

DIGITAL OUTLOOK · Q4 2026 INTO 2027

Digital Outlook 2027.

What changed across design, development, and digital marketing in 2026, and what we would do about it in 2027.

A POINT OF VIEW ON 2026,
ORGANIZED BY DISCIPLINE,
WITH A SCORECARD ON
LAST YEAR'S FORECASTS
AND TEN CALLS WE WILL
GRADE NEXT YEAR.

Contents.

00	How to read this	p. 03
01	The scorecard	p. 04
02	Executive summary	p. 05
03	The macro picture	p. 07
04	Design	p. 10
05	Development	p. 13
06	Search and organic visibility	p. 16
07	Paid media	p. 19
08	Agentic commerce	p. 27
09	The thread that ties it together	p. 35
10	What to prioritize in 2027	p. 32
11	Ten calls for 2027	p. 33
12	How we think about the year ahead	p. 34
13	Sources and methodology	p. 35




A point of view rather than a vendor pitch.

This is a point of view rather than a vendor pitch. It sets out what actually changed across digital design, web development, and digital marketing during 2026, then says where we think 2027 is heading and what we would do about it. We have organized it by discipline because that is how most teams are structured.

One argument sits above the disciplines. Through most of 2026, artificial intelligence read the web, summarized it, and recommended what to buy. By the end of the year it had started doing the buying. That shift, from a layer that interprets your brand to a layer that transacts with it, connects every section that follows.

We have flagged industry data by source so you can weigh it yourself. Where our own numbers sharpen a point we have included them in *From our portfolio* callouts, aggregated and anonymized across our managed accounts, never any single client's identifiable figures. Where the evidence is genuinely unsettled, we say so. There are two such cases this year, and being honest about them is more useful than false certainty.

We have also opened with something most outlooks skip, which is a scorecard on the forecasts the industry published in January. If a document is going to make predictions about 2027, it should be willing to show how the last round of predictions held up.



01 · THE SCORECARD

How January's forecasts held up.

The industry entered 2026 with a set of confident numbers. Most of them moved, and the direction of the misses is instructive: forecasters were too conservative about spend and too aggressive about how fast automation would deliver.

WHAT THE JANUARY CONSENSUS SAID	WHAT HAPPENED BY SEPTEMBER	READ
US ad spend would grow roughly 9.5% (IAB, January)	The IAB raised its full-year forecast to 12.3% on September 10, after a stronger than expected first half	▲ TOO CONSERVATIVE
Social would grow 14.6% and CTV 13.8% (IAB, January)	Revised up to 16.5% and 15.6% respectively	▲ TOO CONSERVATIVE
Meta would overtake Google in net ad revenue (eMarketer, April)	Still a projection rather than a banked result, on a worldwide gap of roughly \$4B	● UNRESOLVED
AI Max would deliver about 14% more conversions (Google, at beta)	Google now publishes an average of 7% for the full feature suite, and independent studies disagree sharply	▼ REVISED DOWN
Ranking in the top 10 drives roughly 75% of AI Overview citations	Later readings cluster far lower, in the range of roughly 17% to 38%	▼ REVISED DOWN

HAPPY COG'S TAKE

The pattern in those misses is worth more than any single number. The market was underestimated because cyclical demand held up better than anyone expected, and the automation gains were overestimated because vendor figures are measured under conditions most advertisers do not have. We would treat platform-published performance claims the way we treat a forecast, which is as a directional input with a named source and a margin, not as a number to plan against.

Five shifts defined 2026.

01 The market ran hotter than the forecasts.

The IAB lifted its US growth estimate from 9.5% to 12.3% in September, with social at 16.5%, connected TV at 15.6%, and commerce media at 13.6%. Paid search, at 8.1%, is now growing at roughly half the rate of social.

02 Automation stopped being a choice.

Google began auto-upgrading Search campaigns using campaign-level broad match and automatically created assets to AI Max on September 1, having closed off the creation of those legacy settings a month earlier. Dynamic Search Ads follow in February 2027.

03 Advertising moved inside the assistants at commercial scale.

OpenAI reported ChatGPT Ads reaching a \$1B annualized run rate in under 200 days, with self-serve buying open across dozens of markets by the end of August.

04 Product data became a distribution channel.

A stack of competing commerce protocols now defines how AI agents read catalogs and execute purchases, and a catalog that is not machine-readable is invisible to all of them.

05 Measurement became the bottleneck.

The IAB found 86% of buyers changing or expecting to change how they measure performance because of conversational AI and agents, and 45% naming the reconciliation of AI-driven and traditional journeys as their hardest problem.

OUR OVERALL READ

2027 is not a year to chase new features. It is a year to make your brand *transactable*.

Structured content an AI can cite, product and service data a machine can parse, clean first-party data feeding the algorithms that now spend your budget, and enough human judgment in the loop to say no to most of what the platforms recommend. Those inputs are what every AI surface rewards, and they are the only part of the stack you still control.

03 · THE MACRO PICTURE

The money tells the clearest story.

The money tells the clearest story, and in 2026 it told a better one than anyone predicted. The IAB now puts US ad spend growth at 12.3% for the year, nearly three points above its January projection, lifted by the Winter Olympics, the FIFA World Cup, the US midterms, and easing macroeconomic anxiety among buyers.

Growth concentrated where it has been concentrating for three years. Social, connected TV, and commerce media took the increases. Linear TV declined again, and paid search grew at 8.1%, which is healthy in isolation and slow relative to everything around it.

12.3%	9.5%	8.1%
US AD SPEND GROWTH, SEPTEMBER REVISION	THE JANUARY PROJECTION	PAID SEARCH, ROUGHLY HALF THE RATE OF SOCIAL

FIGURE 01

SOURCE: IAB

US channel growth, January forecast versus September revision.

CHANNEL	☐ JANUARY ☒ SEPTEMBER	JAN	SEP	CHANGE
Social media		14.6%	16.5%	+1.9 pts
Connected TV		13.8%	15.6%	+1.8 pts
Commerce media		12.1%	13.6%	+1.5 pts
Digital video excluding CTV		9.6%	9.4%	-0.2 pts
Paid search		8.2%	8.1%	-0.1 pts
Digital out of home		7.4%	7.0%	-0.4 pts
Linear TV	DECLINING	-1.7%	-1.5%	+0.2 pts

What that means for planning 2027

The single most consequential fact about 2027 is that the cyclical lift disappears. The Olympics, the World Cup, and the midterms all sat inside 2026, and anyone building a 2027 plan against 2026 comparables is measuring against an inflated base.

Forecasters currently put 2027 global growth somewhere between roughly 5% and 8% depending on the house, and we would show clients that spread rather than a single number, because the honest answer is that the range reflects real disagreement.

The practical implication is unglamorous. If your 2027 targets assume you repeat a World Cup year, rebuild them. Strip the event-driven demand out of your 2026 baseline before you set a growth number, and expect the categories that rode those events hardest to face the toughest comparisons.

5% to 8%

FORECAST RANGE FOR 2027 GLOBAL GROWTH,
DEPENDENT ON THE HOUSE


HAPPY COG'S TAKE

Diversification is still the strategy, and concentration is still the risk. We advise clients to treat social as the backbone of digital reach without making it the whole spine, to keep testing CTV while the cost-to-attention gap is open, and to treat commerce media as a performance layer rather than a brand-building one. The brands that outperform in 2027 will not be the ones with the biggest budgets. They will be the ones who set a clean 2027 baseline, fed accurate data into the right channels, and verified that automation was actually earning its place.

FROM OUR PORTFOLIO · AGGREGATED & ANONYMIZED

Across 56 managed Google Ads accounts, representing roughly \$23.8M in measured spend, we compared channel mix over three windows: July to December 2025, January to June 2026, and July 1 to September 21, 2026. Four things stand out.

SHARE OF MEASURED SPEND

□ JUL TO DEC 2025 ■ JUL 1 TO SEP 21, 2026


i. Performance Max overtook Search. Performance Max rose from 38.7% of measured spend to 45.3%, while Search fell from 52.7% to 44.9%. In the current window it is the largest single channel in our book for the first time, and the crossover happened inside nine months.

ii. Automated campaign types are approaching half of all spend. Performance Max and Demand Gen together went from 42.1% of measured spend to 48.4%. The budget is following the automation, not just the controls inside individual campaigns.

iii. Standard display is collapsing rather than declining. Average monthly display spend fell from roughly \$8,500 to roughly \$2,200, a drop of about 74%, and display now accounts for one tenth of one percent of measured spend. The industry has described display as being relegated for two years. In this book it is nearly gone.

iv. Shopping gained share, from 3.7% to 5.9%. That is the second largest mix shift in the set, and it points the same way as the commerce protocol story: budget is moving toward the formats that run on a structured product feed.

-74%

AVERAGE MONTHLY DISPLAY SPEND

\$23.8M

MEASURED SPEND, 56 ACCOUNTS

\$1.47M → \$1.62M

AVERAGE MONTHLY SPEND

(Average monthly spend rose across the three windows, from roughly \$1.47M to \$1.62M, so these are mix shifts rather than an artifact of a shrinking base. The current window runs through September 21 and is therefore partial. Figures cover Google Ads only, across 56 of 58 active client accounts.)

04

Design.

Last year we described machine experience as the idea to watch, meaning the recognition that a website has a second audience which is not human. In 2026 it stopped being an idea and became a line item.

What 2026 taught

i. Structure became commercial rather than technical.

Strict heading hierarchy, semantic markup, accurate labeling, and clean component logic used to be quality signals that a good team insisted on. They are now the difference between a brand an AI can cite and one it cannot see, which moves them from the engineering budget to the revenue conversation.

ii. The human-authorship reaction sharpened.

As generated polish became free, polish stopped signaling competence. The studios doing interesting work leaned harder into visible evidence of real effort: real photography of real places, raw process, deliberate imperfection, and specificity that a model would not have produced on its own.

iii. Personalization moved into the layout.

Real-time adaptation of messaging, calls to action, and page structure based on visitor behavior shifted from a flourish to an expectation, particularly on commerce and lead generation sites.

iv. Typography kept replacing imagery.

Oversized, viewport-scaled, and kinetic type continued to displace hero stock photography, which is an aesthetic choice that also happens to cut page weight and improve load times.

v. Performance became a design constraint rather than an engineering cleanup.

Core Web Vitals now measure how quickly a page responds to a tap, not just how quickly it paints, which makes interaction cost a decision the designer makes rather than one the developer inherits.

What we expect in 2027

The tension between expressive surfaces and rigid underlying structure will sharpen rather than resolve. The visible web will keep getting bolder and more obviously human-made, while the structure beneath it gets stricter and more machine-legible. The studios that win will hold both at once, producing work distinctive enough to be remembered by people and structured enough to be understood by models.

We also expect trust to become an explicit design problem. When anyone can generate a site that looks like a Fortune 500 company in an afternoon, looking credible stops being evidence of being credible, and design has to carry proof rather than polish.

Design has to carry proof rather than polish.

HAPPY COG'S TAKE

Brand expression and machine readability are two requirements of the same build, and we design for the person and the crawler in the same pass. A site that looks distinctive but cannot be parsed by an AI is an incomplete deliverable, and so is a perfectly structured site nobody remembers. The discipline is doing both without letting either slide, and in 2027 the second half of that sentence is the one most teams will get wrong, because structure is now measurable and taste is not.

05

Development.

If 2025 reorganized development around AI assistance, 2026 was the year the industry started counting the cost of doing it without governance.

What 2026 taught

i. Adoption kept climbing and trust fell.

Developer confidence in the accuracy of AI-generated output dropped year over year even as usage rose, driven by real quality and security incidents reaching production. That gap between how much teams use these tools and how much they trust them is the defining engineering management problem of the year.

ii. Agents moved into the build at production scale.

Coding agents are now opening pull requests in volume rather than suggesting completions, which changes review from a quality step into the primary control surface.

iii. Governance became the thing that separates shipping from stalling.

Gartner expects more than 40% of agentic AI projects to be cancelled by 2027, citing unclear value, cost overruns, and weak risk management. The teams actually getting agents into production are the ones that treated permission boundaries, decision logs, and human approval checkpoints as launch enablers rather than paperwork.

iv. The fundamentals held.

Server-first rendering, TypeScript as the professional default, composable and edge architectures, and security handled inside the framework rather than as a final-stage audit all continued to consolidate.

40%+

AGENTIC AI PROJECTS GARTNER
EXPECTS TO BE CANCELLED BY 2027

What we expect in 2027

The gap widens between organizations that treat AI as a lifecycle capability under human judgment and those treating it as a productivity hack. We expect buyers to start asking about governance during vendor selection, and we expect the first genuinely expensive public failure of an under-governed agent deployment to do more to change industry behavior than any amount of analyst commentary.

Speed without governance is a liability rather than a win.

HAPPY COG'S TAKE

Speed without governance is a liability rather than a win. We are happy to let AI compress timelines, but only behind human review, automated checks, and rollback discipline. The compounding advantage in 2027 is not raw output, it is shipping fast and safely on semantic, accessible, performance-first foundations that serve users and the AI systems now reading every page. When something goes wrong in an AI-assisted build, the question a client will ask is who reviewed it, and a team that cannot answer that has a process problem rather than a tooling problem.



06

Search and organic visibility.

The results page is no longer the destination. It is the answer, and 2026 was the year the industry stopped arguing about that and started trying to measure it.

What 2026 taught

$\sim 2/3$ US GOOGLE SEARCHES THAT END WITHOUT A CLICK	76% BUYERS NAMING AI ANSWERS THEIR TOP PRIORITY	48% MEASURING CITATIONS INSIDE AI TOOLS	45% UNABLE TO RECONCILE THOSE SIGNALS
---	--	---	---

i. Zero-click became the structure rather than a trend.

Roughly two thirds of US Google searches now end without a click, and on queries where an AI Overview appears that figure is far higher. Any content strategy still modeled on organic traffic returning to its former levels is built on an assumption that no longer describes how search works.

ii. The measurement gap started closing, unevenly.

The IAB found 48% of buyers now measuring brand visibility and citations directly inside AI tools, which is real progress from a standing start a year ago. It also found

45% unable to reconcile those signals with the journeys they have measured for years, which is the harder half of the problem.

iii. Optimizing for AI answers became the top priority for

buyers, cited by 76%, ahead of using generative AI inside campaigns. The question shifted during the year from how marketers can use AI to how they reach a consumer who is using it.

iv. Quality became a stated concern.

Anxiety about low-quality AI-generated content was named by 38% of buyers as a media investment challenge, which creates an opening for brands willing to publish things a model could not have written.

FROM OUR PORTFOLIO · AGGREGATED & ANONYMIZED

25 MANAGED GA4 PROPERTIES · JUL 2025 – SEP 22, 2026

4.2x

MONTHLY ASSISTANT SESSIONS, H1 2026
TO NOW

1,404

CHATGPT SESSIONS THIS QUARTER, FROM
68 IN H2 2025

159


GEMINI SESSIONS, FIRST MEASURABLE
MONTH

Assistant-referred sessions ran flat from the second half of 2025 into the first half of 2026, averaging 144 and then 156 sessions a month, before rising to roughly 656 a month between July 1 and September 22. The engine mix inverted along the way. Perplexity supplied most assistant traffic through 2025 and has fallen in absolute terms since, while ChatGPT went from 68 sessions across all of H2 2025 to 1,404 in a partial quarter, and Gemini registered its first measurable referrals in September. Assistant traffic remains well under 1% of sessions, so the level is still small and the slope is the finding.

(GA4 relabelled the ChatGPT source during 2026; figures combine both labels. The current window is partial, running through September 22.)

WHERE THE EVIDENCE IS GENUINELY UNSETTLED

SHARE OF AI OVERVIEW CITATIONS EXPLAINED BY TOP-10 RANKING



Last year we flagged a conflict over how much of AI Overview citation is explained by conventional ranking, with one analysis putting the overlap near 75% and a re-reading putting it closer to 38%. A year later the lower estimates have more support, and some readings go lower still. We are not treating that as

settled, because the studies use different query sets and different sampling windows, and the honest position is that the relationship is real and its magnitude is contested. Our working view has not changed: ranking still matters a great deal, and it is no longer sufficient on its own.

HAPPY COG'S TAKE

This is still SEO in its evolved state rather than SEO replaced, and we do not believe in abandoning proven fundamentals to chase unproven AI-specific tactics. Clean technical foundations, clear entity and topic structure, and genuinely authoritative content are the same things that make a brand citable by a model. On top of that we add structured answers, FAQ and schema markup, content engineered at the passage level, and off-site reputation strong enough that AI systems are confident recommending you. That is why we have evolved our SEO offering into a combined SEO and AEO practice, and why “are we part of the AI answer?” is a reporting line for us rather than an afterthought.

07

Paid media.

Last year we wrote that platform recommendations are designed to increase spend rather than results, and that disciplined teams outperform by treating suggestions as inputs rather than instructions. Ten weeks later Google removed the choice for a large share of Search campaigns, which is a useful test of whether that position was principled or merely comfortable.

What 2026 taught

FIGURE 02

THE AI MAX MIGRATION, IN STAGES

EARLY AUGUST 2026 Creation of campaign-level broad match and legacy automatically created assets closed	SEPTEMBER 1, 2026 Campaigns using those settings auto-upgraded to AI Max, with defaults unreviewed	FEBRUARY 2027 Dynamic Search Ads migration, deferred from 2026
---	--	--

- i. The AI Max migration became mandatory in stages.

Google stopped allowing creation of campaign-level broad match and legacy automatically created assets in early August, auto-upgraded campaigns using those settings to AI Max through September, and deferred the Dynamic Search Ads migration to February 2027. Advertisers who did not change settings before September 1 were upgraded with defaults they had never reviewed.
- ii. The performance evidence is contested, and the contest is the point.

Google publishes an average of 7% more conversions or conversion value for the full AI Max feature suite, revised down from the 14% cited at beta. Independent analyses have pointed the other way, including work on retail campaigns reporting return on ad spend materially below classic match types, and separate findings that search partner placements performed considerably worse than core Search.
- iii. Meta continued automating end to end,

with Advantage+ absorbing more targeting, placement, and creative decisions, and generative creative tools moving into the campaign build itself.
- iv. Advertising arrived inside the assistants as a real channel rather than an experiment.

ChatGPT Ads reached a reported \$1B annualized run rate inside a year, with self-serve buying open across dozens of markets, contextual targeting based on conversation rather than cross-web tracking, and reporting that is still thin enough to make attribution genuinely difficult.

Reconciling the AI Max evidence

Both sets of numbers can be true, and the reconciliation is the whole argument. Automation optimizes toward whatever your conversion data tells it to chase. An account with clean, deduplicated, business-relevant conversion tracking gets a system that pursues the right outcome more efficiently than a human could. An account with duplicate events, low-value actions counted as primary, or a stale goal left switched on gets the same system pursuing the wrong

outcome with total confidence and no hesitation. The variable in those studies is rarely the algorithm. It is the quality of the signal underneath it.

That is why the first thing we did when the migration notices went out was audit conversion actions rather than campaign settings.

14% → 7%

GOOGLE'S AI MAX CLAIM, BETA TO FULL SUITE

\$1B

CHATGPT ADS ANNUALIZED RUN RATE, REPORTED

HAPPY COG'S TAKE

Treat paid media as a precision tool rather than a growth shortcut. We lean into automation where it earns its place, and we keep humans firmly in the loop: clear guardrails on what the AI can and cannot do, brand and non-brand kept separate, auto-apply recommendations reviewed rather than accepted by default, and a relentless focus on feeding the algorithms clean first-party data. When a platform removes a control, the work does not disappear. It moves upstream, into the conversion data and the constraints you set before the system starts spending.

FROM OUR PORTFOLIO · AGGREGATED & ANONYMIZED

GOOGLE ADS SEARCH, SEPT 1-22, 2026

12 of 44

SEARCH ACCOUNTS WITH AI MAX ON

68%

OF SEARCH SPEND IN THOSE ACCOUNTS

6 of 9

ACCOUNTS WHERE AI MAX CPA RAN LOWER

Where we run both, AI Max campaigns posted a lower cost per conversion than the same account's non-AI Max campaigns in six of nine accounts, roughly matched in one, and ran higher in two, including the account that carries 72% of all AI Max spend. The comparison is not a controlled test: AI Max tends to be switched on for stronger campaigns first, and conversion definitions differ by account.

Programmatic moves up the plan

The Google Display Network is shrinking as a line item. As Performance Max and Demand Gen absorb display inventory, advertisers have less say over where standalone display runs and less visibility into what it bought. Reporting on placements, frequency, and audience composition has thinned at the same time the controls have been folded into automated campaign types.

Programmatic buying through a demand-side platform answers most of those gaps. You choose the supply, set the frequency, bring your own audiences, apply your own block lists, and see line-level delivery. The same buy can reach display, native, online video, streaming audio, and connected TV from one plan, with one set of rules for brand safety and one view of reach.

The money leaving GDN should not follow Google by default. It should go where the controls are.

i. Supply quality is now the job.

Made-for-advertising sites, spoofed inventory, and low-attention placements still take a share of open-exchange spend. Curated deals, seller-level block lists, and regular inventory sweeps are the difference between programmatic that performs and programmatic that only delivers impressions.

ii. Audiences travel with you.

First-party segments, contextual targeting, and partner data can be activated across channels in a DSP, which makes programmatic more resilient to the signal loss described on p. 25 than platforms that depend on their own logged-in data.

HAPPY COG'S TAKE

Plan 2027 display dollars as a programmatic budget with a Google component, rather than the other way around. We keep Performance Max and Demand Gen where they earn it, and we move prospecting and retargeting display into DSP buys where we control the supply, the frequency, and the reporting.

FROM OUR PORTFOLIO · AGGREGATED & ANONYMIZED

MANAGED DISPLAY SPEND, GOOGLE ADS AND YAHOO
DSP · JUL 2025 – SEP 22, 2026

84.6% → 98.8%

SHARE OF DISPLAY SPEND RUNNING
THROUGH A DSP

54.4% → 65.8%

VIEWABILITY ON DSP DISPLAY

\$3.73 → \$4.42

DSP DISPLAY CPM

PERIOD	DSP SPEND	GOOGLE DISPLAY	DSP SHARE
Jul–Dec 2025	\$278,896	\$50,816	84.6%
Jan–Jun 2026	\$897,440	\$38,523	95.9%
Jul 1–Sep 22, 2026	\$471,244	\$5,909	98.8%

The shift described on this page has already happened in our own book. Display spend running through a demand-side platform rose from 84.6% of managed display investment in the second half of 2025 to 98.8% in the current window, while average monthly Google display spend fell from roughly \$8,500 to \$2,189. Viewability on the DSP side rose from 54.4% to 65.8% over the same period, and CPMs rose from \$3.73 to \$4.42, which is what buying curated supply rather than open-exchange inventory costs. We paid more per thousand impressions and more of them were seen.

HAPPY COG · DIGITAL OUTLOOK 2027

(Google Ads display across 56 managed accounts; Yahoo DSP display across the managed book. The current window is partial, running through September 22.)

VOL. 02

Connected TV as a performance channel

Connected TV has been bought and judged as an awareness channel: reach, frequency, completion rate. That framing is out of date. Household-level matching and deterministic identity now connect a TV exposure to what happens afterward on a phone or laptop in the same home, and more platforms offer outcome-based optimization and pricing, including cost per visit and cost per acquisition.

The measurement problem is structural, since nobody clicks a TV. The response is to measure what CTV changes: exposed-household site visits, branded

search volume, direct traffic, and conversion rates in other channels. Holdout and geo tests separate real lift from demand that would have arrived anyway, which matters because attribution alone tends to over-credit CTV when search and social are also running.

CTV should be held to the same accountability as search and social, measured on its own terms.

Visits

QUALIFIED SITE VISITS FROM EXPOSED HOUSEHOLDS

Lift

BRANDED SEARCH AND DIRECT TRAFFIC, TEST VS. CONTROL

CPA

OUTCOME-BASED BUYING WHERE THE PLATFORM SUPPORTS IT

HAPPY COG'S TAKE

We buy CTV programmatically so it sits in the same plan, the same audiences, and the same frequency caps as the rest of the media. Before launch we set a baseline for branded search, direct traffic, and conversion rate, then optimize toward qualified visits rather than completion rate. From 2027, every CTV flight ships with a holdout or geo test so the lift is measured rather than assumed. Our 2026 flights did not, which is why the callout below reports delivery and not value.

FROM OUR PORTFOLIO · AGGREGATED & ANONYMIZED

YAHOO DSP CONNECTED TV · JAN 1 – SEP 22, 2026

\$171,573

CTV SPEND, 2026 TO DATE

97.7%

COMPLETION RATE ON 9.1M IMPRESSIONS

0

FLIGHTS THAT CARRIED A HOLDOUT

Connected TV took \$171,573 of managed programmatic spend in 2026 to date, delivering 9.1 million impressions at an \$18.94 CPM with a 97.7% completion rate, and accounting for roughly three quarters of all DSP video investment. What we cannot tell you is what it was worth. None of these flights ran with a holdout or geo control, so the lift is not measurable, and the conversions attribution assigns to CTV are exactly the over-crediting this page warns about. We are reporting the gap rather than working around it, and closing it is the first thing on our own 2027 list.

(Yahoo DSP connected TV delivery, January 1 through September 22, 2026. Attributed conversion figures are deliberately omitted.)

Early read on the August Smart Bidding change

FIGURE 03

THE ROLLOUT, IN STAGES

JULY 6, 2026	AUGUST 17, 2026	AUGUST 27, 2026
Bid Target Adjustment Tool released to review affected campaigns	Budget-limited Target CPA and ROAS campaigns begin bidding to target	Global rollout complete and fully live

Before August 17, a campaign limited by budget on Target CPA or Target ROAS often beat its target, because a capped budget forced the system onto the cheapest auctions. Google’s own example is a \$10 Target CPA delivering \$5. That campaign now delivers closer to \$10. The change covers Search, Shopping, Performance Max, and Demand Gen (including Target CPC on Demand Gen), and Google does not adjust targets or budgets for you.

The early data matches the mechanism. Published median CPC analysis in mid-September showed CPCs rising on budget-limited campaigns and a pool of cheaper clicks moving to campaigns that are not budget-limited. For a clean read, compare August 1 to 16 with August 28 onward and leave the rollout window out.

A target that was never binding is binding now. Stale targets have become the most expensive setting in the account.

HAPPY COG’S TAKE

Every target-based campaign should carry a target set from what the business can afford per qualified lead or sale, not from a number entered at launch and left alone. We reviewed budget-limited campaigns against actual performance before the change, reset targets one at a time, and are now reading post-change results with conversion lag accounted for before making a second adjustment.

FROM OUR PORTFOLIO · AGGREGATED & ANONYMIZED 56 MANAGED GOOGLE ADS ACCOUNTS · AUG 1-16 VS AUG 28-SEP 22, 2026

+18.5%

SEARCH COST PER CLICK

+15.3%

SEARCH COST PER CONVERSION

-4.8%

PERFORMANCE MAX COST PER CLICK

Search cost per click rose 18.5% and Search cost per conversion rose 15.3% comparing the two weeks before the rollout with the four weeks after it. Performance Max cost per click moved the other way, falling 4.8%, which matches the published account of cheaper clicks migrating toward campaigns that are not budget-limited. Shopping was close to flat at plus 2.1%.

(Not a controlled test. The windows differ in length, and cost per click and cost per conversion are rate metrics rather than totals. Seasonality and conversion lag are not controlled for.)

Cookie consent and the changing signal landscape

Third-party cookies are not leaving Chrome. Google dropped the deprecation plan and is retiring the Privacy Sandbox APIs, while Safari and Firefox continue to restrict tracking regardless. The pressure on advertising signal now comes from regulation and enforcement rather than the browser.

In the US, the model is opt-out rather than opt-in, and the mechanism is converging on Global Privacy Control, a browser signal that tells a site not to sell or share personal data. Twelve states required businesses

to honor universal opt-out signals as of August 2026, and new comprehensive state laws took effect on January 1. Enforcement has moved from disclosure to operation: regulators are testing whether the backend matches the banner, and settlements have followed when opt-outs did not work consistently across services and devices.

A consent banner is only as good as the tags that obey it.

i. Consent decides what the platforms can see.

Consent Mode and conversion modeling fill some of the gap when users decline, but modeling needs volume and correct configuration to be reliable. Lower-volume accounts feel the loss first, which feeds straight into how well Smart Bidding performs.

ii. First-party data carries more of the weight.

Server-side tagging, enhanced conversions, and offline conversion imports from the CRM are how you keep a trustworthy signal while honoring the choices users make.

HAPPY COG'S TAKE

Treat consent as part of the measurement stack, not a legal checkbox. We audit whether tags actually respect consent and GPC signals, configure Consent Mode correctly, move key conversions server-side, and import qualified outcomes from the CRM so the platforms optimize on data you are allowed to use and can trust.

Channels to think about for 2027

Google and Meta will hold most of the budget in 2027. The channels below are worth a test budget and a clear question to answer, chosen for audience fit rather than novelty.

i. Reddit Ads

Reddit has become a performance platform. Advertising revenue grew 64% year over year to \$762 million in Q2 2026, performance-oriented revenue passed 60% of ad revenue earlier in the year, and the active advertiser count grew more than 75%. People arrive with specific questions and strong opinions, which suits considered purchases where trust comes before the click. Creative has to sound like the community, and a poorly matched ad gets said so in public.

ii. Assistant and AI answer ads

Ads inside ChatGPT and other assistants are real inventory now (see p. 20). Targeting is contextual and reporting is thin, so tests should be small, with conversion tracking in place before launch.

iii. Retail and commerce media

For brands sold through retailers or marketplaces, retail media offers closed-loop sales data that open-web channels cannot match.

iv. Streaming audio and podcasts

Available through the same DSPs as display and CTV, which makes it straightforward to add to a programmatic plan and measure alongside it.

64%

REDDIT AD REVENUE GROWTH, Q2 2026
YOY

\$762M

REDDIT Q2 2026 ADVERTISING
REVENUE

60%+

SHARE OF REDDIT AD REVENUE FROM
PERFORMANCE

HAPPY COG'S TAKE

New channels earn budget by answering a specific question for a specific audience. We scope each test with a fixed budget, a single success measure, and a decision date, so a promising channel scales and an unproven one is retired without drifting in the plan.

08

Agentic commerce.

The most consequential shift of 2026 is the one fewest marketing teams have budgeted for. Product data stopped being back-office plumbing and became a storefront.

The mechanism is straightforward. When a person shops through an AI assistant, the assistant reads structured product data to decide what to recommend, and increasingly to move the shopper toward a transaction. A catalog with clean, complete, machine-readable data can be surfaced. One without it cannot, and no amount of brand equity compensates, because the system is not making an aesthetic judgment. It is parsing a feed.

What 2026 taught

i. Several competing protocols now define how agents read catalogs, authenticate, and execute payments, backed by different combinations of AI platforms, payment networks, and retailers. No single standard has won, and most brands selling across multiple surfaces will end up supporting more than one.

ii. The rollouts have consistently underdelivered against their announcements. Integrations promised at launch arrived late or partially, and the working settlement through 2026 has been that AI handles discovery while merchants keep checkout and remain the merchant of record.

iii. The card networks moved, which matters more than the protocol announcements. When payment infrastructure starts supporting agent-initiated transactions, the constraint stops being technical feasibility and starts being merchant readiness.

iv. Our own spend data points the same direction. Shopping grew from 3.7% to 5.9% of measured spend across our managed Google Ads book this year, the second largest mix shift in the set. Shopping runs on a product feed, and budget is moving toward the formats that structured data makes possible.

THE AI HANDLES	THE MERCHANT KEEPS	THE CONSTRAINT BECOMES
Discovery	Checkout	Merchant readiness


WHERE THE EVIDENCE IS GENUINELY UNSETTLED

This is the least reliable topic in the document, and we would rather say so than pretend otherwise. Most published material on agentic commerce is vendor content, the sources contradict each other on basic facts including which checkout features are currently live, and the picture changes monthly. Every claim in this section is dated, and we would encourage you to check the current state before acting on any of it. We have deliberately left out several widely repeated figures that we could not trace to a primary source.

Be ready for all of them, which mostly means fixing the feed.

HAPPY COG'S TAKE

The work here is not speculative and it is not exciting. Clean your product and service data, make it complete and structured, and keep it current. That is the same work that improves Shopping performance, feeds retail media, supports on-site search, and makes your pages citable in an AI answer, which means it pays for itself regardless of how the protocol competition resolves. We would not advise a client to build for a specific protocol in 2027. We would advise them to be ready for all of them, which mostly means fixing the feed.



09 · THE THREAD THAT TIES IT TOGETHER

The same shift appears in each discipline.

Read the four disciplines side by side and the same shift appears in each. In design, the crawler is a second audience and structure is how you speak to it. In development, AI builds much of the product and governance is how you keep it trustworthy. In organic search, AI decides who gets cited and authority plus structure is how you earn it. In paid media, AI runs the campaign and your data and creative are the inputs that steer it. In commerce, AI recommends the product and your feed is what it reads.

DIAGRAM A		YOUR BRAND → THE AI LAYER → YOUR AUDIENCE
DISCIPLINE	WHAT THE AI DOES	HOW YOU STEER IT
Design	The crawler is a second audience	Structure
Development	AI builds much of the product	Governance
Organic search	AI decides who gets cited	Authority plus structure
Paid media	AI runs the campaign	Your data and creative
Commerce	AI recommends the product	Your feed

A year ago we described this as AI becoming the layer between a brand and its audience. That was right, and it is now incomplete. Through 2026 the layer stopped only interpreting and started acting, which changes what is being asked of you.

THE 2026 PROBLEM

Being legible to machines.

THE 2027 ONE

Being *transactable* by them.

Being legible to machines was the 2026 problem. Being transactable by them is the 2027 one.

The practical consequence is that the old boundaries between building a website, writing content, and running media have dissolved. A brand's AI visibility depends on a fast, semantic, accessible build, content structured to be cited, an authoritative off-site reputation, clean first-party data, and a machine-readable catalog, all at once. That is why we increasingly bring these capabilities to a problem together rather than in sequence.

And it points to the durable investment for the next several years, which is authority as distribution. AI systems surface the brands they are most confident are credible. A consistent, trustworthy, well-structured presence across the web is no longer just a brand asset. It is how the machines decide whether to put you in front of a customer at all.

There is a counterweight worth stating plainly, because most of this year's commentary leaves it out. As systems take over execution, the brief becomes the product. Every automated system returns what its inputs tell it to pursue, which means constraint setting, data quality, and creative judgment are no longer the preamble to the work. They are the work.

As systems take over execution, the brief becomes the product.

10 · WHAT TO PRIORITIZE IN 2027

Where we would start.

If we were setting the agenda for the year, this is where we would start. None of it requires chasing the newest feature, and all of it strengthens the fundamentals every AI surface now rewards.

01 Rebuild your 2027 baseline before you set targets.

Strip the Olympics, World Cup, and midterm demand out of your 2026 comparables so you are planning against real underlying growth.

03 Instrument your AI visibility.

Track whether and how your brand appears in AI answers across your category, and get a baseline now so the trend is available when it matters.

05 Fix your product and service data.

Complete, structured, current feeds are the prerequisite for every agentic surface and they improve conventional performance while you wait.

07 Put guardrails around ad automation.

Review platform recommendations rather than auto-applying them, keep brand and non-brand separate, and verify incrementality rather than trusting platform-reported lift.

09 Make creative the differentiator.

As targeting and bidding automate, distinctive creative that is human-led and AI-accelerated is where the remaining edge lives.

02 Audit your conversion data before you audit your campaigns.

Automation now spends your budget against whatever your tracking says matters, so the tracking is the single most consequential thing in the account.

04 Make your site machine-legible.

Audit technical health, semantic structure, schema, and accessibility together, because those serve users, search, and AI citation at the same time.

06 Invest in off-site authority.

Earn credible brand mentions rather than links alone, since AI systems weigh reputation built across the broader web.

08 Move budget toward provable performance.

Keep testing CTV while the attention-to-spend gap is open, treat commerce media as a performance layer, and stop funding standard display out of habit.

10 Govern your AI-assisted builds.

Pair faster delivery with human review, monitoring, and rollback discipline, and be able to name who reviewed what.

Ten calls, to be graded.

These are specific enough to be graded a year from now, which is the point.
We will score them in the next edition.

SCORED IN
VOL. 03

-
- | | | |
|----|--|--------------------------|
| 01 | Google completes the Dynamic Search Ads migration to AI Max roughly on the February 2027 schedule, without a further deferral. | ☐ |
|----|--|--------------------------|
-
- | | | |
|----|---|--------------------------|
| 02 | At least one of the major agentic commerce protocols consolidates into another or is effectively abandoned. | ☐ |
|----|---|--------------------------|
-
- | | | |
|----|--|--------------------------|
| 03 | AI-referred sessions remain below 5% of total sessions for most mid-market brands through 2027, while converting at a materially higher rate than average organic traffic. | ☐ |
|----|--|--------------------------|
-
- | | | |
|----|--|--------------------------|
| 04 | AI visibility reporting becomes a standard line in agency reporting rather than a paid add-on. | ☐ |
|----|--|--------------------------|
-
- | | | |
|----|--|--------------------------|
| 05 | A meaningful number of marketers miss 2027 targets set against inflated 2026 comparables, and the misses cluster in categories that leaned hardest on event-driven demand. | ☐ |
|----|--|--------------------------|
-
- | | | |
|----|--|--------------------------|
| 06 | A significant public incident involving an under-governed AI agent changes enterprise procurement language faster than analyst guidance has. | ☐ |
|----|--|--------------------------|
-
- | | | |
|----|---|--------------------------|
| 07 | Assistant advertising becomes a standard line item in mid-market media plans rather than an experimental budget, with measurement still the main complaint. | ☐ |
|----|---|--------------------------|
-
- | | | |
|----|--|--------------------------|
| 08 | Standard display continues its decline and largely disappears from performance-led plans, surviving mainly in programmatic reach buys. | ☐ |
|----|--|--------------------------|
-
- | | | |
|----|---|--------------------------|
| 09 | At least one major platform walks back an automation default after advertiser pushback, in the way Google restored Dynamic Search Ads creation in 2026. | ☐ |
|----|---|--------------------------|
-
- | | | |
|----|--|--------------------------|
| 10 | Conversion data quality becomes an explicit agency selection criterion, because clients learn the hard way that automation amplifies whatever the tracking tells it. | ☐ |
|----|--|--------------------------|
-

React to every announcement, or build what compounds.

The temptation in a year like this is to react to every announcement, and we would push the other way. The platforms will keep shipping, the acronyms will keep multiplying, and the demos will keep impressing. Underneath all of it, the work that compounds is unglamorous and durable. It means building fast and clean, structuring content so it can be understood and cited, keeping your data in order, earning genuine authority, and bringing taste and judgment to the places automation cannot reach.

That is also how we read our own role. AI has made execution cheaper and faster for everyone, which means the value of a partner is no longer doing the work. It is knowing which work matters, having the discipline to skip the rest, and being honest about what the evidence does and does not yet support. That is the lens we bring to every engagement, and it is the lens behind this outlook.

This document is a point of view rather than a forecast we will defend to the decimal. The figures cited are vendor and industry estimates, and we have flagged where they conflict. We would rather be useful and honest than precise and wrong.


13 · SOURCES AND METHODOLOGY

How this was put together.

This outlook synthesizes industry sources published across 2026, including platform announcements, agency analyses, and forecasts from the IAB, eMarketer, WARC, dentsu, Gartner, and others. Statistics are paraphrased and attributed to their originating organizations. All forecasts are vendor-stated models rather than audited results and should be treated as directional. Where sources conflict, we have said so in the text.

The Happy Cog figures in the *From our portfolio* callouts are aggregated and anonymized from our managed accounts. The paid media mix figures are drawn from 56 managed Google Ads accounts and cover July 2025 through September 21, 2026, with the current period partial. Two accounts were excluded for technical reasons and no individual client is identified. Figures are presented as portfolio aggregates.

INDUSTRY SOURCES	PORTFOLIO SET	WINDOW
IAB, eMarketer, WARC, dentsu, Gartner, and others	56 of 58 managed Google Ads accounts	July 2025 through September 21, 2026



PEOPLE FIRST.™

The best work starts with a good conversation. If you are excited about the year ahead, we would love to help you make the most of it.

START A CONVERSATION

happycog.com

hello@happycog.com

WHAT WE DO

Design, development, and digital marketing, working as one team from day one to bring your best ideas to life.

HOW WE WORK

Over twenty-five years of outcomes over outputs, and radical transparency along the way.