

The End of Marketing Campaigns as We Know Them

By [Silvio Palumbo](#), [Yun Lim](#), and [Karl Johnson](#)

ARTICLE JULY 02, 2026 12 MIN READ

This is Part 4 of a five-part series on next-best action.

The previous articles in this series made the case that agent-native next-best action (NBA) is the new imperative in marketing. Empowered by new decisioning science, [AI agents](#) will compose personalized actions from a composable shelf of offers, creatives, and micro-journeys. The technology and the science needed to make this transition possible are already available and constantly improving. Now comes the hardest part: changing the prevailing ways of working.

Companies need to address three key aspects of the broader marketing operating model. The first is a shift in how organizations think about marketing—away from predefined journeys and toward stocking and curating a composable shelf that the NBA can draw from in real time. The second is a reinvention of the execution value chain, redesigning the build process around AI agents that produce content, offers, and experiences at 10 to 100 times the volume, variety, and speed that legacy approaches can support. The third is implementation of enterprise decisioning governance—a cross-enterprise discipline that steers the NBA’s autonomous decision making toward customer-centric, cross-product optimization and ensures alignment with evolving business priorities. Together, these three transformations set up organizations to maximize the value creation potential of NBA and to minimize the risk.

The End of the Campaign as We Know It

For decades, marketing has been organized around codified campaigns and journeys, which involve mapping business objectives to specific audiences and creative, scheduling them on a calendar, or attaching them to a predefined trigger, and deploying them with limited personalization. The campaign is the dominant unit of work, and the calendar is the primary organizing principle. That model was built for a world in which humans made every decision and were responsible for managing complexity—for example, by batching customers into segments and scheduling communications. It wasn't designed for a world in which the system can act on individual customer context in real time.

Agent-native NBA inverts this logic. The system starts with the customer and pulls from the composable shelf to deliver the right action at the right moment. The concept of the calendar as the organizing principle fades away. Predefined audiences yield to real-time contextual targeting. The strategy shifts from “which campaign do we send to which segment this month” to “what is on the composable shelf, how modular is it, and what objective function governs how the system selects items from it for each customer.”

Moving beyond linear campaigns presents the opportunity for a new paradigm that splits marketing into two distinct sides: build and delivery. The build side is where marketers, in collaboration with specialized AI agents, produce and approve content, offers, and experiences, and then load them onto the composable shelf. The delivery side is no longer a manual marketing function at all. Instead, it is executed by fully autonomous agents who handle 1:1 orchestration and personalization to meet the needs of each individual customer, operating within guardrails and business rules set by the enterprise's decisioning governance. (See Exhibit 1.)

This does not mean that every campaign disappears. Brand campaigns, major product launches, regulatory or transactional communications, and seasonal moments will continue to exist as planned pushes. But these become the exception, not the rule. In the mature agent-native model, 70% to 80% of all customer touchpoints shift from predefined campaigns to NBA-driven, real-time interactions sourced from the composable shelf. The campaign as the dominant unit of marketing work is coming to an end.

Reinventing the Execution Value Chain

The more content, offer, and experience variants a company builds, approves, and puts on the composable shelf, the more the NBA can leverage experimentation and personalization. The resulting explosion of targeted content and actions—delivering 10 to 100 times the volume and variety that most campaign-based organizations have today—elevates each unique outreach and touchpoint to its highest relevancy and efficacy. The agentic marketing approach invites the

templating of content and offer constructs, rendering them easy to deploy in real-time for hyper-personalization. Legacy processes and timelines need to evolve for execution at this speed, scale, and flexibility of output.

In most large enterprises, the linear legacy marketing build process takes 60 to 90 days from initial strategy to campaign launch, involves 20 or more people with numerous handoffs, and produces a single campaign with a narrow set of variations. Both the cadence and the throughput are incompatible with today's needs.

The marketer's traditional functional roles and required skill sets change, too. Redesigning around a pairing of humans with AI agents at each step of the process streamlines and accelerates the value chain. The primary focus of the human marketers shifts from doing the manual work to giving strategic direction, using their judgment to refine and revise outputs in response to agent inputs—sharpening the marketing brief, curating the best content variants, validating brand alignment, and setting the learning agenda.

Although marketers can find endless ways to adopt agentic workflows, it's more sensible to prioritize investments with a focus on research and insight definition, strategic and creative brief creation, copywriting and asset production, legal and compliance review (where regulated industries such as banking will benefit the most), and creative build (such as HTML coding).

The velocity gain is dramatic. For example, a brief agent can synthesize past performance and market signals into a draft test plan with targeting criteria, KPIs, and creative hypotheses, collapsing tasks that previously took days into a few hours' work. The marketer then refines the brief rather than building it from scratch. Content and offer agents generate multiple variants, which a legal and compliance companion reviews in real-time rather than as a sequential gate, compressing weeks into hours. Once content is on the shelf, a performance optimization agent continuously identifies top-performing variants and reallocates marketing resources accordingly. Operations that previously required a quarterly review cycle and a dedicated team now happens in real time. (See Exhibit 2.)

Process reinvention aside, the legacy model's organizational structure, ways of working, and agency ecosystem all need revising to support the agentic-marketer pod—a lean, cross-functional team of three to five people, each teamed with AI agents, spanning strategy, content, data, channel activation, and compliance. The pod has little or no dependence on teams outside itself, thus eliminating intake processes, handoffs, and time spent idling in queues or waiting for decisions. The pod owns its domain (such as the life-cycle stage for a set of products), manages its actions on the composable shelf, and operates in rapid iterative cycles of days, not months. This structure accelerates the speed at which marketing can learn, adapt, and respond to what the NBA sees in market. (See Exhibit 3.)

Organizations that have adopted this model have reported reductions in cycle time of up to 80%, with content production compressing from months to days—an exponential increase in content throughput. Agentic-marketer pods can operate with 60% fewer resources than traditional campaign teams require, but the roles that remain critical gain even more value: strategy,

curation, governance, and continuous learning. The continuous learning loop is a core operating discipline, not an afterthought. The pod drives a learning agenda, designs experiments, reads results, and feeds insights back into the shelf and the NBA to compound performance improvements over time.

Establishing Customer-Centric Decisioning Governance

Realizing the NBA's potential to orchestrate across all products and channels to deliver customer-centric experiences requires governance that extends beyond any individual P&L owner. Companies may continue to build most content and offers in a product-centric way, but the NBA engine must make customer-centric decisions, transcending individual product P&L and short-term targets.

If the company's strategy is to maximize profit and prioritize whatever converts best in the short term, that could lead to overpromotion of products from one line of business. This type of product-centric logic risks bombarding shared customers with competing messages, or overemphasizing short-term conversion at the expense of customer lifetime value. The more aggressive the engine parameters are, the faster these problems will show up. Decisioning governance provides the operating model discipline necessary to prevent this outcome.

Decisioning governance has two core responsibilities. The first is to define and continuously recalibrate the enterprise's objective function over time: the hierarchy of KPIs, weights, and valuations that dictate what to optimize for and how to prioritize competing actions. For example, competitive dynamics for a bank might lead it to weigh cross-sell heavily in one quarter but shift toward retention the next.

Decisioning governance is an always-on function, not a one-time exercise. Business priorities shift. Products gain or lose strategic importance. Short-term pressures such as quarterly earnings, investor expectations, and competitive responses create real tradeoffs against long-term lifetime value optimization. A cross-enterprise governance body must review and recalibrate these inputs regularly (at minimum, quarterly) so that every autonomous NBA decision reflects the company's current strategy. Without this discipline, the system will optimize for stale priorities and will widen the gap between what the NBA is doing and what the business strategy requires.

The second responsibility of decisioning governance is cross-enterprise coordination. In most large organizations, multiple lines of business interact with the same customers. In the absence of oversight, each line of business feeds its own priorities into the system independently, resulting in product-centric competition for customer attention rather than a coordinated,

customer-centric experience. The governance architecture must operate across three distinct layers: an executive champion team that sets enterprise-wide priorities and establishes the objective function; agentic pods (whether in marketing, servicing, sales, or customer experience) that build and stock the composable shelf within each priority; and the NBA's air-traffic control layer, which orchestrates delivery against the framework's rules, constraints, and valuations. These guardrails set limits on customer contact frequency, identify channel preferences, define brand standards, and detail business rules that determine which actions take precedence when multiple options are competing for the same customer.

Operationalizing governance requires clarity about where agent autonomy ends and human oversight begins. Agents should autonomously determine action selection, timing, channel, and composition from preapproved shelf assets in cases calling for high-frequency, low-stakes decisions, where speed matters and where the quality controls on the shelf itself limit the downside of any individual decision. Agents should make recommendations for human approval when they identify a gap in the shelf that requires new asset creation, when they detect signs that current objectives may be suboptimal and need modification, or when they encounter high-value customer interactions that raise the stakes of a single interaction to a point that warrants human review.

Humans govern through policy (brand standards, regulatory constraints, fairness requirements, and contact limits) encoded as system constraints that agents must operate within. Policy review and updating occur on a regular cadence by policy, not on a decision-by-decision basis. Humans also retain ownership of decisions that define the system's purpose: strategic objective setting, budget allocation, shelf strategy, and the expansion or contraction of agent autonomy itself. These decisions should never be delegated to AI.

Organizations that set up an appropriate decisioning governance framework create a self-learning system that can maintain its alignment with business strategy. Sooner or later, organizations that deploy agent capabilities without clarifying the autonomy boundaries inevitably face a crisis of trust that impedes scaling and value creation.

A Day in the Life of Agentic Marketing

To bring the contours of agentic marketing to life, consider what a workday looks like for Maya, a VP of customer marketing at a large retail bank, six months into this transformation. She leads a pod of four human marketers who curate the composable shelf, refine objectives, and review agent outputs daily. Her role is to govern the system: setting the boundaries within which AI

agents and her team operate and making judgment calls in situations that require strategic context.

Her morning begins with a governance dashboard summarizing insights from the system's daily performance: tens of thousands of individually composed customer interactions across email, push, in-app, and SMS. The dashboard surfaces three escalations that require her appraisal. The first is a drift signal: an agent has been unduly favoring a cashback offer whose incremental lift is declining. Maya adds a margin efficiency constraint that will push the agent to explore alternatives—a four-minute adjustment. The second escalation involves reviewing new creative variants from her pod's shelf curation work; she approves most, and flags a few for tone misalignment. They are live within the hour. The third escalation arose when an agent detected a behavioral pattern among a target group for premium card upgrades and recommended a new communication strategy. Maya evaluates whether the agent's suggestion aligns with the bank's positioning—and on that basis, she greenlights the recommendation.

By midmorning, she has governed tens of thousands of interactions and made three strategic decisions. Her pod is executing the rest of the day's work. The work differs from campaign management, but it is not necessarily easier. It requires sharper strategic thinking, deeper data fluency, and comfort governing a system that moves faster than any human-only team could.

What This Means for Marketing Leaders

The shift to agent-native NBA demands more of marketing leaders than a technology decision. It requires them to restructure how their marketing organization works—from the strategy that governs what gets built, to the agentic-marketer pods that build it, to the enterprise governance that steers the engine. Marketers who thrive using this model look different from today's campaign operators. They are strategists who set direction for agentic coworkers, content leads who curate shelf quality, data specialists who drive continuous experimentation, and governance architects who calibrate the NBA's priorities against the enterprise's evolving needs. Leaders should invest in developing their best people for these roles, while also being clear-eyed about which legacy roles, processes, and ways of working should no longer exist.

The early results are compelling: up to 80% faster cycle times, exponential increases in content throughput, and 20% to 40% higher customer engagement and conversion rates. The real prize, however, is the compounding effect. Every touchpoint and experiment makes the system smarter. Each governance cycle sharpens alignment with the business. That advantage builds over time, and the gap between organizations that move early and those that wait grows with it.

This is Part 4 of a five-part series on the future of next-best action. [Part 1](#) explores why most NBA programs underdeliver and identifies four structural gaps. [Part 2](#) discusses the shift from marketer-orchestrated journeys to agent-composed actions and the technology architecture that enables it. [Part 3](#) unpacks the decisioning science: from propensity models to contextual bandits to foundation model reasoning. Part 5 addresses measurement: from campaign-level attribution to agentic measurement.

Authors



Silvio Palumbo

Managing Director & Senior Partner
New York



Yun Lim

Associate Director
Chicago



Karl Johnson

Managing Director & Partner
Washington, DC



ABOUT BOSTON CONSULTING GROUP

Boston Consulting Group bridges the gap between ambition and outcomes for the world's leading companies and organizations. We are built for this era of unprecedented change — bringing strategic clarity rooted in over 60 years of deep domain knowledge, combined with applied AI shaped by our practitioners. BCG works shoulder-to-shoulder with CEOs across industries and geographies to deliver transformative impact at scale: stronger returns, transferred capabilities, and change that sticks. For more information, visit bcg.com.