing Reach
- Service accounts: regular delivery of brand and product content
- Video accounts: technical explanations and scenario-based content
- Mini-programmes: integrated product enquiry and service access
- AI supports content production, improving reach efficiency and quality

Private-Domain Engagement
- WeCom-based communities and service platforms tailored to different customer roles and tiers
- AI supports customer insight and sales lead identification
- Evaluation shifted from follower numbers to monthly active users and high-engagement users
- AI helps manage—but does not replace—the industrial sales process

Schneider Electric China contributes
nearly half
of group global monthly active users

China ranks
first globally
in monthly active users

Evaluation prioritises
monthly active and high-engagement users
over follower numbers

COMMERCIAL CHALLENGE: MOVING BEYOND ONE-OFF HARDWARE PURCHASES

B2B Procurement Challenge: From One-Off CAPEX to Ongoing Subscriptions and Operations
For Schneider Electric, technology transformation also required market education.

Commercial Challenge

Traditional Hardware Sales Model
One-Off Capital Expenditure (CAPEX)
- Large, one-off equipment purchases with clear capital timing
- Tangible hardware assets made value easier to see and decisions faster
- Mature procurement processes were often led by purchasing departments

Software and Service Model
Ongoing Subscriptions and Operations (OPEX)
- Subscription models required customers to rethink capital allocation
- Software value emerged through long-term operation and service delivery
- Decision-making involved IT, operations, finance and other departments
- Longer evaluation cycles increased procurement complexity

For Schneider Electric, technology transformation had to be accompanied by **market education**. The company needed to help customers move away from a "hardware-heavy, software-light" purchasing mindset and from one-off purchases towards longer-term operating relationships. This was one of the deeper challenges created by the shift in Schneider Electric's growth logic.

Source: Company interviews and publicly available materials

More broadly, this software- and service-led model also created a commercial challenge for Schneider Electric. Traditional B2B hardware sales were usually treated as one-off capital expenditures: customers bought equipment, paid for it and owned it. Software and system services worked differently. Their value depended on ongoing subscriptions and continuing operations over time.

For Schneider Electric, this meant that technology transformation had to be accompanied by market education. The company not only had to restructure its technical architecture; it also had to help customers move away from a "hardware-heavy, software-light" purchasing mindset. In practical terms, that meant persuading customers to move from one-off purchases towards longer-term operating relationships. This was one of the deeper challenges created by the shift in Schneider Electric's growth logic.

This also explains why Schneider Electric did not present AI as a separate "new story". AI mattered precisely because it was embedded in the company's products, systems and customer relationships. The less it appeared as a standalone layer, the closer Schneider Electric came to the real role of AI in industrial companies: not taking centre stage, but steadily reshaping how value was created.

IV. Case Takeaways: From Hardware Provider to Architecture Shaper

What made the Schneider Electric case representative was not simply that an industrial giant had begun to develop software and use AI. More importantly, through a software-first strategy, software-defined architecture and AI capabilities, Schneider Electric was moving beyond the role of a traditional hardware provider. It was not simply selling hardware, nor was it simply selling software. It was defining how software and hardware worked together, how systems opened up, and how computing power and energy were reconfigured.

The real change brought by AI and digital transformation was therefore not the addition of a "digital business line". It was a change in the company's growth logic. Although software and services still accounted for a smaller share of revenue than hardware, they were increasingly shaping the order of product development, the rules of the ecosystem and the way high-value services were embedded into integrated software-hardware application systems.

In this sense, the Schneider Electric case was not simply a story of a hardware company becoming more software-driven. It revealed a deeper shift: once software begins to define hardware and AI becomes embedded in systems, industrial companies are no longer competing only over individual products. They are competing for architectural leadership.

V. Looking Ahead: From AI-Augmented to AI-Native Systems

As AI technology advances, system design will shift from AI augmentation—layering AI capabilities onto legacy systems—to AI-native design, which starts with AI infrastructure and then rebuilds processes, organisations and systems around it. The goal is to create new business value by making AI part of the system's underlying architecture from the outset.

The automotive industry offers a useful analogy. Traditional internal-combustion vehicles were built around the engine and gearbox. Smart features could be added later, but they remained additions to the same underlying mechanical architecture. Electric vehicles changed the logic more fundamentally: the centre of gravity shifted from the engine and transmission to the battery, motor and electronic control system. That shift did not simply change the product; it helped reshape the structure of the industry and paved the way for smart vehicles.

Schneider Electric's work on alarm design for power distribution systems illustrates the difference. In a traditional system, equipment generates an error code. Operations and maintenance staff then look up the code in a manual, identify the likely fault and decide what action to take. An AI-augmented system makes this process faster by using AI to perform the lookup, but the basic logic remains the same: each type of fault still has to be modelled in advance, and updates are made over weeks or months. AI improves the workflow, but it does not change it. People still have to turn the result into reports, work orders and inspection steps before maintenance can be carried out.

The AI-native approach starts from a different premise. Instead of modelling every possible fault in advance, Schneider Electric can train an AI agent for the power distribution domain, using a large language model and natural-language Skills. A Skill is a standardised task capability, written in natural language, that the agent can call on when needed. Skill documents effectively "teach" the agent how to diagnose and respond to faults. When a new capability is needed, the team can add a new Skill document rather than rebuild the fault model, reducing iteration cycles from weeks or months to hours or days. In this model, the system is organised around large models, agents and Skills from the outset. AI is not added later; it is part of the architecture from day one.

This gives Schneider Electric distinctive advantages in AI-native system design. The company has deep expertise in power distribution and industrial systems, along with extensive fault cases and practical handling experience. These provide valuable resources for training agents. Its expert teams can also convert industry knowledge into natural-language Skills, turning accumulated domain expertise into capabilities that an agent can use. This ability would be difficult for companies without deep industry experience to replicate.

Schneider Electric's EcoStruxure Power Operation, an integrated power management and control platform, shows this AI-native thinking in practice. The team has adopted a development approach centred on prompts and Skills, and the methodology emerging from this work could support further AI-native development.

FROM AI-AUGMENTED TO AI-NATIVE SYSTEMS: POWER DISTRIBUTION ALARM DESIGN

AI Augmentation vs AI-Native: Different System Logic
AI augmentation improves the workflow; AI-native design changes the architecture. Systems move from traditional data models to large models, agents and Skills.

Power Distribution System Alarm Design

AI-Augmented Model:
AI Layered onto Existing Systems—Underlying Logic Unchanged

Use AI to Replace Manual Error-Code Lookup
- Equipment generates error codes; AI performs the lookup previously done by operations staff
- Efficiency improves, but each fault type still requires separate modelling
- New fault types require model updates, creating a heavy engineering workload
- AI remains an assistive tool: reports, work orders and inspections still require human follow-up
- Result: iteration cycles are measured in weeks or months

Architecture Leap

AI-Native Model:
AI as Core Infrastructure from Day One

Large Language Model + Agent + Skills
- Train agents in the power distribution domain
- Skills are natural-language task capabilities that agents can call on
- Skill documents teach agents how to diagnose and respond to faults
- New capabilities are added through new Skill documents, not new fault models
- AI is foundational infrastructure, not a later add-on

Iteration cycles are compressed to hours or days

Iteration Cycle Comparison: AI-Native Accelerates Updates

		Iteration Cycle
Traditional	Manual Error-Code Lookup	
AI-Augmented	AI Lookup, Fault-by-Fault Modelling	Weeks to months
AI-Native	Skill Files Enable Rapid Updates	Hours to days

Schneider Electric's Distinctive Strengths in AI-Native Design

Industry Knowledge:
Deep expertise in power distribution and industrial systems, supported by extensive fault cases and handling experience

Knowledge Conversion:
Expert teams convert industry knowledge into natural-language Skills that agents can use

EcoStruxure Power Operation:
Integrated power management and control platform showing AI-native thinking in practice

VI. Data Centres as a New Growth Engine in the AI Computing Wave

Data centres, another important part of Schneider Electric's business, were pushing the logic of software-defined systems into an even more demanding setting.

Schneider Electric's 2025 results showed that data centres were becoming a core engine of global growth. Demand was strong in North America and parts of Europe, while China also achieved positive growth, supported by demand from the data-centre market. This growth was driven by strong demand for AI computing power.

As single-rack power density moved from the kilowatt level towards the megawatt level, existing power-supply architecture, transformers, cooling systems and power distribution systems all faced major reconfiguration. AI was not simply increasing the number of servers. It was creating an urgent need for deeper coordination and optimisation across the entire data-centre energy and automation system.

Hardware is only one part of the challenge. AI workloads fluctuate sharply between training and inference, which means data centres must keep extremely dense, power-hungry hardware clusters running reliably and efficiently under constantly changing conditions. The real issue is therefore not simply whether the equipment is reliable, but whether power systems, cooling systems and IT loads can be managed as one coordinated system. This is where software-defined data-centre capabilities become critical.

Schneider Electric's role in this field goes well beyond building data-centre facilities. Its capabilities span the full data-centre life cycle, from design and construction to operation. They include AC UPS systems, high-voltage DC power supply, air and liquid cooling, design simulation and energy-efficiency optimisation. This end-to-end capability matters because, once single-rack power reaches several hundred kilowatts, no single element can be optimised in isolation. Power supply, cooling efficiency, space utilisation and operational performance all have to be managed together.

This also explains Schneider Electric's view that data-centre competition is shifting from scale to efficiency. The key question is no longer simply who can build more capacity, but who can support higher-density power architectures and use smarter energy scheduling to reduce both computing and energy costs. Technical competition is therefore moving beyond the performance of individual devices towards the efficiency and resilience of the energy-computing system as a whole.

The economics are clear. Power architecture alone can create significant differences in energy loss: AC and DC power supply systems can differ by more than a factor of two. In an AI data centre running tens of thousands of GPUs, that difference could translate into tens of millions of yuan in annual electricity costs.

This points to a broader lesson in Schneider Electric's growth logic. A software-first approach did not make hardware less important. Instead, it raised the bar for what hardware had to do. As AI advances, data centres need physical infrastructure that is not only efficient and reliable, but also measurable, controllable and able to work in coordination with other systems. The competitive advantage therefore lies in integrated software-hardware capability: combining energy management, automation and software optimisation into one system. Companies that can do this are better placed to turn the data-centre boom into long-term business value, rather than one-off project revenue.

DATA CENTRES AS A NEW GROWTH ENGINE IN THE AI COMPUTING WAVE

Data Centres: AI Computing Demand Raises the Bar for Software-Hardware Integration
Single-rack power density is moving from kilowatts towards megawatts
Power supply, cooling and distribution systems face major reconfiguration

Core Growth Engine in 2025

DEMAND-DRIVEN GROWTH

AI Computing Demand Drives Reconfiguration

- Single-rack power density is moving from kilowatts towards megawatts
- Power supply, transformers, cooling and distribution systems face major reconfiguration
- Strong demand in North America and parts of Europe; China also achieved positive growth
- By 2025, data centres were becoming a core global growth engine for the group

CORE CHALLENGE

Managing Efficiency Under Volatile AI Workloads

- High-power clusters must run reliably and efficiently despite sharp training and inference load fluctuations
- Energy systems and IT loads must be coordinated in real time
- AC and DC power losses can differ by more than a factor of two
- In AI data centres with tens of thousands of GPUs, this can mean annual electricity cost differences in the tens of millions of yuan

SCHNEIDER ELECTRIC'S EDGE

End-to-End Lifecycle Integration

- Covers the full life cycle: design, construction and operation
- Integrates AC UPS, high-voltage DC power supply, air cooling and liquid cooling
- Combines design simulation with energy-efficiency optimisation
- Optimises power supply, cooling, space utilisation and operations together as rack power rises into the hundreds of kilowatts

Data-centre competition is shifting from scale expansion to **efficiency**. The key question is who can support higher-density power architectures and use smarter energy scheduling to reduce both computing and energy costs. **Software does not replace hardware; it raises the bar for integrated software-hardware competition**.

THREE KEY SHIFTS IN SCHNEIDER ELECTRIC'S GROWTH LOGIC

1.
Product Development: From Hardware-Led to Software-Led
Software shapes the development sequence; hardware follows software design.

2.
Ecosystem Rules: From Device Lock-In to Open Architecture
Open automation breaks down closed-system barriers and shifts advantage towards ecosystem leadership.

3.
Competitive Focus: From Product Performance to Architectural Leadership
Competition moves beyond individual products; efficiency and system-level integration become the basis of advantage.

Source: Company interviews and publicly available materials

As software begins to define hardware and AI becomes embedded in systems, the basis of competition for industrial companies is shifting towards architectural leadership. But this is only the next stage, not the end point. As AI applications deepen, system design is moving from AI-augmented to AI-native: architectures are being rebuilt with AI as the starting point. Companies with deep expertise in energy and industrial automation will be well placed to gain an early advantage, because they can turn accumulated industry knowledge into capabilities that AI agents can use.

At the same time, the computing demand behind AI is turning data centres into a new growth engine. As single-rack power density rises, competition is moving from scale to efficiency. This creates demand for software-defined power architectures, AI-driven energy scheduling and integrated software-hardware systems. These are precisely the capabilities on which companies such as Schneider Electric can compete.

CASE 7

Matrix Design's AI Journey: From Industry Know-how to Vertical Models and an Expanding AI Ecosystem

Matrix Design Co., Ltd. ("Matrix Design") was founded in 2010 and listed on ChiNext, the Shenzhen Stock Exchange's market for growth enterprises, in 2022. It is now recognised as a leading player in China's high-end interior design market, with a portfolio spanning residential, workplace, hospitality and commercial projects.

For a mature, creativity-intensive business like Matrix Design, AI was never just another design tool. Its significance lay in the possibility of addressing some of the structural tensions that had long constrained the interior design industry: the tension between bespoke creativity and scaled delivery; the trade-off between costly, time-consuming visualisation and the need for faster client response; and the limits of a traditional service model that still depended heavily on expanding headcount.

Against this backdrop, Matrix Design's response to the generative AI (AIGC) wave was not to build a collection of isolated applications. Over the past two years, the company developed Ark.art, a proprietary vertical AIGC platform designed around the full lifecycle of design delivery.

Through Ark.art, Matrix Design gradually connected early-stage exploration, model training, data cleaning, tool development and business model validation. The aim was to turn AI from a supporting tool into operational infrastructure that could reshape productivity. At the same time, this marked a broader shift in the company's role: from a traditional design service provider into a company capable of exporting AI-enabled commercial capabilities and helping build a wider industry ecosystem.

I. Beyond Generic Tools: Building a Vertical AI Strategy

Matrix Design began exploring AIGC in late 2022. In the early stage, an internal team known as the "Pioneer Society" systematically tested general-purpose image generation models such as Midjourney and Stable Diffusion. The aim was to understand how far these tools could support real design workflows.

The results quickly showed the limits of generic AI in a professional design setting. While these tools were useful for open-ended creative exploration, they could not be converted directly into reliable production capability. In interior design, their weaknesses became particularly clear.

The first issue was the professional usability of the outputs. General-purpose models had not been trained deeply enough on the physical scale of interior spaces or on architectural standards. Their outputs often suffered from distorted perspective, inconsistent material textures, mismatched furniture proportions and, in some cases, design flaws that could not be built in the real world. Designers still had to spend substantial time correcting and refining the images, which weakened the practical efficiency gains from using AI.

The second issue was the barrier to adoption. Open-source models were often built around an engineer's logic, with complex parameters and code-based settings. This made them poorly suited to the working habits of interior designers and left them without a natural, intuitive way to interact with the tools. The randomness of the generation process also made it difficult to maintain a stable professional standard across multiple design iterations.

For Matrix Design, these limitations pointed to a broader strategic conclusion: AI would struggle to enter core business workflows if it remained merely a tool layered on top of existing processes. In June 2023, the company formally launched Ark.art as an internal incubation project, moving beyond the usual resource logic of a traditional design firm and committing to a more ambitious path: building its own vertical AI capability rather than relying on external tools.

The goal was to reduce passive dependence on the evolution of external general-purpose tools. Instead, Matrix Design would build on its own data foundation to develop vertical large-model applications tailored to the aesthetics, standards and workflows of interior design.

II. Building a Technology Moat: Turning Tacit Industry Know-how into Digital Assets

What made Ark.art commercially viable was not simply the use of large-model technology. Its real foundation was Matrix Design's proprietary data lake. Unlike generic models trained on broad public internet data, Ark.art was built on the company's own private data resources.

In the early stages of R&D, however, the company had to work through a period of organisational and technical integration. AI algorithm engineers and interior designers brought different ways of thinking and different standards for evaluating results, creating clear barriers to cross-disciplinary communication.

Matrix Design responded by redesigning its collaboration mechanisms. It introduced product managers with both design and software development experience to act as cross-disciplinary "translators", helping convert meaning between two professional languages.

When the project risked following the logic of a standalone technology product rather than the needs of the business, the company adjusted its technical team structure and made one principle clear: technology development had to remain closely aligned with core business use cases.

The company also needed a way to handle subjective disagreements over the quality of model outputs. To do so, it drew on its long-running design review and voting mechanism. A panel of senior designers acted as an arbitration group, using collective professional judgement to define the benchmarks for model optimisation.

Ark.art's core training data came from more than a decade of Matrix Design's accumulated project experience, including thousands of high-end real-world projects, tens of thousands of high-resolution renderings and the accompanying design texts.

The more important step was making experience explicit. Matrix Design assembled senior designers from across its business units to form a style annotation team, which structured and labelled the company's private image library. The team translated tacit aesthetic judgement—such as spatial circulation, material and lighting effects and colour composition—into parameter tags that AI could read. It also codified industry terminology developed through years of serving leading real estate clients, including labels such as “New Chinese style” and “cream-toned interiors”.

On this basis, the R&D team used LoRA (Low-Rank Adaptation) fine-tuning to continuously feed the model with data from nearly 800 new Matrix Design projects each year. This combination of high-quality proprietary data, precise semantic annotation and targeted fine-tuning turned the company’s long-accumulated core know-how into hard-to-replicate foundational digital assets. It also gave the AI model a form of professional judgement more closely aligned with real business logic.

TECHNOLOGY MOAT ANALYSIS

Beyond Generic Tools: Turning Industry Know-how into a Vertical Model Moat

Ark.art shows how the gap between generic AIGC tools and professional design workflows can push companies towards in-house vertical models. By building on proprietary data, codifying design know-how and using LoRA fine-tuning, Matrix Design turned more than a decade of project experience into hard-to-replicate foundational digital assets.

AI ACROSS THREE KEY AREAS: PRODUCT DEVELOPMENT × CUSTOMER SERVICE × MARKETING

General-purpose AIGC tools (Midjourney / Stable Diffusion) Generic training · engineering-led interfaces · variable outputs	Self-developed vertical solution	Ark.art's vertical foundation model Proprietary data lake · design-specific logic · professional-grade outputs
PROFESSIONAL USABILITY Outputs often showed distorted spatial perspective, inconsistent material textures and misaligned furniture proportions, and sometimes included structural flaws that could not be built. **Low professional usability**		PROFESSIONAL USABILITY Trained on Matrix Design's archive of high-end real-world projects, the model produces design visuals that better reflect architectural standards and realistic spatial proportions. **Commercial usability**
INTERACTION BARRIER Interfaces were engineering-oriented, relying on complex parameters and code-based settings that did not match designers' working habits. **High adoption barrier**		INTERACTION BARRIER Functional modules—including Inspiration Search, Master Style, Visualisation Generation and Partial Reconstruction—are designed around designers' actual workflows and support natural language interaction.. **Lower adoption barrier**
OUTPUT STABILITY Generated results were highly variable, making it difficult to maintain a stable professional standard across iterations and leaving designers with substantial correction work. **Low stability**		OUTPUT STABILITY Senior designers set optimisation benchmarks through a review-and-voting mechanism, while new project data is continuously fed back into the model. **Professional stability**
DEPTH OF DOMAIN KNOWLEDGE Limited training on interior spatial scale and architectural standards, with weak command of industry style labels such as "New Chinese style" and "cream-toned interiors". **Shallow generalisation**		DEPTH OF DOMAIN KNOWLEDGE Senior designers codify tacit aesthetic judgement—spatial circulation, material and lighting effects and colour composition—into structured parameter tags that AI can read. **Vertical depth**

A FOUR-STEP PATH FROM TACIT KNOW-HOW TO FOUNDATIONAL DIGITAL ASSETS

Proprietary data lake × knowledge codification × LoRA fine-tuning × continuous improvement
Turning more than a decade of design know-how into hard-to-replicate foundational digital assets

Building a Technological Moat

01
Build a proprietary data lake
- Thousands of high-end real-world projects accumulated over more than a decade
- Tens of thousands of high-resolution renderings and accompanying design texts
- Private Matrix Design data, not broad public internet data

Data moat: proprietary assets competitors cannot easily replicate

02
Annotate data to codify experience
- Style annotation team of senior designers from across business units
- Spatial circulation, material and lighting effects and colour composition converted into structured parameter tags
- Industry style terms codified, including "New Chinese style" and "cream-toned interiors"

Knowledge moat: making tacit experience explicit

03
Create collaboration mechanisms
- Product managers with design and software backgrounds acting as cross-disciplinary translators
- Senior designer panel setting model optimisation benchmarks
- Technology development kept aligned with core business use cases

Organisational moat: cross-disciplinary collaboration

04
Fine-tune with LoRA for continuous improvement
- Low-Rank Adaptation (LoRA) fine-tuning
- Nearly 800 new projects each year fed back into the model
- Model improvement driven by ongoing project accumulation

Evolution moat: strengthens as the business keeps using it

SCALE OF THE PROPRIETARY DATA FOUNDATION

Thousands of Projects	Tens of Thousands of Renderings	~800 New Projects/Year	June 2023
High-end real-world cases accumulated over 10+ years	High-resolution renderings with accompanying design texts	New project data continuously fed back into the model	Ark.art internal incubation project formally launched

Source: Interviews with Matrix Design (Ark.art); publicly available company disclosures

III. Reshaping Core Business Workflows: From Linear Coordination to Agile Delivery

Ark.art's web platform was officially launched in March 2024. Its modules—Inspiration Search, Master Style, Visualisation Generation and Partial Reconstruction—were designed closely around the actual workflows of interior designers. As AI became embedded in Matrix Design's internal operations, it began to change the resources required for traditional design projects and the time needed to deliver them.

In a conventional interior design process, concept design, scheme development and visualisation production usually follow a linear sequence, requiring nearly 120 hours of coordinated work across multiple teams. Even small adjustments at the early creative stage often lead to substantial rework in later modelling and rendering.

Ark.art significantly shortened the time required to generate interior visualisations. After designers input a floor plan, key materials and style prompts, the system can produce multiple high-quality renderings and material texture options within minutes. In the Partial Reconstruction module, when clients request changes such as replacing a specific piece of furniture or material, the AI can quickly redraw the relevant area and automatically recalculate light and shadow. This directly addresses a long-standing pain point in traditional rendering workflows: the high cost and low efficiency of repeated revisions.

Internal application data showed that, within five months of launch, Ark.art had been validated across 66 large-scale interior-space projects. Across the broader visual production workflow, it reduced the total time needed to generate design visuals to around eight hours, saving designers an average of 75% of their previous output time. In the especially time-consuming final rendering stage, AI shortened the average cycle from five days to two hours—an efficiency improvement of around 98.3%.

In early-stage commercial bidding, Ark.art's ability to generate outputs frequently and in parallel enabled design teams to provide a richer range of visual proposals at lower time cost. This agile response capability increased Matrix Design's project win rate by around 25%.

More broadly, this marked a change in part of the decision-making logic of interior design. In the past, the high cost of producing design visuals meant that design iteration was cautious and infrequent. Now, the process could shift towards high-frequency generation, rapid screening and targeted iteration, significantly reducing the cost of commercial trial and error. In this process, AI mainly took over exploratory and repetitive execution tasks, while final professional judgement and trade-offs remained the responsibility of human designers.

BUSINESS WORKFLOW RECONFIGURATION

From 120 Hours to 8: How Ark.art Makes Design Delivery More Agile

Ark.art cut the overall time needed to generate design visuals from around 120 hours to around 8 hours, saved designers an average of 75% of their output time and increased project win rates by around 25%.

KEY PERFORMANCE METRICS (INTERNAL APPLICATION DATA, VALIDATED WITHIN FIVE MONTHS OF LAUNCH)

98.3%	75%	25%	66
improvement in final rendering efficiency Average cycle: 5 days → 2 hours	average visual-output time saving Overall visual production: 120h → 8h	increase in project bid win rate Agile response improved bid competitiveness	large-scale interior-space projects validated Within five months of launch

TRADITIONAL LINEAR COORDINATION VS. ARK.ART AGILE DELIVERY

Traditional Process vs. Ark.art Process: How Design Delivery Changes

Human-in-the-Loop

AI takes on exploratory and repetitive execution tasks; final professional judgement and trade-offs remain with human designers.

Traditional Linear Coordination   Multi-team handoffs · minor upstream changes trigger extensive downstream rework

Concept Design	Scheme Development	Modelling & Rendering	Client Revisions	Design Finalisation
~20 hours Creative ideation and style direction	~30 hours Spatial layout and material selection Upstream changes → downstream rework	~5 days 3D modelling and final rendering High cost · limited revision flexibility	Multiple rounds Furniture and material changes require re-rendering Light and shadow recalculated each time	Cautious iteration High rendering costs limit the pace and scope of iteration

Total time (multi-team coordination) **~120 hours**

Ark.art Agile Delivery   High-frequency generation · rapid screening · targeted iteration

Creative Direction	AI Generation	Designer Selection	Partial Reconstruction	Design Finalisation
~1 hour Style and core concept defined by designers Human-led creative decisions	**Minutes per option** Floor plans, key materials and style prompts generate multiple high-quality options AI handles exploratory execution	**Rapid decisions** Designers screen multiple options and make the final trade-offs Human-led judgement	**Fast revisions** AI redraws selected areas and recalculates light and shadow, cutting costly rendering rework AI handles repetitive execution	**High-frequency iteration** Rapid generation and screening reduce commercial trial-and-error costs Win rate +25%

Total time (AI-assisted visual production) **~8 hours** 93.3% reduction

A SHIFT IN DESIGN DECISION-MAKING LOGIC

BEFORE

Low-frequency, cautious iteration

- High visual-production costs made design iteration cautious and infrequent
- Minor upstream changes trigger extensive downstream rework
- Bid proposals were limited in number, with visual variety constrained by time
- High commercial trial-and-error costs kept solution exploration conservative

NOW

High-frequency generation · rapid screening · targeted iteration

- Lower marginal cost supports frequent generation and rapid screening
- Partial Reconstruction enables low-cost revisions and reduces rework
- Greater visual variety in bid proposals supports a 25% higher win rate
- Lower commercial trial-and-error costs expand the space for exploration

Source: Matrix Design (Ark.art) internal application data; company interviews (Mar–Aug 2024)

IV. Redesigning the Organisation: New Roles, New Collaboration and New Incentives

As Ark.art became embedded in Matrix Design at the system level, the company also had to adjust the way it organised and managed the work around it. Rather than keeping the platform inside the R&D department, Matrix Design made deeper changes to its organisational structure and incentive mechanisms, turning Ark.art into a collaborative project involving employees across the company.
One of the first changes was the transformation of design-related roles. As AIGC reduced the role of manual modelling in traditional rendering work, Matrix Design began upgrading the skills of the affected employees. Staff with strong aesthetic foundations were guided into new roles as “AI model trainers” and “style artists”, where they contributed directly to Ark.art’s parameter tuning and style development.

The role of core designers also expanded. They were no longer only creative producers; they increasingly became managers of human–AI collaboration. As AIGC lowered the basic technical threshold for producing drawings and visualisations, the commercial value of design shifted more quickly back towards logical thinking, user insight and overall solution coordination. Designers therefore needed to build new capabilities in prompt engineering, data analysis and model tuning. Internal practice showed that, in AIGC projects, designers with cross-disciplinary capabilities achieved markedly better overall performance than was typical under traditional working models.

Matrix Design also changed how it evaluated and rewarded contribution. To support this new form of collaboration, the company moved away from a reward and allocation mechanism centred on a single design function and introduced a gamified DKP (Dragon Kill Points) system for quantifying performance in AI-related projects.

The DKP system measured each member’s contribution to cross-departmental AI projects. It expanded evaluation beyond individual output to include cross-disciplinary collaboration, knowledge sharing and model optimisation. Data showed that, after AIGC was introduced, cross-disciplinary team projects recorded an average increase of around 28.5% in DKP contribution value, alongside improvements in the novelty and feasibility of solutions.

The partner equity incentive plan further reinforced this shift. Together, these mechanisms helped Matrix Design build a stronger organisational foundation for embracing technological change.

V. Business Model Expansion: From Internal Efficiency Tool to Home and Living Ecosystem Platform

The deeper value of enterprise AI is not limited to internal efficiency. Once a company proves that AI can improve core operations, the next question is whether that capability can open up new business boundaries. After Ark.art had validated system-level efficiency gains inside Matrix Design, the company faced a strategic choice: keep the technology internal to protect its design differentiation, or open it externally and build a broader commercial ecosystem.

Matrix Design chose the second path. In early 2024, it established Ark Technology (Shenzhen) Co., Ltd. as a wholly owned subsidiary, marking the start of a new growth curve: from providing design services to exporting AI-enabled commercial capabilities.

1. SaaS for Vertical Design Markets (B2C)

Ark Technology first targeted China's large base of independent interior designers and small-to-medium home renovation studios. It launched Ark.art as a SaaS subscription service, offering professional-grade visual output, proprietary models aligned with domestic aesthetic preferences and a relatively low barrier to use.

This positioning quickly attracted attention in the industry. By early 2025, Ark.art had drawn more than 20,000 registered external designer users. By focusing on core pain points in interior design, the product steadily raised its seven-day new-user retention rate to 22%. It also established early commercial validation and a recurring revenue base in the vertical design tools market.

2. Customisation and Workflow Integration for Enterprise Clients (B2B)

In the enterprise market, Ark Technology offered tiered subscriptions and customised private deployment services for home furnishings and building materials companies, as well as real estate developers. For example, it developed a dedicated AI style model for Dongpeng, a well-known building materials brand.

Enterprise clients could upload high-resolution product assets to a private library. Ark.art could then automatically place these real products into generated, on-trend interior-space renderings. This moved Ark.art beyond the role of a standalone design tool and gradually embedded it in enterprise clients' workflows for new product development, marketing and display.

3. Towards a Full-Chain Digital Industry Ecosystem (Platform)

Ark Technology's longer-term strategy is to evolve into an AIGC ecosystem platform for the home and living industry. The sector faces two structural pain points: design schemes are often disconnected from final product procurement, and real-world product photography is costly.

Ark.art entered the finished furniture and e-commerce marketing segment to address these problems. Home furnishings companies no longer needed to invest heavily in physical photography studios. Instead, they could use Ark.art to generate marketing scenes at scale, incorporating aesthetic preferences from different overseas regions. This reduced the cost of marketing assets for cross-border e-commerce and improved the efficiency of product display.

To accelerate the development of this ecosystem, Ark Technology took a series of steps across both industry and capital. In July 2024, the company showcased Ark.art at the inaugural China Home Furnishings AI Design Conference at the Guangzhou CBD Fair, where it released an AIGC solution for the home and living industry.

In September 2024, Ark Technology secured RMB 5 million in strategic investment from Jifu Fund, which is backed by state-owned capital from Shenzhen's Futian District, through a capital increase that gave the fund a 10% equity stake. This provided funding for Ark.art's ecosystem development and signalled recognition of its industry platform model by external capital and government-guided funds.

In December 2024, Ark Technology partnered with Tencent Cloud and PhotoG to establish the "Design AIGC Joint Lab". Drawing on Tencent Cloud's underlying computing power and data architecture, the lab further advanced the digitalisation of the real estate and home and living industries.

ANALYSIS OF BUSINESS MODEL EXPANSION

From Internal Efficiency Tool to Home and Living Ecosystem Platform

Matrix Design externalised Ark.art through Ark Technology, extending AI-enabled commercial capabilities from internal design workflows into SaaS, enterprise integration and a broader home and living ecosystem.

THREE-STAGE BUSINESS MODEL EVOLUTION: INTERNAL TOOL → SAAS/ENTERPRISE INTEGRATION → INDUSTRY ECOSYSTEM PLATFORM

Stage 1 · Internal Validation
Internal efficiency tool
2023–March 2024 · Incubation phase

VALUE PROPOSITION: COST REDUCTION AND EFFICIENCY GAINS

- Ark.art was incubated internally to support Matrix Design's own design operations
- Overall visual production cycle reduced from 120h to 8h
- Validated across 66 large-scale interior-space projects
- ~25% increase in project bid win rate

Core outcome: 98.3% efficiency improvement

Stage 2 · Market Expansion
SaaS (B2C) + Enterprise Integration (B2B)
Early 2024 · Ark Technology founded

VALUE PROPOSITION: REVENUE GROWTH AND RECURRING REVENUE

- B2C: SaaS subscriptions for independent designers and small-to-medium home renovation studios
- B2B: tiered subscriptions and private deployment for home furnishings and building materials companies and real estate developers
- Custom AI style model for Dongpeng, with products seamlessly embedded in renderings
- Seven-day new-user retention rate reached 22%

Core outcome: 20,000+ registered external designer users

Stage 3 · Ecosystem Development
AIGC Ecosystem Platform for the Home and Living Industry
From H2 2024 onward · Strategic development phase

VALUE PROPOSITION: BUSINESS MODEL EXPANSION AND ECOSYSTEM NODE

- Entry into finished furniture + e-commerce marketing, enabling large-scale generation of cross-border e-commerce visuals
- Reduced reliance on physical photography studios and cut cross-border e-commerce marketing asset costs
- Secured RMB 5 million from Jifu Fund through a capital increase for a 10% equity stake
- Established the Design AIGC Joint Lab with Tencent Cloud and PhotoG

Core outcome: platform model recognised by external capital and government-guided funds

TIMELINE OF INDUSTRY AND CAPITAL STRATEGY (2023–2024)

Key Milestones: From Internal Incubation to an Industry Ecosystem Node 2023–2024
Moved from internal incubation towards an industry ecosystem node, supported by capital investment and industry partnerships.

June 2023
Ark.art internal incubation project launched
Matrix Design formally launched a vertical in-house development strategy to reduce reliance on general-purpose tools.
Strategic decision

March 2024
Ark.art web platform launched
Four core functional modules went live and were validated across 66 large-scale interior-space projects within five months
Product launch

July 2024
Guangzhou CBD Fair: First China Home Furnishings AI Design Conference
Ark Technology showcased Ark.art, released a home and living AIGC solution and began building industry influence
Industry launch

Sep–Dec 2024
Capital and Ecosystem Validation
RMB 5 million strategic investment from Jifu Fund; Design AIGC Joint Lab with Tencent Cloud and PhotoG
Ecosystem development

ORGANISATIONAL AND INCENTIVE ADJUSTMENTS FOR BUSINESS MODEL EXPANSION

Talent Structure Evolution: From Creative Producers to Human–AI Collaboration Managers

- Rendering specialists → AI model trainers and style artists
- Core designers → human–AI collaboration managers
- Design/software product managers → cross-disciplinary translators
- Design value → logic, user insight and solution coordination

DKP-Based Performance System: Driving Cross-Functional Collaboration and Knowledge Sharing

- The gamified DKP system quantified contributions in cross-departmental AI projects
- Evaluation expanded beyond individual output to include cross-disciplinary collaboration, knowledge sharing and model optimisation
- Partner equity incentives strengthened the organisational foundation for technological transformation

28.5% increase in average DKP contribution value across cross-disciplinary team projects after AIGC adoption

Sources: Interviews with Matrix Design (Ark); public disclosures by Ark Technology; Guangzhou CBD Fair publications (2024)

VI. Key Takeaways: AI-Driven Evolution in Creative Services

Matrix Design's incubation of Ark.art shows how AI can reshape how a creative services company operates. The case is not simply about using AI to make one task faster. Its significance lies in how Ark.art created value across several dimensions at once: reducing costs and improving efficiency by sharply shortening visual production time; supporting growth by improving bid win rates and opening up SaaS revenue; and expanding the business model by helping build a wider home and living ecosystem.

The case also shows that this kind of transformation tends to unfold in stages. Matrix Design began by using general-purpose tools to test whether AI could address specific pain points. It then moved into broader user adoption and commercialisation, before reaching deeper changes in the company's underlying management mechanisms.

Three broader lessons emerge.

First, when companies pursue AI-driven transformation, competitive advantage does not come from simply connecting to a general-purpose foundation model. It depends on whether a company can turn high-quality proprietary business data and vertical industry know-how into its own foundational model system. This is what allows AI outputs to be commercially usable while also meeting professional standards of quality and aesthetic judgement.

Second, the harder challenge in moving AI into deeper applications often lies inside the organisation. Companies need to redesign roles, adjust performance mechanisms and renew collaboration systems at the same time, so that employees can move towards new forms of human–AI collaboration. Only when the organisation provides systematic support for AI can companies move beyond local efficiency gains and achieve a broader improvement in overall effectiveness.

Third, externalising internal capability is an important route for companies rooted in real-world operations seeking to expand their business boundaries and explore new business models. Once a company has developed mature digital capabilities internally, it can make those capabilities explicit, modularise them and bring them to market through an independent organisational vehicle. In doing so, it can move beyond the limits of its original physical service model.

Ark Technology's evolution—from an internal support tool into an AIGC ecosystem node in the broader home and living industry—offers a concrete example of this commercial logic in action.

4

TOWARDS AUTONOMOUS INTELLIGENCE: ORGANISATIONS, SYSTEMS, AND TRUST

The business world has moved beyond the initial phase of AI as a tool or copilot and is now entering deeper waters of autonomous intelligence. The “3×3 AI Strategy Matrix” captures a shift from improving individual productivity to optimising system-wide intelligence efficiency, with more advanced applications pointing towards organisational restructuring and the decentralisation of decision-making. As AI agents begin to manage end-to-end business processes, enterprises are no longer dealing solely with tactical questions of deployment, but with deeper strategic and philosophical issues—how they are structured and how they operate, where the boundaries between humans and machines lie, and what ethical principles underpin their activities.

The move towards autonomous intelligence and human–machine symbiosis is reshaping the core operating logic of the enterprise. The early gains from general-purpose large language models (LLMs) are rapidly giving way to highly specialised **multi-agent collaborative networks** and **world models** with real-world perceptual capabilities. This shift in the underlying logic is driving organisations towards M-form structures and, ultimately, towards self-driving enterprises, while prompting a fundamental redefinition of the role of humans in business systems.

This chapter explores the next phase of business evolution. It begins by defining the technological frontier, examining how multi-agent systems and world models are reshaping exploratory knowledge work. It then turns to organisational and human dimensions, exploring the managerial transformations and performance challenges that arise under human–machine symbiosis. Finally, it outlines a digital governance framework grounded in trustworthiness, and considers how this can become a key source of competitive defensibility in scaling AI sustainably.

Technological Frontier: Multi-Agent Collaboration and World Models

4.1

The era in which LLMs function primarily as standalone conversational tools is coming to an end. Enterprise priorities are shifting from generating outputs to delivering outcomes. Producing a marketing report or writing code represents a standardised, verifiable output. By contrast, reshaping a product line based on deep market insight and delivering tangible profit growth reflects a true business outcome—one achieved in conditions of significant uncertainty. Bridging the gap between output and outcome requires two fundamental shifts in the underlying technical architecture: multi-agent collaboration in the horizontal dimension, and the development of world models in the vertical dimension.

4.1.1 From Isolated AI Tools to Multi-Agent Systems: Reconfiguring Exploratory Knowledge Work

The most valuable business activities are rarely linear tasks that can be codified into standard operating procedures (SOPs). Instead, they take the form of exploratory knowledge work, characterised by high levels of uncertainty. Such work typically starts without clearly defined objectives, unfolds through non-linear processes, and depends heavily on real-time judgement shaped by context. The knowledge gained through repeated trial and exploration is also difficult to capture, standardise, and reuse.[1] Traditional monolithic LLMs struggle to handle this type of work. They lack clear role separation and built-in mechanisms to detect and correct errors, and are constrained by limited context windows. As a result, they can gradually lose coherence over time and generate inaccurate or fabricated outputs, undermining their reliability in complex, enterprise applications.

To address these structural limitations, multi-agent architectures have begun to replace monolithic models as the dominant approach. Instead of relying on a single system to handle everything, they break work into specialised roles, allowing different agents to coordinate and cross-check one another. A representative example from leading enterprise practice is the "Generative Enterprise Agent" (GEA) architecture developed by Tezign. Designed for complex, non-linear business exploration, GEA fundamentally abandons the idea of a single, all-purpose model. Instead, it divides cognitive work into modular components, structuring AI as a set of specialised roles that work together in much the same way as a human executive team.

Within this architecture, the Reasoning Agent acts as both the "brain" and the "director." Rather than executing low-level tasks, it focuses on high-level judgement—deciding what to do, how to do it, and whether to continue. It translates ambiguous human intent into structured execution plans, dynamically constructs and updates the overall context, and continuously monitors feedback. When faced with conflicting market signals or incomplete information, it adjusts direction—determining whether to continue exploring or to terminate efforts and cut losses in a timely manner. Complementing this is the Execute Agent, which functions as the system's "operator." Guided by the Reasoning Agent's instructions and context, it executes tasks by invoking specialised modules as needed. These include tools for social media monitoring to capture external signals, and interview modules that conduct multi-turn conversations with digital consumer twins—generating data and surfacing patterns.

This separation of roles makes enterprise applications far more flexible. Organisations can encode the tacit knowledge of domain experts into modular capabilities that can be dynamically orchestrated. Through API-driven Model Context Protocols (MCPs), these agents integrate directly with core enterprise systems such as ERP, CRM, and HRM platforms, enabling action-oriented interactions across functions and systems. For example, when a decline in quarterly revenue is detected, a multi-agent system can automatically initiate a diagnostic workflow. A data retrieval agent extracts regional sales data, an analysis agent identifies root causes, and an execution agent formulates and submits targeted recommendations—such as reallocating marketing budgets—for managerial review. Through this kind of coordinated operation, AI systems evolve from passive question-answering tools into a digital workforce with a degree of operational autonomy and accountability.

Recent breakout applications, exemplified by OpenClaw, have attracted widespread attention and prompted organisations to experiment with agent-based approaches to workflow automation and closed-loop execution. However, caution is warranted. Workflow chains that are quickly assembled around a single tool or within an isolated function rarely constitute true system-level closure. Achieving closure within a single department or task does not translate into enterprise-level optimisation; instead, it risks creating a new "local optimum" trap. Once partial closure is achieved, the critical question becomes: what comes next? Sustainable intelligence efficiency depends on enterprise-wide architectural transformation. Only through unified, top-level design—one that breaks down functional silos and data fragmentation while enabling deep, end-to-end integration across the enterprise—can multi-agent systems fully realise their potential to reshape how the business operates.

4.1.2 Beyond Text: The Rise of World Models and Decision Simulation

While multi-agent systems address how tasks are divided and coordinated within a system, world models tackle a more fundamental challenge: how machines understand the rules governing reality.
In recent years, LLMs based on autoregressive architectures have demonstrated remarkable capabilities in language processing and pattern recognition. However, as enterprise adoption deepens, their structural limitations have become increasingly apparent. At their core, these models predict the next token based on statistical patterns in vast datasets; they do not possess a grounded, commonsense understanding of the physical world, nor can they capture complex causal relationships in business contexts. This lack of real-world grounding leads to "hallucinated" outputs—responses that appear plausible but fail under real-world logic—especially in tasks that require multi-step reasoning, dynamic environment forecasting, or long-term planning.

A paradigm shift is now underway in computational science. The focus is moving beyond scaling model parameters towards building world models capable of internalising space, time, and physical laws. The central idea is to equip machines with an internal "simulation of reality"—a high-dimensional representation of the external world—so they can model and evaluate how environments evolve before taking action or making consequential decisions. This shift from statistical prediction to causal reasoning is giving rise to a new generation of decision intelligence.

In business contexts, one of the most significant breakthroughs lies in the creation of high-fidelity enterprise digital twins and simulation environments. By integrating multimodal data—spanning supply chain nodes, logistics networks, and point-of-sale activity—organisations can move beyond retrospective analysis and adopt a forward-looking, scenario-based approach. Managers are no longer limited to reviewing what has already happened; they can test "what-if" scenarios within dynamic simulation environments. For example, when geopolitical conflict triggers raw material price shocks, or extreme weather disrupts major ports, world models can anticipate cascading effects across global business networks and generate optimised strategies for capacity reallocation and inventory management.

In domains involving intensive physical interaction—such as advanced manufacturing and industrial automation—world models provide the cognitive foundation for embodied intelligence. Physical agents powered by vision–language–action (VLA) models can develop a working understanding of spatial constraints, gravity, and object affordances. They can run large-scale simulations in parallel at minimal additional cost, and interpret three-dimensional spatial data with a level of fluency comparable to language. This enables autonomous navigation in complex, unmanned factory environments, as well as the precise execution of tasks such as high-precision assembly and hazardous inspection.[2]

Beyond physical systems, world models are beginning to operate at the level of strategic foresight. They can learn and internalise competitive dynamics, long-term shifts in consumer preferences, and cyclical macroeconomic patterns. When organisations face major decisions—such as entering new markets or launching disruptive products—these models can simulate competitor responses and evolving customer behaviour within virtual market environments. The result is a set of decision scenarios that combine strategic viability with significantly lower trial-and-error costs.

Ultimately, when multi-agent systems are fully integrated with world models as their foundation of “common sense,” they evolve beyond tools that simply execute instructions. They become autonomous decision engines capable of contextual judgement and long-horizon reasoning. This marks the emergence of a new paradigm—one in which enterprises can unlock system-level intelligence gains and build the next generation of competitive advantage in the age of autonomous intelligence.

4.2 The Future of Organisations: Towards Intelligence-Native, Self-Driving Enterprises

The ongoing shift in technological paradigms is driving a fundamental restructuring of how organisations are designed and operate. As enterprises move from automation towards autonomous, AI-enabled systems, they are reaching a critical inflection point in how they are structured and managed. With multi-agent systems scaling across core workflows, organisations must treat them as part of the workforce—not as tools, but as operational units with clearly defined roles. As a result, the organisation becomes increasingly self-driving, with decision-making and execution distributed across human and machine agents. As machines move beyond their traditional role as tools and become embedded in decision-making processes, conventional hierarchical structures and linear management approaches begin to break down. In their place, more agile and adaptive organisations—built around coordinated human and machine agents—are emerging as the next stage of organisational evolution.[3]

4.2.1 Restructuring: The Contraction of Middle Management and the Emergence of M-Form Organisations

According to classical management theory, hierarchical structures emerged in part to address the limits of human information-processing capacity and constraints on managerial span of control. To translate strategic intent from the top into action and consolidate feedback from the front line, organisations developed complex, multi-layered middle management systems. In the era of autonomous intelligence, however, these constraints are gradually being eroded. As information flows more freely and data processing becomes dramatically cheaper, the need for multiple layers of coordination is diminishing.

Rather than moving towards absolute decentralisation, organisational structures are instead evolving towards more advanced forms of the M-form. On the one hand, multi-agent systems can take on large volumes of coordination work—such as aggregating data, tracking cross-functional progress, and handling standardised approvals. As a result, traditional middle management roles, long centred on passing information up and down the organisation, are being compressed and redefined. At the same time, senior executives can use decision augmentation tools to gain a real-time, enterprise-wide view, significantly extending their span of control. This reduces reliance on layered reporting structures and allows them to directly coordinate clusters of digital workers through real-time data twin dashboards, enabling rapid adjustments to strategy and resource allocation.

At the frontline, meanwhile, the value of roles is being fundamentally re-evaluated. Advances in computing power have dramatically raised the productivity ceiling of individuals, accelerating the emergence of the "one-person company" model in capital markets and at the technological frontier. With tightly coordinated AI agents embedded in automated technology stacks, individuals can now execute entire workflows—from product ideation and coding to global marketing and distribution. In some cases, this can generate revenues in the tens of millions of renminbi at minimal operating cost. This reflects a deeper shift in how business leverage is created: growth is no longer driven primarily by adding headcount, but by building systems and harnessing computing power.

For large, established enterprises, the idea of the “one-person company” goes beyond encouraging individual entrepreneurship. It points instead to the development of employees with strong systems thinking—individuals who can operate with a high degree of autonomy by leveraging AI. This requires breaking down data silos and equipping frontline staff with end-to-end AI capabilities. The core employees of the future will no longer be narrowly specialised operators, but strategic orchestrators who use context engineering to direct networks of AI agents—driving exponential growth with unprecedented leverage.

4.2.2 Human–Machine Symbiosis: Redefining Work Patterns and the Psychological Contract

The restructuring of organisations at the macro level promises a significant release of productive capacity. At the level of individual employees, however, the day-to-day reality of working with AI is introducing new and often unexpected challenges for both performance and wellbeing. A common expectation has been that highly capable AI systems would free up human labour, allowing employees to work fewer hours and focus more on creative tasks. However, empirical research published in Harvard Business Review points to a counterintuitive outcome: rather than reducing workload, AI systems are intensifying both work demands and cognitive strain in new ways.[4]

The study finds that when employees have access to tireless, always-available, high-speed AI assistants, they do not simply finish work earlier. Instead, they use the efficiency gains to take on more work—refining outputs, generating additional options, and exploring the data in greater depth. Over time, this expansion raises expectations for routine tasks, often without employees being fully aware of the shift.

At the same time, the ability of AI systems to handle multiple tasks in parallel places greater demands on human coordination and oversight. Employees’ roles are shifting—from front-line “creators” to “final reviewers” responsible for validating system outputs. While this transition reduces some operational effort, it increases the cognitive burden of reviewing business logic, checking accuracy, and guarding against model hallucinations. Without clear guidance, this high-intensity, multi-task mode of work can blur the boundary between work and personal life, putting traditional psychological contracts under strain.

In navigating this transition, leaders must be clear about what makes humans indispensable in a human–machine system. As algorithms narrow the gap in generating standard solutions and analysing large volumes of data, distinctly human capabilities become more valuable. These include empathy in human interaction, judgement in situations of ambiguity, the ability to build long-term trust, and accountability in the face of complex trade-offs. Accordingly, the role of management must evolve beyond enforcing procedural compliance. It must also involve actively shaping sustainable ways of working—helping employees set clear boundaries in human–AI collaboration and avoid burnout. Leadership in the age of AI will be defined not only by the ability to orchestrate algorithms, but by the capacity to balance system efficiency with human wellbeing, dignity, and fairness.

4.2.3 Process Reengineering: From Task-Based to Outcome-Oriented

To turn human–machine symbiosis at the individual level into sustained organisational advantage, AI frontrunners are moving beyond the traditional “patchwork” approach to adoption. In the past, organisations typically kept their existing linear workflows, adding AI tools to isolated, inefficient steps to make them faster. By contrast, leading companies in 2026 are redesigning end-to-end workflows around the capabilities of multi-agent systems.

In these emerging organisational models, the way value is created is shifting from one-off, project-based delivery to continuous, evolving processes. As a result, enterprises are increasingly adopting an outcome-oriented approach to process design.

First, real-time, system-level feedback loops become the default way of operating. As agents execute tasks, they continuously adjust and improve in real time within live business environments. They track performance indicators—such as conversion rates and customer feedback—and use these to adjust strategy in real time. This allows the business to operate as a self-improving system.[5]

Second, the structure and distribution of human–machine collaboration are being fundamentally reconfigured. In leading technology and financial institutions, the scope of collaboration between human experts and AI agents is expanding across core business processes. Human experts focus on defining intent, maintaining alignment across domains, and setting compliance guardrails. Meanwhile, agent clusters run continuous parallel processing, clean and integrate data, and automatically carry out tasks at scale.[6]

Third, the way organisations evaluate technology investment is changing. AI systems are no longer judged solely by cost savings. Instead, these “digital workforces” are expected to deliver measurable business results, including profit growth. To support this shift, organisations must set up cross-functional units that integrate data and business operations, manage computing resources centrally, and treat them as a core driver of revenue.

Sustained Action: Risk Management and Compliance Frameworks for Scaling Agent Systems

4.3

As enterprises increasingly entrust core resource allocation and decision-making authority to autonomous systems, the nature of risk is shifting. The primary concern is no longer algorithmic performance alone, but the cascading consequences of inadequate governance. From unintended data privacy breaches and gaps in compliance oversight to reputational damage caused by algorithmic bias, AI-related risks are evolving from isolated technical issues into strategic challenges that directly affect long-term organisational resilience.[7] Against this backdrop, establishing a robust, coherent global framework for trustworthy AI governance and risk management has become a prerequisite for scaling agent systems.

4.3.1 From Statements of Intent to Operational Evidence: The New Normal of Global Digital Governance

Over the past decade, corporate approaches to algorithmic governance have largely remained at the level of intent—articulating ethical principles, issuing compliance statements, or establishing ethics committees. However, as regulatory scrutiny becomes more detailed and stringent across jurisdictions, this relatively permissive phase is giving way to a more rigorous, evidence-driven model. Flagship regulations such as the EU AI Act have entered the implementation stage, while major economies in North America and Asia are introducing detailed compliance frameworks and transparency requirements.[8]

This new environment is defined by a shift from declarative compliance to demonstrable operational evidence. Faced with rising expectations from regulators, investors and consumers, organisations must be able to show—at any point—that their AI systems are operating in a compliant and controlled manner. This requires governance across the entire AI lifecycle: from authorising training data and conducting pre-deployment testing and risk assessments, to detecting and mitigating algorithmic bias, and preparing response mechanisms for unexpected incidents. Each stage must generate clear, traceable audit trails. In high-sensitivity domains—such as financial credit assessment, medical decision support and human resource screening—rigorous human-in-the-loop supervision, combined with robust technical safeguards, has become a minimum requirement for managing compliance risks and ensuring business continuity.[9]

4.3.2 Building an Enterprise-Level Lifecycle Framework for Trustworthy AI

As organisations rapidly deploy a growing range of AI tools, a new category of systemic risk is emerging: “agent sprawl.” This reflects an evolved form of shadow IT, where employees independently adopt AI agents, automation tools, external model APIs, data plugins and workflow platforms without unified oversight. Over time, this leads to fragmented, opaque “underground” digital systems that are difficult to govern, audit or even fully understand. In response, forward-looking organisations are reframing the challenge. Rather than reacting defensively to compliance requirements, they are treating data governance and system-level risk control as strategic investments in organisational resilience. From this perspective, a scalable, cross-regional “trustworthy AI” governance framework typically spans four interrelated technical and managerial dimensions.

First, strengthening the foundations of data lineage and privacy protection. The reliability and fairness of model outputs depend on the quality and compliance of underlying data. Organisations should establish automated data quality monitoring and lineage tracking to ensure that all key inputs to decision-making are traceable and verifiable.[10] Where personal data or sensitive commercial information is involved, privacy impact assessments should be built into system design from the outset. In addition, privacy-enhancing technologies—such as differential privacy, federated learning and hardware-level encryption—should be applied where appropriate to minimise the risk of data leakage during model inference.[11]

Second, enhancing model explainability and decision transparency. Trust in AI systems—among both internal stakeholders and external regulators—depends on understanding how decisions are made. Embedding explainable AI (XAI) capabilities into core models helps open up the “black box” of deep learning systems. When AI-generated recommendations inform critical decisions, systems should be able to show key drivers, feature importance and reasoning pathways in a clear and structured way. This level of transparency supports regulatory compliance and builds confidence in AI-assisted decision-making.[12]

Third, establishing full-lifecycle observability and dynamic monitoring. While multi-agent, closed-loop systems improve efficiency, they also increase operational complexity. Organisations therefore need real-time monitoring across the entire lifecycle—from sandbox testing and version control in early development to continuous detection of model drift in production. This enables early detection of performance degradation or logical deviation. In response to emerging threats such as prompt injection, systems should also include safeguards such as anomaly-triggered circuit breakers and degradation alerts to contain risks before they spread.

Fourth, defining dynamic permission guardrails and clear boundaries for human–machine collaboration. Embedding business rules and compliance constraints into system architecture allows organisations to define a secure operating framework for AI agents. When granting API access or control over resources, levels of autonomy must be carefully calibrated. Routine, low-risk processes can be handled autonomously. By contrast, decisions involving core business interests or ethical considerations should trigger mandatory human review, ensuring that ultimate decision-making authority remains appropriately governed.

4.3.3 Trust as Capital: Converting Compliance and Governance into Enduring Advantage

As general-purpose algorithmic capabilities become widely available and foundation models move into the infrastructural layer, “trust” is no longer just a moral expectation—it is becoming a core business asset that competitors will find difficult to replicate.[13]

If governance and compliance are treated merely as constraints on business development, organisations risk becoming reactive when faced with system anomalies or regulatory scrutiny. By contrast, forward-looking enterprises see robust governance frameworks as a “precision steering and braking system” that supports stable, high-performance operations. Built on a solid data foundation and rigorous risk control logic, such frameworks give business teams the confidence to grant calibrated autonomy to multi-agent systems in complex market environments. This, in turn, allows organisations to pursue efficiency gains and business expansion along core value chains in a controlled and sustainable way.[14]

From a longer-term perspective, the ability to demonstrate trustworthy AI operations—through consistent, auditable records—becomes a source of competitive differentiation. When an organisation can credibly assure customers, ecosystem partners, investors and regulators that its AI systems deliver not only strong performance but also adhere to principles of safety, fairness and accountability, this accumulated trust becomes a source of brand premium and a durable competitive barrier.

Looking ahead, the move towards deeper human–machine collaboration will not be a disorderly race for technological advantage, but a systemic evolution grounded in business fundamentals and governance discipline. At its core is a careful balance between system efficiency and human judgement. Ultimately, success in the AI age depends on an organisation's ability to redesign its structure, evolve its management philosophy and build a governance framework that is both trusted and compliant. Only then can enterprises move beyond isolated efficiency gains and achieve true transformation—from incremental optimisation to system-wide evolution.

Endnote

[1] Walter Sun, Sean Kask, and Jonathan von Rueden, “AI in 2026: Five Defining Themes,” SAP News Center, Jan 9, 2026, https://news.sap.com/2026/01/ai-in-2026-five-defining-themes/.

[2] Maria Basso and Michael Römer, "How AI-first operating models unlock scalable value," World Economic Forum, February 12, 2026, https://www.weforum.org/stories/2026/02/how-ai-first-operating-models-unlock-scalable-value/

[3] Microsoft Research, "What’s Next in AI: 2026," Microsoft, December 12, 2025, https://www.microsoft.com/en-us/research/articles/whats-next-in-ai/

[4] Ranganathan, Aruna, and Xingqi Maggie Ye, "AI Doesn't Reduce Work—It Intensifies It," Harvard Business Review, February 9, 2026, https://hbr.org/2026/02/ai-doesnt-reduce-work-it-intensifies-it

[5] Tobias Bock, “Leading in the Age of AI: Why Competence Development Matters More Than Ever,” Emerald Publishing, Jan 14, 2026, https://www.emeraldgrouppublishing.com/opinion-and-blog/leading-age-ai-why-competence-development-matters-more-ever.

[6] Linda Ivy-Rosser, Akshara Naik Lopez, Faram Medhora, and others, “Predictions 2026: AI Agents, Changing Business Models, And Workplace Culture Impact Enterprise Software,” Forrester, Nov 5, 2025, https://www.forrester.com/blogs/predictions-2026-ai-agents-changing-business-models-and-workplace-culture-impact-enterprise-software/.

[7] OriginTrail, “5 Trends to Drive the AI ROI in 2026: Trust Is Capital,” Medium, Dec 23, 2025, https://medium.com/origintrail/5-trends-to-drive-the-ai-roi-in-2026-trust-is-capital-372ac5dabc38.

[8] David Talby, “4 Shifts Defining AI Governance in 2026,” ODSC - Open Data Science on Medium, Feb 10, 2026, https://odsc.medium.com/4-shifts-defining-ai-governance-in-2026-c92c28259139.

[9] Kezia Farnham, “AI Governance: What It Is, Why It Matters and How to Implement It,” Diligent, Jan 27, 2026, https://www.diligent.com/resources/blog/ai-governance.

[10] AI in 2026: How to Build Trustworthy, Safe, and Governed AI Systems,” Keyrus, n.d., https://keyrus.com/us/en/insights/ai-in-2026-how-to-build-trustworthy-safe-and-governed-ai-systems-noram.

[11] Jason M. Loring, “Data Privacy Day 2026: Privacy as the Foundation of Responsible AI Governance,” Jones Walker LLP, Jan 28, 2026, https://www.joneswalker.com/en/insights/blogs/ai-law-blog/data-privacy-day-2026-privacy-as-the-foundation-of-responsible-ai-governance.html?id=102mejb.

[12] “AI in 2026: How to Build Trustworthy, Safe, and Governed AI Systems,” Keyrus, Dec 15, 2025, https://keyrus.com/us/en/insights/ai-in-2026-how-to-build-trustworthy-safe-and-governed-ai-systems-noram.

[13] OriginTrail, “5 Trends to Drive the AI ROI in 2026: Trust Is Capital,” Medium, Dec 23, 2025, https://medium.com/origintrail/5-trends-to-drive-the-ai-roi-in-2026-trust-is-capital-372ac5dabc38.

[14] Beata Socha, “Why Data Governance Is the Cornerstone of Trustworthy AI in 2026,” Strategy, Oct 7, 2025, https://www.strategy.com/software/blog/why-data-governance-is-the-cornerstone-of-trustworthy-ai-in-2026.

Core Team

China Europe International Business School (CEIBS):
Co-founded by the Chinese government and the European Union (EU) in 1994, CEIBS operates campuses in Shanghai, Beijing, Shenzhen, Zurich, and Accra. As China's only business school to originate from government-level collaboration, CEIBS is committed to educating responsible leaders versed in "China Depth, Global Breadth" in line with its motto of "Conscientiousness, Innovation and Excellence". Leaders from the Chinese central government and the EU, respectively, have lauded CEIBS as "a cradle of excellent executives" and "a role model of EU-China co-operation". CEIBS was among the first business schools on the Chinese mainland to receive both EQUIS and AACSB accreditation. In the Financial Times global rankings, its MBA programme has ranked first in Asia for ten consecutive years, while its Global EMBA programme has ranked among the top two worldwide for six consecutive years.

CEIBS X Tezign Generative AI and Business Innovation Initiative:
Launched jointly by CEIBS and Tezign, this initiative connects technological innovation with industrial development. It establishes an operational framework for "AI + Business" to support management research and the commercialisation of AI technologies, contributing to industrial upgrading and economic transformation.

Tezign:
Tezign is a leading enterprise-grade agentic AI company in China. Built on its proprietary Generative Enterprise Agent (GEA) architecture, Tezign develops agent systems that understand business context, support complex decision-making, and generate continuous business outcomes. The GEA architecture is anchored in a "context system" that serves as a single source of truth. It integrates a proprietary Creative Reasoning Model with over 400 agent skills, enabling reasoning, orchestration, and execution within real-world workflows in a closed-loop system. Tezign serves over 180 enterprise clients worldwide and supports more than one million professional users across more than 50 countries.

Guest Editor

CEIBS AI and Marketing Innovation Research Lab

The CEIBS AI and Marketing Innovation Research Lab ("the Lab") is a core part of the CEIBS Research Centre for AI and Management Innovation. It focuses on advancing frontier research and practical applications of artificial intelligence in marketing. The Lab aims to build an open, diverse, and high-quality collaborative ecosystem, supporting continuous innovation and deeper integration in teaching and research in the age of digital intelligence.

Through strategic partnerships with leading enterprises, the Lab brings together diverse resources to respond to a wide range of enterprise needs, conducting research and applied innovation grounded in real-world business scenarios. Its collaboration model is flexible and tailored to each partner's organisational characteristics, strategic priorities, and level of engagement, enabling differentiated collaboration models.

To date, the Lab has established parallel collaborations with companies such as Tezign and DeepZero, each focused on distinct research and application areas. The partnership with Tezign is particularly notable for its strong alignment in research priorities and project design, as well as a level of support that enabled the establishment of a dedicated research fund—namely, the CEIBS x Tezign Generative AI and Business Innovation Initiative—which supports systematic research and accelerates the translation of insights into practice.

By contrast, its collaboration with DeepZero focuses on building a high-level platform for academic–industry exchange. Through jointly organised forums and simulation-based experiments, it promotes deeper integration between teaching and research while strengthening talent development and applied innovation capabilities.

The Lab not only addresses the needs of its current partners but also provides a scalable foundation for broader enterprise engagement. Looking ahead, it will continue to work with organisations at the forefront of AI and marketing innovation, contributing to the evolution of business education and research in the age of digital intelligence. At the same time, it aims to deliver sustained value to its partners and alumni enterprises, working together to navigate emerging opportunities and challenges in an increasingly AI-driven landscape.

Scan to follow the Lab on WeChat
Access the latest research insights and events

Contact
lelaine@ceibs.edu | zhuahua@ceibs.edu

Scan to follow the Research Centre on WeChat
Access the latest AI research insights from CEIBS

Scan to follow the CEIBS x Tezign Initiative on WeChat
Reply “AIBI” to access reports and explore the latest trends in AI-powered business growth

Scan to follow CEIBS on WeChat

Scan to follow Tezign on WeChat