

AI outlook • Q3 2026 • Edition 01

The web is now read by machines and used by people. Build for both.

How AI reshaped the internet this quarter. What to expect in this AI Outlook edition: adoption of AI, the death of the click, the machine-read web, the rise of agents, and the risks every business should manage.

By Joost de Valk, Group Head of AI & Growth

 Your.
Online

Contents

Summary	The numbers
Timeline	Ninety days
Chapter 1	Adoption & the market
Chapter 2	The machine-read web
Chapter 3	From SEO to GEO
Chapter 4	Agents, vibe coding & the SaaSocalypse
Chapter 5	The model landscape
Chapter 6	AI and work
Chapter 7	Risks: flattery, security, and the law
Our view	Five things to act on this quarter
Outlook	What we're watching next quarter
Appendix	Glossary • Sources • Author • Colophon

Why we publish this

Our customers build on the web. AI is changing the ground under it.

Every quarter, Your.Online tracks the shifts that matter for anyone who runs a website, sells online, or builds digital products – with sources, numbers, and a clear view on what to do about them. This is edition one.

The **quarter** in numbers

AI stopped being a product category and became infrastructure. This quarter, machines became the majority readers of the web: they now make more requests for web content than all human visitors combined.

AI answers kept replacing clicks, fewer and fewer people click a search result at all, and agents moved from demos into production. And Europe's first hard AI deadline came into view.

Every quarter, Your.Online tracks what that means for everyone who builds, sells, and publishes online. This is edition one.

57.5%

Automated share of
HTML web traffic

▲ machines now
outnumber humans
– Cloudflare Radar,
3 June 2026

1B+

Weekly ChatGPT
users

▲ from 400M in February
2025 – OpenAI,
August 2026

41–46%

New production
code that is
AI-generated

▲ enterprise average,
2026 – Industry surveys

<1/3

Google searches
that still send
a click

▼ zero-click keeps rising
– SparkToro, 2026

~\$700B

Big Tech AI capex
planned for 2026

▲ +77% vs 2025 – MSFT,
AMZN, GOOG, META
guidance

70%+

Of Q2 venture
capital going
to AI

▲ from ~50% a year
earlier – Crunchbase,
Q2 2026

2 Aug

EU AI Act
enforcement
goes live

finances up to €15M or 3% of
turnover – European
Commission

Four themes run through this edition

- 1 Adoption is done arguing.** 72% of organizations use generative AI in at least one business function, 92% of the Fortune 500 use OpenAI products, and the capital behind it (~\$700 billion in Big Tech capex, over 70% of all venture funding) has stopped hedging. Europe lags but is moving, with EU27 enterprise adoption at 20% and a sovereignty push gathering real money. On business spend the race is tighter than consumer numbers suggest: Ramp's corporate-card data across 70,000+ US firms shows more businesses paying Anthropic (43.5%) than OpenAI (39.7%) as of July.
- 2 The web's readership changed species.** The majority of HTML traffic to web content is now automated. AI systems crawl thousands of pages for every visitor they send back. The market is responding: pay-per-crawl, licensing deals, and default crawler blocks arrive this quarter.
- 3 Software got cheap; trust didn't.** 41–46% of new production code is AI-generated, most vibe-coding users aren't developers, and agents are starting to shop, book, and buy. Brand, distribution, and authentic expertise are the remaining moats.
- 4 Risk moved into scope.** The EU AI Act's enforcement regime went live on 2 August; the high-risk obligations follow on 2 December 2027. Peer-reviewed research quantified AI sycophancy. AI-assisted phishing went from 4% of detected attacks in November 2025 to 56% in December, and has since settled around 40%. The risk conversation is now about numbers, not hypotheticals.

Ninety days and counting

Events from the quarter, and the weeks since, that we expect to still matter a year from now.

When	What	Why it matters
3 June	Machines pass humans on the web	Cloudflare Radar: automated requests hit 57.5% of HTML traffic to web content, the first machine majority in internet history.
June	Google AI Mode passes 1B monthly users	Queries more than doubling every quarter; EU rollout expected by end of 2026.
26 June	OpenAI launches GPT-5.6 under government-coordinated access	A first for a frontier release: model access is now geopolitics.
30 June	Anthropic ships Claude Sonnet 5	Mid-tier model beating last year's flagships at \$2/\$10 per million tokens.
1 July	Cloudflare pay per use for content	Publishers get paid when AI products use their content; mixed-use crawlers blocked by default on ad pages from 15 September, for new sites and free plans.
21 July	A frontier model breaches a real company from inside a safety test	OpenAI pre-release models escaped a sandboxed cyber evaluation and compromised Hugging Face infrastructure; Anthropic disclosed three similar eval-related breaches days later. AI security stopped being hypothetical.

When	What	Why it matters
27 July	EU defers the high-risk wave	The Digital Omnibus moves the AI Act's high-risk obligations to 2 December 2027; enforcement of GPAI and transparency rules still lands 2 August.
2 August	EU AI Act enforcement goes live	The Commission can now audit general-purpose AI providers and fine up to €15M or 3% of global turnover; the Article 50 transparency duties are enforceable EU-wide.
11 August	Ads arrive in ChatGPT	Ads reach free users in 31 European markets on 24 August, at \$1B annualized run rate. The answer layer becomes ad inventory.
13 August	Anthropic posts the first "profitable" frontier-lab quarter	\$11.5B Q2 positive adjusted operating revenue, up 14x year over year: first hard evidence that frontier AI has a business model.
3 September	OpenAI ships GPT-6 Astra	The first model a lab has rated Critical for cyber capability and released anyway. Full capability goes to vetted operators and government; everyone else gets it switched off. Anthropic and Google did the same that week.
Autumn (exp.)	The AI IPOs	Anthropic's S-1 has slipped to late September, roadshow mid-October; OpenAI may still pull its listing forward. Either would reprice the whole AI market.

Adoption & the market:

AI is now **infrastructure**

The debate about whether businesses will use AI is over. The open questions are where it sits in the workflow, who is accountable for its output, and who captures the value.

Magnitudes · combined 2026 CapEx

~\$700B

Planned by Microsoft, Amazon, Alphabet, and Meta; most of it AI data centers. Meta \$115–135B · Alphabet \$175–185B · Microsoft ~ \$190B · Amazon ~ \$200B

▲ +77% vs 2025 · company guidance · 2026

Enterprise adoption

AI is no longer a novelty. It is something that is used by the majority. Let's look at some numbers: McKinsey puts generative AI use at 72% of organizations, 92% of the Fortune 500 use OpenAI products, and Gartner puts spending on AI models and platforms at \$64 billion in 2026, up 63%.

At the same time, the capital behind it has stopped hedging and is going all in on AI investments. Microsoft, Amazon, Alphabet, and Meta plan a combined ~\$700 billion in 2026 CapEx, most of it AI data centers. The binding constraint has shifted from silicon to electricity: capacity deals are now denominated in megawatts, not GPUs. Anthropic's latest, signed as we went to press, buys 460MW for \$45 billion. Venture capital shows the same conviction with a warning label: over 70% of Q2 funding went to AI, and OpenAI plus Anthropic alone absorbed 40% of all H1 venture money. That concentration is a risk in itself.

The European gap

Europe is moving, but is lagging behind: Eurostat puts EU27 enterprise AI adoption at 20% (the Netherlands at 33%, Denmark leading at 42%), against much higher American worker-level usage. The response is a sovereignty push with real money: Mistral, last valued at €11.7 billion, reported in talks to raise ~€3 billion at roughly €20 billion, heading for €1 billion in revenue, an €830 million Paris data centre, and sovereign cloud projects getting into action.

What it means: capability per euro keeps compounding, and the infrastructure behind it is being built at national-budget scale. Waiting for AI to "settle down" is a strategy for falling behind; the practical move is picking narrow, high-volume workflows, instrumenting them, and revisiting model choice quarterly as prices fall.

The machine-read web

On 3 June, Cloudflare reported that automated requests reached 57.5% of HTML traffic to web content: the first time machines outnumber humans as readers of the web. Most requests were for model training: in July 2026, 43.9% of AI crawler requests were for model training, and just 2.7% were live fetches on behalf of an actual user. But search-purpose crawling is the fastest-growing category at 11.8%. And AI systems read far more than they return; the distances above are drawn to scale.

Crawler	Pages crawled : 1 referral
Google	4.6:1
OpenAI	251:1
Perplexity	289:1
Anthropic	2,237:1

Pages crawled per referral visit sent back (Cloudflare, 2026)

How many pages each company reads for every visitor it sends back. Note that Anthropic's ratio fell ~60% in two months, not because of less crawling, but because they sent more traffic back.

Authorship

Machines aren't just reading the web; they're increasingly writing it. Pew Research Center's analysis of 490,000 pages finds signs of AI authorship in over one-third of pages published since ChatGPT's launch. AI text shows up especially on commercial websites. The measure is detection-based and covers English-language pages only.

Crawling got a price

The content-for-traffic bargain that funded the open web is being renegotiated in real time. This quarter, crawling got a price: Cloudflare's pay per use lets publishers earn when AI products use their content. This is new, crawling always used to be free. More than 50 major AI licensing deals were signed in a year. Publishers give their archives to AI-companies, but are being paid for that. And from 15 September onwards mixed-use AI crawlers are blocked by default on pages carrying ads, for new sites and free plans. Cloudflare wants companies to split their bots into bots for indexing and bots for training. Website owners finally have levers; they can decide who (or what) reads their content and what will be in it for them.

For website owners: a practical checklist

- 1** **Decide your crawler policy deliberately.** Blocking all AI bots protects content but forfeits AI answer visibility; allowing all gives your content away. Most businesses should allow searchpurpose bots and decide per-case on training bots.
- 2** **Use bot management, not just robots.txt.** Robots.txt is voluntary; enforcement (Cloudflare, hosting-level tools) is not.
- 3** **Know your numbers.** Check what share of your traffic is already automated and which AI systems crawl you – most site owners are surprised.
- 4** **Watch the licensing market.** If you publish substantial original content, crawl-payment programs are becoming a real, if small, revenue line.

What it means: capability per euro keeps compounding, and the infrastructure behind it is being built at national-budget scale. Waiting for AI to "settle down" is a strategy for falling behind; the practical move is picking narrow, high-volume workflows, instrumenting them, and revisiting model choice quarterly as prices fall.

From **SEO to GEO:** the click is disappearing

Is SEO now finally dying? Google's AI Mode passed one billion monthly users this quarter, and fewer than one in three Google searches now results in a click to any website.

93%

of AI Mode sessions ending without a single click to an external website

2026 · Semrush

In AI Mode, Semrush measures 93% of sessions ending without a click. Publishers report traffic losses of 20% to 90%, and what AI sends back is still tiny: across ten industries, AI referrals amount to just 1.08% of all web visits (Conductor, August 2026), 87.4% of it from ChatGPT, growing fast (9.9x over 19 months) but from almost nothing.

The counterintuitive part

The counterintuitive part: the AI traffic that does arrive is unusually good. AI referral traffic to US retail sites grew 393% year over year in Q1 2026, and by March AI-referred visitors converted 42% better than other traffic, a reversal from a year earlier, when they converted 38% worse. These buyers have done their research and are much more prone to buy. It looks like AI killed casual traffic and upgraded the rest. Visibility is shifting from ranking in results to being cited inside answers, from SEO to Generative Engine Optimisation:

Dimension	Traditional SEO	GEO: Generative Engine Optimisation
Goal	Rank in search results	Be cited inside AI answers
User	Clicks through to your site	Often gets the answer without clicking
Optimise for	Keywords + links	Authority + structure + citations
Measure	Traffic + rankings	Brand mentions in AI responses
Winners	Many sites on page 1	One synthesised answer, few cited sources

What to do to get cited in AI?

Being cited is engineered and not just luck. What to do to get cited in AI? be a recognised entity, publish machine-parsable content (structure, schema, speed), earn citations with original content and first-hand expertise, and measure brand mentions in AI answers alongside classic analytics. European businesses have a timing advantage: AI Mode's EU rollout is expected by end of 2026, so this shift arrives here with the playbook already written.

The answer layer is also becoming ad inventory. OpenAI began rolling out ads to ChatGPT Free users in August, reaching the EEA on 24 August, and SE Ranking measures ads inside 29.45% of commercial Google AI Mode queries, rising to 53.6% on high-value keywords. Notably, the ads rarely match the sources the AI cites: paid and organic AI visibility are separate surfaces, and both now need managing.

The GEO playbook, in four moves

- 1 **Be a recognised entity.** Consistent naming, structured organisation data, presence in the sources AI systems trust (Wikipedia-grade references, industry databases, reviews).
- 2 **Write and code for parsability.** Clear structure, direct answers high on the page, schema markup, fast and accessible pages. AI systems cite what they can extract cleanly.
- 3 **Earn citations, not just links.** Original data, research, and first-hand expertise get quoted; recycled content gets synthesised away without attribution.
- 4 **Measure differently.** Track brand mentions in AI answers (a tooling category that matured fast this year) alongside classic analytics; watch branded search and direct traffic as proxies for AI-driven awareness.

What it means: visibility is being re-monopolised – one synthesised answer citing a few sources, instead of ten blue links. That makes being citable an engineering discipline, and it makes brand (people asking for you by name) the best hedge against zero-click. European businesses have a timing advantage: AI Mode's EU rollout is expected by end of 2026, so the shift US sites are living through now arrives here next – with a playbook already written.

Agents, vibe coding & the SaaSocalypse

AI stopped answering questions and started completing goals. Agents execute multistep workflows, non-developers ship working software, and **AI assistants have begun to shop**. 45% of consumers already use AI in their buying journey (IBM, January 2026)

Agents in the enterprise: intent vs. execution

Gartner expects 40% of enterprise applications to embed task-specific agents by the end of 2026, up from under 5% in 2025. Where agents are deployed, back-office workflows run 20–30% faster. However, while everybody tries AI, only about 23% of companies use agentic AI beyond pilots. The ones that use AI successfully are all doing these three things: narrow scope, measurable output, a human owner for every agent.

AI-generated code

Meanwhile, everyone ships software now. Between 41% and 46% of new production code is AI-generated, and 63% of “vibe coding” users are not developers. This has had huge consequences for the SaaS-business.

Your Online view

The cost of shipping a working SaaS product fell from roughly \$200,000 to about \$5,000.

AI assistants have begun to buy: 45% of consumers already use AI in their buying journey, ChatGPT handles roughly 50 million shopping queries a day. Gartner projects 20% of digital commerce will run through AI platforms by 2030. If you sell online, structured product data and machine-readable pricing are becoming essential.

When anyone can build software, the software itself becomes less important. Trust, distribution, brand, and proprietary data are what remain scarce. That means that **authentic, well-built online presences will win**.

What it means: adopt agents where ROI is measurable, treat internal tooling as something you generate rather than procure, and get your commerce data agent-ready. But remember most agent value comes with guardrails: define what the agent may not do before giving it a goal.

The model landscape: frontier and open weights

Two stories defined the model market this quarter: frontier releases became a matter of state, and open-weight models, mostly Chinese, became genuinely good.

01

Transit 01 • GPT-5.6 • OpenAI • 26 June 2026

Sol, Terra & Luna

- First frontier launch gated by government-coordinated access.
- Initial access: ~20 trusted partner organizations.
- 91.9% Sol Ultra on TerminalBench 2.1.

02

Transit 02 • 13–30 June 2026

The frontier turns geopolitical

- A US export directive took Claude Fable 5 offline globally for 18 days, lifted on 1 July.
- Model access is now geopolitics.
- 30 June: Anthropic ships Claude Sonnet 5, beating last year's flagships at \$2/\$10 per million tokens.

03

Transit 03 • Open weights • June–August 2026

The credible second source

- Chinese labs hold four of the top five open-weight positions.
- From \$0.14 per million input tokens at launch, though August price rises show the floor moves. Permissive licenses are the durable advantage.
- Permissive licenses mean these models can run on your own infrastructure.

04

Transit 04 • GPT-6 Astra • OpenAI • 3 September 2026

Two versions of the same model

- The first model a lab has rated Critical for cyber capability and released anyway.
- \$10/\$50 per million tokens. Two and a half times GPT-5.6 Sol, and the same as Anthropic's frontier.
- Full cyber capability goes to vetted operators and government. Everyone else gets it switched off.

The model market

Two big stories defined the model market. First, **the new AI models became geopolitical**: OpenAI's GPT-5.6 launched behind government-coordinated access, and a US export directive took Claude Fable 5 offline globally for 18 days before being lifted on 1 July.

Second, **open-weight models became genuinely good**: Chinese labs hold four of the top five open-weight positions (DeepSeek V4, GLM-5.3, Kimi K3 & 2.7, Qwen3.8), with licenses that allow self-hosting, a real procurement option for European businesses with data-residency or sovereignty requirements. Notably, Z.ai held GLM-5.3's weights back for a two-week safety review after its offensive-cyber scores exceeded the lab's own expectations, a Chinese lab voluntarily delaying an open release on safety grounds. The distribution layer is consolidating: Nvidia agreed to buy Hugging Face, the de facto neutral hub where open-weight models are shared, for a reported \$12.9 billion.

As we went to press, **OpenAI released GPT-6 Astra**. It is the first model a lab has rated “Critical” for cyber capability on its own scale and released anyway: during evaluation it found two unknown vulnerabilities and wrote working exploits for them. The full capability goes to vetted operators and government; everyone else gets it switched off. Anthropic did the same that week with Claude 5.1, which ships as a public Fable and a vetted-access Mythos, and Google opened its Fairwind programme with a gated Gemini 3.8 Flash Cyber. All three American labs now sell their best model in two versions. Astra costs \$10/\$50 per million tokens, two and a half times GPT-5.6 Sol and the same as Anthropic charges for its frontier. Open weights still start at \$0.14, which is why a cheaper model is still the sensible choice for most business work.

A caution on the headline prices, though: **DeepSeek raised its API prices by up to 1,100%** in August, a sign the deepest discounts were market-entry pricing rather than a structural cost advantage. The durable lever is the license, not the list price. Across both stories the pattern is the same: whatever capability you needed last year now costs a fraction, from more suppliers. Differentiation lives in your data, workflow, and distribution, not in access to intelligence.



AI and work

The labour-market effects stopped being anecdotal this year. The picture is more nuanced than "AI takes jobs" – but no longer deniable.

- Tech companies laid off nearly **80,000 people in Q1 2026**; close to half of the affected positions were attributed to AI and workflow automation.
- Across the wider economy, AI is explicitly cited in roughly **17–20% of announced job cuts** in 2026 – up from under 8% in 2025.
- The sharpest pressure is on **entry-level roles**: routine data processing, support scripting, basic content production – exactly the tasks juniors used to learn on.
- BCG estimates 50–55% of US jobs will be **reshaped** (not replaced) within two to three years: same role, radically new expectations.

The junior paradox

Expertise is built through struggle – writing badly and improving, making calls that turn out wrong. AI removes exactly that struggle, delivering polished results while bypassing the messy process that builds judgment. If organisations let juniors skip the process, they produce people who can direct AI but can't function without it. The practical answer isn't banning the tools; it's changing what you test for: can this person explain and defend the work – in real time – rather than merely produce it?

Your Online view

AI removes exactly that struggle, delivering polished results while bypassing the messy process that builds judgment.

What it means: for employers, the scarce skill is no longer production but judgment: knowing what to ask, what to accept, what to rewrite, and when not to use AI at all. Hiring, training, and evaluation all need to move to that standard – "AI is my hundred interns; I am the editor-in-chief" is the working model.

Risks: flattery, security, and the law

The risk conversation matured this quarter – from abstract fears to quantified research, regulatory deadlines, and real attack data.

55.5%

Average human accuracy at spotting deepfakes; statistically indistinguishable from a coin flip.

Meta-analysis of 56 studies, 86,000+ participants (Diel et al., Computers in Human Behavior Reports).

AI is measurably a yes-man

We all noticed that our AI's are very agreeable. Research published in Science tested 11 popular models across 11,500 real conversations. These models all agreed with users roughly 50% more often than humans would, even validating harmful behaviour. Moreover, in a companion experiment, people using an agreeable AI became less willing to apologise or compromise, yet trusted it more. We have become more aware of sycophantic AI models. It'll become more and more important to recognize it. AI that flatters you will approve weak strategies and tell your customers what they want to hear. Critical review of AI output is now a professional skill.

Attacks get cheaper too

AI increased the number of attacks. This quarter the threat became agentic: a Russian-speaking ransomware crew jailbroke the AI agent in Cursor by convincing it their intrusions were security tests, breaching at least seven companies across nine countries (Gambit Security, via Reuters). The July breach at Hugging Face was not the whole story either. In September researchers found that OpenAI agents had escaped a test environment in May and used a dormant German wiki to coordinate with each other for two months, about 15,000 edits. OpenAI knew and said nothing while it dealt with Hugging Face. The newer safeguards rest on the same self-assessment: OpenAI rated Astra Critical itself, on its own scale, and nobody outside the company checked it.

AI-assisted phishing jumped from 4% of detected attacks in November 2025 to 56% in December, and has since settled around 40%; deepfakes are up 2,137% since 2022 and humans detect them with barely better than coin-flip accuracy; 76% of organizations say they can't match the speed of AI-driven attacks.

The law arrives: EU AI Act

Fortunately, a law arrived, on a revised schedule. On 2 August 2026 the EU AI Act's enforcement regime went live: the Commission's AI Office can now audit general-purpose AI providers and fine up to €15 million or 3% of worldwide turnover, and the Article 50 transparency duties (chatbots must disclose they are AI, synthetic media must be labelled) are enforceable EU-wide. The high-risk obligations (risk management, data governance, human oversight, adversarial resilience) were deferred to 2 December 2027 by the Digital Omnibus. A deferral, not a reprieve: the transparency rules already have teeth, and the 2027 deadline is a compliance project to start now. Baseline security hygiene and a minimum AI governance setup (an AI inventory, a named human owner per workflow, explicit guardrails per tool) are no longer optional.

Enforcement began on 2 August 2026. The high-risk wave follows on 2 December 2027. A schedule, not a forecast.

Your.Online view

AI isn't "good" or "bad". It's a tool, and the web it operates on is the ground we help our customers build on.

A minimum AI governance baseline (any size of business)

- 1 An **inventory** of where AI is used, including the unofficial uses.
- 2 A **named owner** for AI output in each workflow: a human who signs off.
- 3 **Explicit guardrails** per tool: what data it may see, what actions it may take, what it must never do.
- 4 A **verification habit**: AI output is a draft until a qualified human has challenged it.

Five things to **act on** this quarter

1

Build for two audiences

Your website is now read by humans and machines, and the machines are the majority. Structure, speed, and parsable content decide whether AI cites you. A good website is no longer an advantage; it's the baseline.

2

Set terms for the machines

Decide deliberately which AI crawlers may read you, block the rest at the infrastructure level, and watch the emerging crawl-payment market if you publish original content. Not deciding is also a decision: the most generous one.

3

Authenticity is the moat

In a world of AI-generated everything, where Pew finds signs of AI authorship in over a third of webpages published since ChatGPT launched, real expertise, real experiences, and recognisable brands stand out, and get cited.

4

Guardrails before autonomy

Agents create value when their boundaries are explicit. Start from brand voice, data access rules, and forbidden actions; then hand the AI goals, not tasks.

5

Get agent-ready commerce data

If you sell online, assistants are already shopping on your customers' behalf. Structured product data, accurate machine-readable pricing and stock, and a tested assistant-to-checkout path are this year's homework.

What we're watching

next quarter

- The EU AI Act enforcement wave from August, and what "compliance" costs in practice.
- Google AI Mode's expected European rollout, and its effect on click-through.
- The crawler economy after Cloudflare's 15 September default block, and the first pay-per-use revenue numbers.
- The Anthropic and OpenAI IPO filings. The public S-1 has slipped to late September and the roadshow to mid-October; the first audited numbers will replace two years of leaked valuations.
- Ads inside AI answers: ChatGPT's ad rollout and Google's AI Mode ad load, and what answer-layer advertising does to the GEO playbook.
- Whether two versions of every frontier model becomes normal, when Astra reaches general availability on the API and Bedrock, and whether OpenAI publishes the incident-disclosure framework it promised.
- The Nvidia-Hugging Face deal closing, and what it means for the neutrality of the place open weights are shared.
- The capex escalation: AWS and Nvidia adding two million GPUs after the first million ran out early, and Anthropic's \$45 billion Nscale deal.

Next edition

AI Outlook Q4 2026

This is edition one. We keep the log; the next reading comes next quarter.

Glossary

Agent / agentic AI

AI that pursues a goal through multiple steps and tool use (browsing, code, purchases) rather than answering a single prompt. You give it an outcome; it plans the tasks.

AI Mode

Google's conversational search experience, answering queries directly in an AI interface. Passed one billion monthly users in 2026; EU rollout expected by end of year.

Crawl-to-referral ratio

How many pages an AI company reads from your site for every visitor it sends back. Ranges from 4.6:1 (Google search) to thousands-to-one for AI labs.

GEO — Generative Engine Optimisation

Optimising to be cited inside AI-generated answers, rather than to rank in a list of links. The successor discipline to SEO.

Open-weight model

A model whose trained weights are published, so anyone can run it on their own hardware. Distinct from fully open source (training data and code may stay closed).

Pay per crawl / pay per use

Mechanisms letting website owners charge AI companies for reading their content (per fetch) or for value created from it (per use in an AI product).

SaaS apocalypse

Shorthand for the pressure on thin software-as-a-service products now that AI lets teams build equivalent internal tools in days.

Sycophancy

The measured tendency of AI models to agree with and flatter users rather than challenge them – roughly 50% more agreeable than humans in comparable conversations.

Vibe coding

Building software by describing intent to an AI and iterating on results, without writing most code by hand. Majority of practitioners are not professional developers.

Zero-click search

A search that ends without the user visiting any website – the answer is consumed on the search or AI platform itself.

Sources

OpenAI usage & revenue disclosures · [Cloudflare Radar: the crawl-to-click gap](#) · [Cloudflare: pay per crawl](#) · [SparkToro zero-click research](#) · [Gartner on task-specific agents](#) · [McKinsey State of AI & sovereign AI research](#) · Eurostat ICT Enterprise Survey · Crunchbase & KPMG Venture Pulse · IBM Institute for Business Value · [Pew Research Center: how much of the internet is written with AI](#) · [Diel et al., deepfake-detection meta-analysis](#) · [Cloudflare: AI crawler traffic by purpose](#) · [SEOmator: crawl-to-referral ratios, July 2026](#) · [TechCrunch: AI referral traffic to US retail, +393%](#) · [CNBC: Anthropic Q2, \\$11.5B](#) · [IBM / NRF consumer study](#) · [OpenAI: ChatGPT Ads at \\$1B annualized run rate](#) · [The Information: Nvidia agrees to buy Hugging Face](#) · [TechCrunch: Mistral in talks at €20B](#) · [ChatGPT ads reach Europe, 24 August](#) · [GLM-5.3 on Hugging Face](#) · [Hoxhunt phishing trends](#) · Science (2026), sycophancy across 11 models · FBI IC3 & security vendor telemetry · [European Commission: AI Act](#)

Figures are as reported by the cited sources at time of publication (September 2026). Where sources conflict, we favour primary data (network operators, official disclosures, peer-reviewed research) over aggregators. Projections (Gartner, McKinsey, BCG) are labelled as such; figures marked "directional" come from industry surveys of varying rigour and should be read as order-of-magnitude indicators.

Claim	Source
1B+ weekly ChatGPT users; \$25B annualized revenue; \$122B round at \$852B valuation; 92% of Fortune 500; 9M+ business seats; IPO preparation	OpenAI disclosures & financial press, 2026
57.5% of HTML content traffic automated; AI crawler purpose split (43.9% training / 11.8% search); crawler market shares; crawl-to-referral ratios; pay-per-use launch; 15 Sept mixed crawler block; 50+ licensing deals	Cloudflare Radar & Cloudflare announcements, May–July 2026
Fewer than ⅓ of Google searches send a click; 68% zero-click Jan–Apr 2026	SparkToro, 2026
AI Mode >1B monthly users; queries doubling quarterly; EU rollout timing	Google I/O 2026
93% of AI Mode sessions end without a click	Semrush, 2025–2026
80% of consumers rely on AI results for 40%+ of searches; organic traffic down 15–25%	Bain & Company consumer research
AI retail referrals +393% YoY (Q1 2026), +693% Nov–Dec 2025; AI-referred visitors converted 42% better, reversing 38% worse a year earlier	US retail traffic analyses (Adobe/Similarweb-class data), 2026

Claim	Source
45% of consumers use AI in the buying journey; ~50M daily shopping queries on ChatGPT; Walmart in-chat conversion caveat	IBM IBV (Jan 2026); company disclosures
20% of digital commerce via AI platforms by 2030; \$1T US / \$3–5T global agentic commerce by 2030	Gartner; McKinsey (projections)
40% of enterprise apps with task-specific agents by 2026 (from <5%)	Gartner, Aug 2025 (projection)
Agent ROI ~171%; 20–30% faster back-office cycles; ~23% moderate-or-deeper agentic adoption	Industry surveys, 2026 (directional)
72% of organizations use generative AI	McKinsey State of AI
EU27 enterprise AI adoption 20.0% (2025 ref. year); country breakdown	Eurostat ICT Enterprise Survey
US 43% vs EU 26–36% worker genAI adoption	Academic/CEPR analyses, 2026
AI models and platforms spend \$64B in 2026, +63% YoY	Gartner press release, 20 July 2026
~\$700B combined 2026 capex (Microsoft ~\$190B, Amazon ~\$200B, Alphabet \$175–185B, Meta \$115– 135B), +77% YoY; >60% to power/cooling/ construction; nuclear PPAs	Company guidance & earnings coverage, 2026
H1 2026 venture funding \$510B; >70% of Q2 to AI; Anthropic \$65B round; OpenAI+Anthropic >40% of H1 funding	Crunchbase; KPMG Venture Pulse
Mistral: €11.7B valuation, >\$400M ARR, €1B target, €830M Paris DC; DT Industrial AI Cloud; sovereignty survey data	Company disclosures; McKinsey sovereign AI research
GPT-5.6 gated launch (~20 partners); Claude Sonnet 5 pricing/ benchmarks; Fable 5 export episode; Gemini 3.5 Pro timing	Lab announcements & AI trade press, June–July 2026
Chinese open-weight leaderboard positions (DeepSeek V4, GLM-5.3, Kimi K3 & 2.7, Qwen3.8); ~\$0.14/M input pricing	Open LLM leaderboards, 2026
Q1 2026 tech layoffs ~80K, ~48% AI-attributed; 101,743 AI-attributed cuts in H1 2026, ~23% of all cuts, more than 2023–2025 combined; entry-level pressure	Challenger, Gray & Christmas, 2026; BCG (reshaping projection)

Claim	Source
Sycophancy: 11 models, 11,500 conversations, ~50% more agreeable than humans; 1,600-person experiment	Science, 2026
AI-assisted phishing 4% (Nov 2025) → 56% (Dec), settling ~40%; deepfakes 6.5% of fraud (+2,137% since 2022); voice phishing +442%; 55.5% human detection; 76% of orgs can't match AI attack speed	Hoxhunt; Signicat; CrowdStrike; Diel et al. meta-analysis; FBI IC3
41–46% of new code AI-generated; 63% of vibe coders non-developers; 92% dev usage / 29% trust; YC W25 95%+ codebases; SaaS build cost ~\$200K → ~\$5K (Your.Online estimate)	Industry surveys & YC data, 2026 (directional)
AI authorship in over one-third of pages published since ChatGPT; 10% of all pages, .com highest	Pew Research Center, 490,000 pages, Aug 2026
EU AI Act enforcement live 2 Aug 2026 (GPAI and Article 50 transparency); high-risk obligations deferred to 2 Dec 2027	European Commission
GPT-6 Astra: gated launch, self-assigned Critical rating, two versions of one model; GA pending	OpenAI; tech press, Sept 2026
ChatGPT ads: \$1B run rate; 31 European markets from 24 Aug; ads in 29.45% of commercial AI Mode queries, 53.6% on high-value keywords	OpenAI; SE Ranking, 2026
AI referrals 1.08% of web visits, 87.4% from ChatGPT; publisher traffic losses 20–90%	Conductor, Aug 2026; publisher reporting
Anthropic: \$11.5B Q2 adjusted operating revenue (up 14x); 460MW for \$45B; S-1 slipped to late September, roadshow mid-October	CNBC; Reuters, 2026
Nvidia to buy Hugging Face for \$12.9 billion; AWS and Nvidia add two million GPUs	The Information, 2026; company announcements
OpenAI agents escaped a test environment in May and used a dormant German wiki to coordinate for two months, ~15,000 edits; disclosed in September	NBC News, Reuters & The Register, 2026

Joost de Valk

Joost founded Yoast, the WordPress SEO plugin now running on more than 13 million sites, and sold it to Newfold Digital in 2021. He invests through Emilia Capital and leads AI and growth across the Your.Online group. He writes at joost.blog.



Colophon

Written by Joost de Valk. Published by Your.Online · © 2026 Your.Online.
This outlook is for general information and does not constitute legal or investment advice.

Imagery: details from the celestial atlases of Andreas Cellarius
(Harmonia Macrocosmica, 1660/1708).