omd

Rapid Response

AI Mediated Brands

Leveraging AI While Protecting Brand Integrity

Updated: March 6th 2026

WE Create WHAT'S NEXT

AI Moves From “Assist” To “Mediate”

AI has stopped being just a back-office efficiency play and has started becoming the place where consumers discover, evaluate, and buy. This shift has multiple implications for the way that brands are expressed, engaged with, and managed.

Alongside multiple innovations, routes to market, and growth drivers, there are new risks and regulations to navigate. This is why “try the new tools” and “protect the brand” need to be in the same conversation.

Ad platforms are rapidly turning conversational AI into paid media inventory. Google is explicitly repositioning Search as an AI-led discovery experience, with AI Mode framed as a new place for ads to fit naturally into the conversation. We are seeing a structural change as targeting moves from keywords and audiences to multi-turn intent, sentiment, and context.

Retail media is also turning into an agentic operating layer. Walmart Connect states that AI will anticipate, recommend, and act. It has expanded its conversational assistant, Marty, to simplify campaign work.

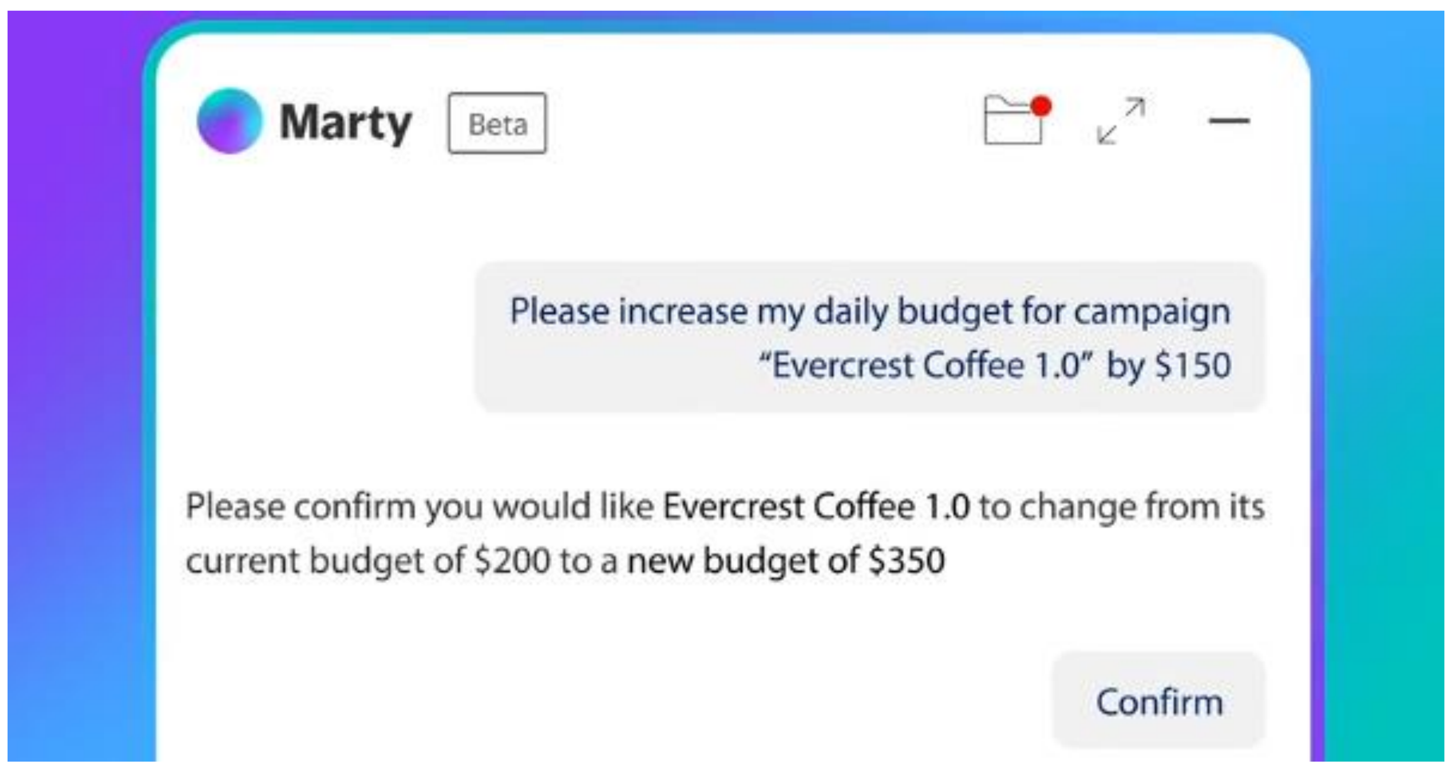

However, consumer trust is scarce. Pew Research Center found that Americans are far more concerned (50%) than excited (10%) about increased AI use, and that concern has risen compared with prior waves of adoption. A growing driver of mistrust is “AI Washing,” where companies overstate claims of AI proficiency.

Low-quality “AI slop” is flooding the open web, often engineered to harvest ad dollars at scale. Verification companies are treating slop as the next evolution of MFA and are building avoidance segments inside DSPs. Brands should define a clear “slop tolerance” like a media policy, treating AI-generated, low-quality environments as a measurable risk category. Use pre-bid avoidance segments where available (e.g., MFA, ad clutter, slop) and supplement with curated supply paths and inclusion lists where appropriate.

The net effect is that while automation increases velocity, it also increases the likelihood that brands become adjacent to noise. When audiences expect fabrication, the old playbook of “more content, more frequency” can backfire.

Taking Control

Regulators and trade bodies are developing their guidance, moving toward disclosure and governance frameworks in advance of tighter policies and tougher punishments.

The European Commission has a Code of Practice on marking and labelling AI-generated content to support EU AI Act transparency obligations, new requirements of which come into play this August.

While in the UK, the Advertising Association has published a Best Practice Guide for responsible GenAI use in advertising, delivering 8 key principles and a playbook for advertisers to follow. The Committees of Advertising Practice's (CAP) monitoring work is a reminder that existing rules on misleading claims and harmful content still apply, even when the asset was generated.

In the US, the Interactive Advertising Bureau (IAB) launched an AI Transparency & Disclosure Framework in January 2026.

In response, platforms are now building “steering wheels” for automation.

Google’s new text guidelines let advertisers constrain AI-generated copy (term exclusions + natural-language guardrails) for Performance Max and AI Max for Search.

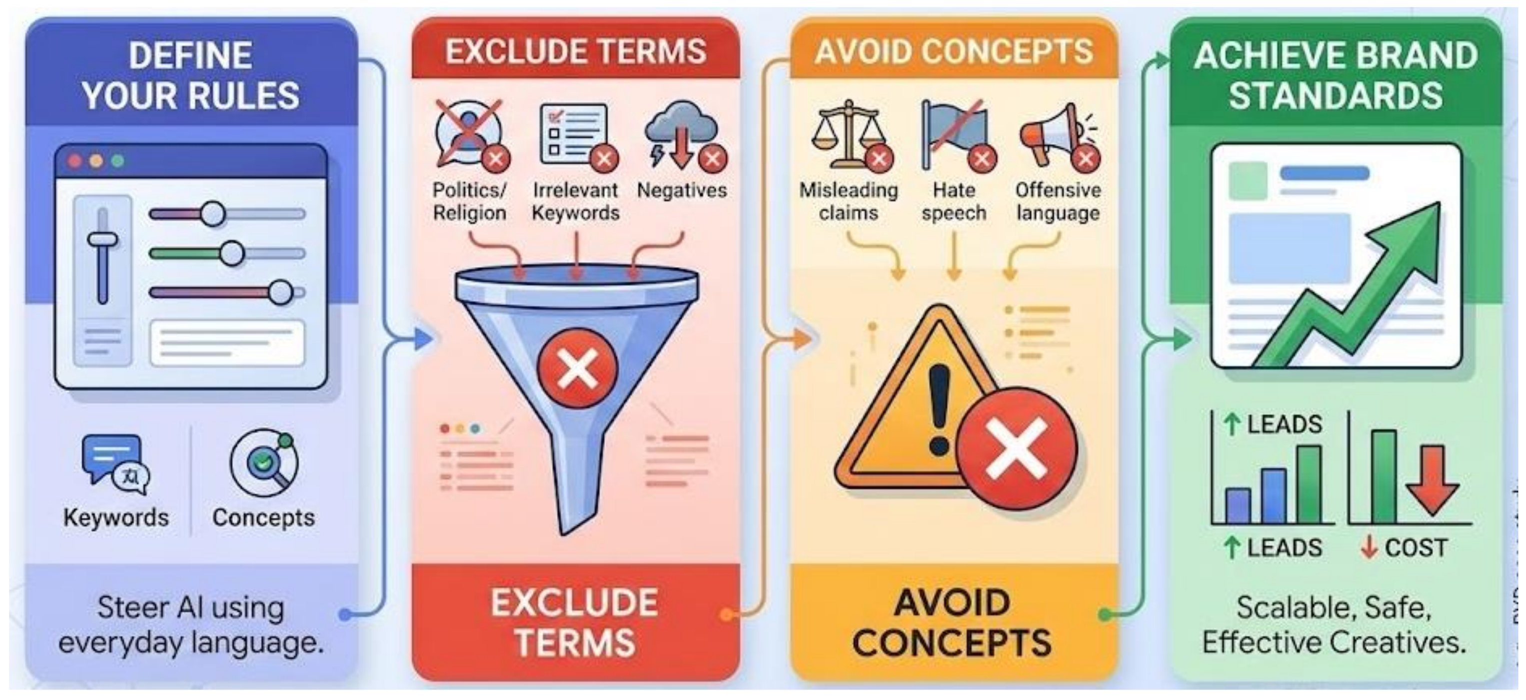

Fully automated copy is not brand-safe by default; policy, tone, and compliance need to be encoded, not “hoped for.” This is a new surface for brand guidelines and values to operate in.

Brands should start with platform controls such as Google’s text guidelines to block banned terms, compliance-sensitive language, and off-tone phrasing. Then extend to a proprietary “prompt library” and approval workflow: what tone is allowed, what claims require substantiation, and what categories require human sign-off.

All of this should be aligned with the various emerging agentic protocols such as UCP, Web MCP, and AdCP to deliver highly personalized yet coherent data, content, and services as brand experiences.

This is an entirely new mechanism through which brands can be expressed. Human-readable brand books are no longer sufficient and must be translated into machine-readable code that can be utilized across a wide and ever-expanding range of contexts.

Stateful AI

Last week, AWS and OpenAI announced a multi-year partnership to co-create a Stateful Runtime Environment in Amazon Bedrock built to run reliable, multi-step agents.

The OpenAI–AWS partnership explicitly positions stateful, governed agents as production-ready building blocks. This increases the value of safety layers and auditability, not just faster outputs at scale.

For advertisers, this matters because the next wave of media tools will increasingly be agent-based workflows to build audiences, generate creative, optimize retail spend, and automate reporting.

It also opens new routes to deliver value and growth.

A significant lock-in mechanism for the AI platforms in recent months has been memory, i.e., the ability for an AI to remember not just the ongoing conversation but also wider history with the user – their preferences, context and data points – meaning it is harder for users to switch between models, as they would need to constantly re-enter context.

A scalable stateful runtime environment means that brands can deliver more sophisticated services based on long term relationships, data permissions, and user context.

Stateful agent environments will push more brands to connect internal systems (PIM, pricing, inventory, CRM) to AI workflows. This means we should treat data usage as a brand promise, not just a compliance box-tick.

Seek to apply strict minimum necessary data rules, isolate sensitive datasets, and demand clear retention and training terms. For example, AWS positions Bedrock Guardrails as a centralized safety layer for GenAI applications.

Build for AI-mediated discovery without losing a human edge. If discovery is increasingly conversational, brand meaning must be both machine-readable and human-believable. Brands like Dove are leaning into trust signals by pledging not to use AI imagery to represent real women.

From Brand To Model Governance

AI technologies are racing forward with ever more sophisticated models and harnesses, delivering agentic capabilities. There will be multiple emerging opportunities for brands to deliver distinct, consistent, and memorable identities that foster emotional connection and trust.

Brands will need to be encoded in new ways so that agentic tools can optimize and create in ways that are unique to them, avoiding the commoditization of a black-box solution. This involves developing protocols for capabilities that are only just becoming clear, with many more over the horizon.

Existing critical areas, such as brand safety, are evolving. Omnicom Media's CASA formalizes the idea that advertisers have a right to know where ads run, who sees them, and whether environments align with brand values. CASA's push for SSP standardization was originally designed to fight opacity and MFA abuse, and it becomes even more important as AI-generated inventory scales.

Brands are now forced to balance hyper-efficient, autonomous AI agents with the growing consumer demand for emotional resonance, cultural authenticity, and transparent business practices.

The key principle is to use AI in a way that earns permission and builds trust across stakeholders. Brands will need guidance to test AI-native tools quickly, while building the governance, data discipline, and quality controls that protect brand integrity at scale.

At Omnicom, the response is to treat AI readiness as an operational discipline: connected identity, connected commerce, and connected governance.

Omnicom's next-generation Omni is positioned as a single operating system that integrates data, identity, and AI to connect strategy to execution. In practice, this enables 'human-in-the-loop' agentic workflows, delivering speed with accountability rather than black-box automation.

The implications of trust as an enabler (and barrier) to AI innovation are a core theme of a new global study due to be released by Omnicom Media this month.

Jean-Paul Edwards
Managing Director, Product
OMD Worldwide