MARKETING MANAGED

AI and Automation Policy

How we use artificial intelligence and automation, and where we stop.

Version	1.3
Date	13 August 2026
Supersedes	V1.2, 3 June 2025
Owner	Phil Bacon, Founder and CMO
Deputy	A Senior Strategist or Director, acting in the owner's absence
Applies to	All staff, contractors, associates and third parties
Next full review	13 August 2027
Status	Published in full to clients and prospects

> **We publish this policy in full.** Clients and prospects can read exactly what happens to their data and their brand before they sign anything.

1. Why this policy exists

Marketing Managed uses AI and automation across most of our delivery. That is not a disclaimer. It is how we run awareness and demand generation for technology and engineering-led manufacturers.

This policy sets out what we use, how we govern it, and where we stop.

We publish it in full. Clients and prospects can read exactly what happens to their data and their brand before they sign anything.

V1.2 governed generative AI only. It assumed AI produced a draft and a human checked it. That is no longer the shape of the work. AI now takes actions across systems, builds working software, and runs governed processes on a schedule. V1.3 covers all three.

2. Scope

This policy applies to everyone working for or with Marketing Managed: employees, workers, contractors, associates, freelancers, volunteers, interns and apprentices, and any third party acting on our behalf.

It applies to all client work and all internal work.

It sits alongside our IT and computer use policy, our data protection policy, and our dignity at work and equal opportunities policies. Where they overlap, the stricter rule applies.

It does not form part of any contract of employment. Marketing Managed may amend it at any time. Where a client's own instructions are stricter than this policy, the client's instructions win for that account.

3. Definitions

Term	Meaning
Generative AI	Systems that produce text, images, audio, video or code from instructions.
Agentic AI	Systems that take actions across tools to reach a goal, rather than returning a single output. An agent that reads a CRM, drafts an email and queues it is agentic. This is the significant change since V1.2.
Automation	Any process that runs on a trigger or a schedule without a person starting it each time.
No-code and low-code	Building working software through visual interfaces rather than written code.
Vibe coding	Producing working software by describing intent to an AI rather than writing the code yourself. Fast, useful, and dangerous when nobody reads the output. Section 10 applies.
Named human	A specific individual who has read the work, understood it, and is accountable for it. Not "the team". A name, recorded.
Live system	Any environment holding real client data, real customer records, or serving real users. A sandbox holding only synthetic data is not a live system.

4. Principles

Six rules. Everything else follows from them.

1. **A named human is accountable for every output.** AI is never the accountable party.
2. **Reading is free. Acting needs approval.** Section 6.
3. **Confidential data goes only into approved tools on business terms.** Sections 7 and 8.
4. **We disclose our method, and we disclose synthetic media.** Section 9.
5. **Every automation has an owner, a log, a stop control and a review date.** Section 11.
6. **Nothing irreversible happens without a human approving it.** Either the specific action, or the automation design that produces it. Section 6 sets out which.

5. How we use AI: the three tiers

Every use falls into one tier. The tier sets the approval.

Tier 1 – Assist. No approval required.

Work that stays inside the agency and touches no live system.

Ideation and brainstorming. Research and summarising. First drafts and outlines. Internal notes and meeting summaries. Proofreading. Data processing on data already held. Prototypes and experiments against synthetic or sample data.

Use approved tools. Apply judgement. Get on with it.

Tier 2 – Produce. Named human reviewer before it leaves the agency.

Anything a client, a prospect or the public will see or receive.

Client-facing content of any kind. Reports and analysis. Campaign assets. Proposals. SEO and web content. Code and automations before they reach a live system. Anything published under a client's name or ours.

The reviewer reads it and understands it. Their name is recorded against the deliverable in Zoho Projects. A skim is not a review.

Tier 3 – Act. Written approval from Phil or the deputy.

Anything irreversible, anything that commits money, anything with legal weight.

- The design of any automation that contacts real people.
- Payments and spend commitments.
- Publishing to a client's live environment for the first time.
- Bulk deletion or overwriting of records.
- Anything touching contracts or employment decisions.
- Any build handling payments, authentication, or the personal data of anyone outside our own staff.
- Adding a tool to the register.
- Synthetic voice or likeness, which is additionally subject to the hard gate at Section 11.

Approval is in writing, recorded in the Governance Log (Appendix C), and given either per action or once at automation design. Section 6 sets out which.

6. Agentic AI: where agents may act alone

Agents read freely. Agents act with approval.

Autonomous. No approval needed.

Read, search, retrieve, analyse, summarise, draft, propose and report. An agent may read across the CRM, the project tracker, the inbox and the file store to build a picture, and hand that picture to a human.

Approved once at design, then runs.

A recurring automation gets full Tier 3 approval at the point of design, then operates without per-run approval, provided it stays inside its approved envelope. The envelope is fixed at approval and recorded: the audience or list it may contact, the message templates it may use, the volume cap, the trigger, and the schedule.

An automation that steps outside its envelope stops and asks. Changing the envelope is a new Tier 3 approval.

This is what lets our outbound sequences run. A sequence approved at design, sending approved templates to an approved list within its cap, does not need a human on each send.

Always human-approved, per action.

- Bespoke outbound composed by an agent for a specific named person, outside an approved envelope.
- Payment or spend.
- First publication to a live client environment.
- Deleting or overwriting records.
- Any action outside an approved envelope that cannot be undone.

> **The test.** If it goes wrong, can we take it back? If no, either a human approves that action, or it runs inside an envelope a human approved. There is no third option.

Standing rule. No agent operates on a client account without that client knowing agents are in use.

7. Confidential information

This section protects Marketing Managed. Section 8 protects our clients. Both apply.

Do not put the following into any tool that is not on the register at Appendix A, and never into a free or personal-tier account:

- Details of our clients, customers, suppliers, agents, partners, employees or shareholders.
- Information about the security of our premises, systems, networks or communications.
- Our trading position, finances, pricing model, margins or commercial terms, where not already public.
- Customer and prospect lists, trade secrets and know-how.
- Anything you have been told is confidential, or that a reasonable person would treat as confidential.

Our own commercial information deserves the same care as a client's. It is easier to leak, because nobody thinks to stop you.

8. Client data, vendor terms and consent

Approved tools only. Client data goes into tools on the register at Appendix A and nowhere else.

Business terms, no training. Every tool holding client data must be on a business or enterprise tier with terms excluding training on our inputs. Free and personal-tier AI accounts never receive client data. This is the most common way agencies leak client information, and it is not negotiable.

Minimise. Use the least data that does the job. Strip identifiers the task does not need. Aggregate where you can.

Tell the client. AI and automation use is confirmed in writing at onboarding, with this policy attached. Restrictions a client sets are recorded on the account and honoured. A client instruction to exclude a tool overrides this policy for that account.

Processor status. We act as processor on client data. The tools at Appendix A are our sub-processors. Under Article 28 UK GDPR we notify clients in advance of any intended addition or replacement, and they may object. Adding a tool is Tier 3 and the client notification is part of that approval, not an afterthought.

International transfers. Several tools at Appendix A are hosted outside the UK, principally in the United States. Transfers rely on the UK Addendum to the EU Standard Contractual Clauses or the UK Extension to the EU-US Data Privacy Framework, as applicable per vendor. The position per tool is recorded on the register and available to clients on request.

Breach and incident reporting. Report suspected data exposure to Phil the same day. There is no penalty for reporting fast and there is one for sitting on it. Where we are processor, we notify the client controller without undue delay under Article 33(2) UK GDPR. Where we are controller, we assess against the 72-hour ICO notification threshold. Phil owns both.

Retention. Automation activity logs are retained for 12 months, then deleted. Prompt and output monitoring records are retained for 6 months. Client data retention follows the client contract.

9. Disclosure and transparency

We do not label every asset. Almost every tool in a modern marketing stack has AI inside it, so blanket labelling tells a reader nothing.

We are transparent about method instead.

Standing disclosure. This policy is published in full. AI and automation use is confirmed in writing with every client at onboarding. We answer questions about our methods specifically and honestly.

Mandatory disclosure. Where an audience would otherwise be misled about what is real:

- Synthetic or cloned voice.
- AI-generated likeness, avatars or presenters.
- AI-generated imagery presented as photography of real people, products, sites or events.
- Automated outreach where a recipient could reasonably believe they are speaking with a named human.

- Chatbots and AI assistants, disclosed in the interface at first interaction.

Never. Fabricated testimonials, invented case studies, fake reviews, or synthetic likeness of a real person without their written permission.

9.1 EU AI Act, Article 50

Article 50 transparency obligations became applicable on 2 August 2026. The Act reaches non-EU businesses where the output of an AI system is used in the Union, so it can apply to work we produce in the UK for EU audiences. We assess this per client at onboarding rather than assuming either way.

Our duties as deployer

- **Deepfakes.** Where we generate or manipulate image, audio or video resembling real people, objects, places, entities or events, and it could falsely appear authentic, we label it. Article 50(4).
- **AI-generated text published to inform the public on matters of public interest.** Disclosed, unless the content has been through substantive human review with a named person holding editorial responsibility for the publication. Our Tier 2 review supplies the review. Editorial responsibility must also be allocated, and where content publishes under a client's name it usually sits with the client. We agree and record who holds it at onboarding. Without that allocation the exemption is not available.
- **Emotion recognition and biometric categorisation.** Where we deploy a tool that infers emotion or categorises people biometrically, we inform the people exposed to it. Article 50(3). Several meeting and video tools ship sentiment features. Check before enabling.

Duties that sit with the provider, not us

- **AI-interaction disclosure** in chatbots and assistants. Article 50(1). We still disclose as a matter of policy, and if we white-label or materially modify a system we may take on provider duties under Article 25.
- **Machine-readable marking** of generative output. Article 50(2). Systems on the market before 2 August 2026 have until 2 December 2026.

Our related duty is not to strip provenance marking. Where our export or editing pipeline removes it, that is a defect to fix. Owner: Phil Bacon. Audit of export pipelines due by 30 September 2026.

Penalties. Article 50 breaches carry administrative fines set by national market surveillance authorities. For SMEs, including micro-enterprises, Article 99(6) caps the fine at the lower of the fixed amount or the percentage of worldwide turnover. Marketing Managed is a micro-enterprise, so the percentage figure governs.

9.2 UK position

There is no UK AI Act. UK AI use is governed through existing law: UK GDPR, the Data Protection Act 2018, PECR, consumer protection and the CAP Code.

SI 2026/425 came into force on 12 May 2026 and requires the Information Commissioner to produce a statutory Code of Practice on AI and automated decision-making. No publication date has been set. When issued it will be admissible in legal proceedings. We will assess and align on publication.

> **Our largest regulatory exposure is PECR, not the AI Act.** Since 5 February 2026 the maximum PECR penalty is £17.5 million or 4 per cent of global annual turnover, whichever is higher, up from £500,000. That is the number that should govern how carefully we run outbound. Section 11.

10. AI-assisted coding, no-code and low-code

We build things. Lovable, n8n, Zoho builders, Vercel, written code. Much of it with AI assistance.

Speed is the point. It is not a reason to ship code nobody has read.

Before anything reaches a live system:

1. A named human has read the code or workflow and understands what it does.
2. No credentials, API keys, tokens or client data have been pasted into any prompt.
3. Secrets live in the platform's own secret store: n8n credentials, Vercel environment variables, Zoho Vault. Never in a build, a repo, a prompt or a document.
4. Dependencies and licences checked. Nothing copyleft in a client deliverable.
5. Tested against synthetic or sample data first. Where the platform provides no separate staging environment, test on a duplicate record set inside the live platform and record that this is what was done.
6. Client IP position per Section 15.

Tier 3 approval and a second reviewer are both required where the build handles payments, authentication, or the personal data of anyone who is not Marketing Managed staff. There is no record-count threshold. One record counts, because one special category record carries more risk than a thousand business email addresses.

Vibe coding needs no approval for prototypes, experiments and sandboxes running on synthetic or sample data. The moment a build reads from or writes to a live system, it is Tier 2 at minimum and the rules above apply in full. An internal tool that reads the live CRM is a live-system build, not a prototype.

Applies equally to written code, n8n workflows, Zoho automations and Lovable builds. The rules follow the risk, not the tool.

11. Automation governance

An automation acts on our behalf, continuously, without being asked. Govern it accordingly.

Every live automation is recorded in the Automation Register (Appendix C) with:

Requirement	Detail
Named owner	One person. Not a team.
Purpose	What it does and why it exists.
Trigger and schedule	What starts it.

Requirement	Detail
Approved envelope	Audience, templates, volume cap, per Section 6.
Activity log	What it did, when, and to whom. Location recorded.
Stop control	How to stop it, and who can. Minimum two people hold the access.
Review date	Quarterly.
Failure mode	What happens when it breaks, and who is told.

An automation with no named owner gets switched off.

Outbound and PECR

Outbound at scale must comply with UK GDPR and PECR. The operative rules:

- **Corporate subscribers.** Unsolicited B2B email to limited companies and LLPs is permitted under PECR. UK GDPR still applies to the personal data of the individuals at those companies. The DUAA 2025 did not change this.
- **Sole traders and partnerships** count as individual subscribers. They need consent or soft opt-in. Do not assume a business address means corporate subscriber. Check.
- **Soft opt-in** applies only where we obtained the details in the course of a sale or negotiations, the marketing is for similar products or services, and an opt-out was offered at collection and in every message.
- **Every message** identifies the sender, gives a valid reply address, and carries a working one-click unsubscribe.
- **Unsubscribes are honoured immediately** and suppressed across every tool, not just the one that sent the message.
- **Volume caps** are set per sequence at Tier 3 approval and recorded in the register. Defaults at Appendix C.

> **Synthetic voice and automated calling remain inactive.** This is a hard gate. It is not available through Tier 3 approval. It opens only when compliance is signed off for that specific market, which requires TPS and CTPS screening, consent position, and call recording notification to be documented and approved.

Third-party platform automation

Some platforms prohibit third-party automation tools in their own user agreements. LinkedIn is the significant one.

Where we propose to use a tool of that kind on a client account, we do not proceed on assumption. The rule:

- The client gives **written permission** before the tool touches their account.
- The permission states plainly that the platform's user agreement prohibits third-party tools, and that the client accepts the associated risk to their account.
- **No written permission means no usage.** There is no implied consent and no grandfathering.

- Permission is recorded on the account and revisited at each quarterly review.

We would rather have that conversation at onboarding than after a restriction.

Quarterly operations review

One review, one date, covering: every live automation, the tool register, open Tier 3 approvals, and any incidents. Owned by Phil. Held quarterly.

12. Prohibited use

You must not use AI or automation to:

- Draft contracts, legal documents or legally binding correspondence without qualified review.
- Make decisions about people. Recruitment, screening, performance management, grievance, disciplinary or task allocation. AI may support the analysis. A human decides. See Section 13 for what this does and does not cover.
- Put confidential information or personal data into a tool not on the register, or into any free or personal-tier account.
- Produce fabricated testimonials, case studies, reviews, statistics or citations.
- Create synthetic likeness or voice of a real person without their written permission.
- Take irreversible action outside an approved envelope without approval.
- Bypass a client's stated restrictions.
- Present unverified AI output as fact.

13. Automated decisions and lead scoring

Section 12 prohibits AI decisions about people. That means decisions with legal or similarly significant effect on an individual: employment, credit, access to a service.

It does not prohibit lead scoring and prioritisation, which is what our outbound does. Ranking which businesses to approach first, and in what order, has no legal or similarly significant effect on anyone. It is prioritisation, not decision-making.

The boundary: a score may order our work. A score may never be the sole reason a person is refused something they asked for. Where an individual queries how they were selected for contact, we explain it in plain terms and suppress on request.

We will reassess this position when the ICO statutory code is published.

When a DPIA is required

UK GDPR sets no volume threshold, and the ICO does not define "large scale" by number. It is judged on how many people are affected, the volume and variety of data, how long processing lasts, and its geographic extent. So we do not use a record count.

We use the ICO's own screening criteria. Where processing meets two or more, we complete a Data Protection Impact Assessment before it goes live. The criteria that realistically apply to marketing work:

- Invisible processing, where contact data is obtained indirectly rather than from the individual.
- Large-scale profiling of individuals.

- Data matching or combining from multiple sources.
- Tracking behaviour or location.
- Use of innovative technology.
- Targeting children or vulnerable individuals.

A DPIA is mandatory regardless of criteria count where processing involves special category data at scale, or systematic and extensive profiling with legal or similarly significant effect.

DPIAs are owned by Phil, held with the Governance Log, reviewed annually and after any material change.

14. Accuracy and quality control

AI is confidently wrong. Plan for it.

Verify every factual claim before it leaves the agency. Statistics, dates, names, technical specifications, quotes and citations get checked against a primary source. AI-generated citations are fabricated often enough that every one must be opened and read.

Technical accuracy matters more here than at most agencies. We write for engineers and technical buyers. A wrong specification does more damage than a dull sentence.

Named SME sign-off. Where content asserts technical specifications, performance claims or engineering detail, a named subject matter expert signs it off before publication. That expert may be ours or, by agreement, the client's. This is a workflow step agreed at onboarding, not a fallback when we are unsure.

Every AI-produced deliverable is checked by a human before it ships. Failing to do this to a satisfactory standard is a disciplinary matter for employees and a contractual matter for everyone else.

15. Copyright and intellectual property

Do not reproduce substantial parts of AI output without checking it. Generative AI can reproduce copyrighted material.

Run plagiarism checks on published content. Check image provenance and licensing before commercial use.

The IP position. Under UK law, purely computer-generated work has an uncertain copyright status, and work with no human author may attract no copyright at all. So we do not promise clients copyright in every asset. What we do promise:

- We assign to the client all rights we hold in the deliverables we produce for them.
- Where a client needs enforceable rights in a specific asset, we involve a human creator meaningfully and record that involvement. Tell us at the outset which assets these are.
- We do not reuse client-specific builds or assets on other accounts.

16. Bias, fairness and brand voice

AI reflects its training data, which skews towards volume and consensus. Two consequences.

Bias. Review output for bias in language, imagery and assumptions. Comply with our dignity at work and equal opportunities policies. Report problems to your manager.

Sameness. AI defaults to generic marketing language. Clients hire us for a specific voice. Edit hard.

17. Security

- Multi-factor authentication on every tool holding client data.
- No shared logins.
- Secrets in the platform's secret store. Never in prompts, repos, chats or documents.
- Access granted on need. At least two people hold admin access to any system running a live automation, so no automation is unstoppable.
- Access removed the day a person leaves.
- Suspected exposure reported to Phil the same day. Section 8.

18. Monitoring

Marketing Managed may monitor use of AI tools and automations for compliance with this policy, including prompts and outputs on agency accounts.

- **Purpose:** compliance with this policy and protection of client data.
- **Lawful basis:** legitimate interests. A legitimate interests assessment is held by Phil and available to any staff member on request.
- **Scope:** agency accounts and agency devices only. Personal accounts and personal devices are never monitored.
- **Retention:** 6 months. Section 8.
- Automation activity logs are retained and reviewed under Section 11.

19. Training and support

Marketing Managed provides training and support on approved tools, including the risks.

If you are unsure whether something is allowed, ask before you do it. That is always the right call and it will never be held against you.

20. Responsibilities

Role	Responsibility
Phil Bacon, Founder and CMO	Policy owner. Tier 3 approvals. Tool register. Breach decisions. Quarterly operations review.
Deputy, Senior Strategist or Director	Acts for Phil in his absence on all approvals under this policy.
Automation owners	Their automation works, is logged, stays in its envelope, and is reviewed.

Role	Responsibility
Everyone	Knows the tiers. Applies them. Asks when unsure.

21. Enforcement

For employees, failure to comply may result in disciplinary action up to and including dismissal.

For contractors, associates, freelancers and third parties, failure to comply is a breach of contract and may result in immediate termination of engagement.

Treated as serious in both cases: putting confidential or personal data into an unapproved tool, taking irreversible action outside an approved envelope, and publishing unverified claims.

22. Review and updates

Full review annually. Next 13 August 2027.

Quarterly operations review per Section 11. Next 13 November 2026.

Reviewed sooner where the regulatory position changes materially, where we adopt a significant new capability, or after any incident.

Changes are communicated to everyone in scope and recorded in the version history.

Appendix A: Approved tool register

Current as at 13 August 2026. Owned by Phil Bacon. Reviewed quarterly.

Category	Tools	Client personal data	Notes
AI assistants and agents	Claude, Cowork	Yes	Business tier. No training on inputs.
Automation and orchestration	n8n, Martech Managed	Yes	Section 11 applies to every live workflow.
CRM	Zoho CRM	Yes	System of record for contacts and pipeline.
Project and delivery	Zoho Projects	Yes	Holds Tier 2 reviewer names.
Build and deploy	Lovable, Vercel	Yes	Section 10 applies. Secrets in platform store.
Content and creative	Canva, Gamma, Descript, Miro	No	Brand assets, copy and creative only. No client contact records, no CRM exports, no commercial terms.
Meetings and notes	Granola	Yes	All participants informed before recording. Emotion and sentiment features off unless Section 9.1 disclosure is made.

Category	Tools	Client personal data	Notes
Email and workspace	Gmail, Google Workspace	Yes	Google Drive is a primary client data store.
Outbound and social	Meet Alfred	Yes	Written client permission required before use on their account. See Section 11, Third-party platform automation. Cap 50 actions per day per account.
Web	Joomla	Yes	First live publication to a client site is Tier 3.
Finance	FreeAgent	No	Internal only.

Not approved: any free or personal-tier AI account, any tool not listed, and any tool a specific client has excluded.

Recorded per tool and available to clients on request: vendor, tier, training exclusion confirmed, hosting location, transfer mechanism, date approved.

To add a tool: request from Phil with the use case and a link to the vendor's terms on training and retention. Addition is Tier 3 and includes client notification under Section 8.

Appendix B: Quick reference

Three questions, in order.

1. **Can it be undone?** No, stop. Get Tier 3 approval, or confirm it runs inside an approved envelope.
2. **Does it leave the agency?** Yes, get a named reviewer and record the name.
3. **Is the tool on the register, on the business tier?** No, stop.

Automation checklist. Owner. Purpose. Trigger. Envelope. Log. Stop control. Review date. Failure mode.

Code checklist. Human has read it. No secrets in prompts. Secrets in the platform store. Licences checked. Tested on synthetic data. IP position recorded.

> **When in doubt, ask Phil.** Faster than fixing it afterwards.

Appendix C: Registers, thresholds and defaults

Governance Log. Zoho Projects. Records every Tier 3 approval: what, who approved, date, envelope where applicable.

Automation Register. Zoho Projects. One entry per live automation with the eight fields at Section 11.

Tool Register. Appendix A, maintained by Phil.

Thresholds and defaults. These are working figures, reviewed quarterly.

Item	Value
Build trigger for Tier 3 and a second reviewer	Any build touching the personal data of anyone outside our own staff, or handling payments or authentication. No record count.
DPIA trigger	Two or more ICO screening criteria. Section 13.
Outbound email, default cap	25 to 30 per day per sending domain
LinkedIn actions, default cap	50 total actions per day per account
Automation activity log retention	12 months
Monitoring record retention	6 months
Automation review frequency	Quarterly

Version history

Version	Date	Change
1.0	06/06/2024	Original generative AI workplace policy, issued as Bacon Marketing.
1.2	03/06/2025	Rebranded to Marketing Managed. Acceptable use revised.
1.3	13/08/2026	Rewritten as an AI and Automation Policy. Adds agentic AI and the approved-envelope model, AI-assisted coding, no-code and low-code, automation governance, the approved tool register, automated decisions and lead scoring, international transfers, breach notification and retention. Replaces blanket authorisation with three tiers. Replaces per-asset labelling with method disclosure plus mandatory synthetic-media disclosure. Adds EU AI Act Article 50, the UK regulatory position and the PECR penalty change. Published externally in full.

No V1.1 was issued.

Marketing Managed. Published in full at every client onboarding. Questions to Phil Bacon.