

Responsible AI & Sustainable Technology Policy 2026

The standards Ember Venture Group applies whenever it designs, develops, recommends, procures or operates artificial intelligence and machine learning systems, and the technology infrastructure that supports them.

ORGANIZATION	VERSION	EFFECTIVE	APPROVED BY	NEXT REVIEW
Ember Venture Group Pte. Ltd., Singapore	1.0	October 2026	Co-Chief Executive Officers	October 2027

1. Purpose

Artificial intelligence earns a place in an institution's decisions only when it is fit for purpose, validated on evidence, overseen by accountable people, and built on technology the institution can control and sustain.

This policy sets out how Ember Venture Group meets that standard. It is a companion to our Sustainability & Responsible Business Statement 2026, which sets out our wider environmental, social and economic commitments.

How to read this policy. Sections 3 to 8 describe how we work today. Section 9 sets out the **2026 to 2027 objectives** through which we are formalizing and recording that practice. We will report progress annually.

2. Scope

This policy applies to every AI and machine learning system that Ember Venture Group Pte. Ltd. designs, builds, recommends, procures or operates, including its Anominal AI platform. It also applies to our advisory engagements on AI adoption, and it binds Ember's founders and the contractors and subcontractors who work on Ember engagements.

3. Principles

- **Fit for purpose.** We select the simplest model that meets the requirement, not the largest available, and record why. In our IMF work for the Central Bank of Somalia, transparent univariate statistical models were used by mandate and large language models were excluded.
- **Evidence before deployment.** Evaluation criteria are defined before modelling begins. Models are validated on held-out and walk-forward data against acceptance gates set in advance, and negative findings are recorded rather than discarded.
- **Human oversight.** AI informs decisions; accountable people make them. Consequential decisions retain a human in the loop, with the authority and information to override the system.
- **Transparency.** We document data sources, model choices, assumptions and known limitations, and explain outputs in terms decision-makers can understand and challenge.
- **Auditability.** Where a system drives decisions, its decision records are retained and protected against alteration, as they are in Anominal AI.
- **Privacy and data sovereignty.** Client data stays on infrastructure the client controls and is processed only for the engagement. It is never used to train models for other clients.
- **Security.** Least-privilege access, private networking, protected credentials, and security testing proportionate to the risk of the system.
- **Fairness.** Where a system affects people, we assess it for bias and disparate impact before deployment and monitor it afterwards.
- **Capability transfer.** We document and hand over systems so that client teams can operate, maintain and extend them without permanent dependence on Ember.

4. Prohibited uses

Ember will not design, build, recommend or knowingly support AI systems intended for:

- unlawful surveillance, tracking or profiling of individuals or groups;

- unlawful discrimination, or the denial of rights or essential services on the basis of protected characteristics;
- manipulating people by exploiting their vulnerabilities, or deceptive impersonation intended to mislead;
- weapons development or targeting;
- any purpose that violates human rights, applicable sanctions or the law.

We decline, or withdraw from, any engagement in which we identify such a use.

5. Sustainable technology

- **Right-sized compute.** Infrastructure is sized to measured workloads. Anominal AI's production platform runs on a single, right-sized dedicated server rather than elastic cloud capacity.
- **Efficient models and training.** We prefer efficient models where they perform as well as larger ones, and run model training as scheduled batches rather than continuously.
- **Data minimization.** We collect and retain only the data a system needs, in efficient storage formats, without redundant pipelines.
- **Open formats and portability.** Open standards, open-source components and portable formats such as Apache Parquet, so clients keep their data, models and documentation and can leave cleanly.
- **Longevity.** Commodity hardware, maintainable architectures and lifecycle costing, so systems are extended rather than replaced, and retired hardware is handled responsibly as set out in our Sustainability Statement.

6. How we apply this policy

STAGE	WHAT WE DO	WHAT WE RECORD
Intake	Establish the purpose, the people affected, the sensitivity of the data and the decisions the system will inform; screen against prohibited uses.	Purpose, risk level and screening outcome.
Design	Select the model approach and infrastructure; design data protection and human oversight into the system.	Model-selection rationale; data-protection and oversight design.
Validation	Fix acceptance criteria in advance; test on held-out and walk-forward data; assess bias where people are affected.	Criteria, results, limitations and negative findings.
Handover	Document the system, train client staff and agree an exit path for data and models.	Documentation, training delivered, exit path.
Operation	Monitor performance and drift; retain decision records; respond to incidents.	Monitoring results, decision records, incidents.

7. Third-party AI services

Ember does not send client data to third-party AI model providers by default. Where a client requests or approves such a service, we use it only with the client's explicit written consent, after reviewing the provider's data-handling, retention and training terms, and we record that decision.

8. Incidents & concerns

Clients, partners, contractors and affected individuals may raise concerns about an Ember AI system with the Co-CEOs through our general contact channel, without retaliation. We investigate every concern, notify affected clients promptly where a system has failed or caused harm, and remediate before the system continues in use.

9. Governance & objectives

The Co-CEOs jointly own and approve this policy. Bartosz J. Solowiej is accountable for architecture, responsible AI and the technology lifecycle; Nicole Solowiej is accountable for governance, contractor acknowledgement and reporting. The policy is reviewed at least annually, and whenever our services or the applicable regulation change materially.

OBJECTIVE	DELIVERABLE	TARGET
Responsible-AI assessment	A written assessment covering the stages in Section 6, completed for every AI engagement.	From Q1 2027
Contractor acknowledgement	Contractors and subcontractors acknowledge this policy before joining an AI engagement.	From Q1 2027
Incident register	A register of AI-related concerns and incidents, with actions taken.	From Q1 2027
Annual reporting	A progress report against this policy, published on embergroup.ai.	By October 2027

10. Approval

Approved by the Co-Chief Executive Officers of Ember Venture Group Pte. Ltd., October 2026.

Bartosz J. Solowiej
Co-Founder & Co-Chief Executive Officer

Nicole Solowiej
Co-Founder & Co-Chief Executive Officer