

HUMAN-AI: Trust, Security and Data Governance

Principles and Commitments

Developed with Legal Counsel North End Law Corporation

HUMAN-AI is built for organisations whose data, knowledge and client relationships are their core assets. This document sets out the principles HUMAN-AI is built on and the operating standards that guide how we protect them.

About this document

HUMAN-AI is a platform in active development, and its architecture will continue to evolve as stronger components emerge. That evolution now includes a deliberate shift in our model strategy: rather than relying primarily on closed, third-party frontier models and hyperscaler AI services, we build on an open-weight foundation model that we fine-tune, customise and operate ourselves, within an architecture over which we hold greater technical control. A conventional security specification, a fixed description of a finished system, would be out of date within months, and keeping it nominally current would invite exactly the overclaiming we refuse to do. What we can state, and commit to, is the standard the platform is built to. This document defines the principles that do not change, the operational standards we work to today, and the way a partner, client or agency (“User”) can check both. It is not a contract, a warranty or an audit report. Legally binding protections sit in the Data Processing Agreement and the commercial agreements that we enter into. This document tells you what to expect from us and how we intend to evidence our approach as the platform matures.

The principles are fixed. They define what HUMAN-AI is, even as the platform matures. The commitments are the operational standards we hold ourselves to now, each tied wherever possible to a mechanism that lets a partner, client or agency verify it. Where a control is still maturing, we say so plainly but report progress on a regular basis.

A note on independence and sovereignty

Our shift toward an open-weight foundation model that we fine-tune and operate ourselves is designed to support greater AI sovereignty: more of the technology that touches user data sits within our own operational control, and less depends on a closed, third-party model provider. We are precise about what this does and does not mean. We do not describe HUMAN-AI as “fully open source”: the foundation model is a starting point that we fine-tune, customise and enhance into a proprietary and differentiated capability, using our own technology, processes and intellectual property. Nor do we describe it as “fully sovereign” or free of third-party dependency, while parts of our infrastructure, including underlying compute, may still rely on third-party providers. Where this document uses language about independence, we use terms such as “designed to support greater AI sovereignty”, “a more sovereign AI architecture”, “reduced dependency on closed-model providers” and “a progressive path toward AI sovereignty”, and we evidence the direction of travel as it is achieved rather than declare it in advance.

The document is structured in five parts: 1- the principles that do not change; 2- the terms we use precisely throughout; 3- the commitments we make by domain; 4- the way we hold ourselves accountable for keeping them; 5- our current compliance posture. This document is updated following material changes and we may provide periodic summaries to partners, clients and agencies where appropriate.

Part 1. Principles

Four principles anchor everything that follows. They are not features. They are the standard every component of HUMAN-AI is held to, and any technology adopted into the platform, including the foundation model at its core, is expected to meet these principles or be subject to remediation, replacement or additional controls.

Isolation by default

Each user's private data is designed to be logically and operationally separated from every other user's. The default state of all data is private to its tenant. The platform is designed so that nothing is shared by accident, by inference, or as a by-product of ordinary use, and anything that falls short of that design is treated as a defect to be fixed, never as a behaviour to be tolerated.

Data ownership

User data serves the user it belongs to. It is not used to train HUMAN-AI's models, or any third-party model, without explicit written consent. Nor is it used in any form (raw, aggregated, de-identified or derived) to improve what HUMAN-AI offers to any other tenant, except through the explicit Contribution mechanism defined in Part 2. A user's knowledge is its asset, portable on exit, and never quietly repurposed. This protection is strengthened by our move to a self-operated model: because core inference no longer depends on transmitting users prompts and outputs to a closed third-party frontier-model provider, there are fewer parties, and fewer paths, through which that data could leave our operational control.

Visible provenance

HUMAN-AI is designed to help users trace answers to their sources, distinguishing private knowledge, approved shared knowledge and model-generated synthesis where applicable, so users can interrogate and challenge any result rather than take it on trust.

Architectural independence

HUMAN-AI is designed to minimise structural dependency on any single external technology provider, including closed frontier-model providers and hyperscalers. We build on an open-weight foundation model that we fine-tune, customise and operate within an architecture we control, progressively developing it into a proprietary and differentiated capability. Where third-party components remain part of the architecture, they are chosen, configured and, where necessary, replaced according to our own standards, not the reverse. Our aim is greater control over the technologies that process user data, evidenced over time rather than declared at a point in time.

What is fixed, and what is not

These principles are fixed. The technologies that implement them are not and should not be. As stronger components emerge, parts of the architecture will change: a database, a hosting region, a foundation model, a model provider, a service. Today that includes moving from primary reliance on closed, third-party frontier models toward an open-weight foundation model we fine-tune and operate ourselves; tomorrow it may include a different foundation model, hosting environment or hardware provider. What does not change is the standard each component is held to. We commit to a set of standards, not to any single vendor.

Every commitment that follows is a standard we hold ourselves to regardless of the technology underneath it. When a component changes, including the model at the core of the platform, the standard remains.

Part 2. Defined Terms

We use the following terms precisely, and in the same sense throughout.

Tenant

The isolated environment belongs to a single user. All of our partner, client or agency's users, data, workflows and agents operate within its tenant, and access is scoped to it.

Private Knowledge Graph (PKG™)

A user's own data, scoped to its tenant. It is accessible only to authorised users, workflows and agents within that tenant. By design, it is intended not to be exposed to, queried by, or surfaced to any other partner, client or agency or Alliance member, and Section 3.1 describes both how that rule is enforced and where enforcement is still being hardened.

Alliance Knowledge Graph (AKG™)

Data that a user has deliberately and explicitly contributed for shared use. It is designed to hold nothing by default, by inference, or as a by-product of ordinary use, and is available only to Alliance members for agreed purposes.

Contribution

The deliberate, governed act by which a user moves specified data from its PKG™ into the AKG™. Contribution is always an explicit decision, never an automatic consequence of using the platform.

Derived Data

Any representation, artefact or measurement produced from user data, including embeddings, vector indexes, caches, fine-tuned model artefacts, evaluation results, aggregations and statistics. Derived Data is user data: it carries the same protections, isolation and lifecycle obligations as the data it derives from, unless the underlying data has been contributed to the AKG™. De-identification or aggregation does not, by itself, remove data from this definition.

Foundation Model

The open-weight base model on which HUMAN-AI's model capability is built. HUMAN-AI does not treat the Foundation Model, on its own, as the product, it is the starting point for fine-tuning, customisation and proprietary enhancement.

HUMAN-AI Model

The fine-tuned, customised and enhanced model derived from a Foundation Model using HUMAN-AI's own training material, techniques, evaluation and configuration. The HUMAN-AI Model, and the proprietary processes used to produce it, are HUMAN-AI's intellectual property.

Model Artefact

Any file or representation produced or used in building or running the HUMAN-AI Model, including weights, checkpoints, adapters, fine-tuning configurations and evaluation results. A Model Artefact built from user data is Derived Data and carries the same protections.

Model Licence

The licence terms under which a Foundation Model is made available (for example, an open-weight licence). We review and track Model Licence terms, including compatibility with commercial use, redistribution, fine-tuning and modification, before a Foundation Model is adopted.

Sub-processor

Any third-party provider that processes user data as part of HUMAN-AI's production architecture: for example, hosting, database, model access, or authentication. Internal tools that do not touch user data are not sub-processors and are recorded separately. A Foundation Model publisher whose weights we download and operate ourselves, and to whom no user data is transmitted, is not a Sub-processor on that basis; it is tracked instead in the model provenance and licence record described in Section 3.2.

Confirmed incident

A security event we have verified to affect user data, or the integrity, confidentiality or availability of a Model Artefact. Confirmation triggers the notification commitments in Part 3. We pursue confirmation without undue delay after first becoming aware of a suspected event, and Section 3.9 governs what partners are told while confirmation is still in progress.

Part 3. Commitments

These commitments describe HUMAN-AI's operational standards, design objectives and maturity roadmap. They describe how HUMAN-AI is being built and run, and we expect partners to hold us to them through the mechanisms in Part 4. They are not intended to create legally binding warranties or contractual liabilities, those are addressed separately in the Data Processing Agreement and the commercial terms.

Each commitment below is tied, wherever possible, to the mechanism that makes it verifiable. Where work is still in progress, the objective is dated and the evidence is named.

3.1 Data Architecture and Isolation

HUMAN-AI is designed around an architecture intended to keep private and shared data strictly separated, with governance applied to each layer.

The PKG™ holds a user's own data, scoped to its tenant and reachable only by authorised users, workflows and agents within it. The AKG™ holds only what a user has deliberately contributed. Movement between the two is a governed action, never a default.

We operate to the standard that:

- **Tenant separation** is intended to be enforced in the data layer, including through row-level security or equivalent controls where applicable. Each row is intended to scope to a tenant, and the system is designed to prevent one tenant from accessing another tenant's data through the application path.
- **Membership and role definitions** are enforced at the data layer, not by application logic alone, so what a user can see and do is governed by the data itself.
- **Sensitive credentials and tokens** are scoped to the individual user who owns them, limiting the blast radius of any single compromise.

The core isolation model has been designed and is progressively hardened across relevant access paths. This work will be evidenced through version-controlled migration history, isolation review and related technical documentation. The objective being to strengthen isolation so that it is enforced consistently at the appropriate technical layers, not only through application logic.

The same isolation standard extends to the model layer. Inference over a user's PKG™ is scoped to that tenant, and Section 3.2 describes how that is enforced now that model inference runs within an architecture we operate, rather than through a shared third-party model endpoint.

3.1a AKG™ Governance

The AKG™ is the platform's one deliberate sharing surface, and it is governed with the same rigour as isolation itself. We operate to the standard that:

- **Contribution is informed and specific**

At the moment of Contribution, the contributing user approves and sees exactly what data will move, who can access it, and for what agreed purposes. Consent to one Contribution does not extend to another.

- **Attribution is the contributor's choice**

Each Contribution carries the attribution level the user selects (named, pseudonymous, or anonymised), fixed at contribution time and not degraded afterwards.

- **Re-identification is prohibited**

Alliance members may not attempt to re-identify a contributor, reverse-engineer a contributor's client base, or otherwise use AKG™ material beyond the agreed purposes. A violation is a material breach of Alliance membership.

- **Competitive conflicts are addressed in membership terms**

The Alliance membership agreement states how contributions are handled where members compete directly, so that no partner, client or agency discovers the answer only after contributing.

3.2 AI Model Governance, Provenance and Data Use

HUMAN-AI's model strategy is built on an open-weight Foundation Model, which we fine-tune, customise and operate within an architecture we control. This section sets out how we select, verify, govern, deploy and evaluate that model, and how user data is, and is not, used in the process.

3.2.1 Model Selection and Approval

No Foundation Model is adopted for production use without documented review and approval. Before adoption, we assess the Model Licence and its compatibility with commercial fine-tuning, redistribution and our intended use; the model's documented training data and known provenance; its security and safety evaluation history; and its technical fit for our architecture and performance requirements. Approval is recorded, including the reviewer, the date and the version approved. Material changes to an adopted Foundation Model, or a decision to change Foundation Model, follow the same process.

3.2.2 Provenance and Supply-Chain Risk

We maintain a record of the origin of each Foundation Model we adopt, including its publisher, release version, and the checksum or signature used to verify the integrity of the weights we download and deploy. Model Artefacts are stored in a controlled registry, version-pinned, and are not modified outside our own fine-tuning and evaluation pipeline. Where a Foundation Model or its dependencies are found to carry a known vulnerability or integrity concern, we treat this as we would any other supply-chain finding. It is triaged, remediated, or the component is replaced.

3.2.3 Fine-Tuning, Customisation and Ownership

Fine-tuning, customisation, orchestration, prompt design, evaluation frameworks, and the resulting HUMAN-AI Model are built using our own technology, processes and intellectual property, and are HUMAN-AI's proprietary assets.

Our standard is that fine-tuning draws on HUMAN-AI's own proprietary and licensed training material, synthetic data and evaluation sets, not a user's raw PKG™ content, unless that user has given explicit written consent to a specific fine-tuning use. Where a user does consent, the scope of that consent is documented in the same way as a Contribution.

3.2.4 Deployment, Hosting and Inference Control

The HUMAN-AI Model is deployed and run within an architecture we control, rather than through a third-party frontier-model API for core inference. Prompts, retrieved context and outputs generated in serving a user's request are processed within our own inference environment and are not transmitted to a third-party model provider for that purpose.

3.2.5 Model Evaluation, Safety and Bias Testing

The HUMAN-AI Model is evaluated on representative workloads before it is deployed and after any material fine-tuning or configuration change, covering accuracy, known failure modes, and bias across use cases relevant to our users. Evaluation results are recorded and reviewed as part of the model approval process in Section 3.2.1.

3.2.6 Portability and Vendor Independence

Because we hold our own copy of the Foundation Model and the fine-tuning pipeline built on it, we are not dependent on continued access to a single third-party model provider to keep the HUMAN-AI Model running, or to keep improving it. This also supports user portability because a user's PKG™ remains exportable regardless of which Foundation Model underlies the platform at a given time, consistent with the standard in Section 3.11.

Data use across the model layer

The same data-use standard applies across the model layer as elsewhere in the platform. User data is not used to train or fine-tune the HUMAN-AI Model, or any third-party model, without explicit written consent.

The same standard applies below the level of training. Derived Data, including embeddings, evaluation results and fine-tuned Model Artefacts, is user data: we do not use one user's data or Derived Data to build features, prompts, heuristics, benchmarks, evaluations or any other improvement that is exposed to, or benefits, any other tenant, unless the user has contributed the underlying data to the AKG™ or given explicit written consent. De-identification or aggregation alone does not authorise such use. Operational telemetry that contains no user content (service metrics, error rates, performance data) sits outside this rule and is used to run the platform.

Because core inference runs within our own architecture rather than through a third-party frontier-model API, fewer parties see user prompts and outputs in the ordinary course of operating the platform. Where we do use a third-party model for a specific, limited purpose, we seek terms and configurations that restrict training, fine-tuning, retention and reuse of submitted data, on the same basis as our own model.

Outputs belong to the user they are generated for. As between HUMAN-AI and the user, outputs produced from a user's PKG™ are the user's, and HUMAN-AI claims no rights in them beyond what is required to operate the platform.

HUMAN-AI is designed to provide provenance so users can understand whether an output is informed by private knowledge, approved shared knowledge or model-generated synthesis. Sub-processor and model-provenance information is available as described in Section 3.6.

3.3 Encryption and Data Protection

HUMAN-AI's data protection approach includes protecting relevant production data in transit and at rest through appropriate technical safeguards. These protections are intended to operate as platform-level defaults where applicable.

- **In transit:** Encryption in transit follows current industry-standard protocols for connections handling user or client data.
- **At rest:** Encryption at rest uses industry-standard encryption provided by the relevant infrastructure or service provider.
- **Key management:** Keys are managed through the infrastructure provider's managed key service, with access enforced at both platform and provider levels.
- **Model Artefacts:** Model Artefacts, including fine-tuned weights, checkpoints and adapters, are held in access-controlled, encrypted storage, version-pinned in the model registry described in Section 3.2, and are not exposed outside the fine-tuning and deployment pipeline.

3.4 Identity, Authentication and Access

HUMAN-AI's access model is designed around least privilege principles so that users and agents receive only the permissions their role requires.

- **Authentication**

The target access model includes controlled authentication and, where supported, domain-restricted provisioning or equivalent safeguards.

- **Authorisation**

Role-based access control with distinct owner, admin and member roles, administered by each partner, client or agency for its own users.

- **Agents**

Workflow agents operate under scoped permissions aligned to the role they act under and are designed not to exceed that role's access.

- **Administration**

Each partner, client or agency provisions, assigns roles to, and de-provisions its own users.

- **Provider access**

HUMAN-AI personnel do not access tenant content in the ordinary course of operating the platform. Support access to a user's tenant requires that user's permission, or a documented break-glass justification in an emergency; it is time-limited, scoped to the need, and always logged. Each user can obtain the record of access to its own tenant.

- **Model administration**

Access to the model registry, fine-tuning pipeline and deployment configuration is restricted to authorised platform engineering roles and logged on the same basis as other privileged access.

The user-facing model above is being implemented and strengthened as the platform matures. The roadmap includes extending equivalent controls to internal platform administration and privileged operational activity.

3.5 Infrastructure, Hosting and Residency

HUMAN-AI's current and intended production configuration uses managed infrastructure in the EU region for core application and database services, subject to configuration and technical validation. Core private tenant data is intended to be hosted in the relevant documented production region, while specific processing locations may vary by service.

HUMAN-AI also relies on graph database, hosting and content-delivery providers, which may operate in different regions. Rather than make a broad residency claim we cannot substantiate for every component, we document the operating region and residency commitment for each provider in the sub-processor register (Section 3.6). Where any sub-processor processes user data outside the EEA or the UK, the transfer rests on a valid mechanism such as an

adequacy decision, Standard Contractual Clauses supported by a transfer impact assessment, or the UK IDTA/addendum as applicable. The mechanism relied on is recorded per provider in the register. Where a user has a specific residency or data-localisation requirement, we can assess the relevant configuration and confirm the applicable position.

Model hosting and inference infrastructure follows the same residency discipline.

Production and local development are separated today. A dedicated, isolated staging tier, used to validate every change before release, is a roadmap near-term priority.

3.6 Sub-processors and Vendor Governance

HUMAN-AI relies on a defined set of third-party providers for hosting, database and graph services, model-related infrastructure, authentication, monitoring and communications. We maintain a vendor due diligence register that records, for providers in the production architecture:

- its role and the services it provides;
- whether it processes user data, and for what purpose;
- the contractual safeguards in place, including data-use, retention and incident-handling terms;
- its primary data location or residency commitment, and the international transfer mechanism relied on where applicable.

We also maintain a model provenance and licence record, described in Section 3.2, covering each Foundation Model we adopt: its publisher, licence terms, version and integrity verification. A Foundation Model publisher to whom no user data is transmitted is tracked in that record rather than the sub-processor register; where a third party is involved in hosting, training or serving the model and does receive user data or telemetry, it is treated as a sub-processor under the standard above.

We distinguish clearly between providers that form part of HUMAN-AI's production architecture, which handle user data and are held to these standards, and internal tools used only for development or record-keeping, which do not. Access to the register follows the nature of the data involved. Where HUMAN-AI processes personal data on behalf of a user, that user receives the current list of sub-processors by name, with each provider's role, location and safeguards, as part of the Data Processing Agreement, and is informed of any intended addition or replacement in time to object before it takes effect, as Article 28 GDPR requires. Where an engagement involves no personal data and the concern is confidentiality alone, no such obligation applies, and sub-processor information may be shared at HUMAN-AI's discretion under appropriate confidentiality arrangements.

3.7 Secure Development

Security and governance are built into the development lifecycle, not added after it.

- All database schema changes are version-controlled as sequential, named migrations, a traceable history of material data model changes, including relevant isolation and hardening steps.
- Deployments run through an automated build pipeline, which standardises release steps and reduces human error.
- Secrets and credentials are expected to be held in managed secret storage and should not be embedded in code or configuration.
- All personnel with access to production systems or user data are bound by written confidentiality obligations, are vetted appropriately to their role, and complete security and data-protection training on joining and at least annually.

- Fine-tuning and evaluation pipelines, and the Model Artefacts they produce, are version-controlled and subject to the same peer-review and vulnerability-scanning discipline as the rest of the codebase, extended to cover model-specific dependencies such as training and inference libraries, adapters and configuration.

We formalise peer review of every change before release, and automated dependency-vulnerability scanning, into documented and evidenced practices, and consolidating serverless functions to the discipline applied to the core services.

The roadmap includes introducing an independent penetration-testing review cadence at the appropriate stage of product maturity, extended to cover the model-serving environment. High-level categories of findings and remediation status may be shared with selected users where appropriate and under confidentiality arrangements.

3.8 Logging, Monitoring and Audit

HUMAN-AI is being designed to log relevant system behaviour and access activity to support security monitoring, operational insight and audit.

- The roadmap includes tenant-scoped audit records for access to user data and significant administrative and agent actions across the platform's main subsystems.
- Logs and metrics are intended to support detection, investigation and operational review.
- Model deployments are versioned, and inference activity is logged to a standard that supports rollback to a prior Model Artefact if a material issue is identified.

3.9 Incident Response

HUMAN-AI's roadmap includes a defined incident-response process covering detection, triage, containment, assessment, remediation, notification and post-incident review.

Where HUMAN-AI acts as processor and becomes aware of a confirmed personal data breach affecting client personal data, HUMAN-AI will notify the relevant controller without undue delay and in accordance with the applicable Data Processing Agreement. The same discipline applies beyond personal data. In the event of any confirmed security incident affecting user data, we will notify the affected user within 72 hours of confirmation and provide the information needed to support the user's regulatory and operational requirements. Where confirmation is still in progress, we inform an affected user of a suspected incident that is reasonably likely to involve its data, so the user can manage its own obligations (including, where it acts as controller, its own Article 33 notification clock) without waiting for our confirmation to complete.

Confirmed incidents include events affecting the integrity, confidentiality or availability of a Model Artefact, for example unauthorised modification of fine-tuned weights, or a confirmed instance of user data being incorporated into a model asset without consent, on the same notification timeline as any other confirmed incident affecting user data. Confirmed incidents should be recorded in an incident register or equivalent security record.

3.10 Resilience, Backup and Continuity

HUMAN-AI's resilience approach includes, or is intended to include managed backups. Backup data is protected through appropriate encryption and provider-level safeguards. Model Artefacts are versioned and backed up so that a prior, validated version can be restored if a deployment issue is identified. Recovery expectations will be communicated externally only once tested and validated.

3.11 Retention, Deletion and Exit

We treat data lifecycle as a first-class governance concern, and we design so that we have no incentive to hold user data longer than the engagement requires.

- **Retention**

User data is retained in accordance with the applicable contractual terms, documented client instructions and applicable law.

- **Deletion**

On request or termination, deletion of user data, including Derived Data, is handled in accordance with agreed contractual terms, including applicable backup retention and expiry processes.

- **Portability**

HUMAN-AI's roadmap includes supporting appropriate data export and portability mechanisms. A user can export its own data at any time. Its PKG™ is its asset, and HUMAN-AI does not retain it beyond agreed windows once the relationship ends.

Because the HUMAN-AI Model is fine-tuned rather than used only at inference time, deletion at the model layer means that fine-tuning draws on HUMAN-AI's own proprietary and licensed material rather than a user's raw PKG™ content, so that a user's ordinary deletion rights are unaffected by the existence of the HUMAN-AI Model.

These controls are being developed and evidenced as part of HUMAN-AI's data lifecycle governance programme: GDPR-aligned data-export and data-deletion processes, a data-subject-request workflow, consent records, and an executable Data Processing Agreement. Relevant materials may be made available to selected users under appropriate confidentiality arrangements.

3.12 Responsible AI and Acceptable Use

HUMAN-AI is built to augment human strategists and operators, not to act unaccountably.

- Significant agent actions are intended to be governed through review, approval or logging workflows appropriate to their risk.
- Strategy outputs and material workflow changes should remain subject to appropriate human review and approval.
- Configuration changes to agents and knowledge sources are logged.
- The HUMAN-AI Model is evaluated on representative workloads before deployment and after any material fine-tuning or configuration change, covering accuracy, known failure modes and bias, as described in Section 3.2.5, and the results are recorded.
- Adoption of a new Foundation Model version, or redeployment of the HUMAN-AI Model following material fine-tuning, follows the approval process in Section 3.2.1 and is not automatic.
- Where a user's own legal or client obligations require AI-assisted content to be identifiable as such, HUMAN-AI's provenance model (Section 3.2) supports that disclosure.

EU AI Act

We track HUMAN-AI's positioning under the EU AI Act as its obligations phase in: our role for each capability (provider or deployer), our reliance on general-purpose AI models, and any transparency obligations that attach to specific features. Because HUMAN-AI builds on an open-weight Foundation Model that we fine-tune ourselves, our position under the Act's general-purpose AI model provisions, including any obligations that attach to substantial modification of a general-purpose model, forms part of that assessment.

We record that positioning in the compliance record described in Part 4. A formal Acceptable Use and Responsible AI policy is part of the wider governance roadmap.

3.13 Confidentiality, Marketing and References

User data, client data, proprietary information, platform outputs, use cases and performance results are not used for marketing, promotion, public reference, case studies, press or sales materials without the user's explicit written permission.

Any case study, testimonial or reference is governed by separate written approval that specifies the content shared, the level of detail, and whether partner, client or agency names, logos or third-party references may be used.

Permission granted for one purpose does not extend to another.

Our fine-tuning methodology, training data selection, evaluation results and model configuration are themselves proprietary and confidential and are not disclosed beyond what is necessary to give user the assurance described in Part 4.

Part 4. How We Hold Ourselves Accountable

A commitment is only as credible as a user's ability to check it. We treat assurance as a standing programme, not a one-off statement, and we commit to giving users appropriate assurance as the platform and control evidence mature.

Ownership and review

Our DPO is accountable for this programme, with security and privacy responsibilities assigned and recorded. Model selection, licensing and fine-tuning decisions are approved through the process in Section 3.2.1 and recorded as part of the same governance programme, under the same named executive owner. This document sits at the top of a defined documentation set: beneath it, the Data Processing Agreement and commercial terms carry the binding obligations, and the Acceptable Use and Responsible AI policy and internal security policies carry the operational detail. It is reviewed at least annually, and after any material change, and its version history is preserved.

A single, current record

HUMAN-AI is maintaining a current internal record of its security and compliance posture, including controls, sub-processors, model provenance and licensing, material changes, and assurance progress.

How partners, clients and agencies stay current.

- On change: material events (new controls, sub-processor changes, a change of Foundation Model, certification milestones, or disclosed incidents) are communicated when they occur.
- Periodic updates may be provided where appropriate, alongside event-driven communications for material changes.
- Relevant assurance materials may be made available to selected users under appropriate confidentiality arrangements and subject to the maturity of the relevant controls.

Co-design and challenge

Early partners, clients and agencies may provide input into HUMAN-AI's governance framework, particularly where their legal, compliance, IT-security and client requirements are relevant.

Part 5. Compliance Posture

We are explicit about where we stand. Misrepresentation would undermine trust and is incompatible with the principles in Part 1.

Established today

HUMAN-AI is being designed and operated with GDPR principles at its core, including data minimisation, purpose limitation, security, retention and data-subject rights. A Data Processing Agreement will be made available where HUMAN-AI processes personal data on behalf of partners, clients or agencies. Data-subject-rights processes are part of HUMAN-AI's privacy governance roadmap and will be documented as they mature. We support partners, clients and agencies' data protection impact assessments with documentation on request, and a named privacy contact is available to partners and their DPOs.

Open-weight model use

Where HUMAN-AI builds on an open-weight Foundation Model, we track the applicable Model Licence and its compatibility with our commercial use and record that assessment alongside our other compliance documentation. We describe our model strategy as designed to support greater AI sovereignty and reduced dependency on closed model providers; we do not describe HUMAN-AI as fully sovereign or free of third-party dependency while any part of our infrastructure, including compute, may still rely on third-party providers, and we will update this statement as that position evolves.

Certification roadmap

Future independent assurance will include SOC 2 Type II, ISO 27001 or equivalent frameworks, depending on partner, client or agency requirements, product maturity and company priorities.

We commit to standards, not vendors. The technology will change but the standard we hold it to will not. You will see the programme evolve as it happens, and you are invited to hold us to your standard at every stage.