The future of discoverability

Why AI visibility is now a strategic marketing priority

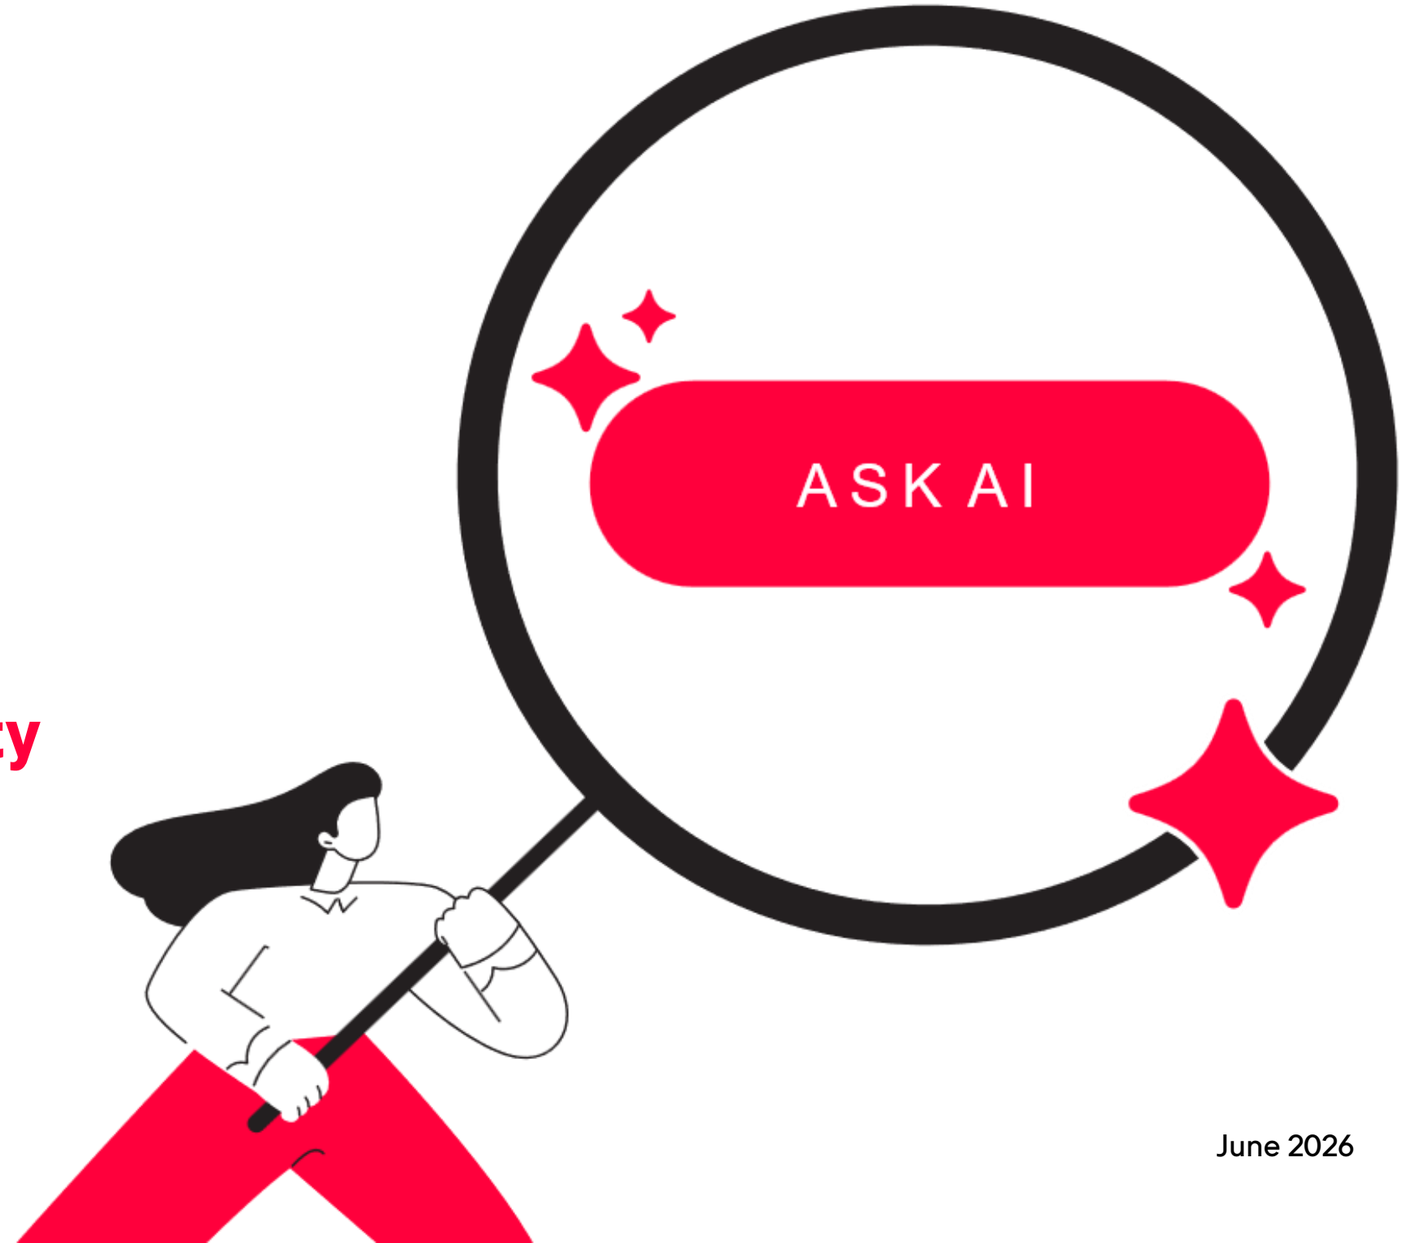

In partnership with:

DEPT®

June 2026

Contents

Introduction 3

Why AI visibility is a more useful framing than GEO optimisation 4

How AI systems form brand understanding 5

The new dimensions of discoverability 7

Governance, ownership and organisational design 11

Measuring visibility in AI-mediated environments 15

Generating internal support for AI visibility initiatives 17

Conclusion 18

Introduction

AI-mediated discovery is fundamentally changing how people use the internet, and no brand is exempt from its impact on consumer search and shopping behavior.

WFA research from 2025 found that 66% of brands expect AI-generated search experiences to reduce organic traffic to their websites. And while 65% are already reviewing their SEO approach in response, only 4% report having made significant changes so far. The uncertainty is understandable. Most existing guidance still frames this as a technical optimisation challenge: improving structured data, adapting content formatting, monitoring prompts, and optimising for AI search results.

Those actions matter, but they only address part of the picture.

Rather than build responses from a single webpage or ranking position, LLMs synthesise information across a broader ecosystem of owned and external sources. This makes discoverability a broader operational and strategic issue instead of an updated SEO workstream. Brand visibility increasingly depends on whether it's consistently understood, reinforced, and validated across the information environment AI systems rely on.

The stakes get even higher as autonomous AI agents become more sophisticated and widely adopted. Capable of completing actions on a consumer's behalf, agents will increasingly handle transactions from booking to buying to reordering with little to no human intervention.

As we move toward an agentic commerce reality, leaders must reframe discoverability as a coordinated system that expands beyond content to include media, governance, creator ecosystems, data and organisational structure.

This guide is grounded in the strategic and tactical questions senior marketing leaders and WFA members are already asking about the future of discoverability and how their brands can remain visible and relevant in an AI-driven customer journey. It explores what these shifts mean in practice, how AI systems form brand understanding, which dimensions of discoverability matter most and where organisations should focus first as AI-mediated discovery continues to evolve.

"AI-mediated discovery is fundamentally changing how consumers find and choose brands. Many marketers understand that this shift is coming, but most are still in the early stages of determining what it means operationally.

This guidance is designed to help CMOs move beyond tactical GEO discussions and start building a coordinated approach to visibility and discoverability in AI-driven environments."

Stephan Loerke, CEO, WFA

"Brands now need to answer to two audiences: humans and AI. AI increasingly shapes what consumers see, consider, and choose. This shift requires a broader approach to discoverability - one that aligns content, data, media, and governance around how brands are represented in AI-driven environments.

This guide helps marketing leaders move beyond tactical optimisation discussions and develop a coordinated approach to building visibility, trust, and influence in an AI-mediated customer journey."

Isabel Perry, Global EVP Strategy, DEPT®

Why AI visibility is a more useful framing than GEO optimisation

One of the risks that has emerged from early generative engine optimisation (GEO) conversations is encouraging organisations to think too tactically, too early.

The instinct is understandable: A new discovery environment appears, so search and SEO teams are expected to deliver a new optimisation playbook with the right formatting techniques, schema, prompt structures, and tools. Much of that work is useful, but optimisation assumes the underlying brand signals are already coherent. And for many organisations, they're not.

In AI-mediated environments, discoverability is increasingly shaped by accumulation and reinforcement.

AI systems appear to build confidence in brands the same way people often do: through repeated exposure to consistent information across multiple trusted contexts.

That has major implications for how visibility is earned. For example, a haircare brand repeatedly associated with dermatologist commentary, Reddit discussions, positive customer reviews, and TikTok tutorials develops a denser, more reinforced signal than a brand that relies primarily on its own website. On the flip side, a software provider with fragmented tooling explanations, inconsistent terminology, or conflicting guidance across touchpoints creates ambiguity that weakens trust and retrievability, even if its traditional SEO footprint remains strong.

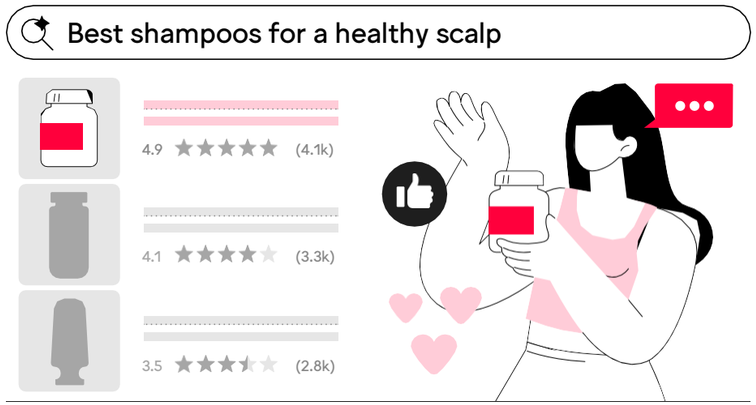

Importantly, the most influential sources for AI aren't necessarily the most polished or highest traffic. In some cases, expert commentary, niche communities, and creator-led explanations carry disproportionate influence because they align more closely to how users frame questions to AI than how a brand describes itself.

GEO tends to frame discoverability as a publishing challenge, whereas AI visibility is closer to a coordination challenge. It forces organisations to consider whether the broader ecosystem around the brand is consistently reinforcing the right signals for LLMs and agents to develop confidence in them.

That includes:

- How claims are substantiated externally
- Whether creators and communities reinforce positioning
- Whether definitions and explanations stay consistent across surfaces
- Whether product and category associations are clear
- Whether important knowledge is structured and maintainable internally
- Whether teams are actively monitoring how the brand is represented in AI-generated environments

This is also why AI visibility cannot sit neatly inside a single function. SEO may own monitoring. PR influences authority. Social and creator teams shape conversation. Product teams influence data quality. Legal and compliance govern consistency and risk. Media investment amplifies which signals gain reach and reinforcement.

In short, the necessary shift is from optimising individual pages toward strengthening the consistency, authority, and coordination of the broader brand ecosystem.

How AI systems form brand understanding

Most LLMs do not fully disclose how responses are assembled or weighted. But while the mechanics vary across platforms and models, the underlying patterns are becoming increasingly clear: AI systems are trying to determine which information appears reliable enough to include in a generated response.

That process is less about identifying the single "best" page and more about building confidence through retrieval, consistency, and reinforcement.

Most modern AI agents combine several layers simultaneously:

- Pretrained understanding
- Live retrieval from the web
- Structured data and metadata
- External validation signals
- Ongoing interpretation of entities and relationships

Because responses are synthesised rather than simply ranked, retrievability increasingly matters as much as visibility itself. A highly ranked page may still be excluded from an answer if the information is unclear, weakly corroborated, difficult to extract, or contradicted elsewhere.

This is also why inconsistency creates outsized risk in AI-mediated environments.

If a brand describes a product differently across its website, social channels, and press coverage, AI systems are left to reconcile conflicting signals themselves. The result can be diluted positioning, inaccurate summaries, weaker category association, or omission altogether.

For regulated or high-trust categories such as finance, pharma, healthcare, and consumer products, those errors can create material brand safety and compliance concerns. AI systems may surface fabricated product specifications, hallucinate or amplify negative claims, confuse one brand's product with a competitor's, or repeat outdated guidance in a context where accuracy matters. These are serious risks that need to be monitored, diagnosed, and routed through the right escalation paths.

The opposite is also true. When descriptions, definitions and positioning are reinforced consistently across multiple trusted environments, systems appear more confident in surfacing and recommending that information.

Discoverability signals at a glance

1. **Clear definitions and explanations**
   → Easier for systems to extract and reuse
2. **Structured information**
   → Helps systems interpret entities, products, and relationships
3. **Consistency across environments**
   → Reinforces confidence in positioning and claims
4. **External corroboration**
   → Strengthens perceived authority and trust
5. **Fresh, maintained content**
   → Reduces risk of outdated or conflicting responses
6. **Expert and community validation**
   → Reinforces relevance and credibility

Where does paid fit into AI visibility?

Sponsored placements in AI-generated environments are starting to emerge, with platforms including Google, ChatGPT, and Perplexity experimenting with different models. But paid placement is not a substitute for building the underlying signals that AI systems use to understand and validate a brand.

As these ad products mature, media teams will have new opportunities to test visibility inside AI-mediated journeys. The bigger shift, however, is that paid, owned, and earned media are becoming more interconnected. Owned content helps establish what the brand wants systems to understand. Earned and shared media help reinforce those claims through publishers, creators, communities, reviews, and expert voices. Paid media can amplify credible signals and support selective testing, but it cannot fully compensate for unclear owned content, weak external validation, or inconsistent brand information.

For now, brands should treat AI-native paid placements as an emerging opportunity, not the whole strategy. The more durable priority is building an information ecosystem that makes the brand easier to understand, trust, and recommend.

PROJECT SPOTLIGHT: UNDERSTANDING WHERE AI SYSTEMS LEARN ABOUT A CATEGORY

For a global beauty brand, the external narrative driven by shared media platforms seemed to have a much heavier influence on AI-generated answers than the brand's owned content. The brand partnered with DEPT® to understand how Reddit and YouTube specifically were shaping its presence in LLM outputs for skincare, haircare, and ingredient-led prompts.

The brand was preparing for a product launch tied to a fast-moving beauty trend, and needed to know which sources, creators, formats, and conversations were already shaping discoverability.

82%

YouTube appeared in 82% of AI responses

40%

UGC accounted for 40% of the top cited content

DEPT® analysed search demand, generated consumer prompts, and used AI visibility tools to understand which sources were being cited across AI responses. The findings showed how much AI discoverability was being shaped outside the brand's owned ecosystem: UGC accounted for 40% of the top cited content, Reddit led citations across the prompt set, and YouTube appeared in 82% of AI responses. YouTube was especially important lower in the funnel, appearing in 95% of bottom-funnel responses, while creator-led channels accounted for 93% of cited YouTube content.

The work showed that improving AI visibility was not just a matter of publishing more or re-optimising owned content for GEO. By identifying where category authority was already being built, they could shape a strategy around the communities, creators, expert voices, and content formats AI systems were actually using to influence their audience. The broader lesson: in AI-mediated discovery, external ecosystems are not just awareness channels. They are part of the evidence base that AI systems use to decide which brands are visible, credible, and worth recommending.

The new dimensions of discoverability

A brand's AI visibility can be understood across five connected factors: **presence, representation, authority, retrieval readiness**, and **influence**. Together, these show whether a brand is appearing in relevant AI-generated journeys, being described accurately, supported by credible sources, easy for systems to cite, and able to shape consideration or choice.

These elements should be assessed in tandem, rather than treated as a step-by-step checklist. A brand can show up often and still be misrepresented, or have strong owned content but lack a presence in third-party media. Looking across all five gives leaders a clearer way to diagnose where brand visibility is strong, where it is fragile, and where teams need to intervene.

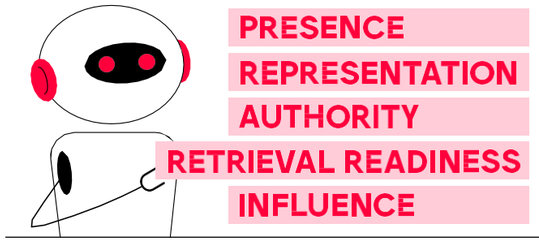

Presence: Understand where your brand is actually showing up

Presence establishes whether the brand is appearing in the AI-generated journeys where it should be eligible for consideration. AI visibility tracking tools can help by running defined prompts across multiple LLMs (such as ChatGPT, Perplexity, Claude, and Gemini) and monitoring brand mentions, citations, source URLs, sentiment, competitor visibility, and changes over time.

A strong baseline shouldn't be limited to a few branded prompts like "What is [brand]?" or "Is [brand] good?". It should cover the journeys where the brand should be eligible to appear: category exploration, problem-led research, product comparison, recommendation, troubleshooting, and purchase-oriented prompts.

The output should show where the brand appears, where competitors appear instead, which sources are cited, and how visibility varies by model or market.

AI visibility data isn't typically retroactive in the same way many SEO teams are used to with keyword tools. Once prompts are selected and added to a tracking platform, the tool can begin collecting results over time. But if the brand was not tracking those prompts last month, there may be no reliable historical record of how often it appeared, which sources were cited, or how competitors were represented.

Where to start: Establish priority prompt groups, current brand presence, competitor presence, cited sources, and gaps. That gives teams a factual starting point instead of relying on anecdotal testing.

Tell me about AI visibility tracking tools

AI visibility tracking is a new and quickly evolving category. Tools vary by model coverage, market and language support, prompt tracking methodology, citation reporting, competitor benchmarking, workflow features, and enterprise governance capabilities. The right fit will depend on a brand's priority markets, internal ownership model, reporting needs, and risk profile.

Examples of tools in this space include:

- Profound
- Peec AI
- Otterly.AI
- AthenaHQ
- Scrunch AI
- Semrush AI Toolkit
- Ahrefs Brand Radar
- Rankscale
- Writesonic GEO

This is not an endorsement list, but an example of the types of platforms emerging in the category.

Representation: Check whether visibility is helping or hurting

Representation examines whether AI systems are describing the brand accurately, consistently, and in ways that support how it wants to compete. Visibility has limited value if the brand is surfaced with vague positioning, outdated claims, missing differentiators, or incorrect product associations.

This requires a structured review of AI outputs against the brand's intended positioning. Are the right product attributes showing up? Are claims accurate? Are competitors being credited with advantages the brand also owns? Is outdated language still appearing? Are AI systems describing the brand in a way that matches how the brand wants to compete?

This is where many issues become visible. A brand may be present in the answer but described generically. A product may be recommended, but for the wrong use case. Or a discontinued claim may still appear because it still lives in legacy content or third-party articles.

Where to start: Instead of trying to "fix the AI answer" directly, trace the representation back to its source. If an outdated or inaccurate product detail is appearing, identify where the AI is pulling it from. That will determine what action to take to remedy the underlying incorrect information.

Authority: Build proof beyond owned channels

Authority shows whether the brand's claims are reinforced by credible sources beyond its own channels. Owned content can explain what the brand wants to be known for, but AI systems also look for corroboration. A claim supported only by the brand's website or press releases is weaker than one backed by relevant publishers or platforms.

The priority is to identify which sources carry weight in the category. This could be shared media platforms (e.g. Instagram, TikTok, LinkedIn, Reddit, Quora, YouTube) or earned media (e.g. digital PR, content partnerships, influencers, UGC, podcasts, customer reviews, product ratings).

The point isn't to try to shape or influence the narrative from every external channel. It is to understand which sources are already shaping AI-generated answers and whether the brand has credible coverage or participation there.

Where to start: Document what claims the brand wants to own, which sources currently reinforce those claims, which sources AI systems cite most often, and where new validation is needed.

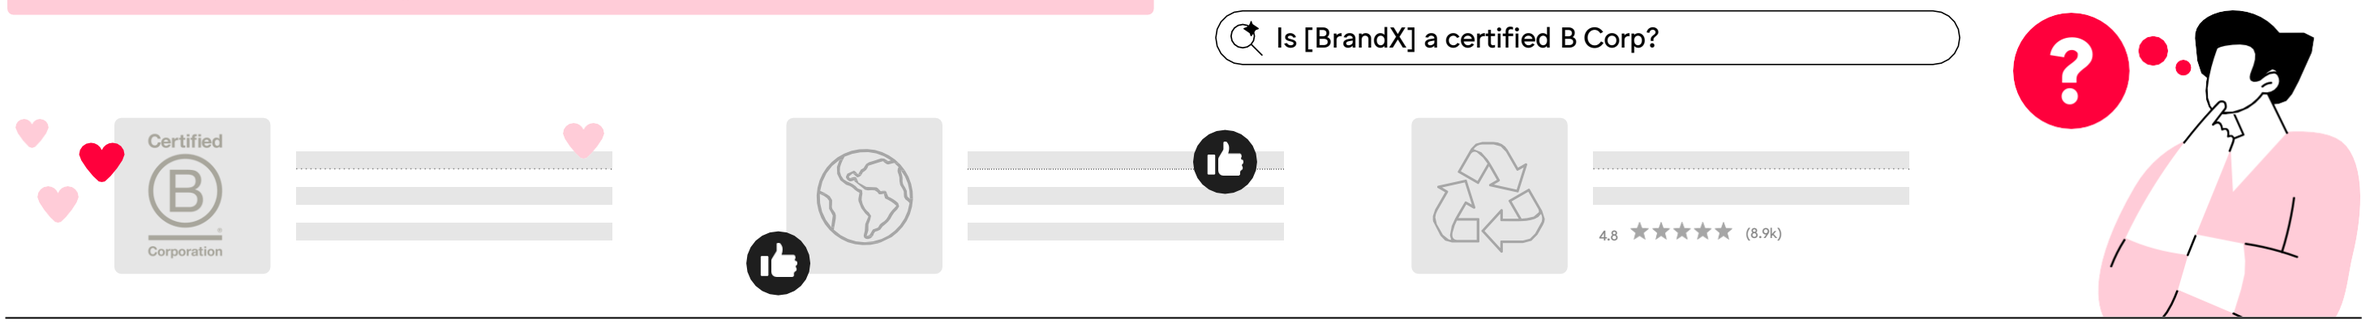

Retrieval readiness: Make the brand easier to use

Retrieval readiness examines whether AI systems can actually extract and reuse the information the brand wants to surface.

This factor relies on content and data quality. A brand may have the right information somewhere, but if it is buried in long copy, phrased inconsistently across pages, missing from schema, stored in PDFs, or different across markets and retailer listings, it becomes harder for systems to retrieve confidently.

High-priority pages should be reviewed for clear definitions, concise explanations, consistent product names, current metadata, and schema markup. In practice, structured data helps label key information for machines: this is a product, this is the price, this is a review, this is an FAQ, this is the brand name, and this is the official description. That makes important information easier to interpret and reduces the risk that it gets missed or misunderstood.

Content should be written and organised in retrievable units, with paragraphs that answer one idea clearly, bulleted or numbered lists, well-formatted tables, and clear Q&A formatting, rather than relying on a long narrative arc. This type of content chunking helps AI systems retrieve small portions of content independently from the rest of a page.

★ **Where to start:** Identify priority content, missing or inconsistent structured data, unclear explanations, duplicate or conflicting claims, and the operational owner for each fix.

Influence: Connect AI visibility to business signals

AI systems increasingly influence what people consider, trust, and choose. They narrow options, frame trade-offs, validate claims, introduce competitors and guide users' next steps, whether a brand is actively managing its AI visibility or not.

For brands, absence is only one risk. AI-generated responses can also shape behavior in ways that work against the brand: for example, recommending a competitor first, matching a product to the wrong use case, or elevating external narratives over the brand's own explanation.

That makes understanding and monitoring influence a critical dimension of discoverability. Without a clear view of how AI systems drive decisions in the category, brands have limited ability to improve the signals, sources and content guiding those decisions.

★ **Where to start:** Combine AI visibility metrics with existing brand and commercial indicators. It will not be perfect at first, but it gives the organisation a way to learn, compare, and adjust investment over time.

Together, these five dimensions turn AI discoverability from a vague visibility concern into a manageable framework. The priority now becomes assigning ownership across the teams and systems that shape each dimension.

PROJECT SPOTLIGHT: TURNING OWNED CONTENT INTO AN AI VISIBILITY ENGINE

A global consumer electronics brand wanted to improve how it appeared in AI-generated answers across priority categories and markets. The work started with prompt research and an AI Overviews gap analysis to identify where the brand had the right to show up, where competitors were being surfaced, and where existing content was not aligned with how people were asking questions in AI-mediated environments.

From there, the team developed a strategy across the brand's owned and external narrative, supported by performance, content and technical audits. The owned content work included 30 LLM-optimised articles across the US, UK, and Australia, plus a Learning Center hub designed to connect related content and make key information easier to find, interpret, and cite.

The results point to what can happen when owned content is treated as part of a broader visibility system:

6.41%
lift in LLM visibility

105%

801%
increase in referral traffic from LLMs.

The key lesson is that content volume alone is not the driver. The impact came from aligning prompt demand, content gaps, technical readiness, and a clearer hub structure around the questions AI systems were already answering.

Who's developing exciting new gadgets?

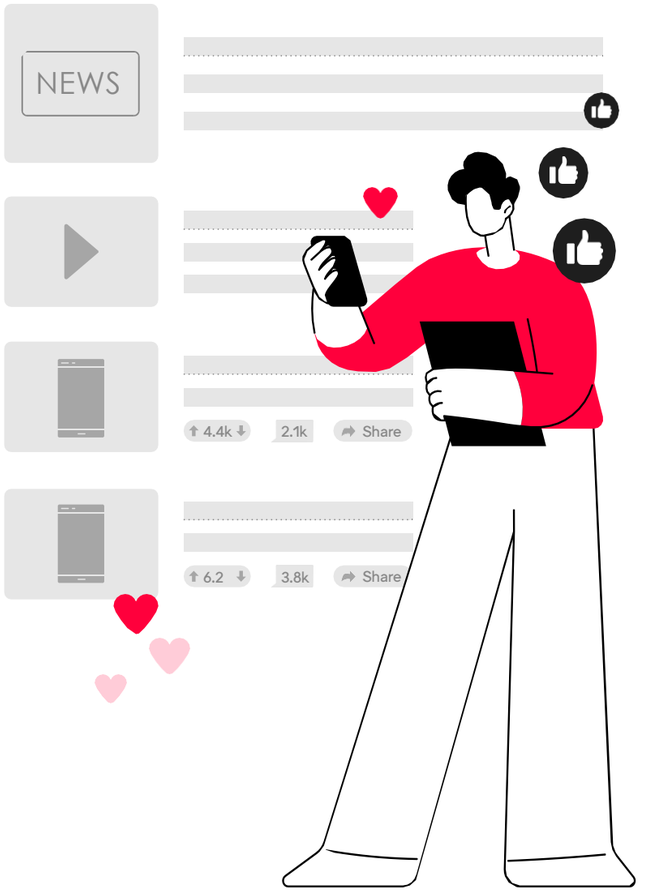

Governance, ownership and organisational design

We've established that AI visibility is shaped by many different signals. The operational challenge is that those signals are usually produced by different teams, managed in siloed systems and measured with varying KPIs.

For many organisations, this is the trickiest hurdle to contend with. The content, data, media activity and technical expertise already exist, but no one is clearly accountable for how those inputs come together in AI-mediated environments.

AI visibility needs a governance model that can coordinate across functions without creating a new bottleneck. This doesn't require building out a large GEO team or centralising every decision. The goal, rather, is to make ownership explicit, establish shared standards, and create a repeatable process for monitoring, improving, and escalating issues.

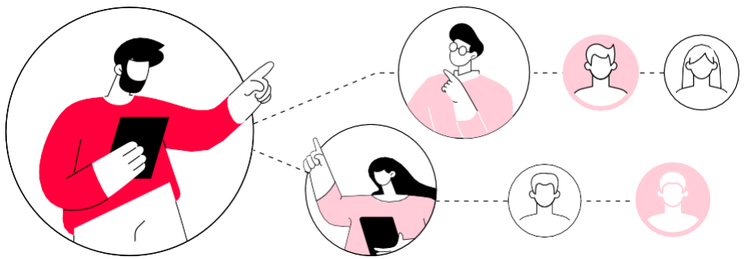

Start by assigning a clear owner

Someone needs to be accountable for AI visibility at the enterprise level, even if execution remains distributed across teams. This person or steering group's role is to coordinate the system: define the measurement approach, convene the right teams, set standards for source-of-truth content, and ensure issues are routed to the right function.

Without clear ownership, AI visibility can quickly become everyone's concern and no one's responsibility.

SEO may track prompts, PR may manage misinformation, legal may review claims, media may test new placements, and brand may update positioning, but no one is looking across the full system and ensuring all messaging is aligned.

For most organisations, the accountable owner will likely sit close to brand, media, digital, search, or marketing transformation. The exact function matters less than seniority, decision rights, and the ability to coordinate across functions.

Establish a core governance model

The most practical structure is a hub-and-spoke model.

The hub sets the shared framework: visibility reporting, prompt tracking standards, source-of-truth messaging, escalation rules, governance cadence and prioritisation. The spokes execute within their areas of responsibility, including content updates, PR activity, paid amplification, creator engagement, retailer corrections, schema implementation, local market adaptation and legal review.

This avoids two common failure modes. A fully centralised model becomes too slow because every content, media, or market decision has to pass through one team. A fully decentralised model creates inconsistent signals because each team optimises its own surface without a shared view of how the brand is being interpreted.

For global brands, the federated model becomes even more important. AI visibility varies by market, language, model, publisher environment, retailer landscape, regulation and local search behavior. Global teams should define the standards. Local teams should monitor, adapt and respond based on what is happening in-market.

Define role clarity across functions

A practical operating model for AI visibility must make each team's responsibilities explicit.



Brand Strategy Core proposition, proof points, category associations, source-of-truth messaging



Content SEO Owned content quality, AI-readable formatting, schema, structured content, source updates

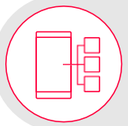

Media Paid amplification, channel signal coordination, selective AI-native media testing



PR Comms Earned authority, publisher relationships, expert validation, misinformation response



Social Creator Creator partnerships, community participation, platform-specific signal building



Analytics Prompt tracking, visibility reporting, sentiment, competitor benchmarking, influence signals



Legal Compliance Claims governance, regulated language, escalation for inaccurate or risky outputs



Commerce Retail Retailer listings, product metadata, reviews, marketplace consistency



Market leads Local prompt tracking, language adaptation, market-specific source monitoring



Tech Data Structured data, knowledge architecture, entity management, workflow integration

Build an escalation process for inaccurate or risky outputs

Monitoring AI visibility has limited value if the organisation does not know what to do with what it finds.

Brands need a clear escalation path for outputs that are inaccurate, outdated, misleading, legally sensitive, or commercially damaging. The process should start with source diagnosis to find out where the AI system is likely pulling information from. Is the issue coming from owned content, retailer metadata, or a broader absence of better information?

Once the likely source is identified, the response can be assigned to the appropriate team. The important point is that AI output issues should not be handled as one-off fire drills. They should move through a defined workflow, with severity levels, owners, response times, and follow-up monitoring.

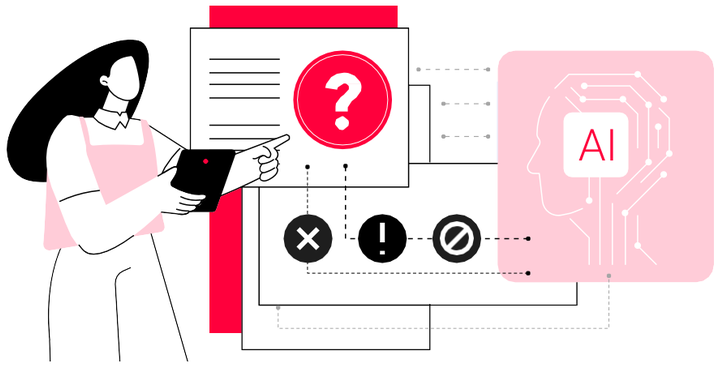

Create a shared source of truth

Governance also depends on having a clear source of truth for the facts and claims the brand wants reinforced.

For some organisations, this may begin as a structured messaging library: approved brand descriptions, product definitions, claims, proof points, disclaimers, naming conventions, and priority category associations. For more complex organisations, especially those with many products, markets, or regulated claims, it may require a more formal knowledge layer or entity-based content model.

The key is that teams should not be rewriting core facts independently across websites, retailer feeds, help centers, PR materials, or elsewhere. That creates drift, which becomes a visibility risk.

A practical place to start is with the information most likely to affect trust, recommendation, or compliance.

This includes:

- Product claims
- Regulated language
- Pricing or eligibility details
- Sustainability claims
- Category definitions
- Expert proof points
- Brand positioning statements

These should have clear owners, approved language, update processes and a record of where they appear.

Establish a governance cadence

AI visibility should be monitored on a consistent cadence, not only when performance drops or misinformation appears.

Teams should regularly review visibility trends, representation issues, priority source gaps, competitor movement, market-specific issues, and progress against agreed fixes. The output should be a short action log, not a reporting exercise: what changed, what matters, who owns the response, and when it will be reviewed again.

This keeps the work from becoming either too tactical or too abstract. It also helps teams learn which interventions appear to shift visibility, representation, or authority over time.

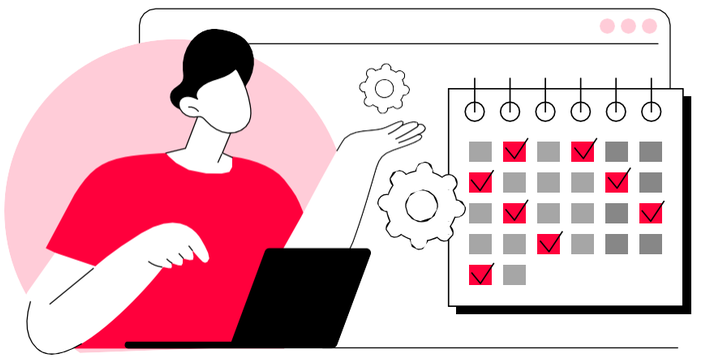

What to operationalise first

Progress happens when AI visibility is treated as a managed operating rhythm. The technical, content, and media work still matters, but it only scales when ownership, governance, and workflows are clear.

Take the first 90 days to focus on the minimum operating model, prioritising the following tasks:

- Name an accountable AI visibility lead and cross-functional working group.
- Define priority prompts, models, competitors, markets and reporting cadence.
- Create or update approved claims, product definitions, and brand proof points
- Define how inaccurate, outdated, or risky outputs are diagnosed and routed.
- Map which teams own content, PR, media, legal, retailer, and market-level fixes
- Establish a monthly review of visibility, representation, risks, and actions

PROJECT SPOTLIGHT: BUILDING THE KNOWLEDGE FOUNDATION FOR AI VISIBILITY

For a leading financial services brand, content volume and quality was not a problem. The organisation had deep expertise, strong traditional search performance, and a large footprint across multiple digital properties. The issue was fragmentation: core knowledge lived across different teams, domains, workflows, and publishing systems, making it harder to maintain consistent, authoritative answers at scale.

01

DEPT® conducted an AI visibility-focused discovery across workflows, governance, tooling, review processes, and content operations. Rather than auditing pages in isolation, the team examined how information was created, approved, updated, distributed, and measured.

02

The findings pointed to a shift from page-based publishing to a governed knowledge system, with canonical entities, structured content, reusable modules, clearer ownership, and centralised governance.

03

The early indicators showed why this matters: nearly 30,000 AI citations, a 58% brand mention rate, an 82.8% average brand citation rate in Google AI Mode, and 91% positive sentiment on non-branded prompts.

04

The broader lesson is that AI visibility is often limited by the operating model behind the content. For complex brands, improving discoverability may require less “more content” and more investment in the systems and workflows that make trusted knowledge easier to maintain, retrieve, and reuse.

Measuring visibility in AI-mediated environments

AI visibility doesn't fit neatly into existing measurement models. Influence can happen before a click, without a click, and across agentic surfaces that never appear cleanly in web analytics.

For example, a new runner tells their AI agent they signed up for their first 5k. The agent recommends a list of three shoe styles for casual running. The user searches the first of the three brands on TikTok and watches an influencer review. Satisfied, they purchase the shoes on Amazon. The AI interaction shaped the decision without creating a simple attribution path, and the shopper never visited any of the brand's owned channels.

In an agentic, zero-click environment, brands' website traffic will decline. The key shift in understanding performance is moving from measuring traffic alone to measuring quality of representation and influence upstream of the click.

Rather than try to replace SEO KPIs or organic traffic metrics with AI-specific equivalents, brands need a model that tracks three things: what AI systems show, how they describe the brand, and whether those patterns connect to downstream behavior.

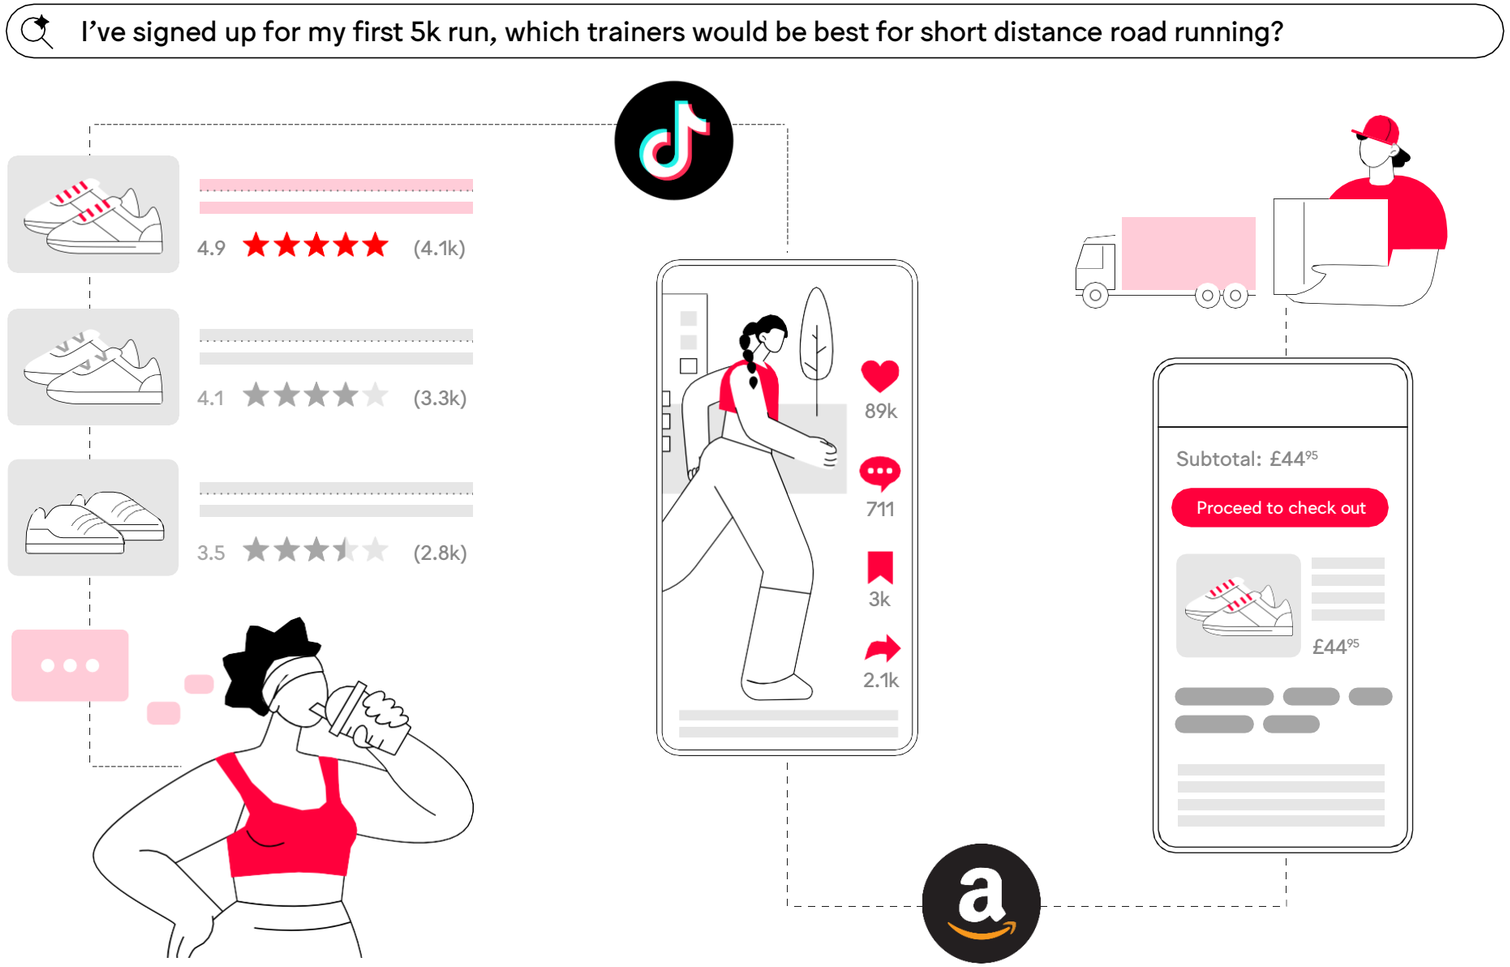

Measurement area	What to track	Why it matters
Visibility	✦ Answer presence ✦ Competitive visibility ✦ Citation share ✦ Source mix ✦ Prominence ✦ Model/market variation	Shows whether the brand is included in relevant AI-generated journeys and how visibility compares to competitors.
Representation	✦ Accuracy ✦ Sentiment ✦ Claim consistency ✦ Category association ✦ Differentiator presence ✦ Misinformation frequency	Shows whether visibility is helping the brand weakening its positioning, or introducing risk.
Influence	✦ Branded demand ✦ AI referral traffic ✦ Assisted conversion behavior ✦ Retailer/marketplace performance ✦ Brand lift ✦ Social/community language	Shows whether AI visibility appears to connect to broader brand or commercial signals.

Direct visibility should be monitored across priority prompts, models, markets, and journey stages. These signals allow you to track where the brand appears, which competitors are being surfaced, and which sources are shaping the answer.

Representation needs a more qualitative layer of review. Automated scoring can help, but human judgment is still required because an answer can be factually accurate and still strategically weak. For example, a brand may be mentioned often but described generically, connected to the wrong use case, or missing the differentiators that matter most.

Influence is the most difficult to cleanly isolate, but it can and should still be tracked directionally. Brands should look for relationships between AI visibility and broader signals such as branded search, share of search, and social or community language. Treat this as evidence-building, not proof of causality—if improved presence and representation coincide with stronger downstream signals, that pattern is useful even if attribution remains imperfect.

Generating internal support for AI visibility initiatives

It may be too soon to frame AI visibility as a fully proven performance channel. A more realistic business case, however, is about building the capability to understand, manage and improve how the brand appears in AI-mediated environments.

The absence of a mature ROI framework is real, but it's not the obstacle it appears to be. Brands have navigated this before, with social media, programmatic, content marketing, etc. In each case, the leaders who built early capability were better positioned when the channel matured. Those who waited for proof were playing catch-up.

Investing in AI visibility goes beyond tooling or technical optimisation. Much of it is time, labor, and coordination: building the right working group, setting standards, experimenting with content and media interventions, and learning what actually moves visibility over time.

That matters because the risk already exists. LLMs and AI agents are already summarising categories, comparing brands, surfacing recommendations, shaping consideration, and, in some cases, transacting for users. In short, they are directly impacting how consumers behave and make choices.

The strongest argument for investing in GEO and broader AI visibility initiatives is that it's low-regret. All of the factors from establishing clearer source-of-truth content to better structuring information to improving cross-team coordination all improve a brand's digital operating model regardless of how AI search evolves.

A practical business case should focus on a phased capability-building agenda:



Diagnose

Baseline visibility audit, prompt and journey tracking, governance mapping (low cost, high learning value)

Improve

Content and data improvements in priority areas, cross-functional ownership, measurement framework

Test and Scale

Selective pilots in AI-native environments, learning agenda, broader rollout based on evidence

This keeps the ask grounded. The goal isn't to secure a large speculative budget for "AI search," but to understand and manage an existing risk—using investments that would be defensible regardless of how the AI landscape develops.

Discoverability becomes an operating capability

AI-mediated discovery is still early, but its implications will expand rapidly over the next 12 to 24 months. As AI becomes increasingly embedded across search, commerce, voice interfaces, and brand-owned experiences, more people will rely on generative systems and autonomous agents to handle decisions and tasks on their behalf.

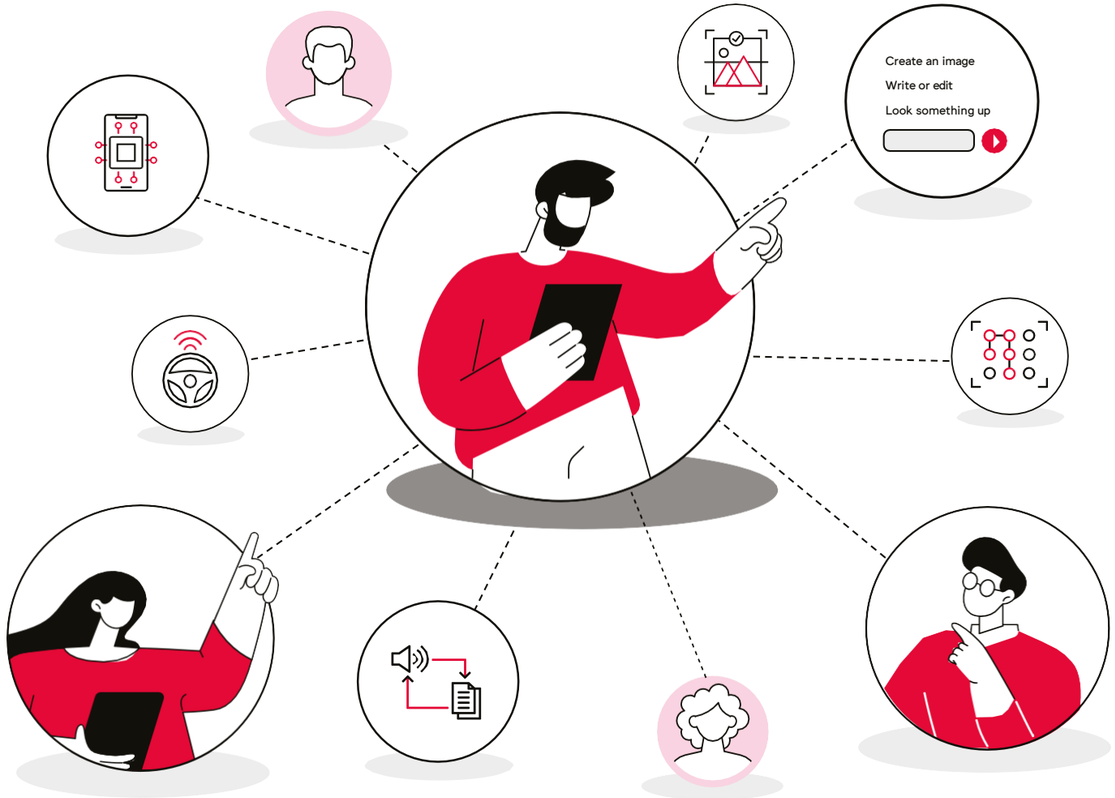

How can our organisation remain relevant and chosen?

To remain relevant and chosen, organisations must treat AI visibility as a business-wide initiative. Establishing and maintaining a presence in AI-generated environments depends on the combined strength of a brand's information ecosystem.

For many organisations, the first step is to assess where the brand fits within the current agentic landscape. From there, the work becomes more practical: improving retrieval readiness, bolstering external authority, and establishing a governance rhythm that helps teams learn what moves visibility over time.

Discoverability is only one part of the job. The more important objective is an accurate, credible, and consistent representation across the environments where discovery is now taking place.

The brands that make progress will be those that build the systems, teams, and evidence base needed to manage discoverability as it becomes more distributed, more synthesised, and more shaped by intelligent intermediaries.

About WFA and DEPT®

About WFA

WFA champions more effective, efficient and responsible marketing communications which help WFA members drive brand growth. It is the voice of marketers worldwide, representing more than 150 of the world's biggest brand owners and national advertiser associations in over 60 markets, which speak for tens of thousands of brands at a local level. Through WFA, they form a global network that offers marketers a unique source of leadership, expertise and inspiration.

If you would like to engage with WFA's AI Community, please contact:



Gabrielle Robitaille
Director of Policy
g.robitaille@wfanet.org

About DEPT®

DEPT® is a Growth Invention company built to help the world's most ambitious brands grow faster. Operating at the intersection of technology and marketing, our 4,000+ specialists deliver growth invention across Brand & Media, Experience, Commerce, CRM, and Technology & Data. We're 50|50 tech and marketing, partner-led, and first to move. Clients include Google, Lufthansa, Coach, eBay, and OpenAI. Certified B Corp and Climate Neutral since 2021.

For more information about DEPT®, please contact:



Jack Williams
SVP Growth, EMEA
jack.williams@deptagency.com



World Federation of Advertisers
London, Brussels, Singapore, New York

wfanet.org

info@wfanet.org

+32 2 502 57 40

linkedin.com/company/wfa
youtube.com/wfamarketers

WFA Competition law compliance policy

The purpose of the WFA is to represent the interests of advertisers and to act as a forum for legitimate contacts between members of the advertising industry. It is the policy of the WFA that it will not be used by any company, or individual, to further any anti-competitive or collusive conduct, or to engage in other activities that could violate any antitrust or competition law, regulation, rule or directives of any country or otherwise impair full and fair competition. The WFA carries out regular checks to make sure that this policy is being strictly adhered to.

As a condition of membership, members of the WFA acknowledge that their membership of the WFA is subject to competition law rules and they agree to comply fully with those laws. Members agree that they will not use the WFA directly or indirectly, (a) to reach or attempt to reach agreements or understandings with one or more of their competitors, (b) to obtain or attempt to obtain, or exchange or attempt to exchange, confidential or proprietary information regarding any other company other than in the context of a bona fide business or (c) to further any anti-competitive or collusive conduct, or (d) to engage in other activities that could violate any antitrust or competition law, regulation, rule or directives of any country or otherwise impair full and fair competition.