

Sustainability of AI Tools & Models

Statement issued: 01 June 2026 · Version 1.0

Atlas Workspace is an AI chief of staff for senior sales leaders. This statement documents the considerations Atlas has made regarding the environmental sustainability of the AI tools and models used to deliver the Service, and the commitments of the sub-processors that power it. It is intended for inclusion in vendor procurement reviews.

1. Architectural choices that reduce footprint

- **Tiered model routing.** The majority of Atlas requests are routed to smaller, more efficient models (e.g. Gemini Flash / Flash-Lite, GPT-5 mini / nano). Larger frontier-tier models are used only when the task genuinely requires them. This materially reduces compute per interaction.
- **No model training on customer data.** Atlas does not train or fine-tune any model on customer content. We do not run training jobs — the most energy-intensive class of AI workload — at all.
- **Data minimisation.** We retain only what the Service requires, and offer export and deletion on request. Less stored data means less ongoing storage, replication and backup energy.
- **Serverless, shared infrastructure.** Atlas runs on shared edge and serverless infrastructure rather than dedicated always-on servers, improving utilisation.
- **No on-prem GPU fleet.** Atlas does not operate its own GPU datacentre. Inference is delegated to hyperscale providers whose datacentres have substantially better PUE than typical enterprise datacentres.

2. Sub-processor sustainability commitments

Atlas's AI inference is delivered through the sub-processors listed below. Their publicly stated sustainability positions, which Atlas relies on as part of its processing chain:

Provider	Stated commitment
Google (Gemini models, via Google Cloud)	Carbon-neutral operations since 2007; targeting 24/7 carbon-free energy across all datacentres by 2030. Google Cloud regions publish carbon characteristics per region.
OpenAI (via Microsoft Azure)	OpenAI inference runs on Microsoft Azure, which is committed to being carbon negative, water positive and zero waste by 2030, and to matching 100% of electricity consumption with zero-carbon energy purchases.
Lovable (hosting & AI Gateway)	Application hosting on shared cloud edge infrastructure operating on 100% renewable energy.

3. What we do *not* currently do

In the interest of being straightforward with procurement teams:

- Atlas does not publish a per-request gCO₂e figure. The underlying model providers do not yet expose this at the API level.

- Atlas does not currently purchase separate carbon offsets — we rely on the renewable-energy and carbon commitments of the underlying infrastructure providers above.
- Atlas does not operate its own datacentre, and therefore does not publish a PUE figure of its own.

We will update this statement as model providers expose more granular emissions data.

4. Engagement & documentation

Sustainability of the underlying tools and models is referenced during vendor engagements and documented publicly at atlasworkspace.app/sustainability, atlasworkspace.app/subprocessors, and atlasworkspace.app/security. Enterprise customers may request this signed PDF statement as part of their vendor onboarding pack.

5. Contact

Questions about Atlas's sustainability position, or requests for additional documentation for a vendor review: hello@atlasworkspace.app.