HARVARD BUSINESS REVIEW

Digital Article / Generative AI

Design Processes to Evolve with Emerging Technology

How industry leaders like BMW, Unilever, and Waymo are embracing real-time visibility, digital twins, and agentic AI. *by Manish Sharma, Lan Guan, and H. James Wilson*

Published on HBR.org / January 30, 2026 / Reprint H09207

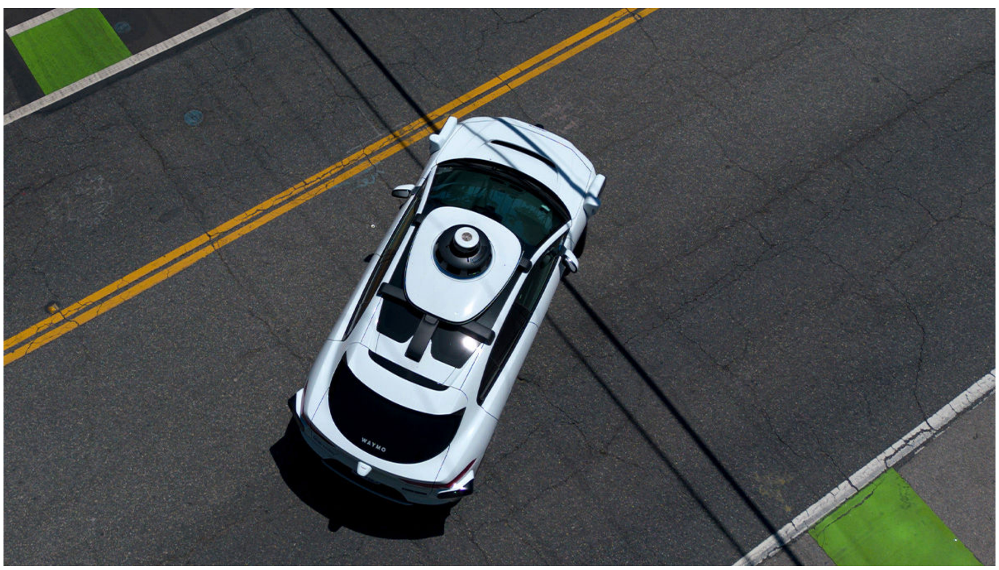

Eric Thayer/Getty Images

For decades, businesses have been trapped in a cycle of painful, episodic change, launching massive re-engineering projects and investing in new IT systems, only to find their organization's fundamental metabolism remains sluggish. Immense transaction costs —the friction of coordinating people, managing information, and aligning complex work—have made deep, continuous transformation prohibitively expensive and risky.

According to our recent unpublished analysis, more than four in ten working hours are related to such transactions within business processes. Now, a new framework for reinvention is poised to collapse these costs by combining three forces: 1) radical process visibility, 2) low-risk experimentation within a digital twin, and 3) autonomous AI agent-driven execution. Together, they transform process reinvention from a disruptive project into a continuous, core capability.

That future is unfolding now. In a 2025 survey of more than 3,000 C-suite leaders we conducted, nearly 9 in 10 (87%) agreed that agentic AI is ushering in a new era of process transformation. Already, two in ten (20%) of those leaders report that their organization is redesigning processes with AI at the core. Yet nearly three in ten (29%) of those leaders are taking an incremental approach, improving existing processes with AI step by step. Meanwhile, industry leaders are building the foundation of a new operating model that makes continuous reinvention practical and offers a roadmap for truly evolutionary process design.

Radical Visibility

Management's mental model of a process often oversimplifies reality. For instance, a single healthcare accounts payable process can have millions and millions of touchpoints that no manager could possibly comprehend. Any improvement initiative based on a flawed model inevitably solves the wrong problem—hence the high failure rate of traditional transformation programs. That's why the essential first step in process reinvention is to codify operations into a dynamic, computable model: a current-state, real-time blueprint of how work actually happens. Achieving this level of process visibility becomes the critical catalyst for every subsequent transformation, laying the groundwork for building a digital twin.

Consider the desire of mining company Boliden to optimize control of the copper grinding process at their Aitik mine in Sweden. Grinding plays a critical part in extracting valuable metals from mineral deposits and remains the primary operational bottleneck at the plant. Even small improvements in throughput and stability of the process can yield big financial gains. Working with multinational giant ABB, the company created a highly accurate as-is model—in effect, a virtual plant. They codified it by creating a direct copy of the plant's actual distributed control system process (DCS) including its customized logic to ensure a realistic replication of the base layer of control.

This control system model was then connected to a process simulation that hosted the process dynamics of the physical mill. The "digital footprints" analyzed were the real-time communication protocols that shared data between the base control layer and the process simulation, mirroring the true "as-is" data flow of the physical plant. This computable model, with its radical visibility into the process, provided the non-disruptive environment of a digital twin needed to test new control strategies against the current state.

The very act of building the new model forced the company to recognize that its old model no longer accurately represented the current state of the process. As ABB and Boliden executives point out, "It is critical to have a simulation that describes the current state of the process . . . Finding and implementing this type of information is a key part of setting up a virtual environment for control development."

Continuous Experimentation in the Digital Twin

Continuous improvement requires continuous experimentation. During the dot.com era of the late 1990s, for example, the emergence of simple A/B testing transformed decision-making by replacing executive intuition with evidence. Instead of guessing which design worked better, companies showed different versions of website designs to random groups of users and let metrics like clicks and purchases decide. In 2012, an engineer had a simple idea to make ad headlines longer by combining two lines of text. Executives rejected the idea until an A/B test proved it would increase revenue by 12%.

Experimentation within a digital twin extends that principle to the core of a company's operations. Organizations can simulate and validate entirely new processes from a clean sheet in a risk-free simulated environment, which assumes that AI agents will handle coordination. This capability replaces episodic, high-stakes process change with an ongoing, low-risk cycle of reinvention.

BMW Group's "Virtual Factory" illustrates this shift. The company built a full digital twin of its production plants, linking building data, equipment data, logistics information, and even 3D simulations of manual human work processes. This virtual twin will also simulate avatars of human workers grabbing parts and tools, and assembling components, to find the best procedure, minimize ergonomic problems, and identify the best number of workers for each job. Engineers can now simulate and validate the integration of new vehicle models virtually, collapsing a costly four-week physical validation process into a three-day process.

Looking ahead, the BMW Group's Virtual Factory is expected to reduce production planning costs by up to 30 percent. BMW's previous validation process for a new vehicle model required manually guiding a physical car body through the production line over several weekends. In the paint shop, this process could require draining and cleaning entire dip coating tanks, incurring enormous costs and production downtime. The new validation process allows planners to run “what-if” scenarios, such as optimizing line layouts for new models in a virtual environment, without disrupting live production.

Unilever applies the same approach to sustainable packaging at its Tinsukia, India, factory. Its improvement digital twin simulates new packaging materials virtually, cutting trial times by 84% and increasing the number of experiments from two to thirty per year. By simulating, testing, and monitoring the performance of new packaging materials virtually, teams can identify the most promising options before committing to physical production, accelerating progress toward the company’s sustainability goals.

The benefits of working with digital twins go beyond direct cost and time savings. For BMW, accelerating the validation cycle, from four weeks to three days, supports the planning of more than 40 new or updated vehicle models across its global network by 2027. For Unilever, increasing its experiment volume from two to thirty trials per year directly drives strategic sustainability objectives: the Tinsukia site has already achieved a 21% reduction in virgin plastic use.

Agent-Driven Execution and Adaptation

Today, roughly one-third of process tasks can be fundamentally reinvented through the effective use of agents, according to our economic analysis. Managers now have unprecedented visibility into their operations and can use digital twins to simulate how AI agent process interventions might perform in practice. These virtual environments provide low-risk sandboxes in which AI agents can calibrate and evolve processes and then transition to live operations with proven performance. Once deployed, these agents orchestrate complex processes in real time, dramatically reducing the transaction costs that have long constrained enterprise agility.

Waymo's fully autonomous driving system offers a glimpse of this future. The company's core process – navigating a vehicle safely and efficiently through complex, dynamic, and unpredictable urban environments—illustrates how agent-driven execution evolves through real-world validation on city streets rather than abstract modeling. This task is one of the most formidable challenges in artificial intelligence: the system must interpret its surroundings, predict how other road users will behave, and act accordingly in milliseconds. The agent must interact safely with countless other uncontrolled factors—human drivers, pedestrians, cyclists, and more in an infinitely variable environment and anticipate the likely future trajectories of every other road user to make proactive, defensive, and socially intelligent driving decisions.

Waymo trains its system on billions of miles of real-world driving data, including more than 100,000 miles of expert human driving in complex urban settings. This allows the AI to establish a strong, human-like baseline for navigating common driving scenarios with consistency, predictability, and social awareness. However, real-world data is inherently sparse when it comes to rare but critical safety events.

To close this gap, Waymo employs reinforcement learning in a massive, high-fidelity simulation environment. In this virtual world, agents repeatedly encounter and resolve millions of challenging and dangerous scenarios—cars running red lights, pedestrians suddenly stepping into the road—that would be impossible and unsafe to practice in reality. This continuous cycle of experimentation with both real-world data and simulated experience allows the system to handle a wider range of scenarios with greater reliability over time.

The programming of the car goes far beyond following static rules of the road. It is developing and refining a process capable of responding intelligently to the complexities of the real world, using a rich suite (lidar, radar, cameras, and external audio receivers) to perceive the dynamic environment and road users, predict their behavior, and plan and navigate a journey from A to B in real time. With advancements in large language models (LLMs) and vision-language models (VLMs), Waymo's next generation AI models will combine Waymo's driving experience and AV-specific agentic AI advancements with the 'world knowledge' and reasoning capabilities of LLMs/VLMs to create models specifically applicable to the driving context.

Embedding that level of agility directly into a firm's live operations requires autonomous AI agents empowered to execute, manage, and adapt entire workflows. This is a significant leap beyond simple robotic process automation (RPA), which typically focuses on discrete, repetitive tasks. Advanced AI agents can coordinate multi-step processes involving professional judgment and manage handoffs between systems and people. By automating these coordination layers, companies can significantly lower the transaction costs of coordination that have long hindered enterprise agility and accelerate the adoption of processes redesigned from scratch for new technical capabilities and human behaviors.

Far from rendering humans obsolete, these agent-driven models elevate their contribution. As AI coordinates and executes complex workflows, employees focus on new and more valuable roles. While AI can map explicit workflows, it cannot capture all the deep, unwritten, contextual knowledge that experienced employees possess, such as the tacit knowledge of BMW workers of the automaker's implicit aesthetic. As processes become increasingly automated and self-adapting, employees become process stewards who ensure that agent-driven operations align with the company's ethical standards, regulatory obligations, and strategic objectives.

...

Industry leaders like ABB, BMW, Unilever, and Waymo show that the vision of a continuously evolving, agent-driven enterprise is not a distant future but an emerging reality. For executives, the goal is no longer to manage incremental change in the traditional, episodic sense, but to institutionalize a continuous cycle of clean-sheet process redesign. This requires building an evolutionary system where process innovation becomes a core organizational capability rather than a periodic disruption. As machines handle coordination, the most successful organizations will redesign roles and elevate distinctly human capabilities like creative problem-solving, innovative judgment, ethical oversight, and strategic direction to guide the systems that increasingly run their operations.

When processes become transparent, testable, and adaptive, companies gain an operating rhythm that evolves as conditions change. This shift collapses internal transaction costs, accelerates innovation, and aligns human creativity with machine-driven precision. The result is a new kind of organization designed for perpetual reinvention and capable of evolving its processes as fast as the world around it.

This article was originally published online on January 30, 2026.

MS

Manish Sharma is the Chief Strategy and Services Officer at Accenture.

LG

Lan Guan is chief AI and data officer at Accenture.

H. James Wilson is the global managing director of technology research and thought leadership at Accenture Research. He is the coauthor, with Paul R. Daugherty, of *Human + Machine: Reimagining Work in the Age of AI,* New and Expanded Edition (HBR Press, 2024).

𝕏 @hjameswilson

