DEFINED.AI

Defined.ai
AI Governance Pack

External Use

May 2026

Table of Contents

Defined.ai AI Policy

Intro

Defined.ai's AI Policy outlines a comprehensive, ISO-aligned governance framework that governs the full lifecycle of data and AI-assisted activities, from ethical sourcing and provenance through secure decommissioning. Positioned primarily as a responsible data partner rather than an AI deployer, Defined.ai emphasizes lawful processing, transparency, privacy, security, human oversight and accountability across all operations. The policy details clearly assigned roles, risk management processes, approved AI tool controls, incident response and continuous improvement mechanisms, aligning with global regulations such as GDPR, CCPA and the EU AI Act, as well as international standards including ISO/IEC 27001, ISO/IEC 27701 and ISO/IEC 42001. Together, these controls ensure that AI related activities are ethical, auditable, risk managed and fit for evolving regulatory and stakeholder expectations.

Policy

What is Defined.ai's approach to AI governance?

Defined.ai operates a structured AI Governance Program embedded within its Integrated Management System and AI Management System. The program governs the full lifecycle of data and AI-assisted activities, including sourcing, curation, annotation, validation, delivery, monitoring and secure decommissioning.

AI governance is operationalized through documented policies, procedures, risk registers, training requirements, audit trails and incident response mechanisms. The program is subject to continuous review and improvement to ensure lawful processing, ethical sourcing, transparency, privacy, security and accountability across all AI-related operations.

To this end, Defined.ai sets measurable objectives that reflect international best practices.

What role does Defined.ai play in the AI ecosystem?

Defined.ai acts primarily as a data partner, collecting, curating, annotating, validating and supplying datasets with documented provenance and compliance safeguards to third parties building AI systems.

Defined.ai may also act as an AI user internally, under controlled conditions, using approved AI tools to support business operations. Defined.ai does not place AI systems on the market, nor does it act as a deployer of AI systems making decisions about individuals.

Does Defined.ai develop or deploy AI systems that make automated decisions about individuals?

No. Defined.ai does not develop or deploy proprietary AI systems that make significant or legally relevant automated decisions about individuals.

Its activities focus on responsible data operations and internal AI-assisted workflows under human oversight. Where AI tools are used internally, they do not operate autonomously or replace human decision-making.

How is accountability structured within Defined.ai's AI governance?

Accountability is clearly defined and assigned across Legal, Privacy, Security, Data Governance and operational teams, with responsibilities aligned to specific decision types.

Oversight mechanisms include:

- Governance committee involvement where required;
- Documented roles, responsibilities and escalation paths;
- Defined ownership for risk identification, mitigation and acceptance;
- Auditability, logging and traceability of decisions;
- Defined remediation and corrective action processes.

Accountability is maintained throughout the full data and AI lifecycle.

Which laws and standards does Defined.ai align with?

Defined.ai's governance aligns with applicable frameworks including but not limited to:

- General Data Protection Regulation (GDPR);
- California Consumer Privacy Act (CCPA);
- EU Digital Services Act;
- EU AI Act;
- Directive on Copyright in the Digital Single Market.

Governance is further informed by international standards and best practices, including ISO/IEC 27001, ISO/IEC 27701, ISO/IEC 42001 and relevant National Institute of Standards and Technology (NIST) guidance. Regulatory developments are actively monitored, and governance controls are updated accordingly.

How does Defined.ai ensure ethical data sourcing?

Defined.ai applies ethical data sourcing practices supported by:

- Supplier vetting and due diligence processes;
- Ethical questionnaires and periodic reviews;
- Contractual safeguards and codes of conduct;
- Dataset provenance documentation;
- Respect for consent, opt-out and lawful basis requirements where applicable.

These measures are designed to support lawful, transparent and responsible dataset creation.

What does “data provenance” means at Defined.ai?

At Defined.ai, data provenance refers to the documented traceability of a dataset’s origin and handling throughout its lifecycle. This includes records of how data was sourced, curated, annotated, validated and delivered, as well as any applicable restrictions, conditions of use or risk considerations associated with the dataset.

How does Defined.ai protect privacy and personal data?

Defined.ai implements technical and organizational safeguards, including:

- Role-based access controls;
- Encryption at rest and in transit;
- Secure storage and transfer protocols;
- Data minimization and purpose limitation principles;
- Incident and breach response procedures;
- Consent and opt-out handling aligned with legal requirements.

These safeguards are integrated into Defined.ai’s broader privacy and information security framework.

How does Defined.ai address bias and fairness?

Defined.ai integrates fairness considerations into its data collection and annotation processes through:

- Bias-aware dataset practices;
- Context-appropriate demographic considerations where relevant and lawful;
- Ethical annotation guidelines;
- Ongoing quality and risk monitoring.

These controls are applied in a risk-based and context-aware manner. Defined.ai does not claim that datasets are free from bias, but applies reasonable measures to identify, document and mitigate known risks where feasible.

What human oversight exists in Defined.ai’s AI activities?

Human oversight is a core governance principle. Defined.ai assigns trained personnel to oversee AI-assisted workflows, with authority to intervene, pause, escalate or stop activities where risks arise.

Human review is mandatory for AI-generated outputs unless a justified exception is approved under documented governance controls.

How does Defined.ai manage AI tools and automation?

Defined.ai maintains an Approved AI Tools Inventory. New tools are subject to a review process that may include:

- Business justification;
- Privacy and security assessments;

- Data Protection Impact Assessments (DPIAs) where required;
- Registration, monitoring and usage conditions.

Use of unapproved AI tools is restricted and may result in corrective or disciplinary actions in line with company policies.

How does Defined.ai manage AI and data risks?

Defined.ai applies a structured risk management methodology aligned with its AI Risk Management Framework.

This includes:

- Risk identification and assessment;
- DPIAs and impact assessments for higher-risk activities;
- Documented mitigation measures and controls;
- Periodic reviews, including quarterly risk reviews;
- Alignment with regulatory expectations and internal audit processes.

Are there prohibited AI uses at Defined.ai?

Yes. Defined.ai does not permit:

- Inputting confidential or personal data into unapproved tools;
- AI-driven decisions about individuals without safeguards;
- AI use without lawful basis or governance controls;
- Deployment of autonomous AI agents without approval.

What happens if an incident occurs?

Defined.ai follows a structured incident response process, which includes:

- Detection and classification of the incident;
- Containment and remediation actions;
- Internal and external stakeholder communication where required;
- Post-incident reviews and corrective or preventive actions.

How does Defined.ai ensure continuous improvement?

Continuous improvement is supported through:

- Internal audits and management reviews;
- Performance indicators and quality metrics;
- Client and stakeholder feedback;
- Regulatory monitoring;
- Periodic policy, procedure and control updates.

How does Defined.ai oversee third-party vendors and data partners?

Third-party providers are subject to governance processes overseen by cross-functional teams. Vendor oversight may include:

- Compliance, privacy and security evaluations;
- Contractual safeguards and audit rights;
- Ongoing monitoring and review;
- Incident notification obligations;
- Ethical and sustainability considerations where relevant.

How does Defined.ai govern the full lifecycle of data and AI activities?

Defined.ai applies lifecycle governance covering:

- Selection and procurement;
- Integration, curation and validation;
- Operation and monitoring;
- Change management;
- Secure decommissioning and retention management.

Lifecycle controls are embedded into operational workflows and governance documentation.

How does Defined.ai engage with stakeholders?

Defined.ai recognizes the impact of AI on clients, contributors, regulators and society. The company supports transparency through clear communication, documentation and governance practices designed to meet stakeholder expectations and applicable regulatory obligations.

FAQ: Ethically Sourced Datasets and AI Model Services

Intro

This FAQ serves as a structured, evidence-based mechanism of how Defined.ai embeds ethical, transparent and governance-ready practices across its data and AI model services. Anchored in internationally recognized frameworks—including the OECD AI Principles, the NIST AI Risk Management Framework and relevant IEEE standards—it provides a verifiable account of the company's approach to responsible data sourcing, consent traceability and end-to-end data lineage.

By documenting processes for bias detection and mitigation, transparency of dataset limitations and rigorous validation and benchmarking, the FAQ highlights Defined.ai's commitment to fairness and accountability. It also outlines operational safeguards such as separation of duties, continuous monitoring, reproducibility and audit readiness.

This framing positions Defined.ai as a trustworthy and governance aligned data partner, capable of supplying ethically sourced, well-documented datasets that support responsible AI development.

Governance and Accountability

Is there a clearly defined governance structure for data and model lifecycle oversight?

Yes. Defined.ai operates under a mature, ISO-aligned governance framework that provides structured oversight of the entire data and model lifecycle. A dedicated team ensures controls, risk assessments, audits and approvals are consistently applied from data sourcing through model decommissioning.

Are roles and responsibilities for ethical decision-making explicitly documented?

Yes. Ethical, privacy and security responsibilities are explicitly documented across our Information Security and Privacy Management System and internal policies. Teams and suppliers follow defined obligations on fairness, accuracy, explainability and responsible data stewardship and employees are trained on these duties.

Are escalation and remediation mechanisms clearly defined?

Yes. Defined.ai maintains formal escalation, audit and proactive remediation procedures under its integrated ISO/IEC 27001 and ISO/IEC 27701 controls. Risks and incidents are escalated to cybersecurity, privacy and governance leads, with structured remediation steps and continuous monitoring embedded in our operational model.

Ethical Data Sourcing and Consent

What is the documented source of each dataset?

Defined.ai requires full dataset provenance documentation for every delivered asset. Suppliers must provide a written list identifying the source of each dataset element, along with the specific assets covered by each consent or release. This ensures traceability and lawful acquisition at the asset level.

How is informed consent obtained and recorded?

Defined.ai secures informed consent through three controlled pathways depending on the origin of the data:

1) **In-house data collections (proprietary datasets)**

   a) Participants are onboarded under fit-for-purpose Data Collection Agreements that explicitly authorize AI/machine learning use, define rights and permitted uses and outline revocation conditions. Consent is formally recorded and stored as part of the internal dataset package, ensuring consistent, contractually sound and fully documented consent across all proprietary data assets.

2) **External supplier-provided datasets**

   a) Suppliers must obtain and deliver "Required Consents"—valid, written model or property releases—before any dataset is accepted. Each consent must map to the specific asset it covers and suppliers must provide a traceability list linking each dataset element to its corresponding consent, ensuring lawful sourcing and full auditability.

3) **Data collected via crowdsourcing**

   a) Collections through Neevo, Defined.ai's proprietary crowdsourcing platform, operate under a dedicated informed-consent framework, within which contributors are explicitly informed of:

      i) The purpose of the collection;

      ii) The intended AI/machine learning uses;

      iii) Any downstream licensing implications and their rights, including withdrawal.

Contributors must provide affirmative consent before participating, and Neevo's workflow records and stores this authorization as part of the dataset metadata. This guarantees transparent, documented and rights-respecting consent for all crowd-generated assets. Together, these pathways ensure a uniform, legally robust consent posture across all data sources used by Defined.ai.

Can consent be traced at a granular (dataset-slice or contributor) level?

Yes. Required Consents match individual assets; this applies both to crowd members and non-crowdsourced data collection participants. At their end, Suppliers must provide a mapping list linking each consent to the specific dataset element it covers. This enables contributor or slice-level verification when needed.

How are data subject rights (erasure, correction, opt-out) operationalized?

Defined.ai operationalizes data subject rights through an ISO-aligned Privacy Management System. It is supported by clearly documented processes that allow the organization to identify individuals, locate their data across the lifecycle and fulfill erasure, correction or opt-out requests efficiently. These procedures are reinforced through continuous governance reviews and dedicated workflows managed by the Privacy team. This framework is strengthened by the Defined.ai Privacy Program, a structured organization-wide governance system that sets policies, assigns responsibilities and ensures consistent implementation of privacy requirements.

The Legal, Privacy and Compliance Team closely monitors the designated Privacy Portal for any requests from data subjects to exercise their rights and resolves them in a timely manner.

Finally, Defined.ai has a dedicated Data Protection Officer (DPO) who oversees compliance, ensures impartial and expert guidance, manages interactions with supervisory authorities and monitors GDPR-aligned rights handling. This combination of structured program governance and DPO oversight ensures that all data subject rights are handled lawfully, promptly and transparently.

Are third-party data providers audited for ethical compliance?

Yes. All data partners undergo a structured Ethical Dataset Compliance review. This includes an Ethical Questionnaire and a partner onboarding workflow assessing legal compliance, consent validity, security posture and responsible data practices. This is reinforced contractually through the Supplier Code of Conduct and Data Licensing Agreement.

Data Provenance, Lineage and Governance

Is end-to-end data lineage documented and auditable?

Yes. Defined.ai maintains full end-to-end data lineage for every dataset—including custom workflow-based collections and off-the-shelf (OTS) datasets—supported by strict provenance and traceability requirements. The Supplier Code of Conduct mandates a written, item-level mapping of each dataset asset to its original source and the corresponding consent or release, ensuring lawful and ethical sourcing across all modalities. The Ethical Dataset Compliance process requires partners to provide comprehensive lineage details—including origin, ownership, permissions, collection workflows, labeling steps and transformation history—as part of onboarding. For OTS datasets specifically, Defined.ai's provenance framework ensures that every file can be traced back to licensed partners or content providers, with clear source-tracking, consent documentation and lineage records recorded in audit-ready form. This creates a complete, independently auditable chain of custody across both bespoke workflow projects and OTS catalog assets, enabling verification of dataset origin, lawful basis and permitted use.

Are data transformations, labeling steps and augmentations traceable?

Yes. Defined.ai ensures full traceability across every stage of dataset development. The Ethical Dataset Compliance Questionnaire requires suppliers to disclose each step in the data lifecycle in detail, including acquisition workflows, labeling procedures, quality-review layers, partial deliveries and any transformations or augmentations applied to the data. Each dataset must clearly document these steps, creating a transparent, reviewable audit trail. This enables Defined.ai—and its

clients—to understand precisely how the dataset was created, validated and modified, ensuring reliability, reproducibility and compliance throughout the workflow.

What controls exist to prevent unauthorized data access or misuse?

Defined.ai enforces strict technical organizational and contractual controls aligned with its ISO/IEC 27001 and ISO/IEC 27701 Information Security and Privacy Management Systems. In addition to secure infrastructure, access restriction and continuous monitoring, all Supplier Data Collection Agreements and contractual frameworks ensure that permitted use is validated and secured across the entire chain, covering source rights, participant consent and alignment with the end-client's intended purpose. These layered safeguards prevent unauthorized access, duplication, misuse or disclosure and ensure that all datasets are handled lawfully, securely and in accordance with approved use conditions.

Is lineage information made available to downstream users?

Yes. All datasets delivered to clients include the required provenance, releases and consent documentation, ensuring downstream users can verify the lawfulness, provenance and composition of the data. Suppliers must also provide a traceability list linking each consent to the exact dataset asset, enabling downstream review where required.

Bias, Fairness and Representativeness

What bias identification techniques are used during dataset creation?

Defined.ai requires suppliers and partners to disclose bias-relevant characteristics during dataset onboarding, using the Ethical Dataset Compliance Questionnaire. This includes documenting dataset composition, acquisition context, labeling methods and any known fairness considerations, forming the basis for bias detection and review during dataset creation.

Are demographic and contextual skews quantitatively measured?

Yes. The Ethical Questionnaire requires partners to provide structured dataset details, including demographic attributes when available, enabling the organization to analyze representativeness, contextual skews and potential imbalance across each dataset. This enables quantitative assessment during the compliance review phase.

What mitigation strategies are applied to identified biases?

Suppliers must ensure fairness and responsible data stewardship, including equitable and lawful collection methods. Defined.ai conducts risk-based dataset reviews, requiring corrective actions, alternative sourcing or re-labeling if fairness or representational gaps are identified. A structured onboarding workflow allows the organization to reject, request modification to or constrain use of datasets that exhibit unacceptable biases.

Are residual risks clearly disclosed to users?

Yes. Dataset provenance packages provided to downstream users include the consent, documentation and dataset-level disclosures required for transparent evaluation. Where suppliers provide information on limitations or constraints, this is transmitted as part of the dataset metadata, supporting clear communication of remaining risks.

Transparency and Documentation

Are dataset documentation artifacts provided (purpose, scope, limitations)?

Yes. Through the Ethical Dataset Compliance Questionnaire, suppliers must provide dataset-level documentation, including purpose, acquisition context, datatype-specific details and any limitations or unknowns. This ensures each dataset arrives with a structured, reviewable documentation package.

Are known limitations and failure modes explicitly disclosed?

Yes. Dataset providers must disclose limitations, unknowns, partial acquisitions and contextual constraints in the dataset tabs of the Ethical Questionnaire. These required disclosures allow Defined.ai to surface relevant risks, incomplete data conditions and model-relevant failure considerations upon delivery. The same logic presides over Defined.ai proprietary OTS data.

Is documentation updated as datasets or models evolve?

Yes. Defined.ai's dataset onboarding workflow requires maintaining versioned dataset documentation. When datasets are updated—whether through new deliveries, transformations or augmented content—the corresponding dataset documentation must be refreshed and resubmitted. The structured, tab-based questionnaire format makes updates repeatable and traceable.

Model Quality, Robustness and Performance

What performance metrics are used and why were they selected?

Defined.ai uses structured, repeatable and domain-appropriate evaluation metrics tailored to the model's modality and intended use. Metrics include naturalness, intelligibility, dialect accuracy, latency, error rates, preference tests and human-in-the-loop quality signals. These were selected because they offer objective, comparable measurements of quality to support ongoing improvement across training cycles.

Are models evaluated under stress, edge cases and distribution shifts?

Yes. Defined.ai frequently conducts adversarial and edge-case evaluations, including red-teaming, attack-vector exploration and stumping exercises designed to expose vulnerabilities and out-of-distribution weaknesses. These methodologies stress-test models under non-ideal inputs to identify failure risks early.

How does model performance degrade under non-ideal conditions?

Through large-scale adversarial testing and human-evaluated benchmarking, Defined.ai identifies patterns such as increased error rates, susceptibility to incorrect outputs, hallucinations and context-loss under stress or atypical scenarios. Results from structured stumping campaigns show how models degrade when exposed to rare concepts, ambiguous prompts or adversarial attacks.

Are retraining and recalibration strategies defined?

Yes. Defined.ai's benchmarking framework incorporates continuous evaluation loops, allowing model owners to retrain after each training cycle or milestone using updated error analyses, regression tracking and crowdsourced human feedback. Fine-tuning, rewriting, demographic tailoring, reinforcement learning with human feedback and direct preference optimization are used to recalibrate and optimize models based on observed performance gaps.

Risk Management and Monitoring

Are risks mapped across intended and foreseeable unintended uses?

Yes. Defined.ai maps risks across both intended and foreseeable unintended uses through its performance-requirements governance framework, which documents approved use cases, escalation paths and intervention thresholds (e.g. precision, recall, fairness, robustness), ensuring models are evaluated under real-world, misuse and edge-case conditions.

Is post-deployment monitoring in place for drift, bias and misuse?

Yes. Defined.ai supports robust post-deployment monitoring through ongoing benchmarking, adversarial and stumping evaluations and continuous human-in-the-loop reviews. These processes enable early detection of model drift, emerging bias patterns and potentially harmful misuse, forming a continuous monitoring loop used by both internal teams and clients.

How frequently are risks reassessed?

Risks are reassessed on a continuous or milestone-based cadence, depending on the evaluation program. Defined.ai performs ongoing regression testing, monthly or cycle-based benchmarking updates and iterative performance reviews aligned with model retraining cycles, ensuring that risk insights remain current throughout the system lifecycle.

Are users notified of significant changes or newly identified risks?

Yes. Defined.ai's governance framework requires transparent communication of any significant changes, including new risks, updated performance requirements and revised safety considerations. Users are informed whenever risk thresholds shift or new conditions arise that may affect safe-to-use expectations, consistent with transparency, accountability and documentation principles.

Auditability and Independent Validation

Can an independent reviewer reproduce key datasets or model claims?

Yes. Defined.ai's governance framework requires complete dataset provenance, consent mappings and detailed acquisition documentation, enabling an independent reviewer to verify and reproduce key claims about dataset origin, composition and lawful sourcing. The Ethical Dataset Compliance Questionnaire provides structured, reproducible evidence for every dataset submitted. The Ethical Questionnaire is part of a broader legal, privacy, compliance and cybersecurity vetting

workflow, as shown in the Supplier Program and Partner onboarding documents, which describe multi-layer evaluations by legal, privacy and cybersecurity teams. These evaluations generate formal written assessments and risk findings.

Is there separation between data creation and validation functions?

Yes. Defined.ai separates data creation, quality assurance and evaluation functions through dedicated workflows such as human-in-the-loop evaluation, benchmarking teams and structured validation layers used in stumping, adversarial testing and quality scoring. This ensures independent verification of data and model quality, separate from those who generated or supplied it. Defined.ai uses fully separated teams for data creation, multi-layer quality assurance and independent validation, ensuring that those who generate data are never the same individuals who evaluate or score it.

Are external audits or peer reviews supported and encouraged?

Yes. Defined.ai undergoes multiple internal and external audits each year, including ISO/IEC 27001 and ISO/IEC 27701 certification audits performed by accredited third-party auditors, ensuring continuous oversight and compliance.

Is evidence retained to support audit findings?

Yes. Defined.ai maintains a documented audit trail, retaining all records, findings, corrective-action tickets and supporting evidence from both internal and external audits.