

Bampot: AI Governance and Data Review Pack

Prepared for organisations reviewing Bampot for use with business information.

More organisations now run a new tool past an AI or data governance review before it can be used with business information. This pack answers, in plain terms, the questions that review usually asks: what Bampot is, which providers handle information, how that information is used and held, and the assurances Bampot can give. If your review needs this on your own paper or under a signed agreement, we are glad to work at your pace.

1. What Bampot is

Bampot is a governed innovation tool for consumer goods. A single business question becomes a set of fully formed product concepts, each with the consumer insight, positioning, format, price and creative direction behind it, and each carrying a Provenance Map that traces every claim and step back to a named source. The governed chain, not a free-text chatbot, is the point: an output can be inspected rather than taken on trust.

2. Vendor and contacts

Legal entity: Bampot LLC, a Connecticut limited liability company, United States. Registered address: 2389 Main St, Ste 100, Glastonbury, CT 06033, USA. General and security contact: support@bampot.ai.

Bampot is registered with the UK Information Commissioner's Office (registration ZC170816). Its UK representative under Article 27 UK GDPR is Data Protection Representative (UK) Limited, trading as DataRep; UK data-protection requests can be sent to datarequest@datarep.com, quoting "Bampot LLC" in the subject. Terms of Service and a Privacy Policy are published at bampot.ai.

3. Which AI providers handle information

The reasoning that builds a brief runs on Anthropic's Claude by default, with OpenAI and Google Vertex AI selectable. Image generation is routed by how faithful the render must be, not by the member's selection: faithful pack renders are generated by Google Vertex AI, and mood visuals and other pack work by OpenAI. Each provider is a named sub-processor engaged under a data-processing agreement, and each is used under its own commercial terms. The full sub-processor list is in section 10.

4. Is information used to train AI models?

For text generation we use paid-tier API and Vertex agreements with Anthropic, OpenAI and Google Vertex AI, under which those providers commit not to use the content we submit to train their models. We use only the paid or Vertex tiers, never any free tier. Each provider retains submitted content only briefly and solely for abuse-monitoring purposes, approximately 7 days for Anthropic, 30 days for OpenAI and 24 hours for Google.

Image generation splits by provider. For pack renders that must stay faithful to an existing design, we use Google Vertex AI on our own Google Cloud project, under the Cloud Data Processing Addendum, under which Google commits not to use the content we submit to train its models. For mood visuals and other pack work, we use OpenAI; we have not obtained, and do not make, any commitment that OpenAI processes this content on a zero-data-retention basis, or that it does not use it to train its own models, and if a no-training or retention commitment is important to you, you should review OpenAI's own published terms. Mood visuals and product impressions are AI-generated concepts rather than photographs of real products.

5. Confidentiality and isolation

Every brief runs in isolation. Each call to a provider carries only your brief, with no shared session across users. Every request for stored work is checked server-side against the requesting account before anything is returned, and there is no direct database access from the browser. One customer's work cannot reach another.

Access to customer data is restricted to personnel who require it to operate, support or secure the Service, currently a single named individual, bound by confidentiality obligations. Administrative interfaces show operational and cost information only and cannot retrieve customer work.

6. Ownership of inputs and outputs

You own what you bring and what Bampot helps you make. Bampot owns its method, never your work or your results.

7. Where information is held

In the United States. Bampot is available in the US, UK and other approved regions today, and is not currently offered in the EU. For our US-based sub-processors we rely primarily on the UK International Data Transfer Agreement, or the UK Addendum to the European Commission's Standard Contractual Clauses. Creem (Estonia) and Fathom Analytics (Canada) are covered by the UK's adequacy regulations.

8. Security

Information is encrypted in transit and at rest, access sits behind authentication, and every request for stored work is checked server-side against the requesting account, with no direct database access from the browser. Bampot does not currently hold third-party certifications such as SOC 2 or ISO 27001; a certification roadmap and a Security Summary with fuller detail are available for review on request.

9. Retention, export and deletion

You can export your work at any time and request deletion at any point. When your plan ends, you have 30 days to export everything before it is removed. Closing an account ends access and removes data, subject to any retention the law requires.

10. Sub-processors

The providers below process information on Bampot's behalf, each under a data-processing agreement, except Creem, which acts as an independent controller for the payment transaction.

Sub-processor	Role and data	Location
Anthropic	Text generation	USA
OpenAI	Text generation, and image generation for mood visuals and pack work	USA
Google Vertex AI	Text generation, and image generation for faithful pack renders	USA
Supabase	Database, authentication and storage	USA (us-east-1)
Vercel	Application hosting	USA
Resend	Transactional email	USA
Creem	Payments and merchant of record; independent controller for the payment transaction	EU (Estonia)

Sub-processor	Role and data	Location
Fathom Analytics	Cookieless web analytics	Canada

11. Governance and auditability

Every Bampot Brief carries a Provenance Map. Each claim and step traces back to a named source, so a reviewer can see how a conclusion was reached rather than accept a black box. For a governance function this is often the deciding difference: the work can be inspected.

12. Documents and agreements

Terms of Service and a Privacy Policy are published at bampot.ai. A data-processing agreement is available on request; requests go to support@bampot.ai. If your review needs this pack on your own paper, a signed DPA, or answers to a specific questionnaire, contact support@bampot.ai and we will work at your council's pace.

Claude is a trademark of Anthropic, GPT of OpenAI, and Gemini and Vertex AI of Google. Bampot is independent, and is not affiliated with, endorsed by, or sponsored by Anthropic, OpenAI or Google; it uses their models under their commercial terms.