


A stylized graphic of a robot head. The top part is a dark blue semi-circle with a black antenna on top. The main body is a large dark blue semi-circle. The bottom part consists of two large, light blue circles representing eyes. The background is a solid dark blue.

THE ROBOTS ARE COMING

A GUIDE TO THE FUTURE OF
AI & MARTECH



The Marketing Cloud

“The future cannot be predicted, but futures can be invented. It was man’s ability to invent which has made human society what it is.”

—**Dennis Gabor**, *Inventing the Future*, 1963

Don't have time for all this reading?
We've got you.

Scan below for a digital summary,
powered by Agent Cloud:



themarketingcloud.com

Table of contents

- Agentic AI, explained 1
- 5 things AI agents can’t do (yet) 3
- How agentic AI will redefine influencer marketing 5
- The AI revolution redefining market research 7
- Break the AI slop loop 9
- When AI is smart enough to write its own
rules for creative optimization 13
- Why “BYOAI” is bad business 15
- How autonomous should AI agents actually be? 17
- Agentic AI: What to expect in 2026 19



Agentic AI, explained

Quite a slippery term for a paradigm shift

“Agentic AI” was 2025’s loudest buzzy phrase, but also a term that a lot of people had trouble explaining. Marketers adopting AI agents have felt the shift already, but the hard part is telling substantive, no-BS capability from shiny relabeling.

An agent isn’t just whatever you say it is

Let’s start with what most experts agree on: Agents aren’t a single thing, they live on a spectrum. Toward the middle of that spectrum are task agents that plan multi-step work and “call” third-party tools; at the far end are semi-autonomous “virtual teammates” that pursue goals, escalate when uncertain, and improve with feedback.

But let’s keep things fairly simple. Agentic AI is software that owns a goal over time and can act toward it without you babysitting every step. It plans its own sub-tasks, uses tools/APIs, and remembers or learns from outcomes to adjust the next move.

The key is autonomy plus feedback loops. Truly agentic AI proposes, executes, checks results, and iterates.

Why marketers shouldn’t sleep on agentic AI

Why should marketers care about any of this?

Because agents shrink the time and distance between insight and action, and they increase efficiency so that your human teammates can focus on the stuff the machines can’t.

In practice, agents are getting good at breaking down messy marketing tasks into steps, calling the right tools (search, analytics, asset libraries, CMS), and documenting their decisions.

“Agentic AI has the agency to operate autonomously, make decisions, and adapt in complex marketing or general environments with minimal human oversight.”

ERIC WALZTHÖNY KREUTZBERG
CO-FOUNDER & CTO, SMARTASSETS



Agents can maintain memory recall across marketing workflows, respect constraints you set (budgets, brand rules, approvals), and run continuous, tiny experiments to improve their results.

Here are a few concrete examples of what a marketing-focused AI agent could do:

- Scan reviews and social chatter on Reddit, cluster themes, and drop a one-page insights brief to be reviewed by your CMO.
- Draft a customer-friendly email based on product update notes, then personalize variants by segment.
- Every Friday, summarize the past week of customer feedback into 5 “what to fix” bullets for product and CX.
- Produce 10 on-brand social captions from a blog post, A/B test, and schedule the winners.
- Monitor social for crisis signals, draft a holding statement in your brand’s tone, route approvals, and post to LinkedIn once approved.

A big reason agentic AI is improving fast involves training in simulated environments, where agents practice multi-step software use with rewards and penalties. Silicon Valley is pouring money into these environments to make agents more reliable.

But let’s not get ahead of ourselves...

Plenty remains overpromised (it’s why critics have flagged “agent washing” as an ongoing problem). When it comes to things like long-term judgments or predictions, or finessing small cultural nuances, marketers still need a human touch.

Integration also needs to improve if marketers want robust execution. That’s why having a centralized agentic AI experience like the one offered by The Marketing Cloud is important, especially if you need a data-secure environment rather than a sloppy, Bring-Your-Own-AI mess.

Compliance is another hurdle. Marketing in regulated categories (pharmaceuticals, finance) still demands human checkpoints because agents will occasionally overstate a claim or cite the wrong basis. Brands and agencies have to weigh the gains in

productivity and efficiency with their own risk appetite, putting guardrails in place sooner rather than later.

That said, agentic AI isn’t just hype, as evidenced by Google’s comprehensive report, “The ROI of AI 2025.” AI agents are swiftly moving into the enterprise core, with potential buyers citing data privacy as their top concern, followed by ease of integration and cost.

How agentic marketing is likely to evolve in 2026

While we’re sure to see ongoing splashy promises and spectacle, the real agentic AI action in 2026 will involve the more mundane work of getting these systems up and running in a way that balances autonomy with risk.

Marketers kicked off this past year with agentic AI front-and-center in the conversation at CES 2025 in Las Vegas, where keynotes leaned into it hard; Nvidia literally framed agentic AI as the next platform moment. The mood back then was bullish, but now that the hype is settling, it’s time to focus on governance, results, and ROI.

That means building systems that are trustworthy, secure, and results-driven. It’s why exploring the potentials of custom AI assistants and AI agents within an environment like Agent Cloud can be so helpful, especially for brands or agencies that don’t have an existing AI infrastructure in place. ■

5 things AI agents can't do (yet)

They won't clean your laptop screen, bring you coffee, or convince your Gen Z coworkers that you're hip. And as of now, you sadly can't count on AI agents to nail anything below.

1. AI agents can't force market-research panel respondents to always tell the truth.

Agents can vaporize bots and speeders, improving the quality of your sample immensely. But they can't guarantee that someone who signed up for the incentives to a niche study for a new line of energy drink is actually a "left-handed orthodontist who golfs biweekly."

2. AI agents can't keep your CEO on message longer than a TikTok.

Stop us if you've heard this one before...You give your charismatic CEO talking points, color-coded Q&A, and three pre-briefs. Yet once they grab the mic, it's off to the races. Your brand has suddenly soft-launched a merger, invented a whole new category,

and quoted a compelling stat that unfortunately does not exist. (That said, AI agents *could* help track and address the resulting fallout on social media.)

3. AI agents won't settle your team's creative differences forever.

Yes, computer vision can score contrast and logo safety. But it can't definitively tell you which Pantone shade is right for a new DOOH campaign, or whether the kerning on a headline looks good-weird or bad-weird.

4. AI agents can't turn "surprise money" into guaranteed results overnight.

Bid models can pace and predict with a fair amount of autonomy. But when fresh budget appears at the end of the month, no agent can conjure perfect audiences, premium slots, and flawless performance on a holiday week. That's not optimization, that's straight-up sorcery. (If you figure out how to do it, let us know.)

5. AI agents can't get everyone to agree on what actually worked in a campaign.

An agent can stitch data together and chart it beautifully, making you the Queen of Decks. It still can't get the brand team, the performance crew, and finance to accept the same narrative. The only "single source of truth" it can definitively promise is a single source of new arguments, with prettier infographics. ■



The next evolution: How agentic AI will redefine influencer marketing

The comms world is on the cusp of another AI evolution. While we've all gotten (somewhat) comfortable with LLMs cranking out branded social media posts and polishing press releases, agentic AI represents something far more transformative.

Agentic AI is artificial intelligence that doesn't just respond to our prompts, but actually thinks ahead, makes decisions, and takes action.

A standard LLM might craft a compelling product launch announcement when prompted. If you want further adjustments, the user prompts and tweaks again, until the task is complete.

An agentic AI system, however, continues working toward its goals even when a human isn't there to babysit. It might monitor industry trends, detect the optimal timing window, automatically segment target audiences, draft personalized pitches for different journalist beats, schedule distribution across time zones, and then track engagement—adjusting the approach in real-time based on what's working.

To be fair, some of this is aspirational. The tech needs to develop, as does our ability to trust agentic AI.

Or, as PRophet CMO Jason Brandt winkingly puts it: “Very few agents—perhaps none—are autonomous at this point. But that’s the end goal, of course...without them destroying us.”

Agents, meet influencers

The scene: Your brand wants to partner with as many micro-influencers as possible in advance of a new product launch. But doing proper discovery and vetting, at scale, isn't feasible.

Agentic AI is already embedded in products like PRophet Influence, which leverages the tech to make basic influencer discovery more efficient. This conversational style interface—the ability to ask the interface to “find influencers that fit this specific profile”—is agentic, but fairly simple, says Brandt.

Greater layers of complexity could make influencer



management much more automated. For instance, a fully agentic system could:

- Find influencers who match a pre-determined profile
- Vet them for brand safety and conflicts of interest
- Contact approved influencers via email to explain project and inquire about rates
- Draft a contract to be finalized by a human supervisor

Obviously, this puts a great deal of faith in the accuracy of your agentic system. When AI agents directly engage with potential collaborators by sending messages and discussing terms, your brand or agency needs to weigh the efficiency gains against potential risks. This is why the current level of agentic autonomy is still modest, despite the field's aspirations.

What's next for agentic AI in comms, social media, and marketing?

Looking ahead to the next six months, expect to see more sophisticated agentic solutions across the comms, PR, and social media space.

That means things like enhanced multimodal monitoring (analyzing images, audio, and video alongside text), or improved integration of standalone AI tools into coordinated "autonomous" functions across the marketing tech stack.

The agentic applications for influencer marketing are pretty diverse. Influencer-focused agents could:

- Rebalance campaign spend mid-flight if certain influencers underperform.
- Autonomously monitor influencer content across platforms, with AI tweaking recommendations (hashtags, posting times, cross-platform amplification) without human prompting.
- Collect immediate audience reactions (comments, shares, sentiment), feed them back into campaign strategy, and update influencer instructions automatically.

The real shocker will be systems that can seamlessly coordinate across multiple channels while maintaining consistent brand voice and strategic alignment.

We're long past the point where it's business-savvy to pride yourself on old school tactics alone. A future-proofed team will learn to "love the machine," and the most successful marketing and comms professionals will partner with these agents as seamlessly as possible.

As the tech is improving daily, being able to focus on the human stuff—strategy, relationship building—offers a good model for how to get the most value out of any tech and people investment. ■

The AI revolution redefining market research

Teams that figure out how to pair human judgment with machine speed are going to run laps around the rest.

QuestDIY surveyed* 219 U.S. market research and insights professionals to see how the field of market research is adapting and evolving with the advent of AI tools. Here are some of the key takeaways.

1) From cool tool to co-analyst

AI has formally graduated from “handy gizmo” to “junior teammate.” Researchers overwhelmingly see AI as a support or opportunity—59% and 36% respectively—and 89% say it’s already improved their work lives.

The future is human-led, AI-supported, with AI doing the heavy lifting while people steer strategy and judgment.

What does that look like on a typical morning? You might find AI drafting first-pass reports, surfacing patterns across messy datasets, and assembling “good enough” hypotheses that a researcher then refines with context, ethics, and business sense.

By 2030, researchers surveyed expect AI to act as an advanced decision-support partner (61%), with stronger generative abilities (56%), synthetic data

generation (53%), automation of core processes (48%), predictive analytics (44%), and “cognitive insights” (43%).

2) Analysis at scale

If you had to pick a single best use case for AI in insights, it’s analysis. Professionals we surveyed cited analyzing multiple data sources (58%); unpacking structured data (54%); automating insight reports (50%); analyzing open-ends (49%); and summarizing findings (48%) as key benefits.

“The center of gravity is analysis at scale—fusing multiple sources, handling both structured and unstructured data, and automating reporting.”

GARY TOPIOL
MANAGING DIRECTOR, QUESTDIY



And it's not simply about churning out this work faster. Researchers report AI improves accuracy (44%), surfaces otherwise-missed insights (43%), and accelerates delivery (43%), with a creativity boost (39%) to boot. In other words, AI helps you get to “the so what” sooner, and sharper.

3) Workflow speedups are real, and really measurable

98% of researchers have used AI at work in the past year, and 72% use it daily or more. Usage is rising (80% say they leverage it more than they did six months ago), and 71% expect to increase usage in the next six months.

Time savings are tangible: 56% report saving 5+ hours per week, while 15% save 10+ hours. That reclaimed time is getting redeployed into higher-value interpretation, stakeholder storytelling, and other human-first tasks.

A concrete platform example: QuestDIY (built by The Harris Poll, and part of The Marketing Cloud) bakes AI into survey drafting, question design, analysis, and deployment—supporting multimedia questions, quick translation, and real-time feedback across 100+ countries. You can field, analyze, and visualize faster than ever. And ISO/IEC 27001 certification is a nice data-security bonus.

4) New guardrails. And new jobs (really!)

None of this works without trust. The biggest brakes on adoption are exactly what you'd expect: data privacy/security (33%), lack of time to learn (32%), training (32%), and integration headaches (28%), with policy (25%) and cost (24%) not far behind. Teams also flag low transparency and the need to validate outputs (those pesky hallucinations).

Practically, that means the “human in the loop” isn't optional. Today, 29% describe their workflow as human-led with significant AI support and 31% as mostly human with some AI help.

“AI can surface missed insights, but it still needs a human to judge what really matters.”

GARY TOPIOL
MANAGING DIRECTOR, QUESTDIY

The market-research team of the near future looks like Research Supervisors and Insight Advocates guiding AI “insights agents.” There'll be (real) people supervising rigor, ethics, and business alignment while AI drafts, cleans, codes, and dashboards.

5) The researcher's role levels up

With repetitive tasks offloaded to AI, researchers can pivot to strategic storytelling, ethical oversight, cultural fluency, and technical confidence.

The net result? More influence in an organization, not less, as researchers move from humble data-gatherers to empowered advisors who shape business decisions.

The leaders we surveyed made it clear that AI elevates researchers into growth drivers and strategic advisors. A successful researcher in the coming years will be a master of AI as well as being adept at all the things the AI can't always manage: human curiosity, the ability to turn data into business-relevant narratives, and subjective judgment. ■

*The research was conducted using QuestDIY among 219 market research and insights professionals, defined as those who reside in the US, ages 18+, are employed full-time, and have a role in market research/insights in their organization. The survey was conducted August 15–19, 2025.

Break the AI slop loop

“AI slop” is the catch-all label critics have given to the tidal wave of low-effort, machine-generated text, images, video, and audio that we’re starting to drown in online.

Coined in the early 2020s, the term conjures images of digital junk food: cheap and plentiful, but nutritionally empty. An influential *New York* magazine essay by Max Read helped push AI slop into the national conversation.

AI slop is essentially content whose primary goal is volume or click-bait utility rather than insight, craft, or originality. Creating quick-and-fast slop can be a tempting proposition, especially for brands with limited resources who need to produce assets at scale.

How big is the slop problem?


A few key stats shed light on how sloppy the current internet is.


- 74 percent of new web pages already contain some AI-generated material. A study by Ahrefs analyzed 900k fresh pages (one per domain) from April 2025 and found that only 25.8% were purely human-written; 71.7% blended human and machine text, and 2.5% were 100% AI.
- Originality.ai’s [ongoing crawl](#) of the top-20 Google results for 500 queries shows AI-generated text peaking at 19.1% in January 2025

- and hovering around 16.5% in June 2025—up from barely 10% a year earlier.
- Marketers are part of the surge. In the same Ahrefs study mentioned above, 87% of 879 content marketers admitted they now lean on generative tools for at least part of their output.
- Audiences notice—and worry. KPMG’s [2024 U.S. consumer survey](#) found 67% list “fake news and fake content” as a top AI concern, beating fear of job loss or algorithmic bias.

Why the slop loop matters

The problem isn’t just that low-quality content is annoying or bad, or that your mom’s Facebook feed is now full of questionable deepfakes.





It's that generative models learn from the very pool they help fill.

The “curse of recursion,” demonstrated by researchers at Oxford and DeepMind, shows that repeatedly training on model-generated data drives output toward bland sameness and factual drift.

If the open web continues to thicken with AI slop, tomorrow’s models will be trained on today’s junk,

accelerating a feedback loop that degrades creativity, accuracy, and even the training data available for future AI systems.

For marketers, that means brand trust erodes, SEO value declines (Google is de-indexing spammy AI sites), and differentiation becomes almost impossible in a sea of formulaic prose and uncanny images with too many limbs. ■

A marketer’s anti-slop toolkit

Just because you’re using GenAI marketing tools doesn’t mean you’re contributing to the slop problem. It’s when marketers get lazy (or lean on the wrong tools) that their output gets sloppy.

1. Start with proprietary raw material

Proprietary information will give you a leg up, rather than letting LLMs pull from the public internet alone. Genuine first-party data is algorithm-proof; it forces the model to build on something fresh rather than remix existing web slurry.

GenAI brainstorming and production tools like those within Agent Cloud can be fed your brand guidelines, brand voice guidelines, and other reference assets, providing helpful guardrails that lead to better results.

2. Tighten the human-in-the-loop

Treat your AI tools as draft engines, not

infallible co-workers. Make sure that any AI-generated content, whether text or visuals, is given a full review by at least one (human) domain expert on your team.

3. Layer multiple models

Don’t rely on a single LLM for all your needs, but rather use what’s best for your current task. This can become labor- and time-intensive, though, as you toggle between Gemini, ChatGPT, and Midjourney.

Purpose-built solutions like Agent Cloud avoid this hassle by integrating access to the most applicable LLMs together, without the user having to determine the best fit on their own.

4. Don't be a lazy prompt engineer

No matter which AI model you’re using, the quality of your prompting makes a huge difference. Generic prompts are more likely to spin up slop. Taking time to add texture, specificity, and style guidelines goes a long way.





When AI is smart enough to write its own rules for creative optimization

by Eric Walzthöny Kreutzberg, Co-founder & CTO of SmartAssets

Most adtech optimizes delivery and targeting. SmartAssets starts earlier, at the creative itself. We've built the AI data and analytics platform that analyzes, predicts, and enhances ad content before it goes live, so brands launch work that hits harder with better ROAS.

Under the hood, SmartAssets surfaces the outputs of roughly a dozen machine-learning models across text, image, and audio. Brands can start with simple checks and progress to sophisticated, multimodal policies. Teams build a library of custom rules that all assets need to abide by, and then the AI is able to vet and score new creative, at scale.

These rules might be things like a textual requirement that the brand is clearly referenced; an image rule enforcing logo presence and specific size; and an audio rule confirming that the brand name is spoken in voiceover. If the goal is consistent across all three modalities (text, image, audio) you compose it once and the platform enforces it everywhere.

Because the models see inside the asset, rules can be object-based and context-aware: specific products, people, colors, scenes, or placements.

Consider a sponsor-heavy brand environment like Formula 1. With just a few custom rules tweaks, the marketing team there could ensure the correct season's car appears, that sponsor marks are current and in-spec, and that nothing outdated slips into an edit. As your creative library changes, your rules change with it.

Coverage, speed, and the edge of flexibility

Two axes matter when comparing solutions: how



much you can check, and how fast you can check it.

Today, SmartAssets evaluates around 300 rules and can process a two-minute video in about two minutes. We've built a common structure for what a rule is, and an engine to assess it at scale.

Now we're pushing control to marketers so they can shape their own policies, and we're paving the way for rules that the system proposes proactively. That's not sci-fi, it's just where the tech is at.

What does that look like in practice? It means SmartAssets could internalize your brand identity well enough to surface new, high-value rules and checks you haven't written or even thought of yet—and it will ensure those checks conform to your own guidelines.

That's where the platform becomes not just a validator of creative assets, ticking off pre-determined boxes, but an actual generator of brand intelligence.

Spoiler alert: The future is agentic

Agentic systems are the logical next step. We don't need more engines; we need more rules, drafted and iterated continuously.

Picture an embedded agentic AI researcher that crawls brand docs, scans the latest behavioral science and neuromarketing literature, and translates insights into potential new rules to optimize your brand's creatives.

The agent explores and recommends, and humans approve with a click after reviewing outputs across a sample of assets. Accepted rules join your live matrix, and rejected ones become lessons that refine the next batch.

But first, an important reminder that agentic does not mean ungoverned! Exploration and recommendation can be autonomous, but implementation stays human-in-the-loop. That balance keeps speed and scale without giving up accountability.

Why this matters now

Generative tools have multiplied output. Without enforcement, volume becomes volatility (and a whole lot of AI slop).

Custom rules in a tool like SmartAssets turn brand guidelines into executable standards, so that no creative slips out the door without being properly buttoned up.

Next, a move toward agentic AI workflows makes sure the rules you have in place for creative can evolve as markets, formats, and consumer behavior shift.

While your human team focuses on what they do best, a choreographed team of AI agents could suggest changes to creative rules based on shifting trends in the real world (preferences for new color palettes, a sudden allergy to overly prominent logos).

That equals a surefire boost to efficiency and quality, at scale. ■

Why “BYOAI” is bad business

by Louis Criso, Director of AI Engineering at The Marketing Cloud

Letting your team run wild with unsecure AI applications is a recipe for disaster. Luckily, there are alternatives.

Everyone’s feeling the pressure to adopt AI at work, but not every business leader is investing in proper AI procurement and training. That’s a real problem, simmering just under the surface.

Employees scramble to show that they’re able to adapt and change to new tech, but they’re not given the training, budget, or tools to leverage the benefits of AI without the potential risks.

In order to navigate this dissonance, we’ve seen the rise of a “Bring Your Own AI” (BYOAI) approach, in which teams and individual employees are furtively implementing their own AI solutions or adopting consumer-grade, non-enterprise versions of tools like ChatGPT.

BYOAI (also referred to as “shadow AI”) introduces serious problems, especially when it comes to data security, regulatory compliance, and even efficiency. It’s the kind of thing that makes your IT and legal teams sweat.

Let’s take a look at why BYOAI is bad for business,

and why purpose-built AI solutions are the way to go.

Lack of skilled prompt engineering

While most AI tools and Large Language Models (LLMs) are intuitive and easy to use, this doesn’t mean that everyone will be successful in leveraging their potential. An AI tool’s outputs are only as good as the inputs—which doesn’t bode well for untrained workers blindly using unverified AI applications.

Prompt engineering isn’t a throwaway skill; there’s actually a lot of nuance to it. Tailoring AI inputs to achieve specific outputs requires experience and patience.

Data security

When employees cherry-pick their own AI tools, they’ll likely turn to consumer-grade applications, which often don’t have robust security protocols. Signing up for a company-wide enterprise account with ChatGPT is relatively safe, provided IT has put



the right settings in place. But many other AI tools are developed by small teams or individuals who aren't focused on stringent data policies.

Anytime someone gives you something for free, there's generally a reason: They want your data.

Allowing unrestricted access to unverified AI applications can result in disaster for a company's data integrity and privacy—especially if employees are blindly dumping proprietary or confidential info into these tools, without reading the fine print about how such data can be used to train the model itself.

Inefficiency

The whole idea of AI tools is to save time. But using these applications incorrectly, or without a cohesive, company-wide plan, can actually introduce fresh roadblocks and inefficiencies. For instance, the consumer version of ChatGPT may be cheap (or even free), and it might be helpful for routine tasks like drafting emails or summarizing meeting transcripts.

But for more intricate operations that demand consistency and accuracy, these tools will likely fall short.

For marketing-specific tasks, having a dedicated application—like PRophet for comms, SmartAssets for creative optimization, or QuestDIY for survey writing—means that you're able to harness the upsides of cutting-edge AI without the trial-and-error of generic LLMs.

Lack of scalability

A homegrown or off-the-shelf AI model might work well when analyzing small data sets, but would cave under the pressure of a massive marketing campaign involving millions of data points.

Integration headaches

A custom AI solution often has to be integrated with your existing tech stack—CRM systems, email marketing tools, and analytics platforms. These integrations can become complex and require

significant IT support, slowing down your operations and leading to delays.

The benefits of purpose-built AI

The products within The Marketing Cloud are purpose-built for marketers. While they may be powered by various LLMs underneath the hood, from Anthropic to Gemini, these tools are integrated into an application that has been expertly designed for specific use cases.

Purpose-built AI applications are meant to scale easily, integrate with your existing tech stack, and uphold high data-security standards.

Here are a few other advantages:

Ongoing evaluation & evolution

Purpose-built solutions come with dedicated teams that actively monitor AI model performance and update their platforms accordingly. This means the organizations using these tools can work comfortably on the cutting edge, without doing the periodic re-evaluations themselves.

The same applies to Agent Cloud, the newest offering from The Marketing Cloud. This solution offers convenient access to all the major multimodal LLMs, with updated versions and new applications being added on a regular basis. You don't have to stay up-to-date on the latest and greatest AI platforms; instead, leave it to the experts, and reap the benefits.

Fresh data

One underlying worry about the future of AI models is that they'll begin to degrade over time. Since there's now such a glut of sloppy AI content online, newer AI models end up "digesting" it—and the results aren't pretty.

Purpose-built AI solutions are less vulnerable to this. And in the case of The Marketing Cloud's suite of products, they're constantly being fed and informed by real, pertinent marketing data from within the Stagwell network of agencies. ■

How autonomous should AI agents actually be?

by Louis Criso, Director of AI Engineering at The Marketing Cloud

Today's best agents are fantastic at speed, scale, and following orders. But when the stakes rise, they still need a human hand at the wheel.

If you listen to the more utopian predictions online, marketers (and everyone else) will soon be able to kick back, relax, and let a choreographed swarm of AI agents handle...just about everything.

As the Director of AI Engineering at The Marketing Cloud, I'm always going to be a bit cautious, especially when there are data security concerns.

I'm comfortable letting agents handle the repeatable, rules-based work where volume and velocity matter most: lead routing, budget pacing, first drafts of ad copy, and campaign reporting.

These are exactly the places where machines shine. Give an agent clear objectives and guardrails and it can autonomously adjust bids within pre-set limits, generate draft social calendars, and surface optimizations from live performance data. That's practical autonomy.

Where the hype outpaces reality

I'm far less bullish on the sweeping promises we keep

hearing, like agents that comb every email, infer your tastes, and book travel or make financial moves without a human's explicit say-so.

The nuance, liability, and personal context in those decisions are still beyond what's responsible to delegate. Brands and consumers shouldn't be asked to trust a "black box" with their money or reputation.

The risks that actually keep me up at night

- **Brand safety:** A single off-tone post or mismatched image can burn equity instantly. Would you really feel comfortable allowing an AI agent to draft a social post, create an accompanying image, and post it to your social feeds, without any human oversight?
- **Bad spend:** Loose controls can misallocate budgets or overshoot caps, crushing ROI. Telling your CFO that "it was the AI's fault" is the 21st century version of "my dog ate my homework."



- **Privacy and compliance:** Mishandled data or silent GDPR breaches aren't forgettable snafus. They're existential crises for large organizations.

None of these are theoretical. They're why autonomy must be paired with strong oversight.

I want mandatory human checkpoints at decisive moments: when spend crosses a defined threshold, when actions touch PII, and before anything goes public.

Adding these failsafes isn't a huge lift in terms of human labor, but it'll seem invaluable once you've dodged a major disaster (like an AI agent who decides to blow the entire Q3 paid media spend on a single Instagram ad).

These safeguards let agents keep the assembly line humming while humans own the sensitive, expensive, or reputationally risky calls.

When does an agent earn more freedom?

I tend to look at four signals:

1. **Success rate:** How often does the AI agent complete tasks correctly end-to-end?
2. **Intervention frequency:** How often do humans need to step in or unwind changes?
3. **Cost per result:** Is efficiency improving, holding, or degrading?
4. **Response time & trust:** Is the team confident enough to let the agent move faster?

If those trend in the right direction, widen the lane. If not, tighten it. And let's be honest: many brands don't have the bandwidth to run rigorous autonomy trials themselves. That's why pre-vetted environments like The Marketing Cloud and Agent Cloud matter.

Test agents like you onboard people

Before any agent touches production, I'd expect a staging environment, shadow assignments, and stepwise permissions.

Start with simulated or historical campaigns.

Graduate to low-risk live tasks. Only then consider expanding scope. It's the same way we coach a new teammate: observe, trial, review, then trust.

Every agent action should be traceable, with a record of what it did, when it did it, and why it thought that was the right move. That means timestamped logs, rationale notes, and explicit escalation flags when it defers to a human. It's how we audit, learn, and improve, and how we maintain compliance.

I'd say these guardrails are non-negotiable, at least for the moment:

- **Budget caps** that cannot be exceeded.
- **Approved tools and data sources only.**
- **Time-boxed access** so agents don't "free-range" at odd hours.
- A big, obvious **off switch** a human can hit instantly.

A realistic near-term future for agentic AI

Let agents do what they're great at, which is fast, repeatable, guardrailed execution.

Keep humans where judgment, empathy, and brand stewardship decide the outcome.

Expand autonomy only when the data says it's safe, preferably inside vetted sandboxes like The Marketing Cloud or Agent Cloud. That's how marketers can move from flashy, hot-air promises to durable value and ROI. ■

Agentic AI: What to expect in 2026

Futurecasting can be difficult when technology changes on a daily basis. But we asked some of the brightest minds within The Marketing Cloud (and the broader Stagwell network) to weigh in on what the field of agentic AI might look like in the coming year.

Reimagining the customer journey

“

The promise of agentic AI ushers in new hope for modern organizations that are truly invested in optimizing the entire customer journey. But how do we unlock the force multiplier of agentic AI?

It'll happen once we can connect disparate consumer data sources—from purchase history and loyalty program behaviors to motivations, affinities, engagement, customer service history, and more—with AI agents that have the rules-based logic and ability to activate channels and make decisions that increase a

consumer's lifetime value in real-time.

This future enables brands to truly learn from and adapt to consumer behavior, while at the same time breaking down silos that can exist across various channel teams that want to control bits and pieces of the consumer journey. Now the consumer—and their behaviors across every touchpoint—can drive the experience instead, and it's up to brands to continually optimize.”



ELSPETH ROLLERT
CEO, THE MARKETING CLOUD



Agents as a way to shape brand narrative

“

Agentic AI is speeding up marketing processes at an unprecedented pace—everything from creative to media to targeting to optimization. Autonomous agents are able to take action with minimal human intervention, make decisions, react to feedback, and plan sequences to complete tasks. Think of how many AI assistants you’ve encountered in the past year. That will only multiply in 2026.

As companies compete to incorporate AI into their business, those who come out as winners in 2026 and beyond will be those who embrace these agents as brand ambassadors. AI is already evolving SEO and will eventually overtake how consumers learn and hear about brands. No longer will companies be able to control their brand narratives and PR in the way they do today. Brands that actively

shape agents and innovate across every aspect of their business will be able to adapt to this wave of technology.

At the same time, as marketing grows more efficient, marketing budgets will begin to expand. Agentic AI will make everything from consumer interface to production of content significantly easier and less expensive. This means clients will spend more because they get better returns on their investment. Just as brands needed digital transformation experts to build websites and apps, they now need them to build agentic AI. There will be a whole new round of agentic AI transformation work, ballooning marketing budgets in 2026.”



MARK PENN
CHAIRMAN AND CEO OF STAGWELL

It’ll take an ecosystem

“

The future of marketing lies in fully integrated, AI-powered ecosystems that unify analysis, audience creation, activation, measurement, and optimization.

By embedding generative AI, 24/7 intelligent agents, and predictive analytics across the marketing lifecycle, we can unlock significant improvements in performance

metrics, drive significant impact, and outpace less innovative competitors. This isn’t just about automation. It’s about rearchitecting marketing as a dynamic, data-driven engine for strategic advantage.”



JOHN KAHAN
CHIEF AI OFFICER, STAGWELL

A framework to support the future

“

We are moving from general buzz around agentic AI to actual agentic enablement. The business reasoning is clear: Enterprises need a reliable infrastructure with security, privacy, and enablement that allows internal users to build agents to solve specific use cases.

In 2026, what we want an agent to do and what it is capable of doing will continue evolving daily. There won't be one kind of agent, but rather an assortment. Slow and fast, simple and complex, this variety of agents will all 'live' and thrive in a sort of agentic hive.

Think of multiple marketing agents,

for instance: Creative agents, Media agents, Optimization agents, and Reporting agents all working in cohesion within an agentic framework.

That agentic framework itself is also in flux. Sure, everyone would like to see a fully autonomous agent. But much like a fully self-driving car, this tech will initially have drawbacks, limitations, and regulations to abide by at first. We should be cautiously optimistic as we enter this exciting new world of AI and marketing.”



MANSOOR BASHA
CTO, THE MARKETING CLOUD

The agentic orchestra conductor

“

In 2026, the narrative of AI in marketing will pivot from a single, powerful monolithic genius (LLM) to the intelligent orchestration of 'AI specialist teams.' The leading marketing platforms won't be powered by a single LLM, but rather by a dynamic, heterogeneous system of agents.

A powerful LLM will act as the strategist, breaking down a task into smaller problems, while the moment-to-moment execution will be handled by fleets of hyper-efficient Small Language Models (SLMs).

These discrete steps or problems—

such as personalizing thousands of ad creatives, A/B testing subject lines or creatives in real time, or generating on-brand video clips—will be handled by these specialized models.

This will transform the human marketer from a creator to a conductor of an 'AI orchestra,' fine-tuning individual SLMs for specific tasks to deliver campaigns that are not only adaptive and personalized but economically scalable.”



ERIC WALZTHÖNY KREUTZBERG
CO-FOUNDER AND CTO, SMARTASSETS



A swarm of specialists

“

General-purpose Language Learning Models (LLMs) have significantly empowered researchers by enhancing efficiency, accelerating insight delivery, and unearthing insights that may have otherwise gone undiscovered.

However, similar to human teams, a shift from generalists to specialists is necessary. Specialists enable deep focus, technical expertise, and the competency to tackle more intricate problems.

In 2026, I anticipate AI agents starting to emerge with specializations in methodologies,

quality assurance, analytics, and reporting. This evolution will further boost our ‘Insight Advocates,’ experienced researchers who link AI-derived insights with organizations to address business-related questions.

In practice, that means the AI agents will operate independently across the research process while maintaining a human ‘in-the-loop’ to ensure ethical standards, governance, and human oversight are upheld.”



GARY TOPIOL
MANAGING DIRECTOR, QUESTDIY

Agents as market researchers, always on-call

“

Agentic AI will transform brand tracking from a tool for static reporting into a proactive decision partner. Rather than simply summarizing data, agents will continuously scan for shifts, surface what matters most to a brand, and deliver decision-ready insights in real time. The result: Clients move from data to action faster and with greater confidence.

In 2026, this won’t replace proven methods; it will amplify them. Agents will streamline how data is

connected, analyzed, and highlighted, giving Insights and Brand teams sharper, cleaner outputs on demand. With QuestBrand, we are actively testing and building these capabilities to compound the value of traditional brand equity measures with agentic intelligence. The pace of the market demands nothing less.”



ZEKE HUGHES
MANAGING DIRECTOR, QUESTBRAND

Now is the time to level up user experience (or get left behind)

“

Google announced its newly reimagined Chrome with AI in September, ushering browsers into the agentic era. By 2026, marketers should expect a new reality where browsers aren't just navigation tools, they're copilots attempting to mediate every brand interaction.

A very big problem for brands is that these AI browsers can function like sales reps from other companies stationed in your store. When a customer lands on your product page, agentic AI could instantly contextualize and present your competitors, side-by-side. The experience you painstakingly designed becomes commoditized into a generic comparison feed.

This accelerates a dangerous trend: the erosion of brand ownership. As agentic systems mature, customer acquisition costs will rise, loyalty will decline, and brands without strong first-party data will find themselves struggling for differentiation.

The good news is that the technology isn't there, yet. Today's AI browsers excel at 'better thinking,' not 'better doing.' They summarize, reorganize and compare, but they don't execute complex, multi-step tasks.

That lag gives brands time to innovate where Google cannot: industry-specific, convergent user experiences that can anticipate customer needs before a copilot intervenes.

In 2026, marketers who seize this window will have the chance to create irreplaceable experiences. Those who don't will wake up to find Google, Perplexity, and others have become the real point of contact with their customers.”



DAN GARDNER
FOUNDER AND CEO OF CODE AND
THEORY (PART OF THE STAGWELL
CHALLENGER NETWORK)

The Marketing Cloud

The Marketing Cloud is a suite of AI-driven products across market research, comms, creative, and media. We're empowering modern marketers with the tech to drive real impact.

Market Research



QuestBrand
Monitor brand health over time.



QuestDIY
Create, target, and deploy surveys.



BERA.ai
Quantify brand investment impacts.



QuestIC
Launch insight communities with engaged customers.

Creative & Media



SmartAssets
Score and optimize creative assets using AI.



TPP Insights
Unlock insights from audience location data.



SET
Manage complex marketing campaigns with AI.



Agent Cloud
Harness the latest LLMs and build custom AI assistants.

Agentic

Communications



Influence
Discover influencers and manage campaigns.



Earn
Plan and develop media pitches with AI.



Media Intelligence
Activate global media and market intelligence.



Discover a solution that's right for you.

Don't have time for all this reading?
We've got you.

Scan below for a digital summary,
powered by Agent Cloud:

