

The Rabbit Didn't Lose the Race, He Changed the Story: A Competitive Framework for the AI Race in AdTech, MarTech, OOH and DOOH

Suraj Prakash
Independent Researcher

Abstract

The fable of the rabbit and the turtle is conventionally read as a story about speed: slow and steady wins the race. This article proposes an alternative reading, that the fable is really a story about what happens when the definition of the race changes, and argues that this reframing is unusually well suited to describing the current state of artificial intelligence adoption across advertising technology, marketing technology, and out-of-home and digital out-of-home (OOH/DOOH) media. The article's central claim is that there is no single AI race underway, but at least four simultaneous and structurally distinct races, AI-assisted advertising, AI-optimized advertising, agentic execution, and agent-to-agent transaction, and that a company's position in one race does not determine its position in another. Drawing on independently verified fiscal year 2025 and 2026 disclosures from The Trade Desk, Salesforce, Adobe, PubMatic, Magnite, HubSpot, and Publicis Groupe, together with the Broadsign and Draft Digital agentic out-of-home campaign of 2026 and the IAB Tech Lab's Agentic Advertising Management Protocols, the article documents a consistent pattern across companies of markedly different core revenue growth alongside markedly faster growth in specifically AI-attributed metrics, and argues that this divergence is itself the more informative signal than either figure considered alone. Rather than ranking companies numerically, an approach this article argues creates false precision and reduces a competitive analysis to a leaderboard, the article proposes a qualitative eight-dimension framework, spanning intelligence, data, decisioning, execution, transactions, measurement, distribution, and feedback, and a four-category typology of competitive posture, encompassing scaled incumbents, persistent specialists, single-layer specialists, and infrastructure builders. The article concludes that as AI capability itself becomes commoditized across the industry, competitive differentiation is likely to move one layer deeper, toward proprietary data, inventory control, transaction ownership, and the protocols other agents are required to use, and that the organizations shaping those protocols may ultimately matter more than the organizations building any single agent.

Keywords: *Agentic AI, AdTech, MarTech, Out-of-Home Advertising, Digital Out-of-Home, Competitive Strategy, AI Monetization, Advertising Technology Standards, Platform Economics, Artificial Intelligence Adoption*

1. Introduction: A Story We Think We Already Know

The fable of the rabbit and the turtle is typically read as a story about speed, or more precisely, about the danger of overconfidence in speed. A rabbit, certain of victory, rests midway through a race and loses to a turtle who never stops moving. The conventional moral, that slow and steady wins the race, is not, this article argues, the most useful reading of the story for a 2026 audience evaluating advertising technology. If the rabbit had simply run the race as intended, he would have won easily, and the race would not be remembered. The turtle did not merely win a contest of speed; the turtle changed what the contest was actually about. This article proposes that the same reframing usefully describes the current moment across advertising technology, marketing technology, and out-of-home and digital out-of-home (OOH/DOOH) media, where nearly every company is racing toward the same word, artificial intelligence, without it being clear that they are running the same race, or that winning the race as currently defined will matter for very long.

This article proceeds as follows. Section 2 argues that the industry is not running one AI race but at least four simultaneous and structurally distinct ones. Sections 3 through 9 examine seven companies individually, drawing on independently verified financial disclosures to illustrate how AI capability and AI monetization diverge in practice. Section 10 places these companies' growth figures side by side. Section 11 extends this argument specifically to out-of-home and digital out-of-home advertising. Section 12 proposes a five-dimension evaluation framework and Section 13 an eight-dimension competitive matrix, explicitly avoiding numerical company scoring. Section 14 proposes a four-category typology of competitive posture. Sections 15 through 17 examine the economic and organizational implications of this shift, including a proposed new productivity metric, the implications for employment, and a description of an emergent agent-to-agent transaction architecture. Section 18 proposes six measurable signals for tracking this transition going forward, Section 19 states this article's reporting-period discipline, and the article closes with a discussion in Sections 20 and 21 returning to the fable itself.

2. There Is No Single AI Race

This article's foundational claim is that the phrase artificial intelligence, as used across advertising technology announcements in 2026, conflates at least four structurally distinct developments, each rewarding different capabilities and posing different competitive questions.

Figure 2. The Four Simultaneous AI Races in AdTech, MarTech, OOH and DOOH
(a maturity progression, not a per-company ranking)




Figure 2. The four simultaneous AI races proposed in this article. Higher-numbered races generally, though not universally, depend on capabilities established in the races below them.

The first race, AI-assisted advertising, encompasses copilots, content generation, recommendations, automated reporting, and conversational interfaces; AI helps a human perform a task the human still directs. The second race, AI-optimized advertising, is more consequential: AI begins making continuous decisions across bidding, budget allocation, audience optimization, forecasting, creative optimization, and pacing. The third race, agentic execution, marks a qualitative shift, in which the system does not merely recommend an action but takes it, creating a campaign, changing a budget, selecting inventory, negotiating, executing a workflow, and monitoring and adjusting outcomes. The fourth race, agent-to-agent advertising, is the least mature but potentially the most consequential: a buyer agent communicating directly with a seller agent, which in turn communicates with inventory, audience, and pricing agents, with humans providing governance and approval rather than performing each step manually. The IAB Tech Lab's Agentic Advertising Management Protocols (AAMP) initiative, covering agentic foundations, protocols, and trust and transparency, together with reference implementations intended to let buyer and seller agents discover inventory, negotiate, transact, and execute advertising workflows, is this article's clearest evidence that the fourth race has moved from speculation into active infrastructure development (IAB Tech Lab, 2026).

This article's position is that once a race is properly disaggregated in this way, a considerably more useful question becomes available: not which company is winning the AI race, but which of these four races a given company is actually running, and whether leadership in one race transfers to leadership in

another. The remainder of this article treats that question as its organizing framework.

3. The Trade Desk: Capability Without Immediate Acceleration

The Trade Desk illustrates why AI capability and revenue acceleration must be evaluated separately. The company's second-quarter 2026 revenue was 715 million dollars, up only 3% year over year, while first-half 2026 revenue reached 1.404 billion dollars, up 7%, with second-quarter adjusted EBITDA of 241 million dollars (The Trade Desk, 2026a). In August 2026, The Trade Desk introduced Kokai Zuma, the latest release of its Kokai platform, adding agentic AI capability to campaign workflows, and the company reported customer retention above 95% and third-quarter revenue guidance of at least 650 million dollars (The Trade Desk, 2026b). This article's assessment is that the more useful question about The Trade Desk is not whether the company possesses AI capability, which it clearly does, but whether that capability can be converted into a new layer of economic leverage on top of an already-scaled decisioning platform, a considerably harder question than capability alone answers.

4. Salesforce: When AI Capability Becomes Measurable Business Activity

Salesforce, situated outside advertising technology proper, offers this article's clearest example of AI capability translating directly into disclosed, measurable business activity. The company reported second-quarter fiscal year 2027 revenue of 11.35 billion dollars, up 11% year over year, while Agentforce annual recurring revenue exceeded 1.5 billion dollars, up more than 240% year over year, and combined Agentforce and Data 360 annual recurring revenue reached nearly 3.9 billion dollars, up more than 210%; the company further reported 7 billion cumulative Agentic Work Units delivered to date, including 3.2 billion in the second quarter alone (Salesforce, 2026). This article treats Salesforce's disclosure practice itself as significant: publishing a specific, audited, AI-attributable revenue metric is a materially stronger form of evidence than a qualitative claim of AI capability, and this article proposes that Salesforce's Agentforce ARR disclosure is the clearest available benchmark for what credible AI monetization reporting should look like across the broader industry.

5. Adobe: The Creative Half of the Battlefield

Adobe reported record third-quarter fiscal year 2026 revenue of 6.76 billion dollars, up 13% year over year, with total annualized recurring revenue reaching 27.50 billion dollars, and AI-first annualized recurring revenue exceeding 650 million dollars, growing more than 150% year over year (Adobe, 2026). This article's position is that Adobe's relevance to the advertising AI race is frequently underweighted specifically because Adobe does not compete in media buying; the contemporary advertising workflow increasingly runs from brief, through audience definition, strategy, creative production, media activation, measurement, and optimization, back into creative iteration, and Adobe is deeply embedded in the creative and strategic stages of that sequence. An analysis of agentic advertising that considers only media-buying platforms is, in this article's assessment, examining only

half of the relevant battlefield.

6. PubMatic: Agentic Advertising From the Supply Side

PubMatic illustrates a structurally different entry point into the same broader shift. The company's first-quarter 2026 results disclosed more than 20 AI agents available on its AgenticOS platform, more than 1,000 AI-powered deals transacted to date, and emerging revenues, a category including AgenticOS, Activate, Commerce Media, and Connect, growing more than 80% year over year to represent approximately 14% of total revenue (PubMatic, 2026). This article notes explicitly that the disclosed figure of more than 1,000 AI-powered deals is a cumulative total since AgenticOS launched in January 2026, not a quarterly run rate, a distinction independent industry commentary has specifically highlighted as important context (PPC Land, 2026a). This article's structural point is that the future of agentic advertising cannot consist solely of a buyer agent communicating with a demand-side platform; it eventually requires buyer and seller agents communicating directly with one another, and PubMatic's supply-side agentic investment is evidence that this architecture is being built from both directions simultaneously.

7. Magnite: AI Embedded Into Supply Infrastructure

Magnite reported second-quarter 2026 revenue of 192.8 million dollars, up 11% year over year, with contribution ex-TAC, the company's primary operating metric, growing 17% to 189.6 million dollars, connected-television contribution ex-TAC growing 36% to 97.1 million dollars, and adjusted EBITDA growing 30% to 70.6 million dollars, a 37% margin; the company subsequently raised its full-year contribution ex-TAC growth guidance to 13% to 14% and identified its agentic offerings as a potential future contributor to that growth (Magnite, 2026a, 2026b). This article's assessment is that the more informative question Magnite's results raise is not whether the company has adopted AI, but what happens once AI capability is embedded directly into the supply-side infrastructure through which media is actually transacted, rather than layered onto a buyer-facing interface alone.

8. HubSpot: The Mid-Market Agentic Platform

HubSpot reported second-quarter 2026 revenue of 911.7 million dollars, up 20% year over year, with total customers reaching 306,446, up 14%, and provided third-quarter revenue guidance of 924 to 925 million dollars alongside full-year 2026 revenue guidance of approximately 3.678 to 3.686 billion dollars (HubSpot, 2026a, 2026b). HubSpot explicitly describes itself as an agentic customer platform for scaling businesses. This article's position is that HubSpot's relevance to this analysis is structural rather than purely financial: it demonstrates that agentic capability is not confined to the largest enterprise platforms, and that a mid-market-oriented company can plausibly use agents to perform work that previously required a considerably larger operational team, a development this article considers capable of altering competitive dynamics well beyond HubSpot's own market segment.

9. Publicis Groupe: The Agency Holding Company's Version of the Race

Publicis Groupe reported second-quarter 2026 organic net revenue growth of 4.8%, accelerating from 4.5% in the first quarter, with first-half 2026 organic growth of 4.7% and a record first-half headline operating margin of 17.5%, up 17 basis points year over year; the company subsequently raised its full-year 2026 organic growth guidance to a range of 4.5% to 5% (Publicis Groupe, 2026). Publicis's own communications describe its strategic direction in terms of connected, agentic-driven capabilities, supported by investment across media, data, and technology. This article characterizes Publicis's position as the agency holding company's version of the same underlying race: rather than attempting to become another demand-side platform, Publicis is attempting to combine media, data, creative, technology, and AI into a single operating model, a strategy this article treats as evidence that the AI race extends well beyond companies that own advertising technology platforms directly.

10. What the Numbers Show When Placed Side by Side

Figure 1 places the core revenue growth and, where separately disclosed, the AI-specific metric growth of the seven companies examined in Sections 3 through 9 side by side, using each company's most recently reported quarter at the time of this article's preparation.




Figure 1. Core or headline revenue growth compared with disclosed AI-specific metric growth, most recently reported quarter, for the seven companies examined in this article. The Trade Desk, Magnite, HubSpot, and Publicis Groupe do not separately disclose a directly comparable AI-specific growth metric and are shown with core growth only. Reporting periods differ by company fiscal calendar; see Section 19 for the specific period attributed to each figure.

Two patterns in Figure 1 support this article's central argument. First, the gap between core revenue growth and AI-specific metric growth, where both are disclosed, is substantial in every case: Salesforce's 11% overall growth against 240% Agentforce ARR growth, Adobe's 13% overall growth against more than 150% AI-first ARR growth, and PubMatic's 13% underlying growth against more than 80% emerging-revenue growth. This article does not read this gap as evidence that AI-specific products are unimportant; on the contrary, it reads the gap as evidence that AI-specific revenue lines are starting from a comparatively small base and have not yet become large enough to move a company's headline growth rate substantially, which is a materially different, and more precise, claim than either an assertion that AI is transforming these businesses or a dismissal that AI monetization is not yet real. Second, The Trade Desk's comparatively low 3% core growth, despite substantial AI and agentic investment, is this article's clearest illustration that AI capability does not automatically translate into revenue acceleration, reinforcing the distinction this article draws throughout between possessing AI capability and successfully monetizing it.

11. Out-of-Home Advertising: A Physically Contextual Battlefield

Out-of-home advertising possesses a characteristic that most digital advertising channels do not: physical context. A given screen is not merely an impression opportunity; it has a location, a venue, a surrounding audience, mobility patterns, time of day, weather, nearby events, retail proximity, traffic conditions, pedestrian flow, screen characteristics, and historical performance. This article's position is that this physical richness makes OOH and DOOH an unusually favourable environment for agentic intelligence specifically, because a traditional system can only sensibly ask which screens are currently available, while an agentic system can, in principle, ask a considerably more demanding question: where the target audience will be, what journey they will take, which inventory combination provides genuinely incremental reach, what it will cost, what outcome should be expected, and how the plan should change if a defined sum of budget is reallocated from one market to another.

This is not a purely hypothetical distinction. In 2026, Broadsign and the agency Draft Digital described an end-to-end agentic OOH campaign in which buy-side and sell-side agents coordinated audience and venue targeting, inventory selection, campaign setup, and execution, with human approval and guardrails retained throughout, delivering more than 830,000 impressions across supermarkets, shopping malls, gas stations, and city streets in the Netherlands for the Lot of Happiness campaign, and moving from brief to booked, human-approved plan in under fifteen minutes, compared with a process the companies describe as historically requiring days to weeks (Broadsign, 2026). Broadsign separately states that its ecosystem powers close to three million static and digital signs globally (Broadsign, company disclosures). This article's position is that framing this opportunity narrowly, as AI media planning, understates its scope; the more accurate framing is autonomous OOH campaign intelligence and activation, or, at full maturity, an agentic OOH operating system.

12. What Does It Mean to Win? A Five-Dimension Evaluation

Rather than ranking the companies examined in this article from strongest to weakest, this article proposes evaluating each along five dimensions of competitive architecture: intelligence, meaning how sophisticated the underlying artificial intelligence capability actually is; data, meaning how proprietary and genuinely valuable the underlying data asset is; decisioning, meaning whether the system can resolve genuinely complex advertising decisions rather than simple ones; execution, meaning whether the system can take real action rather than only recommend one; and distribution, meaning how much actual advertising activity flows through the system in practice. This article's position is that a company can possess extraordinary artificial intelligence capability while lacking meaningful distribution, while a second company can possess comparatively ordinary artificial intelligence capability atop enormous existing distribution, and a third company can own the data both of the first two ultimately depend on; no single dimension, considered alone, determines competitive position.

13. A Competitive Matrix Without a Leaderboard

This article extends the five-dimension framework in Section 12 into an eight-dimension competitive matrix, and deliberately does not assign numerical scores to any company against it. Assigning a company a rating such as eight out of ten creates a false impression of precision that this article's underlying evidence does not support, and converts what should be a research framework into a leaderboard. Table 1 instead poses the central question each dimension raises and identifies what actually matters in answering it.

Competitive Layer	Core Question	What Actually Matters
Artificial Intelligence	Can it reason?	Underlying model plus applied intelligence
Data	Does it know something others don't?	Proprietary signals
Decisioning	Can it choose?	Optimization capability
Execution	Can it act?	APIs and workflow automation
Transactions	Can it buy or sell?	Interoperable protocols
Measurement	Can it prove value?	Attribution capability
Distribution	Can it reach the market?	Advertiser and inventory relationships
Feedback	Does every campaign make it smarter?	A genuine, compounding learning loop

Table 1. An eight-dimension competitive framework for evaluating agentic advertising platforms. This article deliberately does not assign numerical scores to individual companies against this framework.

This article proposes that the durable competitive moat emerging from this framework resembles a flywheel rather than a static asset: artificial intelligence draws on proprietary data, which builds an audience and inventory graph, which powers a decision engine, which depends on transaction infrastructure, which is validated through measurement, which generates feedback, which produces

more data, better decisions, more transactions, and ultimately more revenue. Because the underlying language model can become progressively commoditized, this article's position is that the model itself is not, on its own, a durable moat; the feedback loop connecting these eight layers is considerably harder to replicate quickly.

14. Four Types of Competitors

This article proposes a four-category typology of competitive posture, deliberately framed in qualitative rather than numerical terms.

Figure 3. Four Types of Competitors in the Agentic Advertising Race
(a categorical framework describing strategic posture, not a numerical ranking)



Figure 3. A four-category typology of competitive posture in the agentic advertising race. Company examples are illustrative; a given organization's underlying business and AI architecture can differ substantially even within the same category.

The fast rabbits are companies possessing substantial capital, large existing datasets, established distribution, and meaningful existing transaction infrastructure; their principal advantage is scale already achieved, and their principal risk is assuming that this advantage transfers automatically into the agentic era without further work. The persistent turtles are companies that may lack comparable AI budgets but are quietly building proprietary data, vertical-specific intelligence, specialized agents, and

supply infrastructure; their advantage is focus, and their risk is insufficient scale. The specialists, which this article terms foxes, attack a single layer, such as measurement, identity, audience intelligence, creative, inventory, or optimization, exceptionally well; their opportunity is to become the specialist agent every other participant in the ecosystem needs, and their risk is being absorbed into a larger competitor's operating system. The infrastructure builders are, in this article's assessment, potentially the most consequential category: rather than attempting to build the single smartest agent, they are building the protocols, application programming interfaces, schemas, identity systems, transaction mechanisms, trust frameworks, and interoperability standards that allow agents built by different organizations to interact reliably at all. The IAB Tech Lab's AAMP initiative is, in this article's view, the clearest industry-level example of exactly this fourth category of work (IAB Tech Lab, 2026).

15. A New Productivity Metric: Economic Output per Human

This article proposes that agentic advertising introduces a metric distinct from the revenue growth, cost per thousand impressions, return on advertising spend, margin, take rate, and total advertising spend the industry has measured for years: economic output per human. Consider two hypothetical organizations, one employing one hundred campaign operators to manage ten thousand campaigns, and a second employing twenty-five strategists to manage the same ten thousand campaigns. The second organization has not necessarily eliminated human involvement; it has multiplied the decision-making capacity of each remaining human. This article's position is that this shift could plausibly affect operating margins, achievable campaign volume, customer acquisition capacity, turnaround time, inventory utilization, revenue per employee, and overall profitability, and that this reframing moves the AI conversation from a technology story toward a business-model story.

16. A More Careful Framing of Workforce Impact

This article deliberately avoids the simplified claim that artificial intelligence adoption straightforwardly causes layoffs. A more precise interpretation is that artificial intelligence changes the composition of advertising work: certain repetitive tasks become automated and certain operational roles become smaller, while demand may simultaneously increase for artificial intelligence engineers, data engineers, agent architects, measurement specialists, AI governance professionals, commercial strategists, domain experts, and integration engineers. This article's position is that the more informative organizational question is therefore not simply how many people a company employs, but how much advertising intelligence each employee is able to produce, which this article considers a fundamentally different productivity equation than the one the industry has historically used.

17. A Hypothetical Illustration of the Agent-to-Agent Future

This article closes its forward-looking analysis with a hypothetical, illustrative scenario rather than a description of any currently operating system. A brand specifies a business objective directly to a

campaign agent: a five-million-dollar budget, a target audience of affluent consumers, an objective of maximizing incremental reach, a priority on retail proximity, and a defined frequency ceiling. That agent in turn queries an audience agent regarding where the target consumers can be found, a spatial agent regarding what journeys those consumers take, an inventory agent regarding which available screens can reach them, a pricing agent regarding cost, a forecast agent regarding expected delivery, an optimization agent regarding the most efficient portfolio, a creative agent regarding which execution should run where, and a measurement agent regarding whether the resulting campaign worked, before an activation agent executes it. This article emphasizes that under this architecture, humans do not disappear from the workflow; they move upward within the decision stack, from manually constructing a media plan toward defining the objective and constraints the resulting system of agents operates within, consistent with the human approval and guardrails already present in the one concrete, currently operating example this article has documented, the Broadsign and Draft Digital campaign described in Section 11.

18. Six Signals to Watch Rather Than One Prediction

This article deliberately declines to predict which company will win this transition, and instead proposes six measurable signals worth tracking as this shift continues to unfold.

- 1 AI-derived revenue: whether companies begin disclosing specific, audited revenue or annual recurring revenue directly attributable to AI products, following the precedent Salesforce's Agentforce ARR disclosure has already established (Salesforce, 2026).
- 2 Autonomous actions: how much of the advertising workflow agents can genuinely execute, as opposed to merely recommend.
- 3 Agent-to-agent transactions: whether independently developed buyer and seller agents can actually transact with one another, which this article considers potentially more consequential than the introduction of another AI-branded interface.
- 4 Revenue per employee: whether artificial intelligence adoption is measurably increasing organizational productivity, consistent with the metric proposed in Section 15.
- 5 Margin expansion: whether automation is translating into genuine economic leverage rather than remaining a cost center.
- 6 Data ownership: which organization ultimately owns the proprietary data that makes a given agent more capable, a question this article considers likely to grow more, not less, important as the underlying AI models themselves become cheaper and more widely available.

19. A Note on Reporting Period Discipline

This article treats the conflation of actual results, forward guidance, and informal speculation into a single undifferentiated figure as a meaningful source of error in commentary on this sector, and adopts an explicit convention throughout: a figure is described as an actual, reported result only for the specific period in which the company concerned has reported it, and is otherwise identified explicitly as guidance. The Trade Desk has reported second-quarter 2026 results and provided third-quarter guidance (The Trade Desk, 2026a, 2026b). HubSpot has reported second-quarter 2026 results and provided third-quarter and full-year guidance (HubSpot, 2026a, 2026b). Magnite has reported second-quarter 2026 results and provided third-quarter and full-year expectations (Magnite, 2026a, 2026b). Adobe has reported third-quarter fiscal year 2026 results because its fiscal calendar differs from a standard calendar year (Adobe, 2026). Salesforce has reported second-quarter fiscal year 2027 results, for the quarter ended July 31, 2026 (Salesforce, 2026). Publicis Groupe has reported first-half 2026 results (Publicis Groupe, 2026). This article's research cut-off is September 20, 2026, and no figure in this article should be read as extending beyond the specific reporting period cited for it.

20. Limitations

This article is a conceptual and strategic analysis rather than an empirical study, and its central metaphor, the rabbit and the turtle, is offered as an interpretive framework rather than as a predictive model. The financial figures presented in Sections 3 through 9 and Figure 1 are independently verified against company earnings releases, regulatory filings, and earnings call transcripts, and Section 19 specifies the exact reporting period underlying each figure; readers relying on this article for a current investment or strategic decision should nonetheless verify these figures against each company's most recent disclosures, given the pace of change evident throughout 2026. The five-dimension framework in Section 12, the eight-dimension matrix in Table 1, and the four-category typology in Section 14 and Figure 3 are this article's own proposed frameworks, deliberately presented without numerical company scoring for the reasons stated in Section 13, and should be read as a structured way of asking better questions about this sector rather than as an audited ranking. The hypothetical agent-to-agent scenario in Section 17 is explicitly illustrative and does not describe a currently operating system, though it is grounded in the one concrete, documented example of agentic OOH execution presented in Section 11.

21. Conclusion: Perhaps the Rabbit Never Really Lost

The rabbit in the fable was faster and more capable, and could have won the race as originally defined. He stopped, and the turtle, who never did, crossed the finish line first. But this article's closing argument is that the most important event in the fable was not that the turtle defeated the rabbit; it was that the race became memorable precisely because the expected outcome did not occur. This article proposes that the same dynamic is available to the advertising industry now. If every company builds an AI campaign planner, building an AI campaign planner is no longer a source of differentiation. If every platform can optimize a campaign automatically, automated optimization is no longer a source of differentiation. If every demand-side platform can automate media buying, automating media buying is

no longer a source of differentiation. Under this logic, differentiation moves one layer deeper, toward who holds the proprietary data, who controls the inventory, who owns the transaction, who owns the resulting feedback loop, who establishes the protocol other agents are required to speak, and, ultimately, who builds the track upon which everyone else is running. That, this article concludes, is the more precise description of the actual competitive contest now underway across AdTech, MarTech, OOH, and DOOH, and it is consistent with the reading of the fable this article opened with: the winner is not always the one who reaches the finish line first. Sometimes, the winner is the one who changes what the finish line means.

References

- Adobe Inc. (2026). *Adobe reports record third-quarter fiscal year 2026 results* [Press release, Form 8-K]. U.S. Securities and Exchange Commission. <https://www.sec.gov/Archives/edgar/data/796343/000079634326000147/adbeex991q326.htm>
- Broadsign. (2026). *Broadsign's sell-side AI agent joins with Global Netherlands and Draft Digital to deliver first-ever end-to-end agentic AI-powered OOH campaign* [Press release]. <https://broadsign.com/blog/broadsigns-sell-side-ai-agent-joins-with-global-netherlands-and-draft-digital-to-deliver-first-ever-end-to-end-agentic-ai-powered-oo-h-campaign/>
- HubSpot, Inc. (2026a). *HubSpot reports Q2 2026 results* [Press release]. <https://www.morningstar.com/news/business-wire/20260805185463/hubspot-reports-q2-2026-results>
- HubSpot, Inc. (2026b). *HubSpot (HUBS) Q2 2026 earnings call transcript*. <https://www.fool.com/earnings/call-transcripts/2026/08/12/hubspot-hubs-q2-2026-earnings-call-transcript/>
- IAB Tech Lab. (2026). *AAMP: Agentic Advertising Management Protocols*. <https://iabtechlab.com/standards/aamp-agentic-advertising-management-protocols/>
- Magnite, Inc. (2026a). *Magnite reports second quarter 2026 results* [Press release]. <https://www.nasdaq.com/press-release/magnite-reports-second-quarter-2026-results-2026-08-05>
- Magnite, Inc. (2026b). *Magnite lifts 2026 growth guidance to 13-14% as CTV gains 36%*. PPC Land. <https://ppc.land/magnite-lifts-2026-growth-guidance-to-13-14-as-ctv-gains-36/>
- PPC Land. (2026a). *PubMatic's AI narrative can't hide what the numbers actually show*. <https://ppc.land/pubmatics-ai-narrative-cant-hide-what-the-numbers-actually-show/>
- Publicis Groupe. (2026). *Publicis Groupe: First half 2026 results* [Press release]. <https://www.publicisgroupe.com/en/news/press-releases/publicis-groupe-first-half-2026-results>
- PubMatic, Inc. (2026). *PubMatic announces first quarter 2026 financial results* [Press release]. <https://investors.pubmatic.com/news-releases/news-release-details/pubmatic-announces-first-quarter-2026-financial-results>
- Salesforce, Inc. (2026). *Salesforce delivers record second quarter fiscal 2027 results* [Press release]. <https://www.salesforce.com/news/press-releases/2026/08/26/fy27-q2-earnings/>
- The Trade Desk. (2026a). *The Trade Desk reports second quarter 2026 financial results* [Press release]. <https://investors.thetradedesk.com>
- The Trade Desk. (2026b). *The Trade Desk introduces Kokai Zuma, the latest release of Kokai* [Press release]. <https://www.thetradedesk.com/press-room/the-trade-desk-introduces-kokai-zuma-the-latest-release-of-kokai>