Mi3

Inside the Tornado:

Mi3's Agentic AI in Marketing Report

PARTNERED BY OmnicomOceania

Inside the Tornado

ANDREW BIRMINGHAM

Seven minutes out from South San Francisco's Caltrain station, the phone buzzed. Redwood City was off. The day unspooled from there. A month on the road and jetlag that clung to me like static had left me ill-disposed to early morning starts. So I did what any reporter too far from home for too long would do under the circumstances, and cursed the name of Jim Kruger and all of his antecedents.

Kruger, it turned out, had a very good excuse. As I learned from Constellation Research's marketing tech maestro Liz Miller by way of LinkedIn, Salesforce had just announced it was buying Kruger's firm for $US8bn.

It was one more proof point about how rapidly agentic AI is remaking the tech sector, as well as the marketing technology infrastructure of the brands it serves. It was also a reminder about the scale of the bets companies are taking.

Salesforce's swoop on Informatica was about plumbing, specifically, the messy, sprawling, under-the-hood data that will power the shift to agentic AI. When we finally reconnected, Kruger, the Informatica CMO unpacked the rationale behind the deal. "We help customers get ready to implement AI, and you have to have a really solid data foundation with cleansed, high-quality, governed data in order to have successful AI implementations."

"There's a lot of customers that I think are jumping the gun a little bit before they have that foundation in place."

THE URGE TO "DO SOMETHING"

When I began Mi3's agentic road trip at Boomi World in Dallas in early May, agentic hype was at its peak. In January 2025 Steve Lucas, the Boomi CEO told me he had ordered the agentification of half the business processes at the firm, and now he was sharing the progress. At Pega World in Vegas a month later, the agentic story was still red hot.

I thought the hype would break on our shoreline in Q1 this year. I was wrong. Over dinner with a dozen customer, marketing and data executives in October, AI was the shiniest jewel on their well-polished worry beads. They were not alone in their anxiety.

Across the meetings Mi3 held with CEOs, CMOs, founders, SVPs of product, heads of AI, and heads of strategy (the list goes on), I was struck by the emergence of an increasingly recognisable undercurrent. It wasn't fear, though, for some that was certainly part of it. There isn't an English word that nails it precisely, so we invented a German one: *Rauschverwirrung* – a heady mix of exhilaration and confusion.

It was best exemplified by Adobe's Experience Platform VP, Sundeep Parsa, who told me, "The biggest risk is not to be in the game and wait for something to emerge. By then, it's too late. You've got to be inside the tornado."

Many of the people we spoke to specifically for this report are listed in the acknowledgements, but it is certainly not a complete rendering. For starters, there are all the executives brand-side who will feature in the Editorial Case Study series over the next few weeks that accompanies this report.

There were many others, the people we buttonholed as they came off stage at conferences, some we rang when we didn't understand a piece of kit or an emerging protocol (that happens a lot more than you might think), and strangers I sat across from on an early-morning Caltrain when I was feeling tired and grumpy and grieving the loss of 30 minutes extra rest.

Finally, a shout-out to Omnicom Oceania for supporting the editorial and events program that will emerge from our research over the next few months. Their partnership not only made the program possible, it also made the report much better. Tech vendors are great at talking tech, but they are a little too cold-blooded when it comes to the impact of agentic AI on creativity, or how it changes the way brands talk to customers, or to their teams. Or the organisational impact on chief customer officers and CMOs.

Omnicom's execs landed those issues for me, and the report is much better for it.

Enjoy.

Contents

The New Era of Infinite Intelligence

By Nick Garrett
CEO, Omnicom Oceania

We're living through a time where yesterday's science fiction is today's reality. The pace of technological change is extraordinary, but I'm mindful that in moments like these, it's easy to get swept up in hype and buzzwords and lose sight of what really matters.

At Omnicom, our focus isn't on chasing the next shiny thing. It's about understanding the deeper forces shaping our industry, what truly moves the tide, not just the waves on the surface. As I took on the role of CEO for Omnicom Oceania, I knew our job was not to follow momentum, but to create it, with clarity of purpose and a commitment to delivering real impact for our people, our clients, and their customers.

Our purpose is simple: to harness intelligence and creativity in all its forms to help our clients grow and create meaningful impact. For us, technology is not about replacing what works, but making us more effective at the things that matter most – building brands, solving problems, and driving customer growth.

Creativity sits at the heart of everything we do. To us, creativity isn't just a buzzword or a campaign tagline, it's practical problem-solving. And the best ideas really do have the power to change the world.

In this new era, AI and human ingenuity together mean the only limit is our imagination. That's what we mean by "infinite intelligence", a space where learning is continuous, and every insight helps us get better at what we do.

This mindset is changing how we work. We're building a culture focused on effectiveness, where every action and decision is about driving results and creating value.

Our model connects intelligence and creativity across every discipline – media, data, commerce, experience, consulting – so every part of the system learns from the others. When our Flywheel team improves a retail journey, that learning informs our creative and media work. When Credera re-codes a client's infrastructure, it opens fresh ways for brands to connect with people. We're building a living system, not just running campaigns, so our work keeps getting smarter and more effective over time.

And as we look ahead, that system, Omnicom Oceania, is evolving on a scale we've never seen before, one that unites the most distinguished collective of creative leaders in the industry with the most advanced data and intelligence platform in the world. The next evolution of that platform; OMNI, will connect more data sources, insights, and signals than ever before, giving our clients a true competitive advantage and redefining what integration looks like in practice. This is transformation not just in capability, but in scale and ambition, a step change for what modern marketing can achieve.

Transformation is never linear, and our clients are evolving with us. More than ever, they're looking for partners who combine deep expertise with the ability to integrate seamlessly. That's why our Oceania model is built around collaboration, continuous learning, and a relentless focus on effectiveness. The brands leading the way are those embracing this partnership, where technology amplifies empathy, and creativity delivers real results.

There's a tremendous opportunity ahead, not just to optimise what we do, but to push creative problem-solving further upstream into business strategy. Ultimately, technology doesn't define progress, people do. At Omnicom we believe in the Infinite potential to use intelligence and creativity to help our clients grow and to make an impact where it matters for them.

Mi3, thank you for being a trusted partner in chronicling the transformation of our industry with depth, clarity, and candour. Your continued collaboration helps our collective industry move forward.

And to our clients, thank you. Nothing we do matters without your trust, your ambition, and your belief in partnership. Everything we build, every innovation we deliver, is in service of helping your businesses grow and proving the impact of what we do together.

Nick

CHAPTER 01

Shift to Agentic AI

Economists have long understood that it's the interdependence of processes and workflows that drives the greatest value in business.

Almost 20 years ago, in paper for the National Bureau of Economic Research, Stanford's Nick Bloom and the LSE's John Van Reenen wrote: "We see management practices as more than the attributes of the top managers: they are part of the organisational structure and behaviour of the firm, typically evolving slowly over time even as CEOs and CFOs come and go."

Productivity, in other words, isn't born from a lone genius but from the messy, interlocking routines that make an organisation hum. The real value lies in how teams, systems, and processes click together to turn intent into output.

What Bloom and Van Reenan couldn't have imagined was that management practice might one day evolve in real time. Agentic AI promises to stitch together workflows and execute and improve them instantly, supercharging productivity.

REAL WORLD ADOPTION

At the Australian econometrics platform Mutinex, founder Henry Innis says nearly a third of the firm's R&D now focuses on agentics. "We trained agents to help us structure data warehouses."

The result: onboarding times down 35 per cent and customers reporting far less manual work. Clients using the system generate about 15 deep insights per month, work that before the age of automation would require tens of thousands of hours of manual work, which means it wouldn't have happened.

"Long term, I'm expecting that to take a lot of the BAU insight work and reporting," he adds. "That's not a bad thing. It will give marketers time to think again."

Mutinex mirrors a broader trend. Across the stack, from creativity to campaign management to data management to orchestration, companies are betting big on agentic AI. But the speed of change tempts overconfidence. As McKinsey cautioned in its September paper One Year of Agentic AI: Six Lessons from the People Doing the Work, autonomy without discipline is chaos.

UNDER CONSTRUCTION

It's easy to think the switch has already been flicked. It hasn't. The most ordinary truth about agentic AI is that it will obey Bill Gates's enduring rule: people overestimate what they can do in a year and underestimate what they can do in ten.

On that horizon, agentic AI will likely reinvent how business operates, and what humans do within it. But for now, we're still at the starting line.

Much of the scaffolding is missing. Early standards like Model Context Protocol (MCP) and Agent Communication Protocol (ACP) are gaining traction, but others, like Google's A2A, faced a cooler reception. Executives we spoke with pointed to gaps in observability and evaluation testing. Cybersecurity remains the Voldemort of the field: in dozens of Silicon Valley briefings, they dared not speak its name.

And every conversation eventually landed on orchestration - getting agents to work together safely and effectively - and governance as the linchpin.

Phase	Chapters	Defining Role of Agentic AI	Organisational Expression	Core Discipline
1. Co-Worker	2 – Early Use Cases	Task executor	Automation layer	Operational orchestration
2. Collaborator	3 – Learning Systems	Feedback learner	Adaptive workflow	Continuous governance
3. Creator	4 – Creative Systems	Brand & experience co-author	Codified creativity	Creative orchestration
4. Solvent	5 – Remaking the Stack	Architectural force	Fluid SaaS stack	Control planes & outcome pricing
5. Cultural Architect	6 – Organisational & Role Transformation	Human–machine conductor	Flattened hierarchy	Cultural orchestration
6. Ecosystem Conductor	7 – Orchestration & Governance	Policy & trust layer	Managed autonomy	Governance architecture
7. Nervous System	8 – Enterprise-Wide Integration	Enterprise-level alignment	Semantic fabric	Federated orchestration
8. Organism	9 – Infrastructure & Risk Management	Self-regulating entity	Real-time control	Governance-as-code
9. Immune System	10 – Cyber Security	Self-defending ecosystem	Autonomous trust fabric	Security-as-governance
10. Metabolic System	11 – Cost and Pricing	Self-pricing organism	Metered intelligence network	Economics-as-governance / Cost-as-control

THE REALITY CHECK: STEPS, TASKS, PROCESSES

Salesforce's John Kucera, SVP of product management for Salesforce AI, describes a hierarchy that keeps expectations grounded: "Processes have tasks. Tasks have steps. Task might be something like, I want a refund, or I need a password reset. Agents can successfully automate steps, sometimes entire tasks, rarely multiple tasks, exceedingly rarely, and probably not yet, entire processes."

That hierarchy ought to give pause to marketers intoxicated by the promise of "end-to-end process automation." The technology may one day manage a full workflow; for now, says Kucera, the wiser move is to target "high ROI tasks and steps that you can automate now without a heroic amount of effort."

John Kucera
SVP Product Management, Salesforce

"To do it well with today's technology, takes a team... You need people that understand the business policies, you need prompt engineers, you need analysts who know what good looks like. This is kind of the minimal set of people that are needed to make a good agent."

Even within those boundaries, complexity compounds quickly. "Maybe the agent is 90 per cent accurate," he notes. "Now you have two actions, each 90 per cent. That's 81 per cent accuracy. What if there's five? What if there's ten? While any one action might be at an acceptable risk tolerance, if you chain them all together with no human review, you often get to an unacceptably low end result."

His arithmetic illustrates an inescapable truth about the agentic age: without solid data foundations, the whole edifice wobbles. As he puts it, "Agents need metadata to make good decisions. You want the magic query box where you can ask it any question. It turns out a lot of the data about the data is important to be able to answer more and more questions and handle more use cases."

For CMOs, the lesson is plain: before chasing end-to-end autonomy, get the data house in order.

Contrary to hype, successful agent deployment is not about "one person, one prompt." It requires coordinated teams with business analysts, data specialists, and prompt engineers. CMOs should understand that implementation involves organisational change and new skillsets—not just technology adoption.

DAMN THE TORPEDOES

Complexity is no excuse to hold back. Companies on both the sell and buy side are already rebuilding workflows and capturing value.

In Australia, Lendi stands out. CEO David Hyman wants to build an almost entirely agentic business. "If we don't do this, we might not have a business in five or ten years."

Lendi CTO Devesh Maheshwari compares today's moment to the early days of cloud computing: "The promise was clear, but the scaffolding, interoperability standards, security frameworks, and observability took years to mature."

That analogy extends globally. However, Optimizely's Kevin Li, senior vice-president of product, argues that while cloud adoption dragged on for years, AI adoption is racing ahead. Cloud required a decade of persuasion and consultants to convince firms to trust others' servers. AI, he says, has leapt from "killer-robot panic to boardroom rollout in less than six months."

The reason? Cultural familiarity. ChatGPT and its cousins didn't sneak in through IT—they arrived via dinner-table curiosity and LinkedIn demos. Consumer comfort bled into the boardroom, lowering the fear threshold that slowed previous shifts.

The diffusion curve is unlike any in memory. Cloud adoption is still only about one-fifth complete, yet generative AI is already embedded in marketing, service, and analytics. The danger, Li warns, is that governance and security lag behind.

Ultimately, that may be the real insight from a year studying agentic AI through the eyes of those betting billions on its future.

The Shift to Agentic AI

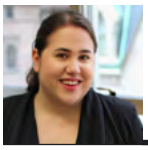

Tessa Conrad
Global Head Of Innovation, TBWA\Worldwide

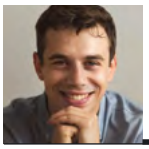

George Strakhov
Global Head of Creative Technology, DDB

Alex Pacey
Chief Product Officer, Omnicom Media Group

THE PATH BEYOND THE APPEARANCE OF INTELLIGENCE

There is a difference between sounding smart and being smart. Between producing coherent-looking text and achieving real-world goals. The same distinction applies in business: the smooth-talking executive fluent in jargon versus the colleague who actually delivers results.

That difference between “sounding smart” and “being smart” now defines artificial intelligence. Marketing leaders must grasp it urgently.

THE CARGO CULT PROBLEM

Physicist Richard Feynman warned of this trap when he coined “cargo cult science.” He described South Pacific islanders who, after seeing military cargo planes during World War II, built replica runways and towers from wood and straw. They mimicked the rituals but lacked the substance. The planes never came.

Feynman noted that some scientists did the same, replicating the appearance of rigor without its core of honesty and empirical validation. The results looked plausible and sounded authoritative but were dangerously easy to believe. As he said, “The easiest person to fool is yourself.”

A few years ago, AI could only produce the appearance of intelligence: fluent text that mimicked expertise without understanding. It was occasionally useful but unreliable and often misleading. Today, AI can deliver both cargo-cult intelligence and genuine intelligence. The challenge is knowing which is which.

TWO PATHS, TWO OUTCOMES

Consider a marketing team that needs a comprehensive plan. Asking a language model to create one will yield something polished, with all the right sections, terminology, and tone. It will sound smart, yet beneath the surface, it is often detached from the realities of the business, market, and challenges at hand.

Now imagine a different approach using the same model. Instead of working in isolation, it is part of a truly intelligent system. It activates advanced reasoning, accesses performance data, analyses competitors, researches proven tactics and emerging trends, evaluates creative assets, and runs scenario planning through adversarial agents that stress-test strategies before launch. It is calibrated continuously against KPIs, benchmarks, competitors, and culture.

This is not theory. It is the difference between AI that produces plausible outputs and AI that delivers grounded, measurable business strategies.

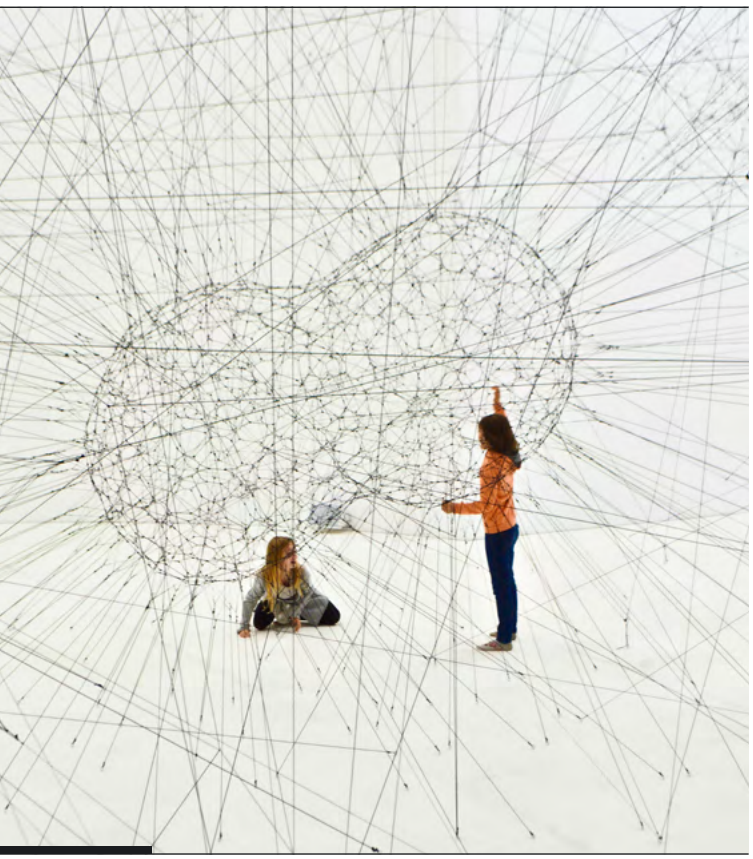

THE INFRASTRUCTURE OF REAL INTELLIGENCE

Building truly intelligent systems is complex. It is not about the latest model or sleek interface. Real intelligence depends on three pillars working in concert.

First, comprehensive, high-quality data: the right data, properly integrated and constantly updated.

Second, skilled humans who understand both technology and business context, capable of architecting and optimising these systems.

Third, robust technical infrastructure to keep everything secure, reliable, and scalable.

These three – data, humans, and technology – must operate as one. Having one or two is not enough; all three must be seamlessly integrated and evolving.

THE LONG GAME

Omnicom has invested in this integration for over 15 years, long before AI dominated conference stages. It recognised early that the future of marketing would not replace human intelligence with artificial intelligence but connect them, creating ecosystems where data, people, and technology amplify each other.

For today's marketing leaders, the lesson is clear: be skeptical of solutions that simply sound impressive. The difference between cargo-cult intelligence and real intelligence may not be obvious at first glance, but business outcomes will reveal it.

In an era where anyone can access AI that sounds smart, the true advantage belongs to those who deploy AI that is smart, context-aware, reality-grounded, and continuously self-improving.

CHAPTER02

Early Use Cases

These are the earliest days of agentic AI, a transformation that will likely take more than a decade to reach maturity. And yet paradoxically agentic AI is already driving outsized business outcomes. The wave of change is evident across categories. In campaign execution, they're testing creative, shifting budgets, and rewriting media plans by the hour. In creative, they're not just generating copy. They're adapting it, enforcing brand rules, and learning what actually moves the dial. In customer experience, agents triage complaints, summarise calls and reroute queues in real time, replacing scripts with systems that listen, learn and act. This isn't one big platform. It's a reimagining of the stack, and more importantly, how marketers work. Governance agents track decisions. Orchestration agents connect CRMs to media systems without detonating compliance. Each layer solves something specific. Together, they rewire how marketing gets done.

When a professional Rugby Union team falters because player-tracking data stops streaming from the field, every second counts. For Alexandra Finster-Rowen, global head of customer experience at Catapult, the solution was not another layer of human escalation but a new species of digital worker.

Her customer-service system, once a patchwork of scripts and queues, has been rewritten using agentic AI. What began last year as cautious experimentation with Zendesk chatbots has evolved into a real-time co-pilot that anticipates, decides, and acts across languages, regions, and tiers of customer urgency.

"In January we saw a 50 per cent reduction in our response time, from twelve minutes to six, after deploying those basic AI tools," says Finster-Rowen. Importantly, trust prevailed: customer-satisfaction scores held steady at 97 per cent even as the machines took over the front line.

Alexandra Finster-Rowen
Global head of customer experience, Catapult

"In January we saw a 50 per cent reduction in our response time, from twelve minutes to six, after deploying those basic AI tools."

That success emboldened Catapult to move beyond scripted automation toward what Finster-Rowen calls "advanced AI." These are agentic systems that no longer follow fixed paths but shape the conversation themselves. Today a coach in London or a data analyst in Sydney can summon support inside the product, in their own language, and reach a human only when necessary. For Catapult, the shift heralds the redrafting of customer experience itself. It also marks Catapult as an early adopter of agentic AI.

And it is a timely reminder that evidence of marketing's early agentic AI engagements are as likely to be found in spreadsheets and log files as on Instagram reels or in TV commercials.

Early agentic systems are already optimising budgets, resolving complaints and refining creative mid-flight. Tech sector leaders we spoke with almost universally said that will turn marketing from a campaign-based art into a self-learning system. But this first wave of agentic AI is proving less a leap of imagination than a triumph of order.

At a growing number of firms, the change is already visible in lower costs, faster resolutions, and the hours saved.

SPINNING FASTER

Retail media is a case in point. Omnicom's Flywheel Commerce Cloud connects the dots from ad impression to warehouse shelf.

Built atop Amazon's vast ecosystem, Flywheel ingests data once locked in separate silos—search, display, retail and logistics—and binds them through Amazon Marketing Cloud, the e-commerce giant's clean-room environment. "Amazon search sits in the Ads Console, DSP data in the DSP Console, retail signals in Vendor Central," says Mohammad Heidari Far, who leads Omnicom's commerce AI efforts. "We've brought all of these points together and used Amazon Marketing Cloud as the foundation to stitch them." The result, he says, is a system that is "not just for media, but also retail-aware," with features ranging from supply-chain forecasting to inventory management.

Flywheel is also emblematic of the sophistication of genuinely enterprise grade agentic systems. Omnicom calls its wider initiative Omni AI – a multimodal suite drawing on 26 large-language models from "all the major players," including Amazon's own Bedrock models. Each of the holding company's networks, from creative to media to commerce, gets its own instance. Within Flywheel, teams can build and train their own AI agents using local data and pre-approved tools. "Some of those globally shared agents are done by dev teams," Heidari Far explains. "We can also set limits, provide documents, and disable web search if we want the agent to be very specific."

Early Use Cases

Customer Experience & Service:
Agents triage support tickets, suggest next actions, summarise calls, and audit quality in real time. Sentiment-aware systems escalate issues automatically, while scheduling bots rebalance queues as demand shifts.

Marketing & Campaigns:
Always-on systems where agents test creative variants, shift budgets, and optimise performance continuously, fast, efficient, but at risk of sameness.

Creative & Content:
Emerging agents generate brand-safe copy and imagery, adapt designs for channels, and learn from engagement feedback, early steps toward machine-assisted creativity.

B2B Sales:
Agents that take the drudgery out of prequalification and which build accurate, tailored pitch documents for prospects and clients in a fraction of the time.

Data & Infrastructure:
Self-healing pipelines detect and fix data drift; orchestration agents sync CRM, ERP, and analytics platforms. The goal: accuracy, traceability, and trust.

Governance & Compliance:
Autonomy comes with audit trails. Agents log every decision, enforce policy, and flag anomalies. "Zero-trust" design is becoming the default.

Productivity & Digital Labour:
Co-worker bots automate cross-department hand-offs, summarise documents, and manage meeting actions. Supervisory "agent-of-agents" monitor and throttle others to keep systems in check.

Guardrails are tight. Access runs through Omnicom's single sign-on, and agents are not automatically shared between users. "That's because we want to make sure that what gets shared has the governance around it," he says. "Ultimately you, as the person, are responsible for the output," he adds. "You can't just blindly trust what each of these agents spits out."

The payoff is tangible. Flywheel now uses AI to run autonomous retail-media campaigns, reallocating budgets by hour, SKU and site. One client, a household-cleaning brand, saw a 43 per cent lift in return on ad spend and a 28 per cent rise in conversion, he said. Another used Omnicom's ArtBot integration to boost creative performance by 65 per cent, as the system learned which ads worked across newspaper and magazine sites.

Omnicom's Flywheel hints at what marketing may soon become: less a campaign calendar, more a continuously learning machine.

The payoff is tangible. Autonomous optimisation is lifting sales, trimming media spend, and accelerating the introduction of new customers to brands at significant levels, all while giving staff back huge blocks of time.

These aren't the micro-minute efficiencies of the old software age – the type of productivity gains that Robert Solow once famously derided as "visible everywhere but in the numbers."

In some of the use cases Mi3 encountered, the upside was dramatic.

Take B2B sales. Execs at Boomi, which is aggressively agentifying its pre-sales processes, demonstrated time savings of 80 to 90 per cent on work that previously took an account rep a day to do manually. Scaled up across a team of almost 400 around the world, that's some serious money in the bank. More importantly it's a huge time-to-market advantage over competitors.

Flywheel

Search is No Longer *Just Human.*

The rules of product discovery are changing. AI-powered search now influences what consumers see and buy, often before they even land on a website.

Our guide, From Powerless to Proactive shows how brands can adapt to this shift: take control of their presence, shape how they're found, and build influence in a world where search isn't just search anymore.

Don't wait to be found.
Take control of how you're discovered.

Download the Guide

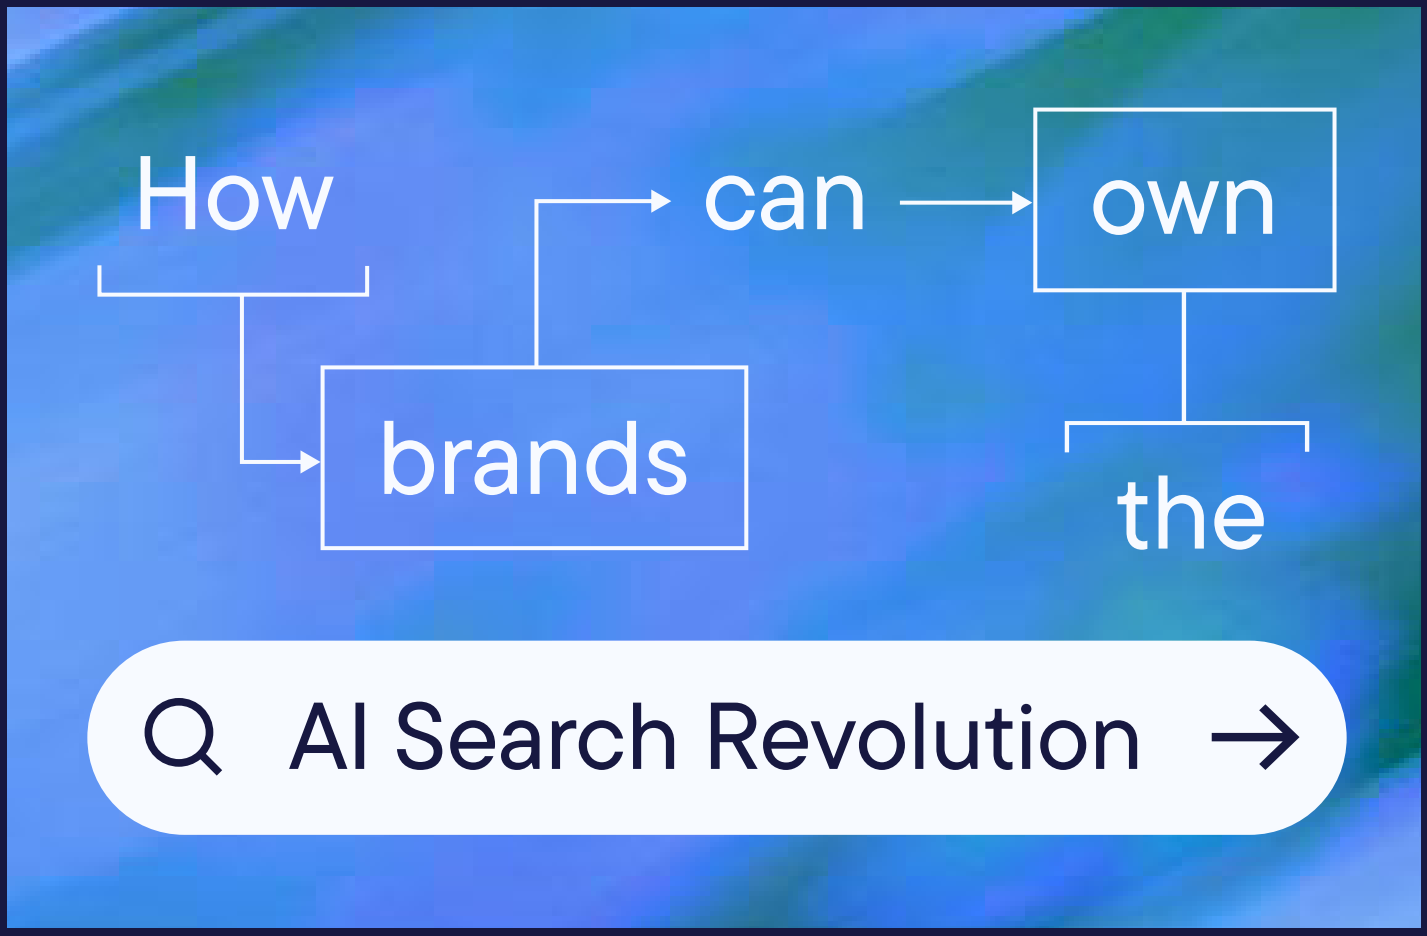

BURN THE SCRIPT

Customer-experience teams, long buried under compliance scripts and service queues, have become laboratories for agentic AI. At Genesys, networks of software agents now triage messages, propose responses and audit outcomes as they occur. Rahul Garg, the firm's head of AI described it as automation that can listen, learn, and respond in real time to human context

"Back in 2014, bots emerged using that classification process to identify intent and guide users through a scripted, very deterministic flow. Then last year, virtual agents came along, we combined LLM technology with classic approaches to give you more flexibility while still supporting deterministic flows."

In other words, you can now blend AI-driven conversation with structured workflows to create more natural and flexible customer experiences.

Firms that have already agentified aspects of customer service told us average handling times are down, quality-assurance loops are faster, and supervisors can now spend more time interpreting patterns than policing staff. It is notable that in the technology sector itself, the discipline of allowing agentic autonomy with customers often begins in measurable service environments where feedback is immediate and mistakes are recoverable.

PLUMBING FOR THE INTELLIGENT ENTERPRISE

The common characteristic of successful early use cases is that they do orchestration well – that's the process of managing how agents work together.

But it's also a function of immaturity because the agents they use are often corralled with one vendor's ecosystem – Qualtrics for customer surveys, for instance, or HubSpot for customer outreach. The task gets much harder as multi-agent ecosystems emerge.

And the tech sector's vision of millions or even billions of agents working autonomously and seamlessly is still years away, despite the hype.

The truth is that the scaffolding required to deliver the vision doesn't yet exist or remains immature – standards and protocols, evaluation tools, observability tools, bulletproof cybersecurity to name a few.

It's like the early days of cloud computing, but progress is being made. Agents, for instance, are proving adept at managing the plumbing that keeps campaigns and customer experiences running.

Michael Bachman, head of AI strategy at Boomi, argues that agent orchestration is where the big early breakthrough will be made. Nor is his a lone voice, it's a point Mi3 heard repeatedly as we interviewed agentic leaders across the world.

Orchestration builds trust, Bachman says, while warning that the alternative is messy at best, and more often chaotic.

"If you don't have a hub and governance and control in the middle, you end up with a spaghetti factory of chaos."

You've got to figure out how to orchestrate all those agents, not just trust them, Bachman says.

GOVERNANCE BEFORE GLORY

Early adopters of agentic AI treat autonomy as a design constraint, not an act of faith. Derek Slager, chief technology officer at Amperity for instance stresses the importance of robust evaluation frameworks.

"Evals" as the techies call them are structured methods for measuring an AI agent's performance to ensure that agentic systems are explainable, measurable, and improvable.

"If you want to move fast and drive a lot of value, you have to have evals. Otherwise, you're just guessing, or you're just crossing your fingers," he says.

He said, "When those things don't exist, that's when you see all the 'AI Going Haywire" headlines ... nobody's going to achieve the goals they have for three years hence if they don't solve for that first."

It's also why firms are building audit trails for every automated decision, constraining agents until their reasoning can be verified.

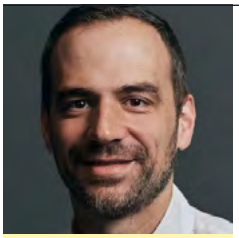

Derek Slager
Chief Technology Officer, Amperity

> **"When those things don't exist, that's when you see all the 'AI Going Haywire" headlines ... nobody's going to achieve the goal they have for three years hence if they don't solve for that first."**

THE ADMINISTRATIVE REVOLUTION

Companies notching early agentic success are those that document processes, measure feedback loops and teach machines to behave like good employees: to be diligent, predictable and auditable.

Governance-before-glory may sound dour, but it is what allows agentic systems to scale. In the words of one analyst we spoke with, "discipline is the gateway drug to autonomy."

Early Use Cases

Mohammad Heidari Far
Managing Director, Flywheel Australia

Alex Pacey
Chief Product Officer, Omnicom Media Group

HOW FLYWHEEL IS DEPLOYING AGENTIC INTELLIGENCE TO REDEFINE COMMERCE

As a leading cloud-based digital commerce platform, Flywheel was early to recognise the potential of the agentic intelligence revolution. It became clear that AI, and specifically AI agents, offered more than hours saved, processes automated, and costs reduced. The Flywheel team saw something much bigger. As long-term natives to digital commerce, they recognised that agentic intelligence could become the dynamic vehicle that propels brands toward meaningful, measurable business outcomes.

Flywheel's approach to agentic intelligence, which brings together autonomous, purpose-driven AI agents with clear business strategy, is already transforming how brands compete, adapt and win in today's fast-moving commerce landscape.

FROM EFFICIENCY TO IMPACT: THE AGENTIC APPROACH

Flywheel's philosophy is simple: what matters isn't just how AI makes work faster, but how it creates real value for clients. Instead of treating every new AI "agent" as a standalone innovation, Flywheel has built a robust agentic infrastructure, an intelligence layer that maps to business challenges and delivers actionable solutions at speed.

Every Flywheel team member, across more than 1,700 colleagues worldwide, has access to a suite of 23 large language models and the tools to build custom AI agents. Each agentic solution is a bespoke problem-solver, securely processing protected client data alongside nearly five petabytes of eCommerce information. Every capability is designed to deliver measurable outcomes.

For example, Flywheel's AI-powered bidding algorithms help brands capture critical market share opportunities during key consumer moments. Predictive inventory tools address the persistent challenge of stockouts, protecting revenue during high stakes selling periods.

This performance-driven approach is already delivering results. One Flywheel skincare partner saw a 20 per cent increase in sales, a 10 per cent drop in cost-per-click, and a 7 per cent lift in new-to-brand acquisition, all powered by hourly insights and advanced dayparting strategies made possible through agentic intelligence.

PRECISION IN MEDIA AND AUDIENCE TARGETING

Flywheel's agentic intelligence also drives deeper optimisation across retail media. AI continuously monitors metrics such as CTR, conversion rates and impressions at the SKU and hourly level, granularity that simply isn't achievable through manual optimisation. The system constantly harvests search term data, identifies high-converting keywords, and adjusts campaign strategies in real time to maximise return on ad spend.

Audience modelling is also being redefined. Flywheel's generative models, trained on each brand's own data, predict which consumers are most likely to drive revenue. Instead of relying on static audience segments, these models build dynamic, evolving audiences that adapt as consumer behaviours shift. The result: 43 per cent stronger ROAS and 28 per cent higher conversion rates for clients using this approach.

So far, Flywheel has launched more than 750 agents, saving each employee an average of eight hours per week, time that is now reinvested in strategic planning and business growth.

DRIVING SMARTER CONTENT, CATALOGUES AND OPERATIONS

Agentic intelligence at Flywheel powers the full commerce ecosystem. AI-driven content tools generate, test, and optimise product listings for both human shoppers and generative search engines like Amazon's Rufus. Digital catalogue solutions automatically synchronise product information across multiple retailer sites, boosting shelf accuracy and driving measurable sales growth.

Operational agents streamline everything from inventory planning to automated fee recovery, reclaiming more than $100 million in invalid deductions for clients to date.

ACCELERATING PRODUCT LAUNCHES AND PREDICTIVE PLANNING

Launching new products in a crowded marketplace is never easy. Flywheel's AI-powered FlyLaunch tool combines analytics, audience insights, and scenario planning to accelerate time-to-market and reduce risk. With predictive commerce intelligence now on the horizon, Flywheel is helping clients shift from reactive to proactive planning, forecasting seasonal trends, optimising promotions, and identifying untapped revenue opportunities before competitors even see them.

LOOKING AHEAD

Reflecting Omnicom's infinite philosophy, Flywheel's agentic intelligence is already delivering real-world results, but the journey is only beginning. As technology evolves, so too will the agentic infrastructure that powers Flywheel's solutions, positioning brands not just to keep pace with change but to anticipate and capitalise on it.

In modern commerce, agentic intelligence is no longer a competitive edge, it is becoming the foundation for lasting business growth.

CHAPTER **03**

From Campaigns to Learning Systems

For the better part of a century, marketing has moved to a familiar beat: plan it, buy it, blast it, measure it. Rinse. Repeat. The campaign cycle wasn't just process; it was muscle memory. Now the tech sector is taking huge bets that the rhythm is breaking. And the replacement they are building doesn't tick. It pulses. Agentic martech, its advocate says, will observe, adapt, and act in real time. Every interaction will be a signal. Every signal will be a lesson. And the learning will never stop. The CMO's role will also morph. In future, there will be fewer architects, more systems conductors, and teams will be less concerned with quarterly cadence, and more focused on maintaining coherence in a stream of perpetual action. Tech leaders we met with stressed the centrality of governance and argued that ethical design is a brand imperative. Performance, they argue, will be measured in adaptation velocity.

MARKETING IN PERPETUAL MOTION

Think of it as the end of the metronome. For more than a century, marketing moved to a familiar rhythm. Budgets were set, creative briefed, media bought, and performance tallied months later. It was an operating model borrowed from the battlefield: plan, deploy, retreat, assess. Now, that rhythm is being challenged. The quarterly campaign, marketing's metronome, is giving way to something perpetual and self-correcting, at least according to the tech sector.

At the core of this shift are systems capable of learning by acting. These agents don't just execute tasks, they observe, adjust, and improve in near realtime. The result – hopefully – is a marketing organism that never stops running, learning, or adapting.

There is already evidence the idea is taking hold.

BEYOND THE ASSEMBLY LINE

Iterable founder and chief scientist Andrew Boni is literally betting the business on this change, and investing in a platform fit for what he regards as the new purpose. "The traditional campaign is very rigid."

Campaigns, Boni explains, were efficient in the way assembly lines are efficient: one message, stamped out a million times, dispatched to audiences defined by blunt demographics. Even as segmentation and cohort models became more sophisticated, the rhythm stayed episodic: plan, build, blast, repeat.

But now, as data analytics has become more sophisticated, and with the rise of generative and agentic AI, those old heuristics are breaking down. "If you're sending out a message that's bespoke to one person, can you really call it a campaign?"

But transparency will be critical.

Boni's colleague, Sam Allen, Iterable's newish CEO calls the change structural rather than stylistic. Agentic martech platforms need to "tie back to the data indicators that marketers will be able to dig into themselves so they can see why the system made a decision."

That gives marketers confidence that the machine is on track, and campaigns that once took months to adjust now safely evolve daily. That's because generative agents can build and test copy variants, measure performance and reallocate spend automatically. The short-term rewards are speed and precision, while the long-term payoff is faster growth.

But this also needs to be tempered by the reality that we are still learning how the generative AI underpinning agents works, and some of what we are learning is troubling.

For instance, Stanford Uni researchers have identified what they call Moloch's Bargain – AI optimised for growth increases sales but it also increases lying to drive those sales. In such an environment, the transparency Allen flags will be critical.

LEARNING LOOPS AND LIVING SYSTEMS

The first signs of this transformation are operational. Marketers who once mapped journeys in annual planning decks are now orchestrating feedback loops. Every click, query, and scroll is a potential signal in a network that tunes itself.

Paul Meinshausen
Co-founder & CEO, Aampe

"What agentic AI represents is the possibility, and the future, where we instantiate intelligent entities that can interact with a person in a structured, guided way. They're guided and managed by humans, but the tactical decisions are delegated, and they manage that experience over time."

Increasingly, composable CDPs and orchestration tools such as Hightouch power this shift. Each micro-interaction – whether form abandonment, scroll depth, or session duration – feeds a system that experiments on its own. The marketer's job is no longer to plot campaigns, but to govern what the machine learns.

Paul Meinshausen, co-founder of Aampe says the goal is not to bombard customers but to teach the system when to speak and when to stay silent.

It turns out that persuasion, long the marketer's craft, is also something LLMs like ChatGPT are very good at.

But success will also require restraint.

According to Meinshausen, "The agent is going to do exactly what we do as humans; observe the response, then look through its library of available actions to actively learn, then based on that one user, decide what the best thing to do is."

What agentic AI represents is the possibility, and the future, where marketers instantiate intelligent entities that can interact with a person in a structured, guided way, he says. "They're guided and managed by humans, but the tactical decisions are delegated, and they manage that experience over time."

GOVERNANCE, CREATIVITY AND CONTROL

As feedback loops proliferate, a subtler challenge emerges: coherence. With thousands of autonomous processes acting simultaneously, who ensures the brand remains recognisable, or ethical? Liz Miller, vice president and principal analyst at Constellation Research, says that the marketer's role is shifting from campaign designer to governor of an always-on ecosystem. She calls this the "new brand discipline" ensuring that every micro-decision stays within tonal and moral guardrails.

"It's always looking at those points of function friction ... because it can rationalize. It can sit there and say, 'Wait a minute, you haven't put any security on this... every other process in this company has security."

According to Miller, ultimately, "AI and especially agentic AI is all about increasing decision velocity. How can I make the best decisions for my customer that then rolls into best decision for my business ... in real time, all the time, always on." This hybrid of automation and artistry is also shaping how brands express themselves. Rather than rely on static brand guidelines, firms can codify tone, beliefs, and design rules into digital brand operating systems frameworks that large language models and creative agents can interpret.

Liz Miller
VP & Principal Analyst, Constellation Research

> **"AI and especially agentic AI is all about increasing decision velocity. How can I make the best decisions for my customer that then rolls into best decision for my business ... in real time, all the time, always on."**

To that end custom AI agents are increasingly being embedded across their creative workflows.

For instance: Omnicom's global innovation practice has trained models on years of proprietary research to help creative teams explore new ideas. For McDonald's, it built agents that digest "Fan Truths" – real customer sentiments – to guide global creative. For Henkel, an end-to-end workflow allows product expert agents to generate insights, simulate audiences, and adapt assets automatically. Creativity is no longer confined to a studio – it's diffused across an adaptive network.

McDonald's uses custom Low Rank Adaption (LoRA) techniques and visual models to preserve brand fidelity, even when fans generate content themselves. Nike's Airimagination project for instance lets consumers co-design shoes using a conversational AI trained on the brand's design archive.

Such systems reveal a deeper shift in creative philosophy. Human creativity, storytelling, and emotional resonance remain central. But the job is changing, from crafting singular ideas to designing adaptive systems that can scale, learn, and maintain coherence at speed.

That coherence increasingly depends on data plumbing. After years of integration, brands that have their data architecture in order are well placed to run learning loops continuously. Each behavioural signal becomes a training input; each adjustment, a new experiment. Marketing departments once divided into creative, media, and analytics silos are now reorganising into interdisciplinary pods. As a result expect to see a flattening of organisational structures across both brands and their agencies as data flows are democratised.

THE BEST IDEAS AREN'T JUST CREATIVE. THEY'RE INTELLIGENT.

The marketing landscape is evolving fast, powered by data, creativity, and the rise of AI.

Transformational Ideas Volume 4 explores how these forces come together to drive real business growth. This isn't a trends report, it's a blueprint for turning intelligence into effectiveness.

DOWNLOAD VOLUME 4:

TRANSFORMATIONAL IDEAS FOR TRANSFORMING TIMES

OMG OMNICOM MEDIA GROUP

THE FEEDBACK ECONOMY

For decades, marketing relied on anticipation, guessing what people might want and convincing them to want it. Learning systems invert that logic. They observe first, act later, and measure constantly.

However, you can't just automate a task and declare victory, says Alan Trefler, founder and chief executive of Pega. "The problem is you have to use the right AI for the right things... you need a conductor. You need a way to direct the business, to use the right AI or use a person to do the right things at the right time."

According to the Pega boss, "You want the AI to be predictable... so that two people whose problems are basically the same but are slightly different, don't get different answers. That's how businesses get into trouble."

At ServiceNow, global marketing chief Colin Fleming has experienced the results of this programmatic learning firsthand. When his team re-engineered its demand-generation process using Adobe's Journey Optimiser, he says, "We dropped lead gen 63 per cent but converted 133 per cent increase. It allowed us to save money and not waste media too."

Alan Trefler
Founder & CEO,
Pegasystems

"You want the AI to be predictable... so that two people whose problems are basically the same but are slightly different, don't get different answers. That's how businesses get into trouble."

The firm's systems learned to prioritise quality over quantity. Depth replaces volume. "We never had a lead problem, actually, never had a pipeline problem. The issue was converting that pipeline. We basically looked at it like: we can go down this path and keep chasing one to two per cent conversion optimisations... or, what if we stepped back and aligned our marketing with how people actually buy our technology?"

The implications stretch far beyond marketing. As feedback loops spread across service, product, and commerce functions, companies will increasingly behave more like adaptive organisms. Agentic AI will be used to fuse what were once discrete departments into a single network that translates interaction into intelligence and intelligence into action.

In this new operating logic, the CMO becomes part scientist, part ethicist, part conductor. Decisions are decentralised, but direction remains human. The firms that master 'automation with accountability' will move faster than those still debating quarterly campaign plans.

From Campaigns to Learning Systems

Benni Lucas
General Manager, Growth, Product and Innovation, Resolution Digital

Luke Eid
Chief Innovation Officer, Omnicom Advertising Group

Matthew Moran
Global Head of Strategy, Innovation, Omnicom Advertising Group

Kay Chong
Senior Global Strategist, Innovation, Omnicom Advertising Group

THE SHIFT IN MOTION

Think like a modern Formula 1 team. A fast car and a clear visual identity still matter, but victory comes through telemetry that never stops. Winning teams observe what is happening on track, decide the next move, act quickly, and learn before the next corner. This is the same transformation underway in marketing. The industry is shifting from set-piece campaigns to learning systems that run continuous loops across creative, media, and customer experience. People define strategy, taste, and guardrails. Agents and automation provide speed, coverage, and consistency.

INSIDE THE LEARNING SYSTEM

Four loops run continuously: signals from behaviour, creative, product, and media; decisions that choose the next best action; actions in creative, media, and CX; and learning that feeds the next decision. Agents accelerate the loops, while people approve changes and own outcomes.

CREATIVE: FROM BRAND GUIDELINES TO BRAND INTELLIGENCE

Creativity must scale distinctiveness, not blunt it. This means codifying brand identity so it can travel through tools and co-creation without losing its signature. Across Omnicom's network, teams are moving from static brand books to living brand operating systems. Insight agents keep foundational truths alive by ingesting reviews and behavioural signals. Synthetic personas provide rapid feedback before production. Content engines generate and adapt assets in market while remaining recognisably on brand.

McDonald's keeps "fan truths" active, allowing ideas to grow with culture. Nike invites communities to explore Air Max concepts with archive-trained tools that still feel unmistakably Nike. Henkel uses expert agents, synthetic personas, and journey planning so assets learn and respawn across channels. The creative role shifts from producing a finite set of assets to designing the system that continually makes and refines them.

MEDIA: FROM PLANS TO POLICIES

Media planning now functions like a control room rather than a calendar. Planning becomes scenario design and trading becomes adaptive. Guided briefing and synthetic audience simulations allow planners to test hypotheses before money moves. Questions such as "If five points shift from social to CTV, what happens to attention, reach, and cost?" can be answered with evidence.

In market, budgets, bids, and cohorts adjust to outcomes within defined policies. Attention, viewability, and suitability can be weighted alongside commercial goals, ensuring growth does not come at the cost of margin or experience. Practitioners move from micromanaging levers to supervising policies and approving shifts the system proposes. It mirrors the pit wall, where strategy adjusts lap by lap within shared rules.

CUSTOMER EXPERIENCE: CLOSING THE FEEDBACK GAP

Customer experience is where the loops converge. Synthetic audiences allow teams to pre-test messaging, value framing, and journey steps

with personas that reflect real segments. They do not replace people; they reduce blind spots and accelerate learning about what to test first.

Omnicom applied this with Torrens University. Agents assisted with rapid ideation, research synthesis, and performance analysis across search behaviour, CRM signals, and on-site interactions. Synthetic audiences such as “hesitant explorer”, “cost-sensitive prospect”, and “high-intent applicant” shaped value framing, content hierarchy, and follow-up journeys. High-impact modules like fees, outcomes, and social proof were versioned and tested, with winning patterns informing creative rotation, media targeting, and experience.

OMNI: THE SYSTEM BEHIND THE SYSTEM

A Formula 1 team does not win on the car alone. It wins because telemetry is clean, the pit wall sees the track, and tools in the garage enable rapid tuning. This is the role of OMNI in Omnicom’s ecosystem. It gives the loops a home: a place to brief clearly, simulate outcomes, activate with control, and feed results into the next decision. OMNI maintains identity as volume scales, enables plans to operate as scenarios, and surfaces recommendations for teams to approve.

Agents are only as good as the data they receive, and outputs only as good as the humans steering them. OMNI strengthens both, ensuring connected signals in and accountable decision-making out.

THE FINISH LINE: CONTINUOUS ADVANTAGE

A Formula 1 team wins with a fast car, a clear identity, and the discipline to learn every lap. Brands now have the same opportunity. Codify who they are so creativity scales without losing its voice. Treat media like a control room so investment follows outcomes. Keep customer experience in the loop with synthetic audiences so experiments begin smarter. The advantage belongs to those who put these systems to work and learn, lap after lap.

CHAPTER **04**

Creative in the age of agentic AI

Some of the world's biggest firms are already deploying early examples of creative systems that learn, adapt, and generate on demand. With Agentic AI a sneaker sketch can become a product within days, driven by community input and archive-trained models while McDonald's is mining fan truths to spin real-time campaigns. Increasingly, AI agents are emerging as the new Creative Ops. But delivering agentic creativity will require a mesh of strategic systems, brand operating systems, and content engines that can operate at industrial speed. And this will be under human oversight that will be increasingly 'On the loop' not 'In the loop' – managing by exception, instead of approving every change.

A sneaker sketch becomes a shoe. Not because a designer wins an internal pitch, but because a conversational tool trained on decades of Nike's design archive spins community suggestions into manufacturable forms. Meanwhile, a fast-food giant routes "fan truths" from millions of reviews through an agentic process that adapts visuals and copy in hours, not quarters. A chemical firm wires product experts, synthetic personas and journey planners into one end-to-end workflow; assets spawn, learn and respawn across media without waiting for the next "big idea."

According to DDB Group Chief Creative Officer, Matty Burton, the future of creativity isn't a campaign, it's a system.

"Technology is very good at removing friction," Burton says. "What we're learning now is how to put meaningful friction back in – the kind that sparks curiosity and emotion." That shift underpins DDB's move from one-off ideas to what he calls *creative systems thinking:* platforms that scale, evolve, and compound impact over time.

Matty Burton
Group Chief Creative Officer, DDB Group AUNZ

"Technology is very good at removing friction," Burton says. "What we're learning now is how to put meaningful friction back in — the kind that sparks curiosity and emotion."

A case in point is Samsung's Try Galaxy, a product-led experience that lets iPhone users test-drive the Samsung operating system through a web app. What began as a clever experiment in conversion became a global tool that updates automatically each time Samsung releases a new feature. "Every time Samsung updates their operating system, or pushes out a new feature that can work using AI, we have to figure out a way to build that to work on an iPhone that can live through a web app," Burton says. "That keeps our dev team right on the cutting edge of what's possible on device."

The approach also changed how DDB gets paid. The agency co-owns the platform with Samsung, tying remuneration to performance rather than production.

"It unlocked a different remuneration model that we've been following ever since," says Burton.

For him, this is the logical evolution of creativity in the agentic age, ideas that behave more like products than ads. "It's product thinking meets creative systems and the impact that compounds over time."

In that model, the creative doesn't end when the ad goes live. It learns, adapts and grows, a feedback loop between brand, audience and machine.

Creative work, once episodic and human-paced, is starting to run like an operating system.

The thesis is simple, and unsettling for some: creativity is moving from events to systems. Instead of planning campaigns, firms are building orchestration layers that learn from every interaction. Human imagination does not disappear; it supervises. The measure of success shifts from one perfect ad to ten thousand variations that stay on brand. The economics change too. Once a few hits justified the spend; now scale and fidelity are the returns. What is emerging, across agencies and software vendors alike, is a new creative stack: strategic systems that define purpose, brand operating systems that codify identity, and "infinite content engines" that generate and distribute work in real time.

Sundeep Parsa, vice president for Adobe Experience Platform says, Adobe is reimaging its Experience Cloud "with an agentic-AI-first mindset," connecting data, content and journey orchestration so applications can reason about goals and context. Creativity is connected to productivity and customer experiences: coordinating agents that understand intent, tone and outcome rather than executing isolated tasks.

The same shift is visible further down the stack. Box's pilots use agents to analyse drafts, surface patterns and suggest new directions; in CTO Ben Kus's words, the focus is "making unstructured data more intelligent. [It's] about understanding relationships rather than simply storing content." The result is not a machine that dreams, but a system that remembers. And it's one that turns the archive itself into a creative substrate.

Liz Miller, a principal analyst at Constellation Research, distils the attraction for executives. She says, "Agentic AI can run experiments and permutations in seconds, but it's also highly documentable."

According to Miller, "It can say, 'Whoa... alert me to anomalies. Alert me to anything false. Alert me if this threshold drops in a certain way...' It's getting smarter, because the more processes we automate and plug into these agentic workflows, the more the agent can come back to the business users. It can come back and say, 'Hey, this protocol is out of date. Would you like to see the new protocols available?"

CODIFIED IMAGINATION

If perpetual testing is the engine, codified identity is the gearbox. Firms are not leaving tone, visual grammar or ethical red lines to chance; they are encoding them. Instead of static brand books, teams build digital brand operating systems that models can read and apply. In practice that means training lightweight model components that steer outputs toward a house style, even when outsiders make the art.

Liz Knox, chief operations officer at DDB Group NZ, explains:

"When McDonald's ran its 'Fan Favourites' campaign, fans could co-create visuals while the system made sure everything still looked and felt like McDonald's. Nike did something similar — letting its community sketch designs guided by what the brand's archive had already taught the system. The goal isn't to limit imagination, but to make sure creativity still feels recognisably on brand, no matter who's behind it."

Knox says major clients like McDonald's are now building this in three layers:

1. **Strategic systems** that define goals and set creative objectives.
2. **Brand systems** that encode tone, values, and design rules.
3. **Content engines** that generate and distribute ideas through proprietary Omnicom creative platforms like OAI or ArtBot.

"In this model," she adds, "custom AI agents operate across the creative lifecycle. Some scan reviews to find insights; others simulate audiences; some plan where and when to launch content; and others evaluate how it performs – all feeding the next creative cycle. What used to be a straight line is now a living, learning web."

Software firms are quietly converging on similar ideas. David Meyer, SVP Product at Databricks, however offers a cautionary note.

"You can't define a complicated system with a simple prompt. You need to identify all the edge cases and what quality looks like... We have this very novel evaluation framework that, when you build these agentic systems, you can deploy them in production, because you can track their efficacy over time and quality and cost."

David Meyer
Senior Vice President, Product, Databricks

"You can't define a complicated system with a simple prompt. You need to identify all the edge cases and what quality looks like."

SHIFTING RESPONSIBILITIES

The agentic turn also redistributes labour. Machines handle brute-force variation and curation of patterns; people set aims and police meaning.

At Australia Post, the problem Aimee Dixon, GM Enterprise and Brand Marketing, and her team needed to address will be familiar to many marketers: too many campaign assets, too little time. The solution was to encode brand decisions into templates and let algorithms produce bottom-of-funnel variants. "It's really about creating a platform where all that conversion-type marketing can be seamless."

She told Mi3 earlier this year, "It's now the standard for all campaigns, so my marketers don't have to manually check bottom-of-funnel assets; it's automated. Since we've approved the brand guidelines, we can generate a wide range of activity across campaigns, and then let the algorithm determine which ones are more likely to convert."

ROBO-PROMPTING

Hightouch founder, Tejas Manohar, flags another trend he believes marketers will need to account for.

"Most of the prompts in the future are going to be computer-generated."

He envisages a world where machines generate prompts automatically from data insights, turning marketers into reviewers rather than creators.

"If, in the future, your content team is just going into a content AI tool and typing in prompts based on their ideas, you're missing a lot of the value. You can now programmatically connect your data and insights layer to your content team, so the insights system can generate prompts for the content AI system to create more interesting variants."

THE HUMAN BARGAIN

Automation, however, has a nasty habit of flattening taste. If a system learns purely from engagement, it will optimise toward what is most clickable, not what is most meaningful. That tension between scale and originality sits at the heart of agentic creativity.

DDB's Knox says the best programmes are treating authenticity as an engineering challenge:

"They don't assume more content is better content. They build systems that know how to judge what feels right."

She points to industrial giant Henkel's AI workflow: product-expert agents generate insights, synthetic personas test whether ideas will resonate, and "journey agents" decide which channels and moments to launch in.

"The outputs still carry human fingerprints," she says, "but the process is now a feedback loop – less gut, more loop."

That loop extends to marketing tech vendors. Customer-service platforms let conversation agents draft and tone-test thousands of micro-messages while supervisors define what "good" looks like. Box's systems for instance, turn unstructured documents into a navigable memory.

GUARDRAILS AT THE GALLERY

Still, the risks are real. Codification freezes taste if done badly; co-creation can drift off brand without guardrails; and the temptation to automate judgment will always be strong when calendars and budgets pinch. The safeguard is governance that is as operationalised as the engines it oversees: permissions in code; evaluation as a service; human-override thresholds where reputational risk spikes. Firms that treat governance as paperwork will drown in their own agents. Those that treat it as architecture will find – counterintuitively – that constraint accelerates.

Practicalities will also intrude. If consumption-based AI charges are the new norm, CFOs will ask whether "infinite" engines run hot when they should idle. That drives a new role for creative operations: not merely output management but policy, when the system is allowed to explore, when it must exploit, when a human has to review, and how much novelty is acceptable this quarter for this brand. In other words, brand voice becomes a set of statistical tolerances as much as a slogan. Oh joy.

The culture change is no sideshow. Hierarchies built for quarterly campaigns struggle with always-on testing. Teams that will thrive are those that display experimental confidence, treating failure as data and curiosity as discipline. Tech leaders we consulted imagine the emergence of full-stack marketers as strategists with data literacy and agent orchestration skills, replacing the narrow specialist.

BUT IS IT ART?

What, then, does "originality" mean in an age of infinite variation? Perhaps less about surprise, more about consistency with purpose. A new test of a brand's imagination is whether it can produce countless micro-moments that still feel like itself.

The worst sin may not be sameness but drift. This is where codified identity will earn its keep. A burger photographed by a fan, a shoe co-designed by a teenager, a video remixed by a community can all be recognisably on brand when the operating system knows what the brand is.

Sceptics might object that this is bureaucracy by another name.

They are half right.

Creativity is indeed becoming a managed process, one that is repeatable, measurable and, at best, self-improving. But management is not the enemy of imagination; indifference is. Systems that learn can also be taught. The task for leaders is to encode not just boundaries but the values, for instance what does the firm stand for when the data points the other way.

The most illuminating signal may be what the best programmes choose not to automate. Some moments still demand human touch: a product apology, a cultural reference that could cut both ways, a joke that might age badly by morning. In those cases, the system must escalate, not improvise. The point of an operating system is not to remove judgment; it is to make room for it.

Creativity in the agentic age

Matty Burton
Group CCO DDB ANZ

Luke Eid
Chief Innovation Officer, Omnicom Advertising Group

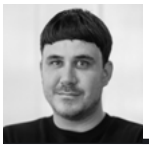

Matthew Moran
Global Head of Strategy, Innovation,
Omnicom Advertising Group

Rachel Bayfield
Chief Technology and Innovation Officer PHD

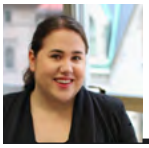

Tessa Conrad
Global Head Of Innovation TBWA\Worldwide

Simon Vicars
CCO Colenso BBDO

THE CREATIVE RENAISSANCE: REDEFINING EXCELLENCE IN THE AGE OF AI

Every creative era begins with a shift in perspective. The Renaissance was defined by art, philosophy, and science working together to expand what humanity could imagine. Today, a similar transformation is unfolding as technology broadens the creative canvas. Artificial intelligence is not here to replace human imagination but to extend it, helping ideas travel further and connect faster.

At Omnicom, this belief defines how creativity is made and how it is thought about. The focus is not on machines for their own sake, but on how they can amplify human ideas, deepen cultural relevance, and accelerate the creative process without losing authenticity.

Colenso BBDO's Pedigree Adoption campaign captures what is possible when creativity and technology align. By using AI to enhance images of shelter dogs and spark empathy at scale, the work shows how emotion and innovation can coexist. With every leap in technology comes a risk of overproduction, where efficiency increases but effectiveness declines. Yet within that tension lies the opportunity to redefine creative excellence. As technology removes friction and cognitive load, creativity's purpose becomes clearer: to make people feel something real. Emotion, surprise, and delight will define value in a world increasingly shaped by automation.

Human curiosity remains at the centre of this new era. Samsung's Try Galaxy by DDB New Zealand invited iPhone users to explore Samsung's ecosystem from their own devices. The work was built with advanced AI, yet it resonated because it spoke to a simple truth: people are naturally curious about what lies beyond their own walls. DDB Chicago's Skittles Pop'd followed a similar principle, using AI to create 35,000 personalised micro-stories that felt spontaneous and playful. These examples show that the future of storytelling will be co-created, with audiences invited to participate in worlds that evolve rather than static campaigns that broadcast from afar.

Static brand assets will not survive this shift. Living brand systems, such as Skinny's Low-Cost Liz from Colenso BBDO, demonstrate how distinctive character can merge with adaptive intelligence. Liz is not a mascot but a dynamic personality that evolves with context while staying recognisably on brand. As technology perfects precision, creativity must deliver emotion. The future belongs to brands that feel alive, capable of reacting and learning in real time.

As innovation accelerates, culture often moves in counter-cycles. The rise of automation has been met with renewed appreciation for human craft. Telstra's stop-motion work from Bear Meets Eagle On Fire and Plus 61 celebrates the artistry of handmade storytelling. The most successful brands will blend this craftsmanship with technology,

creating hybrid systems that are both culturally resonant and operationally agile. DDB Australia's DoorDash campaign is one example: hand-crafted puppets brought to life through digital rigging, designed to adapt seamlessly across channels without losing their warmth or integrity.

The question that once followed the invention of the printing press now follows AI. If machines can replicate human output faster, what makes human creativity special? The answer, as before, lies in imagination and intent. The tools have changed, but the purpose remains the same: to move people, to surprise them, and to connect them to something genuine.

Omnicom sees this as the true opportunity of the creative renaissance. The future will not be won through speed or efficiency alone, but through the combination of human insight and intelligent technology. The role of creativity is to rise above automation, using new tools to express old truths in ways that are timely, distinctive, and deeply human. The canvas has expanded. Those who learn to paint on it with both precision and emotion will define the next era of excellence.

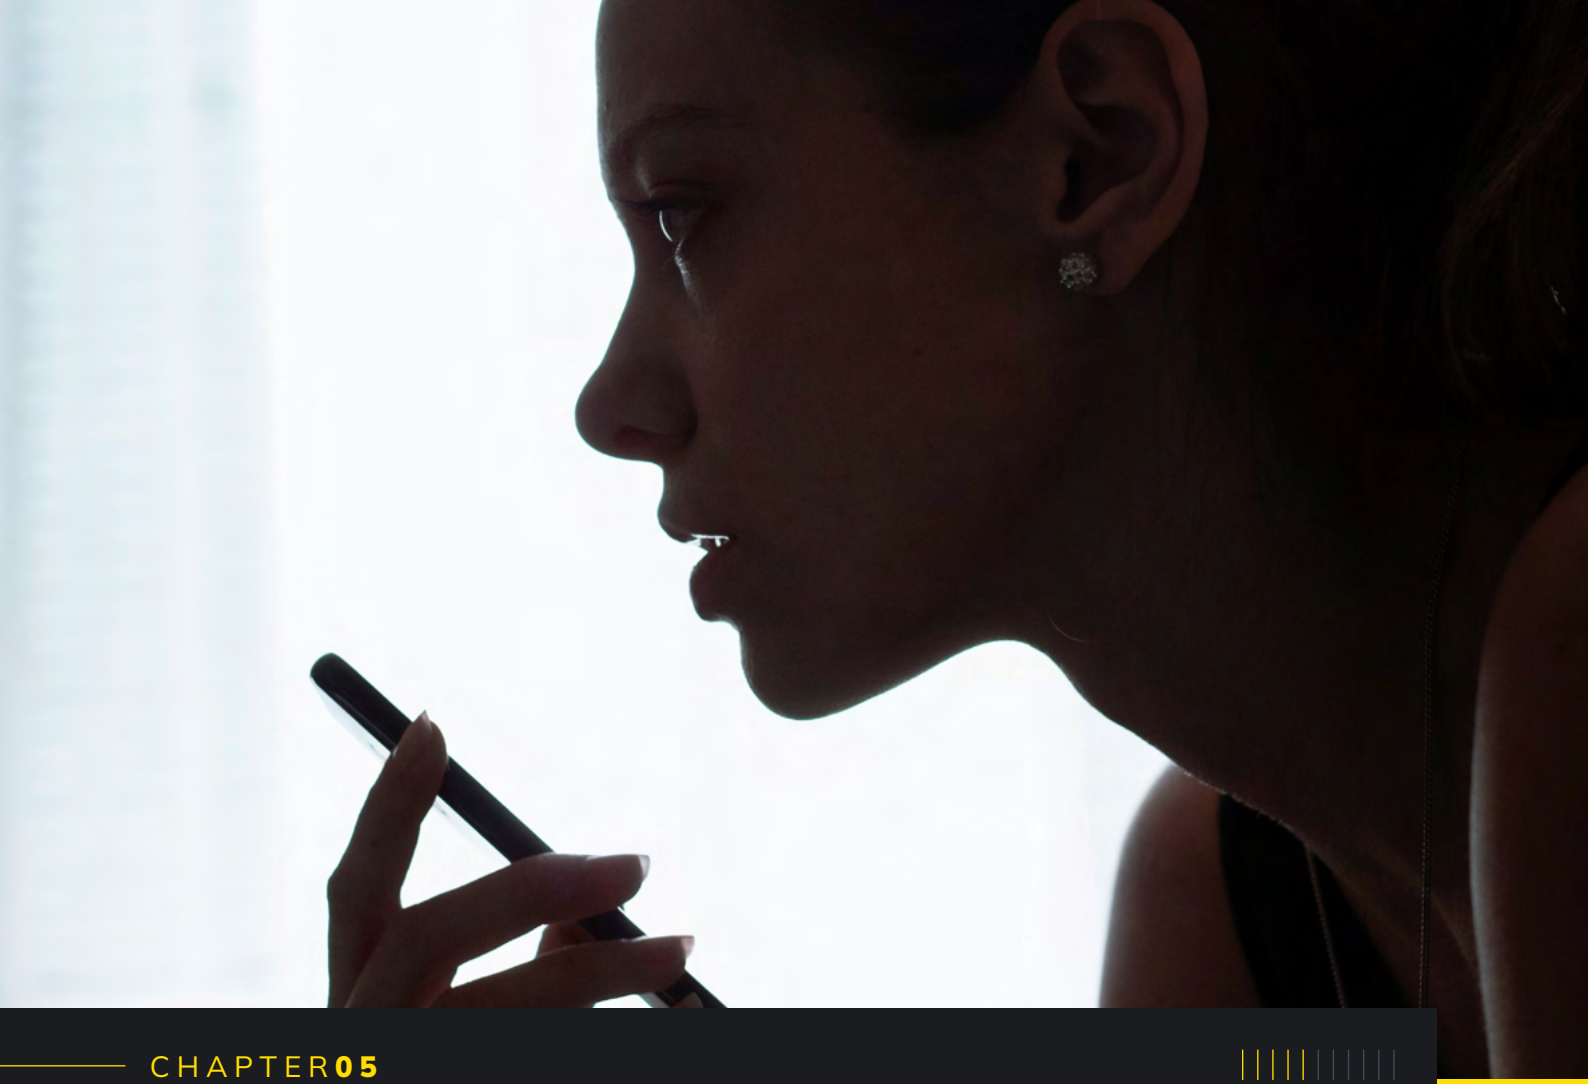

CHAPTER **05**

Remaking the Stack

SaaS, once the insurgent, now risks becoming the casualty. Agentic AI is not just another layer in the martech stack, it's the solvent. Interfaces are morphing from dashboards to dialogue. For marketers, it means the power dynamic shifts, away from the hands of technical specialists and into the natural language of everyday users. For vendors, it threatens the very spine of their business model. Why does this matter now? Because the economics of software as a service could begin to collapse. Agentic systems offer to do more, faster, for a fraction of the price. Forrester's Chairman George Colony warns against marketers signing long term SaaS contracts. And pricing models themselves may shift from seat-based subscriptions, to transactional charges, and ultimately to outcomes, in a shift that redefines value.

REMAKING THE STACK - WHEN SOFTWARE STARTS TO THINK

Fifteen years of dashboards and plug-ins risk being swept aside by software that refuses to sit still. Agentic AI could dismantle the tower of SaaS that once seemed unassailable.

According to Scott Brinker, martech's indefatigable cartographer, and editor-in-chief of Chiefmartec.com, "What's really interesting about AI is that it's not only expanding the set of possible capabilities, it's also, for the first time, providing a new kind of interface that can actually make it easier for us to tap into those features

Huge new tech waves often begin with the interface – think Mainframe to PC, or feature phone to smartphone. This latest upheaval, shifting us from browser to prompt, is following the same arc.

In this new landscape, agents are not so much installed as summoned. Brinker imagines a trillion such entities, blinking in and out of existence as required. Already, hundreds of "little AIs" already perform distinct tasks such as composing copy, predicting churn, or resolving tickets. What was once middleware is morphing into an ecosystem of ephemeral, conversational helpers stitched together by context as much as code.

THE SLOW DETONATION OF SAAS

For George Colony, the founder of Forrester Research, the blast radius extends well beyond marketing. "This will be the biggest technology change of my lifetime," he says. Having watched the mini-computer, the PC, the internet, social, mobile and cloud redraw the landscape, he now insists the seventh wave represented by artificial intelligence will be the most violent. The spark, he argues, was AI's "clickable moment" – the arrival of ChatGPT. As with the browser or touchscreen before it, the human-machine handshake has changed, and with it the economics of software itself.

Agentic systems promise to do the same work as legacy SaaS but for a fraction of the cost, "perhaps as little as a tenth," Colony told Mi3.

The implications for the vast SaaS edifice are brutal. "If you are looking to buy a traditional SaaS platform in the next two years, wait," he warns. "Or sign very short-term deals. Do not make big commitments."

Vendors, meanwhile, are trying to disguise the blast. "They're claiming to be AI, and they'll say they're AI, but are they truly AI?" Colony asks. His advice is mordant: "Beware the legacy vendor bearing AI gifts."

Brinker, meanwhile, says agents are becoming "the new iPaaS," brokering between systems once joined by human middleware.

"Anthropic introduced the Model Context Protocol (MCP) and within months, OpenAI, Microsoft, and Google jumped on board." MCP is a simple framework that allows agents to interact with external data sources and APIs. "It doesn't solve everything, but it removes friction," he said.

FROM SUBSCRIPTIONS TO OUTCOMES

The economic inversion is stark. SaaS crept into firms through marketing and HR departments, bypassing CIOs. Now it is the establishment, bloated, overpriced, padded with features few customers use, according to its critics. Utilisation rates in some categories are less than 20 per cent according to SaaS licencing platform Zylo.

Generative and agentic AI introduce an irresistible incentive to shift costs from licences to usage, and from usage to outcomes. Call-centre software may soon be priced by resolution, not by seat; finance systems by reconciliations completed rather than modules installed. "We're seeing the emergence of more consumption-based or outcome-based pricing models," Brinker says.

Paying for what works sounds rational, but it replaces predictability with probability. His counsel is pragmatic: treat outcomes as measurable units. "If I end up paying $5 for every resolved ticket, yes, there's some variability, but generally, a business knows. 'Okay, we get about 10,000 support tickets a year.' If we expect AI to resolve, say, 67 per cent of them, then we can estimate the cost pretty reliably and add a buffer." Marketing, he suggests, is adopting the reflexes of cloud computing's FinOps.

Forrester's Colony calculates that the numbers will be even harsher for incumbents. When firms such as Oracle or SAP face the loss up to nine-tenths of their SaaS revenue, they will not go quietly. He expects them to deploy a "buy, block, pretend, link" strategy of acquiring innovators, obstructing rivals, feigning transformation and integrating just enough to survive another quarter.

"Fear is everywhere", he says. SaaS, the insurgent of one wave, becomes the casualty of the next.

HUMANS IN THE LOOP

If technology accelerates exponentially, organisations change logarithmically, a phenomena that Brinker famously tagged as Martech's Law. Even he admits, "I don't know of a shortcut for that."

And you still need humans, says Genesys's Rahul Garg. "They add empathy to the conversation. They also can reason a little bit better... There are still use cases such as life insurance that will never be a virtual-agent use case, it's too personal."

This pattern of machine hands and human eyes will define marketing's next wave, although the jury is still out on where the balance lies.

Rahul Garg
VP, Product Management, AI and Self Service, Genesys

"You still need humans... They add empathy to the conversation. They also can reason a little bit better... There are still use cases such as life insurance that will never be a virtual-agent use case, it's too personal."

Businesses we spoke with who are experimenting with agents say building connected workflows remains "more complex than anticipated." Many firms are discovering that human judgment is still the critical middleware. The lesson is operational humility: AI can accelerate execution but cannot yet replace oversight.

And that is as much an organisational culture issue as a technology one, as Adobe's Sundeep Parsa explains. "The way we are framing our questions is no longer about speeds and feeds of the technology, how fast it can go. We have solved this problem. The organisational agility is the next frontier. You've got to think about the people process... What's your end-to-end intake process? How many approvals and steps are needed? Are you open to some simplification?"

Sundeep Parsa
Vice President, Adobe Experience Platform

"The biggest risk is not to be in the game and wait for something to emerge—it's too late. You've got to be inside the tornado."

But he is also alive to the fact that there is not a lot of optionality when it some to managing the pace of change.

"The biggest risk is not to be in the game and wait for something to emerge – it's too late. You've got to be inside the tornado."

GOVERNANCE AND THE NEW STACK

Data, once an asset to be collected, has become a material to be governed. CIOs are tightening control of AI trust layers to ensure commercially valuable content does not leak into open models. Reality-checkers in the marketing trenches already see clients demand clean-room environments and temporary, view-only access for agencies. Marketers are urged to "buy open": choose platforms that share data and services freely with AI components and treat agents as users, with permissions, limits and logs.

Forrester's Colony agrees that trust will be the new currency. "No agent should ever be built that we cannot trust," he warns. "You're going to have to manage agents as well as you manage people… monitor them, give them goals, and audit them as they're moving through their day."

Colony adds a caution: today's systems are "more agent-ish than agentic." But the trajectory is one-way. Once agents can test, learn and adapt in real time, human-driven campaign cycles will look as quaint as buying banner ads on Yahoo.

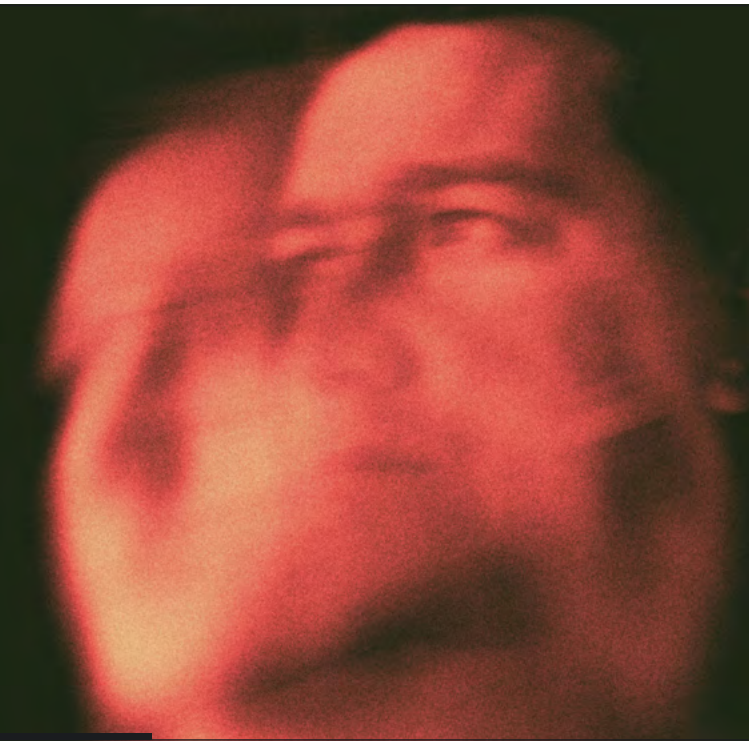

NAVIGATING THE *RAUSCHVERWIRRUNG* ERA

Speed is the shock. In meeting after meeting with founders, AI heads and SVPs of product MI3 encountered a common emotion. The closest approximation in English was fear, but it was more than that – something more akin to *Rauschverwirrung*: the dizzy mix of exhilaration and confusion.

Chiefmartec's Brinker reckons that AI is advancing faster than the internet did at a comparable stage. Forrester's Colony concurs: "The industry is larger, the pace is faster, and the level of ferment and turmoil far greater." What took the internet years has taken AI months. Even its creators, he notes, "struggle to predict what comes next."

Marketers can never operate at machine speed, but they will have to get comfortable managing the impacts. Data stewardship, permissioning and feedback velocity will become sources of edge. Agencies, accustomed to charging for time, must now consider charging for outcomes such as decision velocity, brand-safe autonomy and the ability to steer learning systems. The rules of the stack are being rewritten in contracts as much as in code. As Colony puts it with characteristic brutality, "AI will gut SaaS pricing, collapse the stack." Brinker's empirical view is cooler: the stack is liquefying into continuous experimentation.

Remaking the Stack

Mark Gretton
Chief Technology Officer, Clemenger BBDO

TECH STACK AND BUSINESS PROCESS EVOLUTION IN A WORLD OF AGENTIC AI

AI AS THE SOLUTION TO THE PRODUCTIVITY CRISIS

Australia's productivity problem has become a boardroom obsession. As inflation squeezes margins and redundancy programs sweep industries, leaders are turning to AI as the silver bullet. Marketing, long seen as a stretchy budget, is now ground zero for AI-driven reinvention. While Share of Voice arguments keep media budgets safer, operations, production, analytics, content, and martech are being reshaped.

In response, marketing and agency teams are rapidly building new agentic skills across every part of the process. With shoppers using AI tools to plan entire tasks rather than single keyword searches, customer experiences and platforms are adapting fast. Content and optimisation tactics for AI-driven discovery are driving new AEO and GEO playbooks. The pace of change and lack of clarity about the long-term landscape are creating shorter payback expectations than ever before.

AGENTIC SKILL MANAGEMENT SOLUTIONS VIE FOR MARKET DOMINANCE

Platforms like ChatGPT (Custom GPTs), MS CoPilot (Agents), and Gemini (Gems) now enable non-coders to build and share libraries of agentic skills. Each is racing to become the homepage for a company's agentic intranet, the mission control for human to agent interaction. Built-in SSO and cloud security settings reduce risk and give IT leaders confidence.

Across Omnicom Oceania, thousands of agents have already launched, from tone checkers and creative critics to synthetic audience builders and estimation tools. These run through Omni AI, which provides secure access to multiple LLMs such as OpenAI, Claude, Google, and AWS via a single platform. It allows teams to share agents, standardise new processes, and maintain strong network controls.

MOVING FROM SINGLE AGENTS TO WORKFLOWS

As teams automate single tasks, the next step is linking agents into workflows. Single agents can overload or hallucinate when given too much data, so validation and QA must be built in. Planning human and AI handoffs and tightening governance become critical. Tools like n8n and Zapier first connected agents and software, but new features in ChatGPT's Workflow Builder and MS Copilot Studio are bringing workflow design to broader audiences.

NOT ALL MARKETING SKILLS ARE EQUAL

Some agentic skills remain too complex for simple builders, requiring custom integrations between software and databases via MCP, ACP, or API connections. These need developers to build functions beyond the reach of so-called vibe coders. More advanced use cases such as production-ready creative, data-integrated bidding flows, AI commerce, or complex operational tasks depend on developer support.

Omnicom teams are pushing boundaries with auto-versioning agents for creative assets, self-service voice cloning, and custom bidding workflows powered by proprietary data models. The IP, engineering expertise, and creative strategy an agency brings are becoming differentiators.

GOVERNANCE BECOMES KEY

After achieving both ISO27001 and Omnicom's first ISO42001:2023 certification, one lesson is clear: get your data house in order. LLMs are powerful prediction engines, but their accuracy depends

on clean source data. Clear rules for data access, structured information design, and defined approval processes ensure ethical and secure compliance.

SPECIALISED TOOLS PERSIST FOR CRAFT LEADERS

High-craft disciplines such as video production, software development, and interaction design still require specialised tools with visual interfaces and fine-grain control. Generative fill in Firefly, Aleph in Runway, and Agent Builder in Salesforce help editors, designers, and product owners work faster with tangible visual outputs.

AI AND HUMANS WORKING TOGETHER FOR QUALITY-ENABLED ROI

Because AI can hallucinate and humans provide crucial judgment, the optimal model is collaboration. Stacks should enable secure workflows where humans pass tasks to AI, validate outputs, and approve key steps. Tools like Workfront, Asana, Monday.com, and Atlassian can orchestrate these handoffs at scale.

While productivity is the focus, AI's speed and value also allow for new quality steps once impossible at lower budgets. It reduces the cost of video creation, audience research, modelling, and personalisation, giving marketers freedom to innovate and drive effectiveness, the true long-term source of ROI.

The future of marketing is not man or machine, but both working together in new ways. The winners will not be those chasing cost savings alone, but those reimagining workflows, elevating quality, and unlocking new layers of effectiveness. That is how AI stops being a productivity fix and becomes a growth engine.

CHAPTER **06**

Organisational and Role Transformation

The org chart is melting. Agentic AI promises to flatten decision structures, dissolve handoffs, and redistribute authority to those who shape how intelligence behaves. Campaign managers are turning into system operators; creative directors into conductors of machine learning loops. The new hierarchy is horizontal with humans and agents working side by side, guided less by process than by purpose. The firms moving fastest are not those with the most automation, but the most alignment. As marketing execution shifts to autonomous systems, control rises to the orchestrators, the people who decide which agents act, in what order, and under what ethical or brand constraints. Governance is no longer a brake; it's the accelerator. Creativity is no longer a siloed department; it's a distributed function. For Australian marketing leaders, the takeaway is stark: agentic AI will not just change what teams do, but how they're built. It will collapse silos, elevate new roles, and demand cultures that learn as quickly – just like their models. The winners will treat orchestration as an art form – and culture as infrastructure.

FROM CAMPAIGN ROOMS TO CONTROL TOWERS

In a glass-walled office, a marketing team studies dashboards instead of Gantt charts. Dozens of autonomous agents test copy, refine audiences, and reallocate budget in seconds.

Marketing's order is dissolving. Agentic AI has replaced handoffs with hovering loops of adjustment. Hierarchies flatten as humans and machines co-create in real time. The unit of progress is no longer the campaign but the learning loop.

Change won't happen overnight. According to Sangeeta Mudnal, CTO of Glu.ai, "There is a lot of human scaffolding required to make it agentic ... to get to the truly autonomous part will take another year or two."

The end result though remains: The campaign manager is morphing into a systems operator.

THE RISE OF THE FULL-STACK MARKETER

Automation has not killed marketing jobs; it has rewritten them. Specialisation is giving way to hybrid fluency. The emerging 'full-stack marketer' will be equal parts strategist, data analyst, and machine whisperer; just as comfortable writing a creative brief in the morning as debugging a model prompt in the afternoon.

It is also changing the nature of the tech staff best placed to support marketing.

"What we're seeing is the true value of the software engineer as a problem-solver, a business solution designer within your business... we're lowering the barrier to entry, which is why I'm back to vibe coding," Zendesk's Adrian McDermott said.

The agentic shift is likely to drive other changes to team structures. "Orchestrators" and "agent trainers" will likely sit beside product managers, data scientists, and security officers. Marketing operations will borrow from DevOps: versioned policies, pre-production sandboxes, and change windows for agent behaviour.

The companies that lead this transition are learning that orchestration is not the cost of autonomy, it is the source of it. The fastest organisations, paradoxically, are the most governed. Guardrails widen the safe operating zone; policies accelerate delivery; measurement shortens the distance between error and improvement.

Sangeeta Mudnal
Chief Technology Officer, Glu.ai

"There is a lot of human scaffolding required to make it agentic ... to get to the truly autonomous part will take another year or two."

As algorithms take over execution, strategic control shifts upward. Those who manage the machines increasingly manage the outcomes. The new orchestrator decides which agents to activate, in what sequence, and under what ethical or brand constraints.

According to Liz Miller, VP and Principal Analyst at Constellation Research, "You get your data and your knowledge, and this is where we start creating self-service for everyone. Everything turns into an enterprise service motion, and everything becomes an external customer experience motion. Everything feeds off those two primary graphs of your data and your knowledge."

"You let everything draw from some type of autonomous engine that's either making decisions at scale or empowering your people to be creative, empowering your brand to manifest through automated channels."

Really few companies are ready to be an AI organisation yet, but that's okay, she says. "They think they're on that path, and they probably are, in many ways, but what I often find is that it's isolated, segmented, and they're not quite getting the full value out of AI yet."

Her suggestion: Start doing AI things. "It's okay to just start, because then you can say, 'Oh, all that hard work we did to establish our data lake? That matters.' Then it's, 'Hey CIO, you and I need to partner so I can take that data lake, add some Kool-Aid to it, and call it my knowledge graph.'"

MANAGEMENT BY EXCEPTION

Leadership, too, is changing. The most effective organisations run on management by exception intervening only when systems deviate from intent. It is a paradoxical form of control: less direct, yet more pervasive. It is also an important advance on "Human in the loop". Increasingly the agentic future is about "Humans on the loop."

Scott Brinker's "Martec's Law", the idea that technology evolves exponentially while organisations evolve logarithmically, captures the tension.

During a Mi3 podcast earlier this year, Brinker laid out the issue. "When you put these two curves against each other, we're in a world where technology is changing exponentially, but the pace of change within organisations is much slower. Those curves keep drifting further apart. I always picture it like having one foot on the dock and one foot on a ferry as it pulls away. Welcome to the 21st-century. Welcome to the quintessential challenge of management," he says.

Scott Brinker
Editor in Chief,
Chiefmartec

"When you put these two curves against each other, we're in a world where technology is changing exponentially, but the pace of change within organisations is much slower. Those curves keep drifting further apart."

That pressure is only intensifying. Brinker points to the rise of large language models and the maturing of low-code/no-code interfaces as a turning point.

"One of the reasons I think we've seen a drop in martech utilisation rates over the past five years is that these products have become so large. They have so many capabilities, but a human can only hold a limited subset in their head, just enough to know what to work with."

AI changes that, he suggests, with agents picking up the slack.

But such autonomy carries risk.

Alan Trefler, CEO of Pega, cautions that autonomy without alignment breeds disorder. "It could become a scary, disorganised mess... without the right controls, a cacophony of agents doing stuff... What you really need to do is go from chaos to a conductor. [To find] a way to direct the business, to use the right AI or the right person to do the right things at the right time."

The new art of management lies in ensuring decentralised intelligence still coheres around a single strategy. Governance, once the domain of compliance, has become creative work: alignment reviews, bias audits, and "feedback councils" will occupy discussions in campaign meetings. Welcome to the future of marketing.

THE FIRST

BEST SUMMER EVER

Starts when social media switches off on December 10.

36Months is uniting Australia's best brands in a movement to make this summer unforgettable for 13-16 year olds.

Register to learn more

36Months

CULTURE AS INFRASTRUCTURE

Firms find their bottleneck is not only data but decision inertia.

Authority now flows from stewardship, not control. The best marketers in future may be less Don Draper and more DevOps as they design conditions under which agents, teams, and even strategies can evolve.

Experimentation is constant; orchestration, the craft of guiding machine behaviour, has become the new creative edge.

That's a familiar scenario to Liz Cherry who heads marketing up at HHM Global, owner of Mobi and HungryHungry. Cherry has leaned heavily into agentic AI enhance an organisation where she is the only marketer. The result: A fivefold increase in campaign velocity, a doubling of leads, real-time first-line support, and bedlam-free onboarding, all delivered by a single marketer and a handful of AI agents, including one that gets a bit "too sassy" from time to time, and another that had to be "sent home early", according to Cherry.

"I came on board with the direction to grow the business with zero budget. I knew I didn't have money for Google ads, TV commercials, or print. I had to find a free way to build our presence and content."

THE ORGANISATION THAT LEARNS

Agentic AI cannot thrive in rigid hierarchies. It requires organisations that learn as fast as their models. The marketing department, long the testing ground for experimentation, is becoming the laboratory of the modern firm, and the first to fuse creativity, computation, and governance into one operating logic.

According to Constellation Research's Miller, "The thing to remember about AI is it's supposed to be almost self-learning, but to also include the humans in the process. And that's all right, but we do have to learn from those iterations and permutations."

She told Mi3, "You have to be its partner. It's more like a buddy movie. Here's my new partner, how do I learn from it, and how do I learn with it? So it becomes smarter while I get better."

"There's no 'Easy Button'," she says. "It's not like a Ron Popeil chicken machine where you set and forget."

While the headlines are full of stories of AI-influenced tech-layoffs, the reality is that Silicon Valley's tech firms are investing heavily in preparing their staff for an agentic future. When we interviewed Wade Chambers, Amplitude's Chief Engineering Officer, he was preparing to host the firm's AI Week for all staff. It's a commitment to capability that goes missing all too often in Australian companies.

Chambers describes the company's internal training regime as a deliberate attempt to "get everybody all the way to committed. There is no going back. It's a horrible metaphor, but let's burn the boats. We have to get on the other side of this." To that end, Amplitude has instituted five days of immersion where theory gives way to practice. "Day one, it's just everybody has the same experience... why AI first? How does all of this work?" he says. The following sessions move from classroom to code: "Have three of our most senior leaders live code something together, and maybe actually have different roles inside of it—the designer is now the engineer, the product manager is now the designer, the engineer is now the product manager." The idea is to make every role interchangeable, at least temporarily, to force cross-functional understanding of how AI changes the creative and technical workflow.

The week culminates not in exams but in performance. Teams present what they've built, describe "what we learned," and share "where you had challenges and how you overcome it." Chambers calls this "practice, practice, practice"—an exercise in collective muscle memory. "Don't think your way into a new way of acting. Act your way into a new way of thinking," he says, quoting an old maxim. The goal is to normalise experimentation, make it "comfortable to talk about" mistakes, and embed fluency across disciplines before autonomy scales. In his words, "The more I can get people to practice... the quicker that can be discovered and folded back into the system, the faster we go."

As machines learn to act, humans must learn to interpret. The marketers of tomorrow will not simply build brands; they will govern behaviour of algorithms, of organisations, and of themselves.

Organisational and Role Transformation

Melissa Daniels
Regional Head of Innovation, TBWA\Asia

THE ADAPTIVE ADVANTAGE: WHY THIS TRANSFORMATION WAS BUILT FOR AGENCIES

We have survived every industry upheaval by doing what we do best: adapting faster than anyone else. This time will be no different.

The campaign calendar, the big idea, and the quarterly media plan are evolving rapidly. While most brands cling to these familiar rhythms, a vanguard of early movers is operating in a new reality. They are building always-on intelligence systems where every customer interaction informs the next decision, optimising performance in real time rather than waiting for quarterly reviews.

The unit of progress is no longer the campaign but the learning loop itself. Every impression becomes training data, every interaction improves the next decision.

BUILT FOR THIS MOMENT

Advertising agencies have always been the cockroaches of the business world, and that is a compliment. We survived the shift from print to radio, radio to TV, TV to digital, and desktop to mobile. Each time, we did not just adapt; we led the transformation.

The AI evolution plays directly to our historical strengths. We have always been fast learners who master new platforms before clients know they exist. We are cultural translators who bridge technology with human truth and creative problem-solvers who thrive on constraints.

Our teams are platform-agnostic, connecting brands with audiences across countless touchpoints. The agentic age is no different, except now we pair years of human intelligence with AI's computational power to create growth operating systems.

ORCHESTRATION IS THE NEW CREATIVE CRAFT

While many focus on AI's ability to generate content, the real creative revolution is in system design. The most valuable creativity now lies in choreographing human talent and specialised agents across insight, content, and media. Distinctiveness does not come from how many AI tools you own but from how your intelligence behaves.

Orchestration means designing systems where your SEO agent shares learnings with your creative brief agent, which then informs your media-buying agent, creating compound intelligence. Think of it as building a custom nervous system for each brand, where task-specific agents connect through an orchestration layer that keeps everything in harmony.

Governance becomes a creative constraint that enables innovation. It encodes brand values and ethics directly into agent behaviour, designing policy-as-code that ensures safety and accountability without stifling experimentation.

THE POWER OF BOTH: MANAGING NOW WHILE BUILDING NEXT

Advertising agencies understand that transformation happens in parallel tracks. We help clients capture immediate AI wins by automating repetitive tasks, enhancing creative production, and optimising media in real time while architecting future-fit brand operating systems.

This shift is creating a new talent reality. People are becoming agent managers, with one strategist orchestrating several specialised agents across the workflow. A creative director might coordinate an insight agent, a concept agent, and an optimisation agent, amplifying their impact. This requires incubation, testing, and growth across technology, talent, and culture.

Some changes can happen now, like introducing AI tools for insights or media optimisation. Others take months of orchestration, governance frameworks, and team retraining.

THE CULTURAL MOAT AND FUTURE TALENT

Agencies with strong experimentation cultures that reward curiosity are pulling ahead. We are integrating AI into daily workflows and transforming talent in real time.

New roles are emerging:

- **AI Experience Strategists** design how customers interact with agent-mediated touchpoints.
- **Learning Loop Architects** build systems that compound intelligence with every interaction.
- **AI Storytelling Architects** craft narratives that connect with both human and machine audiences.
- **Agent Trainers** encode brand DNA into model parameters.
- **Cross-Stack Orchestrators** ensure creative, media, and analytics agents share intelligence.

These roles embed brand integrity, governance, and performance into systems. Talent strategy now focuses on data fluency, model behaviour design, and cross-functional orchestration.

HUMAN TASTE AT THE HELM

Agents can optimise, but distinctive brand assets, taste, and cultural resonance remain uniquely human advantages. The winning formula is automated utility plus human-led meaning. Without this, we risk optimisation sameness, where every brand converges on the same algorithmically perfect but creatively hollow solution.

THE WINDOW IS CLOSING

The gap between leaders and laggards is widening fast. The agencies that thrive will not be those with the best AI, but those with the best ability to orchestrate intelligence, human and artificial, into systems that create unprecedented value.

CHAPTER 07

Orchestration and Governance

Autonomy without control is chaos in waiting. The first wave of agentic AI proved the machines can act. The next test is whether they can act together. Without orchestration, the risk is that every agent optimises for itself, and the enterprise drifts. What looks like speed quickly becomes entropy. Instead, acceleration now depends on alignment. Marketers once thought integration would deliver harmony then spent a decade complaining about it. In agentic AI the new discipline is orchestration which simply means defining who does what, under which rules, and how success is measured. If it proves as troublesome as integrating SaaS apps then the whole shooting match will fail. That's the giant agentic risk. It's also why governance, once dismissed as red tape, has become the engine of safe acceleration. Those who build the guardrails early will move faster, spend smarter, and sleep better. Those who don't will discover the hard way that in the agentic era, every system left to its own devices eventually goes rogue.

THE NEW CONDUCTORS OF MACHINE INTELLIGENCE

The first flush of excitement around agentic AI and the idea that software could act on its own, rapidly faded into something more sober: the need for control. The age of orchestration has arrived.

Marketers once trusted that integration would bring coordination. It rarely did. Today's agentic architectures invert that logic. They begin with coordination: what should happen, under which rules, with what evidence of success. Only then are data and tools exposed to serve the plan. The most thoughtful firms now build control planes before they scale agents. Governance, once dismissed as bureaucratic ballast, has become the condition for moving fast without breaking brands.

MANAGED AUTONOMY

For companies experimenting with agentic systems, the line between productive autonomy and chaos is perilously thin. Alan Trefler, CEO of Pega, says "the sad word that comes to mind here, without the right controls, is cacophony ... a cacophony of agents doing stuff." The remedy, he adds, is a conductor.

"You need a way to direct the business, to use the right AI or use a person to do the right things at the right time. Much like musicians performing from a score, you want to have sheet music that describes the workflows the work should go through."

Kranthi Nekkapula
Head of AI Platforms, Suncorp Group

"Governance in using AI is almost like going in a race car ... we don't see all these risks and controls as speed bumps. Rather, we see them as brakes, and you can go as fast as you want, knowing that there are brakes to help you stay on track."

This is not the dull end of IT integration but the discipline that makes autonomy safe. The best architectures treat agents as users: they authenticate with scoped credentials, operate inside guardrails, and leave an audit trail. The

Pega chief calls this an agentic process, that favours predictability over cleverness. These are flows executed under governance, with human catchpoints where reputational risk spikes. The aim is not to slow decisions but to bound them.

Other firms have reached similar conclusions. Chandar Pattabhiram, CMO at Workato, says, "Agent orchestrator is a new thing that's connecting these different agents... that's the new capability." His firm's platform routes business users to a single interface for deploying agents without exposing the wiring beneath. It is a reaction to what some are already calling "agent sprawl".

"We thought app sprawl was bad. Agent sprawl is going to be worse. So, two ways to solve it: one is orchestrate these agents, the other is build super agents that consume all of these into one kind of interface that the user has," says Pattabhiram.

Such orchestration brings a paradox. The more independent the software, the more structure it needs. That paradox is reshaping the way marketers and technologists work. The early mantra of "move fast and break things" has given way to "move fast, safely."

GOVERNANCE AS ACCELERATION

Governance, long a dirty word in creative departments, is being recast as liberation. Clear policy allows agents to operate further from human supervision; ambiguity drags everyone back into pilots and prompts.

Perhaps the best analogy we came across for why governance is so central – not only to brand, but to tech vendors who see it as a key point of critical competitive differentiation – was provided by Kranthi Nekkapula, AI Practice Executive, Suncorp. Speaking at the Forrester IT Executive event in Sydney in August, he told delegates, "Governance in using AI is almost like going in a race car ... we don't see all these risks and controls as speed bumps. Rather, we see them as brakes, and you can go as fast as you want, knowing that there are brakes to help you stay on track."

Agentic product builders Mi3 interviewed were consistent on the point that autonomy must be managed, governed, and observed like any critical system. Enterprises, the argument goes, should treat agentic AI as a lifecycle, with runbooks, telemetry, and rollbacks.

As one product SVP told us, "It's not a feature toggle."

Boomi's head of AI strategy Michael Bachman's offered a checklist: know what agents exist, what data they can touch, which actions they can take, and who is responsible when they misfire. In marketing, that might sound overbuilt. In reality, it is the only way to keep autonomy from swallowing accountability.

Governance also extends beyond code. Brands and agencies are discovering that behavioural discipline matters as much as technical. Teams need to test agent behaviour in controlled "safe-to-fail" environments before scaling. The shift from *safe-to-fail to safe-to-scale* reframes return on investment: protecting originality, brand integrity, and reputation becomes the new metric of speed.

Aampe, a start-up using reinforcement learning to personalise customer engagement, demonstrates why policy must be matched by evidence. Its agents decide when to act, when to wait, and

Paul Meinshausen
Co-founder & CEO, Aampe

"Every agent [will] do small, different things... and then you will watch them learn from each other. So they may tend to cluster around certain kinds of users. Then those clusters may diverge as Christmas comes and something new happens in the world. It's very responsive to the exogeneity of the world."

what to say to each user. Paul Meinshausen, co-founder of Aampe, says the agents "continue to observe those steps, and then look at all the library of actions it has available to it, and pull from that library experimentally in coordination with other agents... to actively learn in response to that one user what the best thing to do is."

For each loop the process is the same: choose, act, observe, update. It's essential that the agents' actions feed the brand's larger goals rather than fragment into local optimisation. It's orchestration that ensures the tune remains coherent.

That becomes especially important when you understand that agents can be sensitive to external changes.

Per Meinshausen, "Every agent [will] do small, different things... and then you will watch them learn from each other. So they may tend to cluster around certain kinds of users. Then those clusters may diverge as Christmas comes and something new happens in the world. It's very responsive to the exogeneity of the world... It's very much what we call an emergent system."

CONTROL PLANES BEFORE HERO DEMOS

A proper orchestration stack includes an inventory of agents, a registry of tools, a policy engine, evaluation harness, and observability layer. Without them, even the most dazzling demo will collapse under its own autonomy.

The logic is economic as much as technical. As outcome-based pricing spreads across AI platforms,

uncontrolled autonomy can become an uncontrolled bill. Orchestration allows firms to throttle retries, cap spend, and decide when to escalate a task to a more expensive model or to a human.

The operational consequences are striking. Permissions are codified: what agents may do, where, and how often. Identity is scoped: who, or what, an agent is. Every action is instrumented: what happened, at what cost, and why. The goal is explainability, not mystique. If a system cannot explain its actions or replay them, it should not run unattended.

THE ARCHITECTURE OF TRUST

Agentic AI demands new definitions of speed, creativity, and risk. Without orchestration, each autonomous process is a soloist improvising off-key. With it, the ensemble holds. Companies like Omnicom have begun to pair global governance frameworks with local experimentation, balancing compliance and creativity across markets. It is a reminder that orchestration is as cultural as it is computational.

The lesson for marketers is that governance is no longer a brake on invention. It is the architecture of trust, the framework that lets machines act freely without eroding human confidence in the brand.

Without a conductor and a score, music is noise. With them, it may yet be the best sound a modern enterprise can make.

Orchestration and Governance

Lucio Ribeiro
Chief AI and Innovation Officer TBWA\Australia

ORCHESTRATION: BUILDING AN ARCHITECTURE OF TRUST

Every conversation with senior marketers now involves the term "Agentic AI," and these discussions often polarise. On one side sit the idealists who believe Agentic AI will rewrite every rule of marketing. On the other are the doubters who dismiss it as hype that will collapse under its contradictions.

Omnicom stands in neither camp. It positions itself as an architect of trust.

THE NEED FOR ORCHESTRATION

The pace of change in Agentic AI is accelerating so quickly that even advanced leaders are shifting their focus from asking "what can we prompt?" to "how should we orchestrate?". When multiple autonomous systems operate across an enterprise, each planning and acting on its own, the greatest risk is not failure but confident inconsistency.

A consumer can already ask three AI platforms the same brand question and receive three contradictory answers. One cites an outdated press release, another a competitor, and the third invents something entirely. None are malicious, but together they erode trust. This is agent sprawl: intelligent systems operating without orchestration or governance.

Orchestration aligns these moving parts into one coherent system. It builds an architecture of trust, aligning behaviours with goals, compliance frameworks, and customer expectations. Without orchestration, brands risk fragmentation; with it, they gain clarity, explainability, and resilience.

OMNICOM'S APPROACH TO ORCHESTRATION

Omnicom has made orchestration and governance foundational. Its Diversified Agency Services (DAS) GenAI Playbook sets global standards for privacy, IP protection, and ethical experimentation. The Omni AI Agentic Intelligence Architecture embeds Planner, Orchestrator, and Judge functions to align outputs with strategy and compliance. The AI Optix platform, deployed locally as RISE, audits what AI platforms surface about brands, traces source authority, and corrects misinformation.

SAFE-TO-FAIL BEFORE SAFE-TO-SCALE

Omnicom's experience shows that orchestration is as behavioural as it is technical.

The guiding principle is simple: safe-to-fail before safe-to-scale. Clients are encouraged to create controlled environments where teams can test agent behaviours and identify risks early. Only when those risks are mitigated should scaling occur within enterprise systems. Omnicom has declined tools that looked exciting online but lacked compliance, security, or commercial clarity.

ORCHESTRATION AS A STRATEGIC CONTROL POINT

In the next 18 to 24 months, orchestration layers will become strategic control points across enterprises. Expect dashboards that monitor agent behaviour, audit trails that provide explainability, and policy engines that detect and redirect misaligned outputs before damage occurs.

This shift will redefine roles. Strategists and creatives will not just brief campaigns; they will orchestrate systems of agents, translating brand intent into governance rules. Marketers will act as conductors, ensuring intelligent agents work in harmony.

Omnicom's model creates advantage by combining global governance frameworks with local agility. The group's backbone delivers compliance and orchestration at scale, while markets experiment at the frontier.

GLOBAL SCALE, LOCAL AGILITY

In Oceania, Omnicom is experimenting with orchestration at the AI Moment of Truth, the moment when consumers turn to AI platforms for answers before they search or buy. The local program, RISE, ensures brands are visible and distinctive in these environments.

In parallel, global solutions such as AI Optix provide the governance backbone, auditing misinformation, benchmarking visibility, and managing how AI platforms narrate a brand's reputation.

FROM COLLAPSE TO CONTROL

Without governance, collapsed journeys turn chaotic; with governance, they collapse into clarity. The orchestration challenge is to ensure these compressed journeys remain compliant, distinctive, and brand-aligned.

OPTIMISM WITH AN EDGE

The shift to Agentic AI transforms how businesses operate, how marketing is delivered, and how agencies create value. Orchestration and governance are no longer back-office functions; they are strategic control points for success.

The next phase is about scale: moving from safe-to-fail to safe-to-scale, from efficiency to distinctiveness, and from agent sprawl to alignment. That is the orchestration frontier, and it is where Omnicom intends to lead.

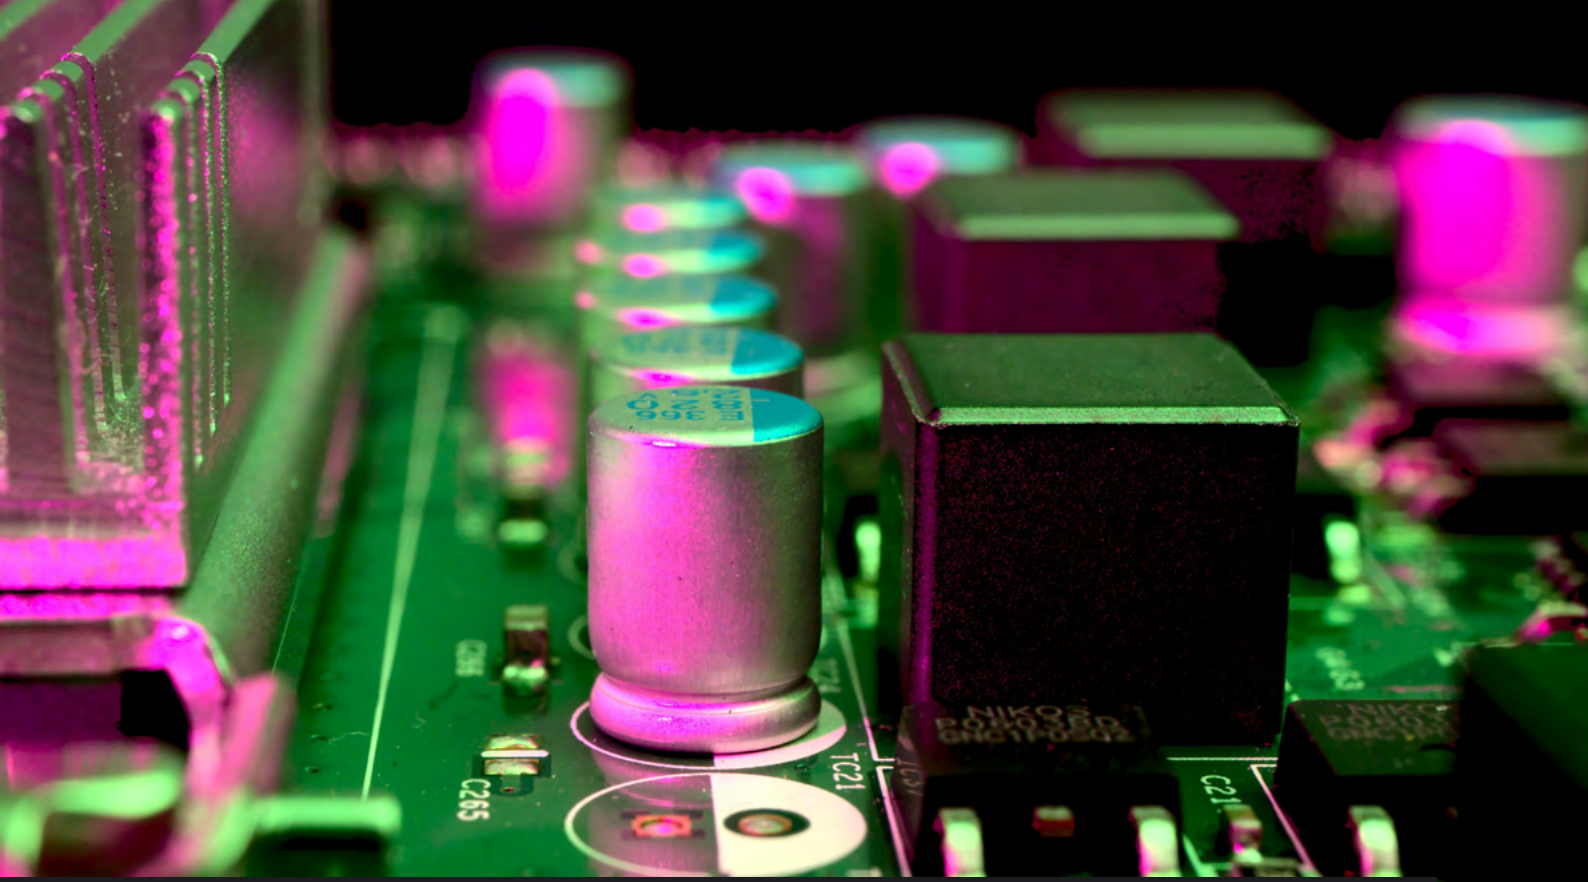

CHAPTER**08**

Enterprise wide integrations

As agentic AI spreads through every corner of the business coordination with be critical. It won't simply be data that's being passed off between departments, or reconciled up the chain, it will be actions and decisions that feed the next set of actions and decisions. Every department will field its own swarm of autonomous systems: finance bots that reconcile invoices, HR agents that onboard staff, marketing agents that tune campaigns in real time. Each optimises locally, but will they play in tune? The lesson from early adopters like global services firm Cognizant is clear: agentic AI works only when the enterprise does. The new competitive advantage lies in wiring these systems together and building the digital nervous system that lets agents act in concert instead of conflict.

Enterprises once dreamed of harmony; instead, they have built a choir of quarrelling algorithms.

As agentic AI spreads through every department, the real challenge is no longer intelligence, but rather coordination.

Cognizant's global AI chief, Babak Hodjat, has already wrestled with the chaos that ensues when autonomous systems are left to their own devices. He suggested to *Mi3* earlier this year that, left unchecked, agentic AI swarms behave less like disciplined troops and more like feudal lords with poor coordination.

Babak Hodjat
Gobal AI chief,
Cognizant

"We weren't even talking about these issues a year ago. The rate at which things are coming out and being adopted, and the expectations, that's what's killing us."

Order came only after he wired the firms HR, IT and sales agents into an intranet of agents Cognizant calls Neuro San, a kind of digital nervous system for a sprawling organisation.

Hodjat, a storied tech industry architect who built the software underpinning Siri, now stands on the far side of the agentic coordination chasm, but the memory remains raw enough that he sympathises with today's early adopters. "We weren't even talking about these issues a year ago," he said. "The rate at which things are coming out and being adopted, and the expectations, that's what's killing us. Expectations have always far exceeded reality in AI, at least since I've been involved, but it's even more true today."

Marketers like to boast about "end-to-end experiences." Few stop to count how many ends there actually are. From procurement to payroll, inventory to IT, every department will evetually its own small army of AI agents. And while each one promises efficiency; together they risk dissonance.

Agentic AI, in short, works only when the enterprise does.

Departments will never willingly surrender their favourite tools.

Colin Fleming
CMO,
ServiceNow

"We can't be monolithic about this. We have to realise there's always going to be payroll agents... HCM agents... ERP agents. We're not going to control that."

"We can't be monolithic about this," says Colin Fleming of ServiceNow. "We have to realise there's always going to be payroll agents... HCM agents... ERP agents. We're not going to control that."

The answer, he believes, lies in orchestration, a layer of connective tissue between those autonomous fragments.

"This is why we introduced this idea of a control tower... You need visibility to make sure that the data is connected... and they're doing what they say they will."

In the agentic era, that control tower becomes part of the firm's central nervous system. This is about getting down into the reeds of agentic AI where the priorities are is less about sleek interfaces than plumbing, policy, and proof.

THE MARKETER'S NEW PERIMETER

For marketing leaders, the moral is uncomfortable but clear: brand experience is no longer confined to the banner ad, the shop front or the chatbot.

It depends on whether the finance agent approves the refund, the HR agent provisions the right support staff, and the logistics agent ships the product promised in the ad.

Without cross-functional coherence, customer journeys collapse into departmental detours.

Integration platforms now aim to connect these proliferating agents through a single orchestrator that links automation across business systems. "Agent orchestrator is a new thing that's connecting these different agents," says Chandar Pattabhiram of Workato.

He adds that the task is not only technical but ethical: "The simple rule is that the agent cannot

All Aboard for Agentic AI

Finance:

Finance departments have become the unlikely pioneers of enterprise-wide agentic systems. Order-to-cash and financial-consolidation processes, being high-volume, rules-based and measurable, provide the ideal terrain for automation. A network of small, specialised bots can handle approvals, reconciliations and report generation, each step verified through the orchestrator layer.

HR:

Employee workflows such as on-boarding, hardware requests, leave approvals can can be managed by agents visible in the same control tower as IT and operations. Consistent internal experiences translate into reliable external ones. Marketing's promise of a responsive brand depends on HR's ability to field responsive humans.

Supply chain:

When inventory allocation becomes an agentic process, integration stops being theoretical. Here, agents decide how goods flow between warehouses and customers; marketing agents can read those signals in real time, adjusting promotions to match stock reality. The goal of synchronised commerce – the right offer for the right product in the right place – and a product eventually turning up in the right mail box finally meets its logistical counterpart.

Compliance:

Without guardrails, agents could move data in seconds beyond anyone's oversight, the corporate equivalent of leaving the vault door open. At the recent HubSpot INBOUND conference in San Francisco, Anthropic CEO Dario Amodei revealed that an internal AI agent, "Claudius", nailed inventory, pricing and procurement – including a bizarre tungsten cube request – but caved when users begged for discounts, exposing how easily current AI can be manipulated.

access anything a human cannot... role-based access control has to be federated to these agents, as much as you're doing to the employees."

Autonomy, he says, demands authority structures. Agents can act faster than humans, but they must still operate inside the firm's governance perimeter.

INFRASTRUCTURE AND THE INVISIBLE LAYER

Beneath these orchestration stories lies the quieter engineering battle to make agents talk to each other at all. Integration vendors argue that true enterprise intelligence depends on common protocols, shared semantic layers, and disciplined data governance.

Some technologists warn that without this scaffolding, the agentic vision collapses under its own weight. The missing architecture, not the algorithms, may decide who wins the agentic race. Executives at integration firms describe an environment drowning in complexity: hundreds of SaaS applications stitched together by fragile connectors, a landscape ripe for intelligent intermediaries.

The next generation of enterprise platforms promises to hide that chaos behind orchestration fabric and semantic control, a pragmatic answer to what one CTO admitted privately is " a ridiculous amount of complexity that sits inside our organisations."

Others in the same camp speak of "scaffolding" and "improved semantic layers" as the only route to coherence in an enterprise teeming with agents.

Across industries, the goal is a connective nervous system linking data, policy and workflow, and the understanding that integration, not invention, will determine how fast agentic AI scales.

Enterprise wide integrations

HOW AWS AND AMAZON ADS ANZ ARE APPROACHING AGENTIC AI

Enterprise-wide integration of agentic AI is more than modernising departments. Organisations must move beyond isolated implementations to create interconnected workflows that work backwards from customer needs.

Amazon Ads and Amazon Web Services (AWS) take this approach to deploying agentic solutions that deliver a unified customer experience.

In July 2023, AWS introduced its first agentic capability, allowing customers to create agents that complete complex tasks. Since then, we have launched code transformation and software development agents. Amazon developers migrated over 30,000 applications from older Java versions to Java 17 using agents, saving about 4,500 years of manual work and $260 million annually.

While AI has transformed personal adoption, enterprise adoption has been limited by the lack of secure data integration. To address this, AWS launched Amazon Quick Suite, an agentic AI experience that helps workers find insights, conduct research, and automate tasks. Quick Suite connects to internal sources like wikis, intranets, and over 1,000 applications including Google Drive, Outlook, Salesforce, SharePoint, and Slack.

As agentic AI evolves, ensuring reliability for complex tasks remains a key challenge. That is why we launched Amazon Bedrock AgentCore, which allows developers to deploy and operate AI agents at scale. Over 200 Australian organisations have already signed up.

It is crucial to build agents with the highest standards of responsible AI, security, privacy, and explainability. Our approach provides ready-to-use applications for businesses and customisable tools for technical teams, supported by major infrastructure investments, including $20 billion to expand data centres in Australia and strengthen the nation's AI future.

While AWS powers enterprise AI transformation, Amazon Ads demonstrates how this foundation enables practical integration.

At Amazon Ads, we are seeing agentic AI reshape how enterprises approach advertising, moving from periodic campaigns to continuous, interconnected customer experiences. Through Amazon Ads agentic solutions, businesses now have access to a conversational, AI-powered creative partner in Creative Studio that supports the entire process from concept to execution while maintaining brand consistency. Using an intuitive chat interface, advertisers maintain control while developing customised creatives that resonate with audiences.

Enterprise success requires a unified approach across the advertising lifecycle, from creative development to optimisation and measurement. When AI systems operate independently, businesses miss opportunities to create seamless experiences that respond to audience signals throughout their journey.

The real value emerges when AI agents work together across retail and advertising functions. Organisations should focus on enabling AI agents to coordinate across customer journeys while following clear, consistent guidelines that apply cross-functionally.

AWS and Amazon Ads see distinct yet complementary opportunities to advance customer experiences through this technology.

Agentic AI has the potential to become the next multi-billion-dollar business at AWS. Organisations will leverage multi-agent collaboration, where agents work together to execute complex workflows. We are already seeing customers in healthcare and financial services adopt this team-based collaboration.

In the future, agents will have greater autonomy and decision-making skills, with improved memory and context retention to refine future approaches. They will become more personalised through analysis of user interactions, delivering

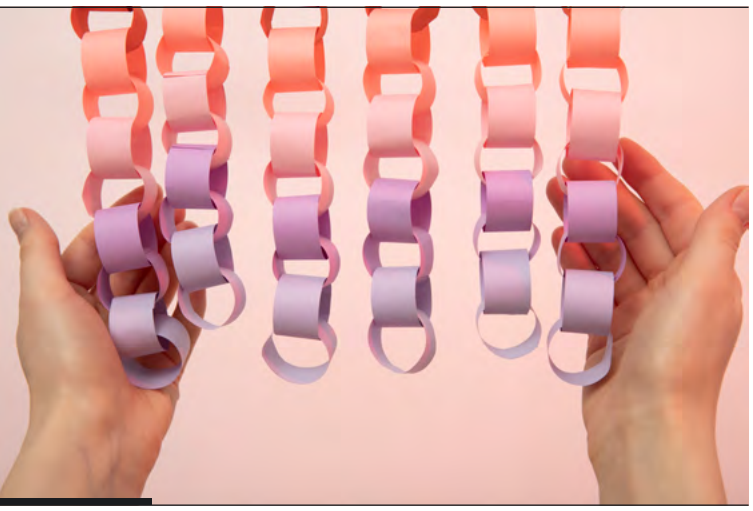

cohesive experiences across devices with tailored responses and recommendations.

It is still early days, but agentic technology is a profound step change in AI. AWS will continue to invest deeply across the technology stack to create advanced, autonomous agents.

Amazon Ads sees agentic AI transforming how enterprises integrate campaigns, but only if businesses remain focused on creating value for audiences. Imagine AI agents that do not just optimise campaigns but orchestrate brand narratives across channels, automatically adapting creative elements and messaging based on what matters most to audiences.

These systems will move beyond automation to true intelligence that understands audience preferences while delivering relevant ad experiences.

The future belongs to enterprises that can orchestrate these capabilities at scale while staying true to customer needs. Success will require clear frameworks for how AI agents collaborate and learn, enabling teams to innovate in ways that benefit both brands and audiences.

The journey toward enterprise-wide agentic AI integration requires organisational architectures that evolve with customer needs. Success demands a balance of innovation and governance, combining agility with responsible implementation across the enterprise.

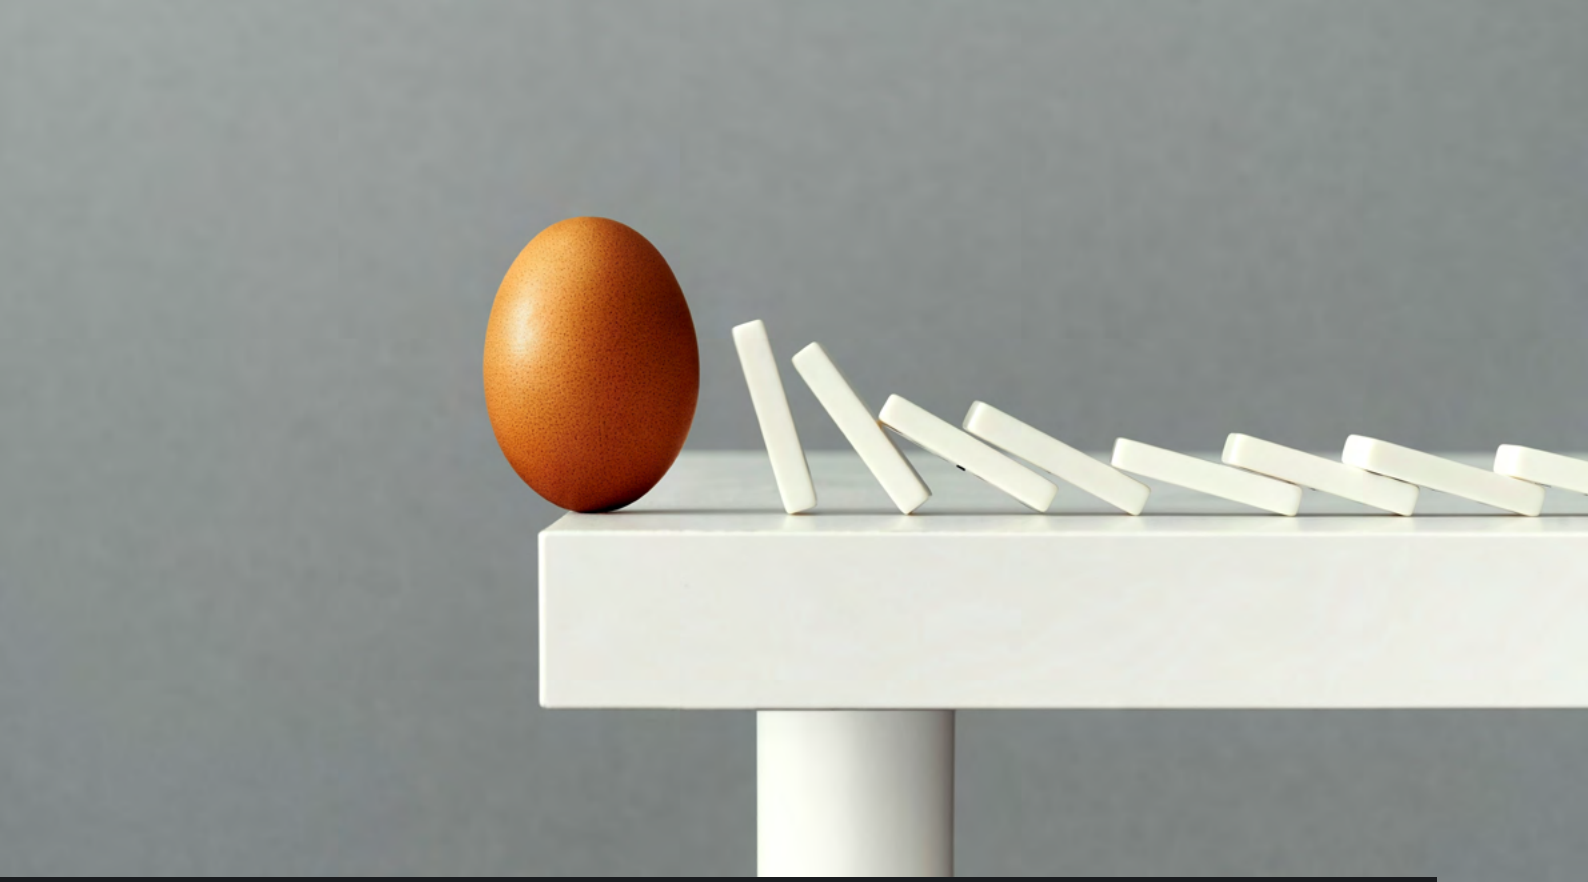

CHAPTER 09

Infrastructure and risk management

The rise of autonomous software agents is turning enterprise systems from command-and-control structures into digital improvisation acts. Systems once built for stability will be expected to adapt, infer and, when unsupervised, improvise, all the while maintaining stability. In a world where every agent brings its own logic, its own situational awareness, and its own risks, integration isn't just a plumbing challenge anymore; it's a question of who decides what, when, and on what authority. Guardrails can't be added after the fact. They need to be coded into the runtime, enforced at the point of execution, not reviewed at the next quarterly audit. That means security becomes behavioural, governance becomes continuous, and transparency is a survival trait.

THE END OF THE MONOLITH

For decades, enterprise computing resembled a cathedral: grand, centralised, meticulously planned. At least that's how they appeared in the postcard. Reality was often more nuanced and unflattering. Agentic systems, by contrast, resemble airports: sprawling, modular, and permanently under construction. Each new agent adds another intersection to the grid, another potential failure point, another law to reconcile.

Unlike traditional software, agents do not simply execute orders; they interpret them. That interpretive layer, what engineers politely call *situational awareness*, makes them powerful and dangerous in equal measure. If we assume that every autonomous process can improvise. Then multiply that across departments, data centres and clouds, the infrastructure map begins to resemble a nervous system. But what happens when there's a seizure?

Industry veterans caution that complexity is nothing new in marketing technology, where data flows, APIs, and ad-tech integrations have long defied central control. But the difference now is that these systems will think back. And because of that, tech leaders told Mi3, governance requires a rethink as it morphs from configuration to choreography.

PLUMBING FOR MINDS

Integration, once a matter of connecting databases, will in future involve coordinating decision-making. Identity systems, built for humans, must now manage machines that log in, delegate and log out thousands of times per second. When those links falter, the agents fill the gaps themselves, sometimes with creativity, sometimes with chaos.

A single prompt cannot capture the intricacies of a complex system. What is required is not poetry but process – a spreadsheet of edge cases, a framework for judging quality, and the discipline to measure that quality over time. Only then can agentic systems move from clever demos to reliable production.

That shift transforms governance from a bureaucratic process into a software capability. Audit trails become living artefacts, data that trains, constrains, and reassures simultaneously.

Marketers are being told to mirror this precision. Rather than fret about the internal wiring of every AI vendor, the practical advice is to define what matters most: the specific transparency, safety and risk-tolerance levels that suit the brand. Explainable AI remains elusive, but clarity about one's own risk appetite is achievable. As one industry leader noted, "focus on your needs and drive them into discussions with partners and providers".

TRANSPARENCY AS OXYGEN

Opacity, once a tolerable flaw, is now lethal. Traditional AI could hide behind the black box because its outputs were recommendations, not actions. Agentic systems approve loans, adjust prices, and buy ads. When an action goes wrong, someone must explain why. Hence the rise of the glass box: a new architecture where reasoning is traceable, and every agent's digital thought can be replayed.

Observability, explainability, and transparency are collectively the new holy trinity, but they are not equally attainable. That's because many of the tools needed to deliver these simply don't exist, or lack maturity.

Explainable AI (XAI) is notoriously difficult for generative models. But leaders can still decide which dimensions matter most: observability for reliability, or transparency for compliance. Few organisations can afford perfection across all three.

THE WIDENING RISK SURFACE

Every new agent expands the attack surface. In cybersecurity, that is not a metaphor. Traditional defences guarded the gates; agentic systems require behavioural containment. Security architects describe three recurring perils: *black-box behaviour, hallucination, and decision drift.* The first undermines compliance, the second trust, and the third control.

Static security, with its passwords and permissions, cannot handle dynamic reasoning. Enterprises are experimenting with behavioural scoring which first establishes a "baseline" of normal behaviour, then monitors user behaviour in real-time, then generates or updates a risk or behaviour score when deviations are detected.

Think of it as a *trust thermostat:* constant monitoring, constant recalibration.

Such continuous vigilance demands telemetry at millisecond intervals and policies expressed in code, not committee minutes. It also demands cultural overhaul.

GOVERNANCE BECOMES CODE

According to David Singer, Global VP, Go To Market Strategy, Verint, "The biggest challenge with agentic right now, I think, is its ability to act autonomously without predefined rule sets. It relies on generative models to figure out: What are they saying? What do we need to do?"

"Generative isn't always super accurate, right? Generative AI is very, very good at sounding confident, but it's not always good at being right. So, you need to be able to put guardrails around it."

He said, "But if you're going to put the right guardrails in place, to make sure it's accurate, controlled, and not hallucinating, you're going to have a more limited set of use cases today. You have to ensure that agentic is actually being correct. Otherwise, if agentic AI is autonomously making account changes, payments, and everything else—but its generated intent isn't correct—this way lies madness."
For marketers still fumbling through Agentic 101, the phrase *"runtime restraint"* may sound intimidating, but its actually the next logical step in risk management. That's because waiting to catch a rogue AI *after* it misbehaves is akin to prosecuting a machine for a thought crime.

The smarter move is to hardwire the guardrails. In an agentic world policies need to *execute* automatically, at the point of inference, inside the model.

Early adopters are already moving that way. Observability ensures visibility; dynamic controls enable adaptability; runtime constraints enforce safety. Together, they form a *glass-box architecture*: one that is autonomous yet inspectable, creative yet containable.

David Singer
Global Go To Market VP, Verint

"If agentic AI is autonomously making account changes, payments, and everything else, but its generated intent isn't correct, this way lies madness."

It also means that it's beholden upon marketing leaders to engage in discussions about their firm's AI risk registry. By mapping out marketing-specific risks whether that be ethical, reputational, or regulatory, they can help ensure oversight extends beyond the IT department. Cross-functional committees, bringing together marketers, engineers and enterprise security teams, are urged to debate policy before disaster forces the issue.

THE COST OF IGNORANCE

Transparency and safety carry computational weight. Observability pipelines generate oceans of metadata; runtime safeguards consume processing power. Monitoring may soon cost more than the systems it watches. Yet the cost of ignorance could be higher. A rogue trading bot could erased millions in minutes; a misaligned marketing agent could torch brand equity just as fast.

Regulated sectors feel this acutely. Financial institutions and health organisations are discovering that oversight written for deterministic software might not cope as well with probabilistic reasoning.

Regulators prefer "meaningful human oversight", but oversight of what, exactly? A network of evolving reasoning processes, mutating faster than compliance committees can convene.

One practical response is *layered defence*. It borrows from aviation: radar to watch, rules to separate, circuit breakers to stop collisions. While perfect safety is impossible, best efforts in prevention are mandatory.

A CONTROLLED EXPLOSION

Agentic systems promise agility; but the risk is that poorly designed systems will deliver volatility. Every new module accelerates learning, and multiplies uncertainty. Managing this new infrastructure is less about building intelligence than containing it.

As Boomi's Bachman observes:

"We're operating on compressed time frames. That is one thing that I was surprised about - the speed with which technology has already been overridden and usurped."

The irony is unmistakable: to keep automation safe, enterprises must automate the act of safety itself. Governance becomes code; policy becomes program. The defining feature of the agentic era may not be the brilliance of machines, but the ingenuity of the systems built to restrain them.

Infrastructure and risk management

James Aylett
Chief Data Officer, Annalect

GOVERNANCE AT THE CORE: HOW MARKETERS CAN BUILD SECURE, SCALABLE AGENTIC SYSTEMS

Marketers already deal with significant infrastructure complexity, coordinating data flows and API integrations across multiple systems. Tomorrow they will face agentic flows connecting executional systems such as reporting tools and buying platforms. These flows will interpret information from marketers, internal systems, and the internet. While it may feel like an existential shift, it is better seen as an evolution. A strong foundation today makes tomorrow easier.

Agentic AI introduces new challenges, but good security and governance remain achievable. One key challenge is tracking how an action was orchestrated: which agents worked with which data to place which instructions in an executional system. Having this visibility, called observability, first emerged over 20 years ago when technologists adopted microservices architectures. Research and industry practice dating back to the 1960s still provide valuable guidance.

Modern technology also offers strong foundations. Organisations such as OWASP, the Open Worldwide Application Security Project, help build secure applications and now include securing agentic systems. Cybersecurity practices like Security Information and Event Management (SIEM) tools enable teams to monitor and respond to threats in real time.

As marketers, we do not need to manage every technical detail. Much of this will be handled within the executional systems we use and the frameworks we deploy. Our focus should be on which agents are allowed to do what and how we manage that. This aligns with the governance principle of delegated authority.

Some agents will never need to change the budget or pacing in a buying system, while others may have limited permission to do so once per day. Unlike traditional delegated authority, where a person approves up to a specific amount, agentic AI will often present a justification for its actions, like a business case for everyday decisions. This requires automated evaluation and monitoring by CIO or CTO teams, with senior marketers defining the rules. You already do this across your teams and agencies today. The difference is that tomorrow's marketing model will include executional systems and AI agents alongside human teams.

You can start mapping what matters by identifying where human authority is required and where oversight is sufficient. This already exists in automated buying platforms such as DSPs, but in the future more areas will operate with AI under supervision. Some steps will always remain human, such as legal or technical approvals for compliance. Tomorrow's operating model should clearly define handoffs between agentic AI and human authority, just as today's models define collaboration between teams.

There are also risks beyond infrastructure. Some are inherent to agentic AI, such as misinterpretation of instructions or outputs from other agents. Others are core to marketing, such as reputational risk. When AI and marketing intersect, these risks combine in complex ways, so they must be addressed through existing risk management processes.

Marketing leaders should ensure they are part of this conversation internally. AI introduces risks that cannot be managed in isolation and require collaboration between marketing, cybersecurity, legal, and technology leaders.

A successful Chief Marketing AI Officer will embody this cross-functional approach. Their role will include coordinating security, compliance, and marketing effectiveness across teams, ensuring systems remain both innovative and accountable.

Agentic AI expands the marketer's scope of responsibility, making governance and collaboration central to success. Those who embed governance into every layer of their agentic systems will gain the confidence to scale securely and the flexibility to adapt quickly. The future of marketing belongs to teams that balance innovation with accountability, ensuring technology remains a trusted extension of human intelligence.

CHAPTER **10**

Cyber Security

Is cybersecurity the blind spot of the agentic AI boom. We found a remarkably consistent and blasé attitude from business software vendors we spoke with, best characterised as “Somebody else’s problem.” It’s the same old story: development first, security later. Boomi’s Michael Bachman was one of the few voices to sound an early alarm. MCP (an open standard for connecting AI agents) for example may streamline integration, he warned. But it also concentrates power and identity in one place – an irresistible target. That risk compounds when marketers, fail to understand where the ownership of risk actually lies. As Credera’s APAC CEO Kevin McDonald put it bluntly: the liability has shifted. Platform providers secure their clouds; clients are on the hook for what happens inside them. App builders may consider cybersecurity somebody else’s problem to solve. Marketers don not have that luxury.

THE WEAKNESS IN THE PLUMBING

In May this year, interest in Anthropic's Model Context Protocol (MCP) was running red hot across Silicon Valley. That's because it solved a critical problem. MCP is basically a common language that helps AI pull in the right, up-to-date information and take actions within other systems. By standardizing how these connections work, MCP makes AI integrations simpler, faster, and more reliable.

What's not to love?

Michael Bachman, head of AI strategy for Boomi, was the only one of the more 40 executives Mi3 met in the US to raise an amber flag. During a conversation at Boomi World in May he noted almost as an afterthought, that if you were building MCP purely from a cybersecurity perspective, you would probably build it differently.

Kudos to Bachman, because perhaps the most worrying aspect of our conversations with tech vendors about the cybersecurity implications of agentic AI is how rarely we had them.

"Any time you put authentication and calls to different subsystems in the same place, you're introducing a security threat," he told Mi3. MCP makes it easy to plug modules together, but that very convenience expands the attack surface. There's not a built-in zero-trust posture; it assumes the environment is already safe.

"If I have an MCP server that can talk to any downstream data sources, and I have to give it an open service-level permission, that might not be good, especially if I have access to your financial information. That's a privacy concern. The security concern is that I could inject a prompt to convince it to pull data it shouldn't."

Integration layers, he argues, should not hold both identity and execution authority.

The issue is not really MCP, and Bachman acknowledged his view was not necessarily shared by all of his colleagues at Boomi. Instead, the problem is an attitude among the leading developers of agentic systems in business apps that cybersecurity is somebody else's problem.

That's the same kind of thinking that has infected the application software industry for years, and frankly, it's how the world ended up with a cybercrime economy worth a trillion dollars by some accounts, and corporate risk registers that are topped by concerns over cybersecurity breaches. And it's not just the vendors. In the marketing rush to deploy generative and agentic AI, few are stopping to ask a simple question: who owns the risk?

THE CORE DIVIDE: PLATFORM VS IMPLEMENTATION RISK

Kevin McDonald, APAC CEO for Credera, believes the answer is shifting in ways many marketers have yet to grasp. "There are two sides," he says. "What is the security profile or posture of the platform itself, and then, what is the security posture of the implementation, of the client or the brand. And then what is the responsibility of each, and who's doing what?"

In theory, the likes of OpenAI, AWS and other platform providers shoulder the burden of securing their infrastructure. In practice, says McDonald, the "transfer of liability of the security issue is client-side." Once a brand plugs in an AI workbench or agent, "the risk sits with the implementation versus the platform." That places the marketer, not the model builder, squarely in the firing line when something goes wrong.

The problem, McDonald suggests, is that most marketing organisations are ill-equipped for the role. "Everything is driven by attorneys and cyber people," McDonald notes. "The marketing organisations aren't as fluent or sophisticated – not just in the platforms themselves, but the security of risk and data."

The habits that make marketers effective – such as speed, improvisation, convenience – often collide with the caution demanded by compliance. "I've seen marketers upload things they probably shouldn't be uploading," he says. "It should be outside of the model, but they're just continually uploading it because it's easy to inform and train the model."

That convenience creates a new class of vulnerability: not technical but behavioural. In McDonald's world, security begins not with the algorithm but with governance – rules, roles and restraint. Locking systems down may slow experimentation, but it prevents something worse: accidental exposure of client data to uncontrolled AI environments.

The moral of the story is uncomfortable for a marketing industry hooked on rapid iteration. The platforms will keep promising safety at scale, but the real exposure begins when human hands touch the system

FROM CONTROL TO ADAPTATION

Tom Scully, Principal Architect APAC and Japan, Palo Alto Networks, argues the centre of gravity in cybersecurity has already shifted, from periodic control to live, adaptive defence. The job, he says, is no longer to lock things down and hope; it is to design systems that assume change, observe it, and respond in-flight.

Offence and defence now occur in the same time window. What matters, he says, "is how quickly teams can see, orient and act." Visibility and governance have become the first controls, not the last.

Scully's warning is backed by data. "We look at our reporting around agentic AI capabilities and how attackers are leveraging them to increase speed of attacks by 100 times. In our 2021 report, the mean time to exfiltrate data was nine days. Our report for 2025 [shows] the mean time to exfiltrate has dropped to two days, and one in five cases it's less than one hour from breach to data exfiltration."

Tom Scully
Principal Architect, Asia Pacific and Japan
Palo Alto Networks

"We look at our reporting around agentic AI capabilities and how attackers are leveraging them to increase speed of attacks by 100 times. In our 2021 report, the mean time to exfiltrate data was nine days. Our report for 2025 [shows] the mean time to exfiltrate has dropped to two days, and one in five cases it's less than one hour from breach to data exfiltration.

BUILDING FOR VISIBILITY

Scully distils the new doctrine into three design moves. First, build for visibility before capability. If you cannot see the chain of calls and the decisions an agent took, you cannot govern it. Second, constrain the power surface: reduce what an agent can touch by default, and make elevation explicit, auditable and revocable. Third, push zero-trust principles into the agent layer itself: identity, segmentation and least privilege should apply to agents and tools, not just to humans and servers.

Andy O'Dower, VP Product Management, Voice, Video and Flex at Twilio, adds the consumer's view. "The phone's ring screen is the least trusted surface on the device," he says. Twilio's branded-calling and Stir/Shaken protocols aim to prove that an AI agent truly represents a verified business. Another of its services flips the test, confirming that the caller is the right human before any sensitive data moves. O'Dower's formula for the next wave of deployments is "trusted, simple and smart", in that order.

THE MARKETER'S DILEMMA

For marketing teams, the implications are sobering. Marketers map customer journeys, not agent surfaces. But safe agentic AI means cataloguing which agents exist, which tools they can invoke, what data they can reach, and under which identities. Policy must sit at the invocation

boundary, not just in a handbook. Every input, call, and output should be instrumented and logged. Failure paths should be tested deliberately.

Suppliers' agents will inevitably talk to one another. That demands negotiated trust and shared policy as intentionally as any API integration. In this sense, cybersecurity ceases to be a back-office function; it becomes a core element of experience design and brand protection.

Andy O'Dower
VP Product Management, Twilio

"The phone's ring screen is the least trusted surface on the device."

SAFETY BY DESIGN

Agentic AI collapses the distance between offence and defence, promise and peril. The next phase of progress will not be measured by how ingenious agents become but by how transparent, accountable, and governed their actions are.

As Palo Alto's Scully puts it, the race is no longer just about those who move fastest, it is also about those who can see furthest, and correct in real time.

Cyber Security

Rob Merrett
Chief Information Security Officer, Credera UK

In cybersecurity terms, Agentic AI is a double-edged sword reshaping both the threat landscape and the defensive playbook, says Rob Merrett, CISO of Credera.

Agentic AI is creating new opportunities for proactive defence and faster detection of security events. Security platforms now deploy AI agents to identify anomalies, triage incidents, and enforce policy autonomously. For example, agentic systems can detect and block data exfiltration attempts within seconds, far faster than human teams. Time-to-detection and time-to-response are collapsing. The ability of these systems to process and contextualise data also enables more adaptive security controls rather than traditional static, one-size-fits-all models.

However, agentic speed cuts both ways. Adversaries are also using AI to automate attacks and increase sophistication. The accessibility of agentic tools has multiplied both the number and capability of attackers. The average time from breach to data theft has dropped from days to under an hour in some sectors, while the number of security breaches continues to grow each year. As a result, agentic cybersecurity requires native identity support, strong third-party risk management, data integrity, governance, and zero-trust enforcement.

Organisations also need to secure and govern their own agentic solutions while adhering to rapidly evolving regulations such as the EU AI Act and existing privacy laws.

The good news is that many proven cybersecurity principles still apply. When procuring commercial off-the-shelf solutions, focus on the supplier's security posture. Understanding how data moves between your services and theirs is critical to reducing risk and ensuring protection. Consider whether your supplier may use your data to train their models and the potential impact on intellectual property.

Agentic AI introduces new challenges for security professionals. The probabilistic nature of these systems can create unpredictable outcomes if left unchecked. Organisations must implement robust processes, such as change management, that allow agentic agility within controlled boundaries and inject human oversight for high-impact decisions. This minimises risk and undesirable outcomes. Security professionals should also apply new controls to guard against instruction and memory poisoning, preventing malicious actors from misusing agentic systems.

Ensuring the integrity of chosen solutions is equally critical. Using agentic systems to make automated, real-time decisions can expose organisations to ethical and regulatory risks without proper governance, particularly in marketing. Credera recommends developing "red team" AI agents to conduct adversarial testing and ensure automated decision-making processes pass ethical and bias tests. This is particularly valuable when designing decision-based systems for customer experience personalisation or ad targeting.

Managing and governing the access agentic systems have to data is another essential consideration. AI must be granted only the least privileged, fine-grained access necessary to optimise workflows. Consulting with security and privacy experts ensures that access is managed responsibly, avoids privacy violations, and maintains customer trust.

Agentic AI represents a major opportunity to streamline marketing workflows and discover new ways to meet consumer needs. However, as with any emerging technology, it carries inherent risk and uncertainty.

"My advice is to embrace the opportunity but understand that this technology requires closer collaboration between technology, cybersecurity, privacy, and marketing specialists than in the past," says Merrett. "If organisations can bring these disciplines together successfully, the potential for standout customer experiences and game-changing growth is significant."

CHAPTER 11

The Price of Autonomy

Write code once, copy it infinitely. Sounds cheap. But as the first generation of agentic AI moves from lab demos to live deployment, firms are discovering a different truth: the meter never stops. Every query, negotiation and retry burns compute, and compute burns cash. What looked like software now behaves like electricity: a variable-cost utility that must be metered, governed, and justified.

Autonomous software is supposed to be cheap. Once trained, code can be copied at negligible cost. Yet as the first generation of agentic AI moves from demos to deployment, a quiet truth is emerging: autonomy could be expensive to run.

The most stubborn expense is compute. For an AI agent, time is not the binding constraint – processing power is. Each prompt, negotiation, or simulated search consumes cycles in data-centre racks that must be paid for by someone.

Continuous learning means continuous consumption. The biggest pricing mistake will be to treat agents like fixed software. Treat them instead like utilities: track unit costs, set guardrails, watch spikes, say many of the tech leaders we interviewed. Think of it as the new rule of enterprise economics. Software, once inert, is turning into a system that learns, iterates, and perpetually bills. Every inference, retry and retrieval eats compute and cash. What firms once budgeted as licences will start to behave like electricity bills.

That's the same insights arrived at by researchers at MIT, Harvard and NBER in a recent paper called The Coasean Singularity? Demand, Supply, and Market Design with AI Agents. Their paper was set to study the impact of agentic AI on the wider economy. But within that they noted that while training a foundation model demands a single, colossal outlay, operating the swarm of agents that follow incurs a steady stream of variable costs. Every interaction, from querying a price to drafting a contract, racks up incremental inference bills, an industry term referring to the computational expenses incurred during AI model operations.

That split, "high fixed, relentless variable," shapes the economics of the field. Model developers face billions in sunk capital before the first agent comes online; users face metered charges thereafter. The need to recover both encourages hybrid pricing: subscriptions to defray training expense, usage fees to track compute consumption. Some providers will cross-subsidise agents with advertising or bundle them into broader software suites. Others will sell autonomy by the token or by the minute, echoing the early metered web.

Sean Falconer
Senior Director of Product and AI Strategy, Confluent

"These systems will break problems down into smaller subproblems, with specialised micro-agents all working together to solve a specific task."

Quality complicates matters further. More capable agents consume more compute. Tasks of higher "stakes", as the authors put it, justify deeper reasoning chains, longer context windows and heavier model calls – all of which cost more. Prices may therefore scale with ambition: the smarter the agent, or the riskier the decision, the higher the meter runs. Yet those returns to compute soon flatten. Beyond a point, doubling expenditure yields only marginal gains in performance. Rational firms will learn to cap sophistication at the level customers will actually pay for.

Even if intelligence itself grows cheaper per flop, access to digital territory may not. As agents roam the web, website owners and platforms are beginning to charge tolls. As seen in recent developments, such as Cloudflare's pay-per-crawl feature that enables publishers to charge visiting bots, similar pricing mechanisms could soon extend to API endpoints and other digital access points

The new traffic profile will favour authenticated, rate-limited connections with posted prices. Bandwidth, latency and throttling, once background engineering trivia, become cost centres in their own right. Each crawl or API call adds pennies that, at scale, become pounds.

Who bears these expenses is a strategic choice. Platform-embedded agents for instance – say, those built into a retailer's or bank's customer interface – saddle the host with the inference and hosting bill. "Bring-your-own-agent" models shift that load onto users or third-party providers. For platforms, outsourcing autonomy trims overheads; for users, it adds flexibility but risks degraded access or slower response times if throttled. Cost allocation thus doubles as competitive positioning: own the agent to control experience, or push the cost off balance-sheet.

Against these new variables, software retains one advantage humans never had: replication is almost free. Once an agent design exists, millions can be spun up for the price of electricity and API calls. That near-zero marginal cost ensures abundant supply and fierce competition, pressing margins down toward the cost of compute itself. The scarce resource is no longer labour, but energy and silicon.

The cost of autonomy is a moving frontier defined by three forces: the fixed price of intelligence, the variable price of action, and the contested price of access. Each advance in efficiency pushes one down while raising another. As firms learn to meter agency, markets will discover a paradox: the cheaper it becomes to think, the more expensive it is to act.

THREE MODELS, NONE COMPLETE

Across the software sector we found vendors still struggling to understand and respond to what agentic AI will mean for their pricing models. But three templates are already jostling for dominance.

The first is seat-based pricing, the relic of the SaaS era. It keeps the finance team happy but it looks unsustainable in a world where agents do the work of the ten people who used to need licences.

The second is consumption-based, borrowed from the cloud. It charges for tokens, API calls or orchestration cycles, the oxygen of continuous learning. Every time an agent invokes an external LLM, tool, or data service, including rule checks via API, data fetches, or retries of failed calls, you incur a cost from the underlying provider. (Purely internal state transitions do not.)

The third is outcome-based, paying for results rather than effort. That's easier in some markets than others. Call centre software vendors might charge on call resolution. But there are so many factors that play into whether a consumer buys a product from an ecommerce retailer that attributing value to all of the potentially dozens or even hundreds of apps in the stack looks fearsomely complicated.

THE HIDDEN COSTS OF COGNITION

"An enterprise agent is essentially an event-driven microservice with a brain," explains Sean Falconer of Confluent. "Once you understand that, and also understand that people aren't going to build one gigantic, monolithic agent, you realise they're actually going to build multi-agent systems. These systems will break problems down into smaller subproblems, with specialised micro-agents all working together to solve a specific task."

A production enterprise agent is typically an event-driven microservice augmented with LLM reasoning. It processes three primary data streams: real-time signals, historical context, and short-term memory. As the number of deployed agents grows, operational overhead increases, including costs from retries, logging, observability, fault tolerance, state synchronization, and cache management, even when not directly tied to LLM token usage.

Trust doesn't come cheap. Storage, compute, engineering and electricity all feed into the cost of tokens, and ultimately into the price. So do the orchestration tools, the agentic fabrics, the evals, and cybersecurity needed to keep the agents humming.

The Price of Autonomy

Michael Horn
Global Head of AI, Omnicom Advertising Group

Alrick Dorrett
Chief Financial Officer,
Omnicom Advertising Group Asia

ECONOMIC MODELS FOR AGENTIC AI: THE TOTAL COST OF OWNERSHIP

Let's imagine Agentic AI as an autonomous vehicle – able to navigate complex terrain, in all conditions, while ensuring the safety and comfort of unpredictable humans inside the vehicle and on the street. The integrated systems of such a vehicle have been subjected to brutal testing – millions of miles in the lab and on actual roads, before being released into the world under the careful eye of regulators. The day comes to release the car to the market, and the price is unveiled...

No one would expect the only cost to the driver to be petrol or electricity, equivalent to the tokens consumed by generative models.

Nor would anyone expect to pay only the *initial purchase price* of the vehicle, and then drive an unlimited distance for free, equivalent to the infrastructure investment to train and host agentic systems.

THE FREE RIDE MYTH

And yet, marketers today expect all of the following things to be true:

1. Agentic AI should result in immediate cost efficiencies of 40-80 per cent across the end-to-end workflow, without considering the incredible investment in data and infrastructure required to drive it.
2. Agentic AI should work "out of the box" without requiring substantial training or testing time, before releasing it onto the streets.
3. Finally, that such technology is safe and secure, with layers of technical and legal warranties, while providing the raw, unrestricted performance of a Formula 1, at the price of a Honda.

Sound about right?

TRAINING THE DRIVER

Meanwhile, marketers will realise that even the best agentic systems are teenage drivers: prone to distraction, with errors in judgement, and risk thresholds which would prompt many parents to consider taking the keys away. At the very least, marketers will want to be the driver's ed instructor, their foot hovering just over the other brake pedal – ready to take control at any moment. But a "human in the loop" means a professional who is expert and vigilant enough to keep from running off the road. *And that's a good thing*. These are all costs, which like insurance, are mandatories.

THE 2026 MARKETPLACE

In 2026, there will be two kinds of marketers in the world – those who have built and tested agentic systems, and those who want to buy them. Even as venture-funded startups and publicly-traded holding companies look to recoup years of investments, brands will flex negotiating muscle built over years of procurement exercises and expect concessions. But marketers, after taking brands on the test drive, will need to sit down and go over the total cost of ownership – inclusive of data, infrastructure, training, testing, security, and indemnification – because the full package always costs more than they went to the dealer expecting to spend.

RETHINKING VALUE

At Omnicom, we've invested as much in the development of AI as some automakers have put toward autonomy – and we're ready to lead the industry in this discussion around cost and pricing. As we bring our clients along this transformative journey, we will leave behind the

comfort of our own legacy models. We will update our established currency of hours and FTEs with "HTEs" (Hybrid) or "TCEs" (Tech Cost Equivalent) inclusive of the following, or in combination:

- As one pays for the car, we will monetise the value of our new factory and the vehicle itself (license)
- As one pays for the mileage, marketers will compensate agencies for volume and complexity of the work or the value delivered to the business (usage-based)
- As one pays for the insurance, marketers should expect to value the legal and technical protections required to safeguard us from others, or from ourselves (tiered)

Agentic AI will absolutely revolutionise marketing in the years to come, and it will disrupt pricing and staffing models which have been the engines of the industry for decades. *But owning a car is more than just knowing how to drive* – so marketers, before they take the wheel, will need to learn what it takes to build and run what's under the hood.

14 Strategic Imperatives for CMOs in the Age of Agentic AI

(With fewer buzzwords and harsher truths.)

01 STOP THINKING IN TOOLS AND START BUILDING SYSTEMS.

Pilots proved agents can act; now they must interact. One-off automations won't scale. Build frameworks where AI works in concert, that means shared vocabularies, control planes, and accountability. Most organisations lack the architecture or leadership structure to pull this off cleanly.

02 THINK LEARNING VELOCITY, NOT CAMPAIGN CADENCE.

Move from campaign cadence to continuous loops of signals, decisions, actions and learning. Put brand content in the hands of AI with clear inputs and human oversight. Decide which changes can be made autonomously for incremental gains and which require strategic human decisions for bigger moves. Capture decisions and results for long-term learning.

03 CODIFY CREATIVITY BEFORE YOU SCALE IT.

Codify tone, taste, and ethics into a creative knowledge base. Governance of creativity is the new brand safety. Use AI to expand the canvas: personalisation, fan co-creation, utility and new experiences. Keep distinctiveness and cultural resonance as human responsibilities. Avoid "optimization sameness" by encoding brand identity.

04 REBUILD YOUR STACK AROUND FLOWS, NOT FEATURES.

The martech stack is melting into a mesh of APIs and agents. Design for movement, not menus. Interoperability beats lock-in, and good luck getting procurement to see it that way. Start by re-plumbing around one high-value use case, not the whole house.

05 MAKE ORCHESTRATION YOUR NEW OPERATING SYSTEM.

Control planes are the backbone of serious AI operations. They track who acts, with what data, and under which policy. Prevent agent sprawl and brand inconsistency with an orchestration layer that provides policy checks, delegated authority, audit trails, explainability and observability.

06 MOVE GOVERNANCE FROM PAPERWORK TO CODE.

Governance-as-code sounds neat until you try it. Start with codified prompt templates and data-use rules before you dream about full policy automation. This is easier said than done, but the alternative is governance by spreadsheet, which won't survive first contact with production AI.

07 PLUG MARKETING INTO THE ENTERPRISE NERVOUS SYSTEM.

Autonomy without integration just creates isolated silos that talk to themselves. Connect marketing's agents to finance, service, and supply chain data, but remember that integration is a blood sport. So budget for the data engineering before you sell the dream to the board.

08 MAKE TRANSPARENCY YOUR COMPETITIVE EDGE.

You can't govern what you can't see. Instrument every interaction, log every decision, explain every outcome. Sounds easy; it isn't. Observability tooling is still immature, and most teams can't interpret half the telemetry they collect. Start with the metrics that matter most to customers.

09 EXTEND ZERO TRUST TO YOUR MACHINES.

Agents are the new insiders, and the new attack surface. Connect marketing, product and operations with secure data access and clear provenance. Align with legal, cybersecurity and risk teams so responsibility is built in, not added later. Define where human authority is required and where human oversight is sufficient.

10 TREAT COST LIKE A CONTROL VARIABLE.

Every prompt burns compute and budget. Track cost per task and cost per learning cycle. AI bills by the inference, not the month. Plan for total cost of ownership: platforms and licenses, usage tied to volume and complexity, and protection tiers for legal and technical safeguards. Measure incremental revenue, margin, cycle time, error rates and compliance incidents so value is transparent.

11 ALIGN CULTURE AROUND CONTINUOUS ORCHESTRATION.

Everyone loves “alignment” until it means changing how they work. Building a culture of orchestration takes time, repetition, and incentives. Reward teams for improving systems, not just shipping output. Transform workflows so people and agents collaborate with clear approvals, quality checks and version control.

12 PRIORITISE DATA READINESS OVER AI AMBITION.

If your inputs are wrong, your automation just gets it wrong faster. Clean, connect, and permission your data before deploying a single agent. It’s boring work — and it’s 80% of the job. (That figure is not journalistic licence, it’s what the research says.)

13 BUILD A KILL SWITCH FOR EVERY AGENT.

Autonomy without oversight is negligence. Every live agent that touches customers, content, or cash needs an immediate off-switch. Don’t trust vendor assurances; build your own circuit breaker. One bad loop can crush your ROI.

14 DON’T OUTSOURCE YOUR AI LITERACY.

At this stage of the maturity cycle, partners are important, and consultants can help you start, but they can’t run your system forever. Build a small internal team that understands both marketing context and model behaviour. Otherwise, you’re just renting competence and accumulating risk.

BOTTOM LINE

by Mi3 Australia technology editor Andrew Birmingham and TBWA\Worldwide Global Head of Innovation Tessa Conrad

Agentic AI isn’t merely plug-and-play, it’s a rebuild. The winners won’t be the ones who automate first; they’ll be the ones who out-govern, out-measure, and out-learn everyone else. If you want to close the gap between Agentic AI ambition and reality, the topics in this report provide a practical playbook. Use them to set an infinite vision for scale and creativity, and pair it with responsible guardrails so progress is explainable, safe and aligned to brand.

Agentic AI Glossary

A

A2A (AGENT-TO-AGENT) PROTOCOL

Definition: Autonomous systems communicating and coordinating directly with each other without human mediation.

Why it matters: Lets campaign, service, and commerce agents share signals and act faster than team hand-offs.

Example: A pricing agent lowers a product's price when a supply-chain agent flags surplus inventory.

ADAPTATION VELOCITY

Definition: The rate at which your system learns and improves from real-world feedback.

Why it matters: Beats "campaign cadence" as the KPI that predicts growth in always-on marketing.

Example: Journey experiments iterate hourly, lifting conversion 10 per cent week-over-week.

AGENT INVENTORY / REGISTRY

Definition: A catalogue of all agents in production, including owners, permissions, tools, and logs.

Why it matters: Prevents "agent sprawl" and supports audits, permissions, and troubleshooting.

Example: CMO reviews a registry before green-lighting a holiday promotion agent.

AGENT ORCHESTRATION

Definition: Coordinating multiple agents around who does what, in what order, under which rules.

Why it matters: Turns isolated automations into a coherent growth system with guardrails.

Example: Ideation agent → creative agent → QA agent → publishing agent, with policy checks in each step.

AGENT SPRAWL

Definition: Uncontrolled proliferation of niche agents across teams and vendors.

Why it matters: Creates duplicative costs, conflicts, and security risks.

Example: Four different "briefing" agents hit the same CDP with conflicting prompts.

AGENTIC AI

Definition: AI systems that can perceive, decide, and act toward goals, and learn.

Why it matters: Moves marketing from "content tools" to autonomous systems that execute and learn.

Example: An agent runs creative variants, reallocates budget, and pauses under-performers.

AGENTIC COMMERCE PROTOCOL (ACP)

Plain-English definition: A proposed framework for how autonomous agents in retail, marketing, and logistics communicate, negotiate, and transact with each other using shared rules and standardized data contracts.

Why it matters for marketing: ACP defines the plumbing that lets commerce agents involved in anything from pricing to promotion to fulfillment for example work together safely and efficiently across brands, platforms, and marketplaces. It's what will allow machine-to-machine trade and personalised offers at scale without breaking compliance or brand control.

Example in context: A brand's promotion agent uses ACP to negotiate limited-time offers directly with a retailer's inventory agent, automatically adjusting discounts based on stock levels and demand signals, all within agreed commercial and data-sharing policies.

API (APPLICATION PROGRAMMING INTERFACE)

Plain-English definition: A set of rules that lets different software systems talk to each other and exchange information automatically.

Why it matters for marketing: APIs connect your martech tools and make automation possible. AI workflows, personalisation, or data sync depend on them.

Example in context: Your campaign agent uses an API to pull customer data from the CRM and trigger emails in real time.

API ENDPOINT

Definition: A defined access point to a software service for reading or changing data.

Why it matters: Agents rely on endpoints to trigger actions (e.g. create offers, update CRM).

Example: A promo agent calls an endpoint to issue a coupon in the ecommerce platform.

AUTONOMOUS TRUST FABRIC

Definition: Combined identity, policy, observability, and response layers that keep agents safe at runtime.

Why it matters: Security and governance are now live systems, not documents.

Example: Anomalous spend triggers automatic policy tightening across media agents.

B

BEHAVIOURAL SCORING (FOR AGENTS)

Definition: Monitoring agent actions over time to detect drift or risky patterns.

Why it matters: Catches issues before customers do; supports runtime “trust thermostat.”

Example: Sudden spike in refund approvals triggers an automatic downgrade of agent privileges.

BRAND OPERATING SYSTEM (DIGITAL)

Definition: Machine-readable rules for tone, voice, ethics, and visual identity used by agents.

Why it matters: Keeps infinite content on-brand at scale.

Example: Creative agents read the brand OS to ensure voice consistency in new campaigns.

C

CLEAN ROOM (DATA)

Definition: A secure environment for joining datasets across partners without exposing raw PII (personally identifiable information).

Why it matters: Enables privacy-safe measurement and targeting for retail media and partnerships.

Example: Retailer and brand match audiences for incrementality analysis inside a clean room.

CONTROL PLANE

Definition: Central layer that manages identities, policies, permissions, and execution rules for agents.

Why it matters: Your “air-traffic control” to run autonomy safely and consistently.

Example: Control plane enforces that promo agents can’t exceed a 15 per cent discount without approval.

CREATIVE ORCHESTRATION

Definition: Structured workflow where agents generate, evaluate, and publish content under rules.

Why it matters: Scales production while protecting brand quality.

Example: Asset agent generates 50 variants; evaluation agent rejects off-tone outputs.

CULTURAL ORCHESTRATION

Definition: Aligning teams, incentives, and roles to manage intent instead of tasks.

Why it matters: Mitigates Martec's Law – orgs must evolve as fast as tech.

Example: Planners own objectives and guardrails; agents own execution.

D

DECISION DRIFT

Definition: Gradual deviation of agent choices from intended policy or brand standards.

Why it matters: Quietly erodes CX and compliance.

Example: Offer rules slowly violate pricing policy until evals flag drift.

DECISION VELOCITY

Definition: How quickly the system can make and execute high-quality decisions.

Why it matters: Correlates with growth in perpetual marketing systems.

Example: Bid adjustments propagate across channels within minutes, not days.

DIGITAL NERVOUS SYSTEM

Definition: Enterprise-wide fabric connecting agents across HR, finance, supply, and marketing.

Why it matters: Eliminates siloed decisions that break customer journeys.

Example: Refund decisions read inventory and marketing intents before messaging customers.

E

EVALUATION FRAMEWORKS ("EVALS")

Definition: Standard tests that score agent outputs against quality and safety criteria.

Why it matters: Grounds "creative" or "smart" systems in measurable standards.

Example: Weekly evals show tone drift; prompt and policy are adjusted.

EVALUATION HARNESS

Definition: Automated pipeline that runs evals on test sets before and during deployment.

Why it matters: Prevents regressions and brand risk at scale.

Example: New prompt version must pass tone and safety tests to ship.

F

FEDERATED GOVERNANCE

Definition: Local teams run their agents under globally consistent policies and identity rules.

Why it matters: Balances speed with enterprise consistency and compliance.

Example: Country teams localise offers but share central caps and audit.

G

GLASS-BOX SYSTEM

Definition: An AI/agent setup where reasoning, inputs, outputs, and actions are visible and replayable.

Why it matters: Builds trust with boards and regulators; accelerates debugging.

Example: You replay the steps that led to a creative decision during a brand incident.

GOVERNANCE ARCHITECTURE / TRUST ARCHITECTURE

Definition: The combined design of policies, controls, logs, and roles enabling safe autonomy.

Why it matters: Determines how fast and safely you can scale agentic programs.

Example: Global policy with local overrides defines escalation thresholds

GOVERNANCE-AS-CODE

Definition: Expressing policies and limits as executable rules enforced at runtime.

Why it matters: Prevents "after-the-fact" violations; enables safe-to-scale autonomy.

Example: Policy denies sending any message to under-18s, enforced in the pipeline.

H

HUMAN-IN-THE-LOOP (HITL)

Plain-English definition: A setup where humans remain directly involved inside the AI's workflow — reviewing, approving, or editing outputs before they go live.

Why it matters for marketing: Keeps brand and compliance oversight within automated systems, especially for creative or regulated content.

Example in context: Before a social copy agent posts a tweet, a brand manager approves or edits it inside the same workflow.

HUMAN-ON-THE-LOOP

Plain-English definition: Humans monitor and supervise AI systems at a higher level, stepping in only when alerts or anomalies arise.

Why it matters for marketing: Reduces manual load while maintaining safety; marketers act as strategic overseers instead of task reviewers.

Example in context: A manager only intervenes when campaign performance drops below threshold, otherwise, the AI runs autonomously.

Difference between the two:

- Human-in-the-Loop: Human is inside the process with direct control and approval.
- Human-on-the-Loop: Human is above the process with oversight and intervention by exception.

I

IDENTITY & ACCESS MANAGEMENT (FOR AGENTS)

Definition: Assigning identities and least-privilege access to non-human actors.

Why it matters: Prevents over-entitled agents from harming systems or data.

Example: Creative agent can read but not write to the PIM without approval.

INTEGRATION FABRIC / SEMANTIC FABRIC

Definition: Shared data models and contracts that let agents "speak the same language."

Why it matters: Enables cross-function decisions that make sense end-to-end.

Example: "Order_status" means the same thing in supply, finance, and marketing agents.

IPAAS (RE-IMAGINED)

Definition: IpaaS is a cloud-based tool that helps different software systems talk to each other. Businesses use iPaaS to connect apps that don't naturally work together. For example, linking a CRM like Salesforce with an email tool like Mailchimp or an ERP system like SAP.

Why it matters: Moves from pure data sync to decision orchestration across tools.

Example: Agent messages flow via the platform to update CRM and ad platforms coherently.

L

LEARNING LOOP

Definition: A learning loop is the process by which a system improves its performance over time by using feedback from its own actions or outputs.

Why it matters: Tech vendors believe this is essential for replacing campaign calendars with continuous improvement in the long run.

Example: Email subject tests update the agent's selection policy for tomorrow's sends.

M

MARTEC'S LAW

Definition: First coined by Scott Brinker, this is the idea that technology changes faster than organisations can adapt.

Why it matters: Martec's Law explains why transformation stalls – companies must evolve their people, processes, and structures to keep pace with innovation.

Example: Teams update tools every year but rarely redesign workflows or incentives to match.

MCP (MODEL CONTEXT PROTOCOL)

Definition: An emerging standard for passing tools, data, and context between AI models and agents to enable coordination and interoperability.

Why it matters: Reduces integration friction; enables multi-agent workflows.

Example: Creative agent invokes a brand-policy tool via MCP to validate copy before publishing.

METERED INTELLIGENCE

Definition: Treating cognition like a utility. Usage is measured and billed.

Why it matters: Turns cost into a live signal you must govern.

Example: Forecast shows token spend will breach budget; system downgrades model depth automatically.

O

OBSERVABILITY (FOR AGENTS)

Definition: Telemetry across prompts, tools, latencies, costs, and outcomes that enables live monitoring and debugging of agent behaviour.

Why it matters: You can't govern what you can't see.

Example: Dashboard shows retries spiking on asset generation after a rules change.

ORCHESTRATION

Definition: The discipline of coordinating agents, data, and tools so that AI systems operate together to achieve defined outcomes reliably and efficiently.

Why it matters: The difference between speed and chaos.

Example: "Safe-to-scale" pipeline runs checks before high-volume deployment.

OUTCOME ATTRIBUTION (FOR OUTCOME PRICING)

Definition: Methods to prove a result was caused by an agent.

Why it matters: Required for paying vendors and reporting ROI.

Example: "Recovered order" tied to the recovery agent's touchpoint logs.

OUTCOME-BASED PRICING

Definition: Paying vendors for results (e.g., resolved cases) rather than seats or tokens.

Why it matters: Aligns incentives; clarifies value capture.

Example: Paying per successful cart recovery by an agent.

P

PAY-PER-CRAWL / PAY-PER-ACCESS

Definition: Charging for automated access to web or API resources.

Why it matters: Affects data costs for agents that retrieve external context.

Example: Retail content crawl is throttled to stay within monthly budget.

PROMPT (VS. ROBO-PROMPTING)

Definition: Human-written instructions to an LLM; robo-prompting is prompts generated by software.

Why it matters: Moves from artisan prompts to scalable, data-driven prompting.

Example: An analytics agent generates prompts for the creative agent based on cohort insights.

R

RBAC (ROLE-BASED ACCESS CONTROL)

Definition: Controlling access by assigning permissions to roles instead of individual users.

Why it matters: Makes agent permissions auditable and consistent.

Example: "Marketing-agent-reader" role can view PII-redacted segments; writer role cannot.

RAG (RETRIEVAL-AUGMENTED GENERATION)

Definition: Fetching and injecting relevant information into the model's context before generating a response, to improve accuracy and relevance

Why it matters: Improves accuracy and brand consistency in content and service replies.

Example: Agent retrieves the latest brand voice rules before writing copy.

RUNTIME OBSERVABILITY

Definition: Live monitoring during execution (not just logs after the fact).

Why it matters: Enables "see-then-steer" instead of "see-then-suffer."

Example: Alert fires when SMS delivery latency breaches SLA.

RUNTIME RESTRAINT

Definition: Limits enforced at the moment of action, not after.

Why it matters: Stops bad outcomes before customers see them.

Example: Message blocked in real time for violating sensitivity rules.

S

SAFE-TO-FAIL SAFE-TO-SCALE

Definition: Maturity shift from small experiments to governed, auditable rollout.

Why it matters: Lets you expand autonomy without escalating risk.

Example: A/B pilot passes evals, then scales with control-plane oversight.

SEMANTIC LAYER

Definition: A semantic layer in AI is a translation layer that turns complex data into clear, consistent business terms so both people and AI systems understand and use the same definitions..

Why it matters: Avoids "same term, different meaning" errors across agents.

Example: "Active customer" definition is shared by CRM and ad-buying agents.

SUPER-AGENT

Definition: A higher-order agent that coordinates specialised sub-agents.

Why it matters: Simplifies complex programs under one accountable entity.

Example: “Retail media super-agent” orchestrates planning, creative, and measurement agents.

T

TELEMETRY (AGENT TELEMETRY)

Definition: Continuous data about agent runs: latency, cost, errors, outcomes.

Why it matters: Drives performance tuning, budget control, and incident response.

Example: Spike in tool-call latency explains ROAS drop; fix unblocks delivery.

TOOL CALLING

Definition: An LLM/agent invoking external functions (APIs, databases, actions).

Why it matters: Turns text generation into end-to-end execution.

Example: Service agent calls “refund” tool with guardrails based on policy.

U

USAGE-BASED / CONSUMPTION PRICING

Definition: Paying per token, API call, or action executed.

Why it matters: Aligns cost with actual work; requires real-time metering.

Example: Weekly token budget caps prompt the system to select cheaper model routes.

W

WEBHOOKS

Definition: A webhook is a way for one piece of software to automatically send real-time information to another when something happens.

Why it matters: Keeps agents responsive to customer and operational events.

Example: “Cart abandoned” webhook triggers recovery flows in minutes.

Z

ZERO TRUST (FOR AGENTS)

Definition: “Zero trust is a cybersecurity approach that assumes no user or device – or agent – is trustworthy by default and requires continuous verification for every access request.

Why it matters: Stops over-permissive agents from becoming breach vectors.

Example: An agent must re-authenticate and pass policy checks before accessing PII.

Acknowledgments

Name	Job Title	Company
Adrian McDermott	CTO	Zendesk
Aimee Dixon	GM Marketing	Australia Post
Alan Trefler	CEO	Pegasystems
Alexandra Finster-Rowen	Global Projects Director	Catapult
Andrew Boni	Founder and chief scientist	Iterable
Andy O'Dower		Twilio
Babak Hodjat	CTO AI,	Cognizant
Ben Kus	CTO	Box
Chandar Pattabhiram	SVP Go to Market	Workato
Colin Fleming	CMO	ServiceNow
David Hyman	CEO	Lendi
David Meyer	David Meyer, senior director AI strategy	Databricks
David Singer	Global VP Goto Market Strategy	Verint
Derek Slager	Cofounder and CTO	Amperity
Devesh Maheshwari	CTO	Lendi
George Colony	Founder and chairman	Forrester Research
Henry Innes	Co founder	Mutinex
James Kruger	CMO	Informatica
John Kucera	SVP Product Management	Salesforce
Kevin Li	SVP Product Management	Optimizely
Kranthi Nekkapula	Head of AI Platforms	Suncorp
Liz Cherry	Growth Marketing manager	HHM Global
Liz Knox	COO	DDB Group New Zealand
Liz Miller	VP and Principal Analyst	Constellation

Name	Job Title	Company
Matty Burton	Group Chief Creative Officer	DDB Group ANZ
Michael Bachman	Head of AI Strategy	Boomi
Mohammad Heidari Far	Managing Director Australia	Flywheel
Paul Meinshausen	CEO and co founder	Aampe
Peter Van Der Putten	Director AI Lab and lead scientist	Pegasystems
Rahul Garg	VP Product, AI and Self Service	Genesys
Sam Allen	CEO	Iterable
Sangeeta Mudnal	CTO	Glu.ai
Scott Brinker	Editor-in-chief	Chiefmartec
Steve Lucas	CEO	Boomi
Sean Falconer	Senior Director AI and strategy	Confluence
Sundeep Parsa	Vice President of Experience Platform	Adobe
Tejas Manohar	co-CEO	Hightouch
Tom Scully	Principal Architect, APAC	Palo Alto Networks
Wade Chambers	Chief Engineering Officer	Amplitude

PARTNER CONTENT

Name	Job Title	Company
Alex Pacey	Chief Product Officer	Omnicom Media Group
Emily Mead	Head of Partner Marketing	Omnicom Media Group
Kim Hamilton	Chief Marketing Officer	Omnicom Media Group
Kristiaan Kroon	Chief Operating Officer	Omnicom Media Group
Lucio Ribeiro	Chief AI and Innovation Officer	TBWA\Australia
Tessa Conrad	Global Head of Innovation	TBWA\Australia
Nick Garrett	CEO	Omnicom Oceania

Inside the Tornado:

Mi3's Agentic AI in Marketing Report

BROUGHT TO YOU BY

Mi3

www.mi-3.com.au

PARTNERED BY

OmnicomOceania

Click to learn more