

Artificial Intelligence Governance

This policy applies the seven principles of the Ethical Technology Charter to every system in the group that learns from data or generates text, code, images or recommendations. It establishes an AI register, sorts each use into one of four classes, prohibits the fourth class outright, requires approval by the Ethics & Technology Council before any use goes live, fixes the provenance, inference and transparency standards for the permitted classes, and guarantees every client a human decision. The group makes no automated decision about any person, and this policy is the mechanism by which that commitment is kept.

CODE	CATEGORY	VERSION
IGC-DT-005	Data and technology	2.0
EFFECTIVE	REVIEW	OWNER
2024-03-01	Annual, by the Ethics & Technology Council. Next review due March 2027.	Head of Technology and Ethics

1. Purpose and the Charter

The Board adopted the Ethical Technology Charter in 2023 because the group intended to use learning systems in research, screening and operations and wanted the limits fixed before the first system went live. The Charter's seven principles are human accountability, proportionate automation, provenance of data, restraint in inference, welfare in biological capital, transparency to clients and the right to a human decision. This policy is the Charter's operating text for artificial intelligence.

The group's position is settled and is stated on every relevant page of its public site: no automated client decisions, no sale of client data, no facial recognition. Vendors are held to the same standard through the Third-Party Risk and Outsourcing policy. Nothing in this policy relaxes a requirement of the Model Risk Management policy; a learning system is a model and carries a tier, an owner and a validation cycle in addition to the requirements here.

The Head of Technology and Ethics owns the policy. The Ethics & Technology Council, which the Head chairs and on which the Chief Risk Officer, the General Counsel, the Group Data Protection Officer and the independent director Priya Vantongeren sit, approves every use and reviews the policy each year. The Council reports to the Board.

2. Classes of use

Every use of a learning or generative system is assigned to one of four classes according to what the system is permitted to do with its output. Class 1, assistive, covers drafting, summarising, translating, searching and code assistance where a named person reads, edits and issues everything the system produces. Class 2, analytical, covers pattern detection, anomaly flagging, screening prioritisation and

research signal generation where the output is an input to a person's analysis and is never shown to a client as a conclusion.

Class 3, decision support, covers systems that present a recommendation alongside the evidence for it to a named approver, who may accept, amend or reject it and whose decision is recorded with reasons. Class 3 is permitted only where the approver has the time, the information and the authority to decide differently, and the approval rate of every Class 3 system is monitored so that a system whose recommendations are never rejected is re-examined.

Class 4, autonomous action, covers any system that would take an action affecting a client, an employee, a counterparty, a position or a regulatory filing without a human decision in between. Class 4 is prohibited. No exception may be granted by the Council, the Executive Committee or the Board without an amendment to the Charter itself. At 30 June 2026 the register held 26 approved uses: eleven in Class 1, nine in Class 2, six in Class 3 and none in Class 4.

- Class 1, assistive: a named person reads, edits and issues every output. Eleven uses.
- Class 2, analytical: output is an input to human analysis, never a client-facing conclusion. Nine uses.
- Class 3, decision support: recommendation with evidence to a named approver who records the decision. Six uses.
- Class 4, autonomous action: prohibited without amendment of the Charter. No uses.

3. The AI register and approval

No learning or generative system is used for any group purpose until it is entered in the AI register and approved by the Ethics & Technology Council. The register entry names the class, the purpose, the accountable executive, the system owner, the datasets used for training, tuning and prompting and their provenance, the population the system acts on, the inferences it is designed to draw, the inferences it is prohibited from drawing, the method by which those prohibitions are tested and the tier assigned under the Model Risk Management policy.

Approval proceeds in stages. A Class 1 use may be approved by the Head of Technology and Ethics and notified to the Council. A Class 2 or Class 3 use requires the full Council, a completed privacy assessment under the Data Protection and Privacy policy and a pre-approval test report. Every approval is for a stated scope and a stated period of not more than twelve months, after which the use is re-approved or withdrawn. Six proposals were declined in the twelve months to 30 June 2026.

The accountable executive for every use is a member of the Executive Committee. That person answers to the Board for the use, whether or not they understand the technology, and is expected to understand it well enough to explain what it does and what it cannot do. This is the first Charter principle, human accountability, and it cannot be delegated below the Executive Committee.

4. Provenance and the group boundary

Every dataset used to train, tune, ground or prompt a system is traced to an entry in the processing register and carries a documented origin and lawful basis. Data whose origin cannot be stated is not used. Client data may be used only within the scope of the mandate under which it was collected and never to train a system that will serve another client. Data received from a counterparty under a confidentiality undertaking is not used for any learning purpose.

Public generative services are blocked at the group boundary under the Information Security policy.

Approved systems run inside the group's own environment or inside a contracted vendor environment that is bound by Charter terms: no retention of group prompts or outputs, no training on group data, no onward

disclosure and audit rights for the group. A vendor that cannot accept those terms is not approved, whatever the capability of its system.

Outputs of Class 2 and Class 3 systems are stored with a record of the inputs and the system version that produced them, so that any output can be reproduced and examined. This record is a record under the Records Management and Retention policy and follows the retention period of the decision it supported.

5. Restraint in inference

No system may be designed, tuned or prompted to infer a person's health, beliefs, political opinions, sexual orientation, ethnic origin, family circumstances or any other characteristic that a mandate does not require. No system may be used to identify a person from an image, a voice or any biometric feature. No system may score a person's propensity to any behaviour. These prohibitions apply whether the inference would be an intended output or a by-product.

Before approval, every Class 2 and Class 3 system is tested for prohibited inference by the model validation unit, using held-out cases designed to elicit the prohibited outputs. A system that produces a prohibited inference in testing is not approved. A system found to produce one in use is withdrawn on the day of discovery by the system owner or the Head of Technology and Ethics, and the withdrawal is reported to the Council and the Board.

The fifth Charter principle, welfare in biological capital, has a specific application: a system used to assess an animal-health, aquaculture or agri-science position must include the welfare standards in the Animal and Biological Welfare in Investment policy among its inputs, and may not be tuned to optimise a return measure that excludes them.

6. Transparency and the human decision

Every client report, valuation, recommendation or communication in which a Class 2 or Class 3 system contributed carries the assistance statement: a short note naming the class of system used, what it contributed and the person accountable for the content. The statement is a fixed form approved by the Council and is not varied by division. A client can therefore see which decisions were model-assisted, which is the sixth Charter principle.

Any client, at any time, may require that a decision to which a system contributed be reviewed by a named person who did not make it, and may require that a future decision be made without system assistance. The review is completed within ten business days and the outcome given in writing with reasons. The request is recorded in the AI register against the use concerned. Eight such requests were received in the twelve months to 30 June 2026, and in two the original decision was changed.

Staff whose work is affected by a Class 3 system, for example in screening or in operations, have the same right to a human review of any decision about them. The group does not use learning systems in recruitment, performance assessment or remuneration.

7. Incidents, monitoring and reporting

An AI incident is any output that breaches a prohibition in this policy, any use outside its approved scope, any client complaint concerning a system-assisted decision and any failure of a Class 3 approver to record a decision. Incidents are reported to the Head of Technology and Ethics within one business day and logged in the AI register. Two incidents were recorded in the twelve months to 30 June 2026, both scope breaches by Class 1 tools, and both uses were suspended until retrained.

The Council reviews the register, the incident log, the approval rates of Class 3 systems and the human-review requests each quarter. It reports to the Board annually with the full register, the declined proposals and its own assessment of whether the Charter is being kept. Internal Audit reviews the register and a sample of approvals every two years.

Key controls

- AI register of 26 approved uses naming class, purpose, accountable executive, data provenance and prohibited inferences.
- Ethics & Technology Council approval before any Class 2 or Class 3 use goes live, and re-approval within twelve months.
- Absolute prohibition on Class 4 autonomous action affecting any client, employee, counterparty, position or filing.
- Provenance trace of every training, tuning and prompting dataset to the processing register.
- Pre-approval testing for prohibited inference and same-day withdrawal of any system found to produce one.
- Boundary block on public generative services and Charter terms in every vendor contract.
- Assistance statement on every client output to which a Class 2 or Class 3 system contributed.
- Human review of any system-assisted decision within ten business days, and incident reporting within one business day.

Related policies

- Model Risk Management (IGC-DT-004)
- Data Protection and Privacy (IGC-DT-001)
- Responsible Innovation (IGC-SE-003)
- Third-Party Risk and Outsourcing (IGC-DT-006)

Colophon

IGUAKO Capital is a work of institutional fiction by the Iguako Institute for Applied Unreality (iguako.tech). This document is part of that work. It creates no rights or obligations, describes no licensed entity and is not an offer, a contract, a client record or advice.