AI agents inside the marketing team

Who publishes, whether each claim has a source, and whether AI-made content is labelled

IN BRIEF

01 Marketing is the department where an agent's output is published. A page. A post. A campaign email. It goes out in the company's name, to people who cannot ask the agent what it meant. That is why the questions to ask first are not about where an agent runs. They are who publishes, whether each claim has a source, and whether AI-made content is labelled.

02 The public material differs. One vendor's help pages say its content agent "doesn't publish automatically", and that "you will review and approve all social posts before they are published or scheduled to post". Another vendor's training module describes one approval of a plan. After it, "the agent takes over, creating and publishing the campaign on its own". Public workflow templates write a post each morning and publish it with no approval step. Few of the pages we opened say where the agent runs.

03 The record explains why the questions matter. The United Kingdom's advertising regulator says that when an AI tool produces something problematic, "the advertiser remains responsible for any output used in their ads". The United States regulator has a rule against fake reviews and testimonials, which it says will "deter AI-generated fake reviews". It also acted against a company that claimed its AI product was "98 percent" accurate without the evidence. One technology news site had to correct 41 of the 77 articles its AI system wrote. One test found that platforms labelled fewer than a third of the AI-made posts put to them.

04 This paper is about the marketing team's own work, from the brief and the draft to campaigns, claims and reporting. It sets out what an agent can read, how it reaches marketing systems, what it may do without a person, and what the rules already say. It ends with seven questions for any supplier, and a place to start.

05 The position throughout is simple. An agent drafts and checks. A named person approves and publishes. No agent publishes a page, posts to a social network, schedules or sends a campaign, or replies to a comment. None places or changes an advert, or changes a budget, a bid, an audience or a consent record. No agent writes a review or a testimonial. Every factual claim in a draft carries its source. Research is about companies and markets. No agent scores a marketer or reads mood. The agent runs on infrastructure the company controls.

IN THIS PAPER

01 Where marketing teams are today 2
02 The work nobody has packaged 3
03 Three doors, one gate 4
04 What an agent may do 6
05 What the rules already ask 7
06 A worked example: one software company's marketing team 9
07 Seven questions before an agent writes in the company's name 10
08 Where to start 10

PART 01

Where marketing teams are today

41 of 77

of the 77 articles a technology news site had published from a trial of an AI engine carried corrections, as a news report of January 2023 put it

30%

of 516 posts with AI-made images and videos on five large platforms were correctly labelled as AI-generated (“just over 30%”), in a test by a specialist publication published on 23 October 2025

The suppliers have arrived, and many sell the publishing itself. Vendor pages describe agents that write articles, build campaigns and post to social channels.

They disagree on who publishes.

- HubSpot’s help article for its content agent says: “You’ll review and edit all drafts before publishing, as content agent doesn’t publish automatically.” Its page for a social media agent says “you will review and approve all social posts before they are published or scheduled to post” [1].
- A Salesforce training module on campaigns built by an agent describes one approval, of the plan. Then: “Once it’s approved, the agent takes over, creating and publishing the campaign on its own” [2].
- One published workflow template has the trigger “Every day at 6:45 AM”. A model writes a caption for each channel, and the next steps post them to the company’s pages on two social networks [3]. We found no approval step among its steps.
- A second template turns one keyword into a long article, and its last step is “Automatically post as Draft or Published” [3].
- A third does the opposite by default. “Posts are created as drafts on purpose”. The same page tells the builder how to switch “from draft to publish for fully automated posting once you trust the output” [3].
- Hootsuite’s page asks whether its agent can publish content automatically. Its answer: “No.” It can draft, “but people remain in control of reviewing, approving, and publishing content”. Microsoft says of an agent that builds customer journeys: “Nothing goes live until you approve it” [4].

So approval can mean every item, one plan, or a setting someone may switch off.

Few pages say where the agent runs. Writer’s trust page lists “Managed, multi-cloud options, including dedicated private cloud deployment”. Buffer says of its writing assistant: “the text you enter into the assistant will be shared with OpenAI” [5]. On the other supplier pages we opened, we found no statement on the point. So where it runs is still a question to ask every supplier. It is not the question that tells the products apart.

A regulator has said who answers for the output. On 11 June 2026 the United Kingdom’s Advertising Standards Authority published four points on AI for advertisers. One reads: “If an AI tool produces something problematic, the advertiser remains responsible for any output used in their ads.” Another says that “even when ads are generated or distributed through automated AI platforms, responsibility does not shift” [6].

Fake reviews are banned by rule, AI-made ones included. On 14 August 2024 the Federal Trade Commission announced a final rule on reviews and testimonials. Its announcement says the rule will “deter AI-generated fake reviews”. The rule took effect on 21 October 2024 [7].

A claim about an AI product needed evidence it did not have. On 28 April 2025 the same regulator announced a proposed order against Workado, a company that promoted its AI content detector as “98 percent” accurate. The regulator said that “independent testing showed the accuracy rate on general-purpose content was just 53 percent”. Under the order the company may make such a claim only if it “has competent and reliable evidence to support the claim at the time it is made” [8]. A figure in marketing copy needs a source before it is published.

PART 01 · CONTINUED

A publisher had to correct what its AI system wrote. In January 2023 Engadget reported on the technology news site CNET, which had published articles written in a trial of an "internally designed AI engine". Of the 77 articles, "41 feature corrections" [9].

Platform labels are not a safety net. In a test published on 23 October 2025, the publication Indicator put 516 posts containing AI-made images and videos on five large platforms. "Only 169 posts, or just over 30% of the total, were correctly labeled as AI-generated" [10].

We found no named, dated, independently measured result for a marketing agent, and this paper quotes none as proof.

PART 02

The work nobody has packaged

We listed the jobs a marketing team does and looked for each among the templates and prebuilt agents we opened. They are thick around writing and posting. We did not find them where the work is checking a draft, a claim or a campaign for a named person.

WHERE TEMPLATES OR PREBUILT AGENTS EXIST	WHERE WE FOUND NONE
A caption written each morning and posted to the company's pages	**A list of every factual claim in a draft, with its source or a note that none was found**
An article written from a keyword and posted to the website	**A check that a quotation has a named source, the person's consent and the original wording**
A campaign created and published after one approval of the plan	**A list of the assets in a campaign made with AI, and whether each is labelled**
A customer journey built by an agent and held until a person approves it	**A check before a campaign goes out for the unsubscribe link, the sender and the audience**
Web content drafted for the editor to review and publish	**A list of pages that state something the product facts no longer support**

PART 02 · CONTINUED

Two findings matter more than the list. First, in the templates that post, the words reach the public before anyone on the marketing team has read them [3]. Whether a person reads them is a setting, and whoever chooses it has decided what the company will answer for.

Second, the careful products show the better pattern. A draft lands in front of the editor, and nothing is published until a person publishes it [1][4].

The gaps sit where the checking is slow and the decision already belongs to a named person. That makes them the safer place to begin.

PART 03

Three doors, one gate

An agent is only as useful as what it can reach. In a marketing team that means three kinds of thing.

Documents.

The brand and style guide, approved messaging, product facts, the claims register, past content. These are indexed once as knowledge sources, with each reader's permissions kept, so an agent answers with the passage it relied on and says when the documents are silent.

Systems of record.

Pages, campaigns, assets, events, registrations. Most are reached through an interface the company already runs. The company gives the platform its description, chooses which operations become tools, and decides which need approval.

Public sources and figures.

Company registries, filings and published research are reached through web search. Campaign results, spend and plan come from a reporting database that is read only.

All three pass one gate: every tool call is checked against the company's policy, which can allow it, record it, hold it for a person, or block it.

Where the company's interface checks the caller, the agent acts as the marketer using it.

SIX WAYS INTO A MARKETING TEAM'S SYSTEMS

	ROUTE	WHAT IT REACHES	WHAT TO KNOW
01	**A file the user uploads**	A brief, a draft, an interview transcript, a survey export	Nothing is connected. The safest start
02	**A knowledge source**	The brand and style guide, approved messaging, product facts, the claims register, past content	Permissions are kept per reader
03	**An imported description of an interface**	Pages, campaigns, assets, events, registrations	You choose the operations. Changes wait for approval
04	**A reporting database**	Campaign results, spend, plan	Read only
05	**A server from the tool catalogue, run in your cluster**	The data warehouse, web search, document conversion	It runs inside. The search words go to the search provider
06	**A company registry**	Legal name, registered address, sector, filings	Public. Through web search or the registry's own interface

145

public interface descriptions read, from 26 publishers; 142 import

2,942

of the 5,554 operations in the 142 are sent as changes

What is published. We read 145 public interface descriptions from 26 publishers and ran each through the platform's importer. 142 import, none is refused, and 3 are over the size limit. 31 more were left out before the run because their format would not import. The 142 carry 5,554 operations, of which 2,942 are sent as changes, and the method alone cannot tell which only read [11].

What that exercise taught us. In these files a change is often a page published, a campaign sent or a post scheduled.

What is not published. Four of the 26 publishers account for 111 of the 145 files. The largest advertising platforms and social networks publish none a file address reaches. For those, the agent works from a file a person exports.

Web search. The search tool runs on a server inside the company's cluster. The words of the search go to the search provider. So an agent searches for a company or a market, not a person, and puts nothing from an unpublished plan into the search.

Text only. The agents in this paper write and check text. They generate no image, audio or video. Where a campaign uses assets made elsewhere, an agent lists what the asset library records.

So choose the operations at import, start with the ones that only read, and add nothing that publishes or sends until the approval step has been seen working.

PART 04

What an agent may do

Two rules run through every agent in this paper. It drafts and checks, and a named person approves and publishes. And every claim carries its source.

01	**It reads** what the marketer using it could read, and no more.
02	**It drafts** the brief, the article, the page, the post, the email, the report, with each factual statement tied to its source.
03	**It checks and lists** what departs from the style guide, what has no source, and what is missing or out of date.
04	**It stops** before anything that changes a record or reaches the public. A tool that changes something waits for approval.
05	**It does one job.** An agent is given the tools for one duty. The agent that checks a campaign holds no tool that sends one.
06	**It scores no marketer and reads no mood,** a colleague's or the public's.
07	**Approval is per item.** A person approves each piece that is published. Approval of a plan is not approval of what the plan produces.

That gives 35 agents in nine areas: content; website and search; social media; email and campaigns; events and webinars; product marketing; claims, assets and approvals; market and accounts; operations and reporting. The full list is a separate document. Four lines are drawn tighter than the rest.

Publishing.

An agent may draft an article, a page, a post, an email or a reply to a comment. It publishes none of them. The editor, the channel owner or the campaign owner approves each one and publishes, posts or sends it. An agent may draft a posting plan. It schedules nothing. It may set spending against the plan. It changes no budget, no bid and no audience, and it places no advert.

Labels and assets.

An agent lists the images, audio and video in a campaign, whether the asset library records each as made or changed with AI, and whether it is labelled. It does not decide what a label says or whether one is needed. The brand lead decides.

Claims, reviews and testimonials.

An agent lists every factual claim, comparison and superlative in a draft, with the source that supports it or a note that none was found. The marketing lead and legal approve or change each one. No agent writes a review or a testimonial. It may check that a quotation the company plans to use has a named source, the person's consent and the original wording.

People.

Research is about companies and markets. About a named person an agent uses three sources only: what the company already holds, what a company registry publishes about a company's officers, and what a buyer's own website says. It takes nothing from social network profiles, and the company buys no contact lists for it. An account list is a list of companies. Reports about audiences are about groups, and reports about results are about campaigns and channels, never marketers.

PART 05

What the rules already ask

This part needs an advertising lawyer's read, and a privacy officer's, before it is relied on. It reports what the texts say; it is not advice. It does not say how any agent, or any piece of content, would be treated under any law. That is for counsel. Nothing in it says that a platform, or an agent, makes a company's marketing lawful.

Labelling AI-made content

The maker of the system marks its output. Article 50(2) of the AI Act: "Providers of AI systems, including general-purpose AI systems, generating synthetic audio, image, video or text content, shall ensure that the outputs of the AI system are marked in a machine-readable format and detectable as artificially generated or manipulated." The paragraph has stated exceptions [12].

The company that uses it discloses in two cases. Article 50(4) says deployers of a system that makes "image, audio or video content constituting a deep fake, shall disclose that the content has been artificially generated or manipulated". On text it says: "Deployers of an AI system that generates or manipulates text which is published with the purpose of informing the public on matters of public interest shall disclose that the text has been artificially generated or manipulated." The text duty does not apply "where the AI-generated content has undergone a process of human review or editorial control and where a natural or legal person holds editorial responsibility for the publication of the content" [12].

What this paper does with that. Here a named person reviews, approves and publishes every piece. Whether a piece informs the public "on matters of public interest", whether that review meets the exception, and what a label should say are for counsel.

These duties apply from 2 August 2026. That is the Act's general date of application [12]. It is not the date for high-risk systems. Regulation (EU) 2026/1744 moved those: the high-risk duties "shall apply from: (i) 2 December 2027 as regards AI systems classified as high-risk pursuant to Article 6(2) and Annex III; and (ii) 2 August 2028 as regards AI systems classified as high-risk pursuant to Article 6(1) and Annex I". The same regulation gives providers whose systems "have been placed on the market before 2 August 2026" until "2 December 2026" to meet Article 50(2) [13].

A code of practice exists. The European Commission's page on it says: "Even though adherence to the code is voluntary, the transparency requirements under article 50 of the AI Act are legal obligations" [14].

Reviews and endorsements

European Union. Annex I of the Unfair Commercial Practices Directive lists, at point 23b, stating that reviews come from consumers who used the product "without taking reasonable and proportionate steps to check that they originate from such consumers". Point 23c: "Submitting or commissioning another legal or natural person to submit false consumer reviews or endorsements, or misrepresenting consumer reviews or social endorsements, in order to promote products" [15].

United Kingdom. Schedule 20 of the Digital Markets, Competition and Consumers Act 2024 lists "Submitting, or commissioning another person to submit or write" a "fake consumer review". It defines one as "a consumer review that purports to be, but is not, based on a person's genuine experience". The paragraph has been in force since 6 April 2025 [16].

United States. Under 16 CFR 465.2 it is an unfair or deceptive practice "for a business to write, create, or sell a consumer review, consumer testimonial, or celebrity testimonial that materially misrepresents", among other things, "That the reviewer or testimonialist exists" [7]. This is why no agent here writes a review or a testimonial.

PART 05 · CONTINUED

Email

In the European Union the starting point is consent. Article 13(1) of the ePrivacy Directive says "electronic mail for the purposes of direct marketing may be allowed only in respect of subscribers or users who have given their prior consent". Article 13(2) allows a seller to email its own customers about similar products, provided they "clearly and distinctly are given the opportunity to object, free of charge and in an easy manner", when their details are collected "and on the occasion of each message" [17]. Each country has its own law, and we opened none.

A person can say stop. Article 21(2) of the GDPR: "Where personal data are processed for direct marketing purposes, the data subject shall have the right to object at any time". Article 21(3): after an objection, "the personal data shall no longer be processed for such purposes" [18]. This is why the Campaign Check Agent looks for the unsubscribe link before a campaign goes out, and why no agent changes a consent record.

Who owns AI-made material

The United States Copyright Office reported on 29 January 2025 that the outputs of generative AI "can be protected by copyright only where a human author has determined sufficient expressive elements". It names "creative arrangements or modifications of the output, but not the mere provision of prompts" [19]. What that means for a company's material is for counsel.

Marketers

Evaluating workers is on the high-risk list. Annex III of the AI Act lists, under employment, "AI systems intended to be used to make decisions affecting terms of work-related relationships, the promotion or termination of work-related contractual relationships, to allocate tasks based on individual behaviour or personal traits or characteristics or to monitor and evaluate the performance and behaviour of persons in such relationships" [12]. Its date is 2 December 2027 [13].

Reading emotion at work is prohibited. Article 5(1)(f) bans "the use of AI systems to infer emotions of a natural person in the areas of workplace and education institutions", except for medical or safety reasons [12]. By its words that article is about the workplace and education. This paper does not say it covers a consumer's mood. The agents here read no mood, the public's or a marketer's, as a design choice.

MARKET	WHAT WE READ	WHAT IT ASKS
European Union	AI Act, Article 50(2) and 50(4)	The provider marks AI-made content; the deployer discloses deep fakes and some public-interest text [12]
European Union	Unfair Commercial Practices Directive, Annex I, points 23b and 23c	No false reviews or endorsements; reasonable steps to check that reviews are real [15]
United Kingdom	Digital Markets, Competition and Consumers Act 2024, Schedule 20, paragraph 13	No fake consumer reviews, submitted or commissioned [16]

PART 05 · CONTINUED

What the rules have in common

01	**The company answers for what is published in its name,** whoever or whatever wrote it.
02	**A claim needs its evidence before it is made,** not after.
03	**A review or a testimonial comes from a real person's experience,** or it is not used.
04	**Some AI-made content is marked or disclosed,** and a person who reviews a text and answers for it is part of how the rule is written.

PART 06

A worked example: one software company's marketing team

Take a made-up company: a maker of accounting software for small firms, with a marketing team of 12. Nothing here is a customer's story, and every figure is illustrative; it is section 4 put to work.

01

An editor asks for a brief and a draft of a guide to a change in invoicing rules. The agent drafts the brief, and once the editor approves it, the guide. Each factual statement carries its source.

02

The Brand and Style Check lists nine departures from the guide, each with the rule. The editor accepts seven and overrides two.

03

The Claim Check lists 14 factual claims. Twelve have a source in the product facts or the claims register. Two do not: a line that the product is "the fastest in its class", and a figure for hours saved. The marketing lead cuts the first, and legal asks for the evidence behind the second.

04

Before the launch email goes out, the Campaign Check reports that one of three versions has no unsubscribe link, and that the audience includes a segment with no recorded basis for contact. The campaign owner fixes both and sends the campaign herself.

05

The Asset Label Check lists the six images in the campaign. The asset library records two as made with AI by the design team. One of the two carries no label. The brand lead decides the label.

Every one of these reads sources the company already has, changes no record by itself, publishes nothing, and leaves a record of who asked and what was shown.

PART 07

Seven questions before an agent writes in the company's name

01 **Who publishes?**
Which pages, posts, emails or adverts go out without a person reading them first? Is approval per item or per plan, and who holds that setting?

02 **Does each claim have a source?**
Does a draft show where each fact, figure and comparison came from? What happens to a line with none?

03 **Is AI-made content labelled?**
Can we list which assets were made with AI? Who decides the label?

04 **What may it change?**
Can it change a budget, a bid, an audience, a consent record or a live page? Does each of those wait for a named person?

05 **Does it write reviews or testimonials, and what does it use about people?**
Our own records, a registry, a buyer's website, a scraper or a bought list?

06 **Where does it run?**
On our infrastructure, in our private cloud, or in the supplier's service? Where is the model? Where do the words of a web search go?

07 **What is scored and recorded, and can we stop it?**
Does anything score a marketer? Can we switch off one agent, one tool, or all of them, at once?

PART 08

Where to start

Pick work where reading is slow, the decision stays with a person, and nothing reaches the public on the agent's word.

Brand and Style Check
The agent lists each departure from the guide with the rule. The editor decides.

Claim Check
The agent lists every claim with its source or the lack of one. The marketing lead and legal approve.

Site Content Audit
The agent lists pages that are out of date or contradict the product facts. The web editor corrects them.

One team, one channel, invented data first. Then beside the manual process, measured against it.

About Katonic AI

Katonic AI builds the platform that organisations use to build, run and govern AI agents on their own infrastructure: in their own data centre, in a private or sovereign cloud, or with no connection to the internet at all. The platform is in production with governments, banks, telecom operators and utilities. Katonic is certified to ISO 27001.

KATONIC.AI/SOLUTIONS/MARKETING

SOURCES

Read on 22 November 2026. Figures are the publishers' own and were not audited by us.

01 HubSpot Knowledge Base, "Create content with content agent", "Last updated: November 18, 2026" on the page. https://knowledge.hubspot.com/website-and-landing-pages/create-content-with-content-agent ; HubSpot, product page for its social media agent (frequently asked questions), undated, read 22 November 2026. https://www.hubspot.com/products/marketing/social-media-ai-agent

02 Salesforce Trailhead, module "Agentic Campaigns", unit "Create and Monitor Your Agentic Campaigns", undated, read 22 November 2026. https://trailhead.salesforce.com/content/learn/modules/agentic-campaigns/create-and-monitor-your-agentic-campaigns

03 Zapier template "Publish daily tailored social captions for your brand", undated, read 22 November 2026. https://zapier.com/automations/marketing/social-media-marketing/social-posting/publish-daily-tailored-social-captions-for-your-brand ; n8n workflow template 8027, "Fully automate blog creation with GPT & publish to WordPress via Google Sheets" ("Last update 9 months ago" on the page). https://n8n.io/workflows/8027-fully-automate-blog-creation-with-gpt-and-publish-to-wordpress-via-google-sheets/ ; n8n workflow template 6117, "AI SEO Blog Automation for WordPress with featured images-End-to-End" ("Last update 3 months ago"). https://n8n.io/workflows/6117-ai-seo-blog-automation-for-wordpress-with-featured-images-end-to-end/ ; n8n workflow template 5773, "Generate & schedule social media posts with GPT-4 and Telegram approval workflow" ("Last update 9 months ago"). https://n8n.io/workflows/5773-generate-and-schedule-social-media-posts-with-gpt-4-and-telegram-approval-workflow/

04 Hootsuite, product page for its Wisdom agent (frequently asked questions), undated, read 22 November 2026. https://www.hootsuite.com/platform/owly-writer-ai ; Microsoft Learn, "Journey Creation Agent in Customer Insights - Journeys (preview)", page metadata dated 17 November 2026. https://learn.microsoft.com/en-us/dynamics365/customer-insights/journeys/journey-creation-agent

05 Writer, trust page, page metadata published 10 June 2024, modified 25 October 2026. https://writer.com/trust/ ; Buffer, "AI Assistant" page (frequently asked questions), undated, read 22 November 2026. https://buffer.com/ai-assistant

SOURCES continued

06 United Kingdom, Advertising Standards Authority and Committee of Advertising Practice, “AI and Deepfakes: Four Things Advertisers Need to Know Before They Hit “Run””, CAP News, 11 June 2026. https://www.asa.org.uk/news/ai-and-deepfakes-four-things-advertisers-need-to-know-before-they-hit-run.html

07 United States, Federal Trade Commission, press release “Federal Trade Commission Announces Final Rule Banning Fake Reviews and Testimonials”, 14 August 2024. https://www.ftc.gov/news-events/news/press-releases/2024/08/federal-trade-commission-announces-final-rule-banning-fake-reviews-testimonials ; 16 CFR Part 465, section 465.2(a), eCFR current text (source note: 89 FR 68077, 22 August 2024). https://www.ecfr.gov/current/title-16/part-465 ; Federal Register record for the rule (published 22 August 2024, effective 21 October 2024). https://www.federalregister.gov/documents/2024/08/22/2024-18519/trade-regulation-rule-on-the-use-of-consumer-reviews-and-testimonials

08 United States, Federal Trade Commission, press release “FTC Order Requires Workado to Back Up Artificial Intelligence Detection Claims”, 28 April 2025 (a proposed order). https://www.ftc.gov/news-events/news/press-releases/2025/04/ftc-order-requires-workado-back-artificial-intelligence-detection-claims

09 Engadget, “CNET had to correct most of its AI-written articles”, 25 January 2023. https://www.engadget.com/cnet-corrected-41-of-its-77-ai-written-articles-201519489.html

10 Indicator, “Tech platforms promised to label AI content. They’re not delivering.”, Alexios Mantzarlis and Nasha Dutta, 23 October 2025. https://indicator.media/p/tech-platforms-fail-to-label-ai-content-c2pa-metadata

11 Templates and prebuilt agents of agent and workflow platforms, as opened for sources 1 to 4; and 145 interface descriptions from 26 publishers fetched 22 November 2026 and run through the platform’s importer (Katonic research).

12 Regulation (EU) 2024/1689 (AI Act) of 13 June 2024, Article 5(1)(f), Article 50(2) and (4), Article 113, Annex III point 4(b). http://publications.europa.eu/resource/celex/32024R1689

13 Regulation (EU) 2026/1744 of 8 July 2026 (Official Journal 24 July 2026), the new Article 111(4) and the replaced Article 113 point (c). http://publications.europa.eu/resource/celex/32026R1744

14 European Commission, “Code of Practice on Transparency of AI-generated Content” (policy page; its timeline gives 10 June 2026 for the final code), read 22 November 2026. https://digital-strategy.ec.europa.eu/en/policies/code-practice-ai-generated-content

15 Directive 2005/29/EC of 11 May 2005 (Unfair Commercial Practices Directive), Annex I points 23b and 23c, consolidated version of 28 May 2022. http://publications.europa.eu/resource/celex/02005L0029-20220528

16 United Kingdom, Digital Markets, Competition and Consumers Act 2024 (c. 13), Schedule 20, paragraph 13(1) and (5)(b), with the commencement note “in force at 6.4.2025 by S.I. 2025/272”, text as shown on 22 November 2026. https://www.legislation.gov.uk/ukpga/2024/13/schedule/20

17 Directive 2002/58/EC of 12 July 2002 (ePrivacy Directive), Article 13(1) and (2), consolidated version of 19 December 2009. http://publications.europa.eu/resource/celex/02002L0058-20091219

18 Regulation (EU) 2016/679 (GDPR) of 27 April 2016, Article 21(2) and (3). http://publications.europa.eu/resource/celex/32016R0679

19 United States Copyright Office, NewsNet Issue 1060, “Copyright Office Releases Part 2 of Artificial Intelligence Report”, 29 January 2025. https://copyright.gov/newsnet/2025/1060.html