

OECD AI Transparency Report

Organization: EuroAIGuard Conseil (FRA)

Reporting Period: Q2 2026

Published: October 2, 2026

Section 0 – Your role in the AI system lifecycle

Responding as (please select all that apply):

- Deployer

If any of the fields above are selected, please indicate your organization size:

Micro enterprise: fewer than 10 employees

Section 1 – Risk Identification and Evaluation

1. How does your organization define and/or classify different types of risks related to AI?¹

- Maintain an internal taxonomy, classification system, and/or set of risk area definitions
- Align with specific external frameworks, standards, reports, and/or guidance documents such as taxonomies, classification systems, and/or sets of definitions for risk areas (e.g. the NIST AI Risk Management Framework, the ISO/IEC 42001, the OECD AI Principles, the EU General-Purpose AI Code of Practice, the International AI Safety Report). (please specify): The organisation's information-security risk process is aligned with ISO/IEC 27001. Personal-data risk is governed under the EU GDPR. For the AI systems it deploys, the organisation applies the EU AI Act risk categories as its classification reference.
- Other (please specify): As a deployer that develops no AI model, the organisation classifies AI-related risk at two points. At adoption, each third-party AI tool is subject to a documented supplier evaluation before use, and the residual risk of that reliance is formally recorded and accepted by the direction. In use, a data-minimisation rule governs what may be processed through AI tools, no client engagement data and no directly collected contact data are processed through them. The organisation records its own AI systems in an AI inventory maintained on its own platform, with EU AI Act risk classification applied by the same rule-based engine it provides to clients. Its two internal agents support marketing and commercial research on publicly available sources only, they are internal, take no autonomous outbound action, and every output is reviewed and validated by a named human before it leaves the organisation.

How does your organization define thresholds at which severe risks, unless adequately mitigated, may be considered "unreasonable"? Please describe the framework(s) most relevant to your organization's role and context. If your organisation does not use formal thresholds, you may describe instead the policies, decision criteria, or governance processes you apply when assessing severe AI risks. You may also describe: • Specific examples of contexts in which your models or systems could, in your assessment, pose unreasonable or intolerable risk; • How mitigations are graduated by capability or

propensity level; • How your organisation considers the specificities of different AI technologies, such as generative AI and agentic AI, including any specific risk assessment or governance mechanisms applied to agentic AI capabilities.

Formal thresholds : We apply a documented risk methodology (ISO/IEC 27001 §6.1.2–6.1.3, in force in version 1.0). Criticality = likelihood × impact on a 4×4 matrix, retaining the highest impact across confidentiality, integrity and availability. Four bands graduate both the deadline and the conditions for accepting a risk : 1 to 3 acceptable as it stands , 4 to 6 treatment within 12 months or reasoned acceptance by the direction, 8 to 9 treatment within 3 months, acceptance only with a documented compensating control , 12 to 16 immediate treatment, no acceptance in any state. That top band is our operational definition of an unreasonable risk.

Limits set at zero : No client engagement data and no directly collected contact data pass through AI tools, only professional information already public on the open web, no client data entered into our platform ever passes through an AI tool, no AI tool performs an outbound action.

Answered by design : Our compliance software embeds, calls and trains no AI model, it is a deterministic rule engine, recorded as such in our ISMS scope. A compliance determination that is not reproducible is a risk we judged unacceptable to place in a client's hands. If any AI component were ever integrated, our secure development policy requires a dedicated AI management perimeter to be opened before production release.

Agentic AI : Two internal agents support research on public sources, recorded in our own AI inventory with EU AI Act classification, per-action permission gating, no autonomous outbound action, validation by a named human before any output leaves.

Generative AI : A production charter dated 11 July 2026 binds every AI-assisted session, every regulatory figure verified against a primary source before it is written, never from memory, and human validation before anything leaves.

2. Both before and after deployment, what practices does your organization use to identify and evaluate vulnerabilities, incidents, emerging risks and misuse throughout the AI system lifecycle?

- Internal benchmarks, standard metrics (e.g., frequency, severity scores)
- Qualitative assessments (e.g., expert judgment, narrative scoring)
- Vulnerability reports
- Automated monitoring / remote status reporting (e.g., telemetry pipelines, automated alerts when risk thresholds are breached)
- Specific risk identification practices for agentic AI capabilities, e.g., autonomous action-taking, tool use, multi-step planning (please specify): Before adoption, each AI tool is assessed for its agentic surface, which connectors and tools it can reach, and what it could do without human action. This is reviewed in practice, not assumed. In operation, agentic risk is contained by design rather than measured, per-action permission gating, no autonomous outbound action, and validation by a named human before any output leaves the organisation. The account used to access AI tools is covered by the quarterly access review.
- Other(s) (please specify): Supplier evaluation before adoption , each third-party AI tool is assessed against our security policy before use, covering contractual terms, model-training settings, data retention and third-party connectors, with the residual risk formally accepted by the direction. Our security incident procedure

and register cover AI-related events on the same basis as any other detection sources include monitoring alerts, audit logs, CVE watch, and reports from clients or third parties.

Please briefly describe your organization's decision-making process for deploying an AI model or system including: who has decision authority; what evidence and criteria are required (e.g., performance, safety/security, robustness, fairness, privacy, compliance, independent external testing, evaluation, validation and verification (TEVV) or assurance); contextual caveats or limitations required in interpreting risk evaluation metric results; what approvals/sign-offs are needed; and what safeguards (e.g., mitigations, monitoring, or phased release) must be in place to move into deployment.

Evidence : Before adoption, each third-party AI tool is assessed in writing against our security policy, contractual terms, model-training settings, data retention, international transfers, connectors, account security. The residual risk is stated explicitly.

Sign-off : Adoption and its residual risk are recorded in a direction approval memo.

Caveat on metrics : As a deployer we have no visibility into a third-party model's internals and run no model-level evaluation. Our assurance comes not from what a model scores but from bounding what it can reach and what it can do.

Safeguards : Data-minimisation rule fixing what may never be processed per-action permission gating, no autonomous outbound action, validation by a named human before any output leaves. Permanent restrictions, not a release phase.

After deployment : AI events fall under our security incident procedure and register, the AI-tool account is covered by the quarterly access review risks are re-examined at the quarterly management review.

Following deployment, how do you use incident reports — including those shared by other organizations — to identify emerging issues, update risk assessments, and trigger corrective actions.

Incidents are governed by a documented procedure and recorded in a single living register. Detection sources are defined and include **reports from clients or third parties** through a published contact channel, alongside monitoring alerts, audit logs, CVE watch and direct observation.

Every suspicious event is logged **at detection**, even if later requalified as a false alarm, the requalification itself is part of the record. Entries are append-only, never edited away.

3. Does your organization have mechanisms to receive and/or share reports of risks, incidents or vulnerabilities by/with third parties?

Yes, both receiving and sharing

If yes, please indicate which stakeholders reporting mechanisms may be accessible to, recognising specifics will depend on the nature of the report and risk level. For each, indicate whether access is on a standing basis (Always), activated based on context (Ad hoc/as relevant), or not currently accessible.

	Always	Ad hoc / as relevant	Not currently accessible
Relevant employees (security, legal/compliance, etc.)			☐
Public authorities (e.g., public reporting offices for security incidents or data breaches)		☐	

	Always	Ad hoc / as relevant	Not currently accessible
Supervisory authorities (e.g., data protection authorities)		☐	
Industry partners or industry fora		☐	
Users (end users, enterprise customers, developers using APIs)		☐	
Business customers / deploying organizations (if different from “users”)		☐	
Supply-chain partners (cloud, data providers, etc.)		☐	
Academia / researchers			
Independent auditors			
Standards bodies / multi-stakeholder initiatives (OECD.AI, ISO/IEC, IEEE, etc.)			
Specific civil society organizations			
Media / general public (people or communities indirectly impacted)			
Other			

Please describe how your organisation determines when and how reporting mechanisms are made accessible to different stakeholders (e.g. based on legal or regulatory requirements, internal policies, or risk level). Please also provide public URLs or details on how to submit reports. If none, write “Not applicable”

How access is determined. Three logics, none of them discretionary. Our published channel is open unconditionally to anyone, with no filtering by stakeholder type. Notification to authorities is triggered by legal obligation rather than by our own judgement of who should be told. Notification to clients is triggered by our incident procedure. Severity governs what must be assessed rather than who may report. From grade three upwards, the personal-data dimension is systematically assessed and the conclusion is written to the register even when it is negative.

Receiving. Our incident procedure names a published channel, contact@euroaiguard.com, for reports from clients or third parties, alongside monitoring alerts, audit logs and CVE watch on our dependencies.

Sharing. Where an incident affects personal data, we notify the French supervisory authority within 72 hours under GDPR Article 33, and affected clients within 48 hours under the commitment set in our incident procedure for our processor role. The notification decision, including a reasoned decision not to notify, is recorded in the incident register.

On the table above. Our single published channel is open to anyone, which is why most rows are marked as available ad hoc rather than as standing arrangements. The organisation is founder-

operated and has no employee, so that row is marked not currently accessible because the stakeholder group does not exist, not because a channel was withheld.

Limits, stated plainly. We operate no dedicated vulnerability disclosure policy or bug bounty, and we participate in no incident-sharing community. At our size, reporting flows through the named channel above and through our supplier and client contractual chains.

How does your organization share relevant information as appropriate?

- Reporting to regulatory authorities as relevant

Please describe what type of information is shared.

To the supervisory authority : where personal data is affected , the content required by GDPR Article 33(3), nature of the breach, categories and approximate volume of data subjects and records concerned, likely consequences, and the measures taken or proposed.

To affected clients: where we act as processor , the nature of the incident, which of their data is concerned, the timeline, and the measures taken.

Along the supplier chain : notifications received from suppliers, including changes of sub-processor, are relayed to affected clients the same day, with the notice and objection periods we owe them.

4. If your organization has programs to incentivize risk(s) disclosure, incidents and systemic vulnerabilities, please share what type of programs you use?

- Not applicable

5. Does your organization use or contribute to the development of international technical standards, best practices or tools for the identification, assessment, and evaluation of risks?

Uses

If contributes or uses, please select all that apply:

- ISO/IEC standards (e.g., ISO/IEC 27001, ISO/IEC 42001, ISO/IEC 23894, ISO/IEC 42005)

6. Does your organization collaborate with relevant stakeholders across sectors to assess and adopt risk mitigation measures to address risks that could have broad and potentially cascading socio-economic impacts? Please indicate which stakeholders you collaborate with, and whether collaboration is on a standing basis (Always-on), activated based on context (Ad hoc/as relevant), or there is no current collaboration.

	Always-on	Ad hoc / as relevant	No current collaboration
Relevant employees (security, legal/compliance, etc.)			☐
Public authorities / regulators / AI safety institutes / other public sector bodies and inter-governmental coordination structures			☐

	Always-on	Ad hoc / as relevant	No current collaboration
Industry partners or industry fora		☐	
Users (end users, enterprise customers, developers using APIs)			☐
Business customers / deploying organizations (if different from “users”)		☐	
Supply-chain partners (cloud, data providers, etc.)		☐	
Academia / researchers			☐
Independent auditors			☐
Standards bodies / multi-stakeholder initiatives (OECD.AI; ISO/IEC; IEEE; AIID etc.)		☐	
Independent civil society organizations (e.g., for evaluations, audits, or policy review)			☐
Media / general public (people or communities indirectly impacted)			☐
Other(s) (please add any other relevant stakeholder groups):			
Not applicable			

Please describe examples of how your organization collaborates. Examples may include: participating in multi-stakeholder or industry initiatives; partnering with academia, civil society, or the public sector; establishing internal or joint governance structures with external experts or in partnership with other efforts to share knowledge and jointly address risks.

Pilot with industry partner : we run a pilot with a partner organisation. The collaboration is explicitly risk-oriented.

Advisory engagements with deploying organisations : Where we work with an organisation deploying AI, the engagement is itself the collaboration, we assess their AI systems against the EU AI Act and ISO/IEC 42001 with them, and the mitigation measures are adopted by them, not prescribed by us.

Supply chain : We completed a full review of our suppliers' data processing agreements in 2026, recording per-supplier points of attention rather than filing them unread, and we hold a commitment to relay any change of sub-processor to affected clients the same day, within the notice periods we owe them.

Multi-stakeholder. This reporting framework is our participation. We joined it as a small deployer to document, at the smallest organisational scale, what the Code of Conduct looks like in practice.

Any further comments and implementation documentation.

Section 2 – Risk Management and Information Security

7. Across the AI system lifecycle, what steps does your organization take to manage risks and vulnerabilities? Please indicate which specific practices, tools or metrics your organization uses at each lifecycle stage. You may refer to tools or metrics present in the OECD.AI Catalogue of Tools and Metrics or suggest additions. Respondents may complete only the lifecycle stages relevant to their organisation's role. Stages that do not apply may be left blank.

Plan and design

Collect and process data

Build, integrate and interpret model

Verify and validate

Deploy

Operate and monitor

Please provide a consolidated description of your organisation's approach to risk management across the AI system lifecycle, including how the practices identified above operate together and any cross-cutting governance, review or escalation processes. Please also indicate how often these practices are reviewed and what types of events trigger reviews and updates. Where relevant, please distinguish between different types of AI systems (e.g. generative or agentic AI).

We deploy AI, we do not build it, with no model metrics to rely on, our assurance comes from containment rather than measurement. Two boundaries, what may enter an AI tool, and what may leave it. Every other practice serves one of them.

Our chain : A written supplier evaluation feeds a direction decision, which fixes the limits before use data-minimisation in, per-action gating and named human validation out. In operation, anything anomalous enters an append-only incident register at detection. Each entry carries a root cause structural causes are carried into the risk treatment plan rather than closed locally, and the affected risk rating is re-examined.

Cadence : Management and access reviews quarterly, restoration test half-yearly, policy review and awareness session annually.

Triggers between cycles : Any recorded incident, a new AI tool, a supplier changing sub-processor, a structural root cause, a regulatory change, as with the recent revision of the EU AI Act calendar.

By type : Generative use, the severe risk is an incorrect regulatory statement reaching a client or the public, controlled by primary-source verification and human validation. Agentic use, the severe risk is autonomous action, controlled by per-action gating, no outbound action, and periodic review of the connector surface.

Please describe any measures used to ensure data quality and fairness across the data collection and curation stage (e.g., dataset audits, filtering training data to remove harmful content or personal

information, fine-tuning on curated data, use of moderation tools, assessment of limitations to model generalisability).

Not applicable at the training stage : We collect and curate no training data, run no fine-tuning, and operate no moderation tooling. Dataset audits, training-data filtering and generalisability assessment belong to the model developers whose systems we deploy, not to us.

On fairness : our compliance platform is a deterministic rule engine. It cannot acquire bias from data. Identical inputs always yield identical classifications, and every result is traceable to the rule that produced it.

8. What security measures are used in testing AI systems?

- Not applicable

9. How does your organization protect intellectual property, including copyright-protected content?

9.A. Safeguards that promote lawful data use and transparency in AI training:

- Not applicable

9.B. Safeguards that reduce the risk of dissemination of outputs that may infringe copyright:

- Not applicable

10. How do you protect privacy and prevent your systems from disclosing confidential or sensitive data across the AI lifecycle?

- Make information about data privacy, practices, and controls accessible to a diverse set of stakeholders
- Use security controls, including encryption at rest and in transit
- Provide organizational customers with controls to configure and implement privacy choices (e.g., data retention, logging, access, or processing settings)
- Other(s) (specify): Our strongest privacy control is architectural rather than technical, a data-minimisation rule fixes what may never enter an AI tool. No client engagement data, no data entered by clients into our platform, and no directly collected contact data is processed through any AI system , only professional information already public on the open web. AI-assisted deliverables are produced on anonymised templates completed outside AI tools. Our compliance platform embeds, calls and trains no AI model, so client data never reaches a model by design.

Describe key privacy measures and data-protection standards applied.

Framework :GDPR in both of our roles, controller for our own processing, and Article 28 processor for the data our clients enter into our platform. Security follows ISO/IEC 27001. We work from EDPB and French supervisory authority guidance.

Records and agreements. Article 30 records maintained for both roles, controller and processor. A model Article 28 data processing agreement for the processor role. Our suppliers' agreements were fully reviewed and archived.

Technical measures : Encryption in transit, TLS 1.2 minimum with HSTS. Multi-tenant isolation enforced by systematic filtering on every data query and covered by dedicated test suites. Role-based access control with server-side guards. Systematic audit logging recording identity, organisation, address and before/after detail.

Breach handling : Supervisory authority within 72 hours under Article 33, affected clients within 48 hours under our agreement, with the notification decision recorded either way.

For AI specifically: the measure is exclusion rather than protection, client data never enters an AI tool, and our platform embeds no model, so it never reaches one.

Limits, stated plainly : Encryption at rest is delegated to our infrastructure provider , application-level field encryption is under evaluation, not in place. Automated purge of audit logs at twelve months is decided but not yet implemented.

11. How does your organization assess cybersecurity risks and implement policies to enhance the cybersecurity of advanced AI systems across the AI system lifecycle? Respondents may complete the checkbox subsections in 11.A–11.D or describe their overall approach under 11.E.

11.A. Governance, policy and assessments

- Maintain or update security policies aligned with recognized standards (e.g., NIST CSF, ISO/IEC 27001, SOC 2, CIS Critical Security Controls), and review them regularly in response to emerging risks or regulatory guidance
- Maintain vulnerability management processes (scanning, channels for external reporting, patching, remediation)
- Implement incident response, disaster recovery, or business continuity plans, including escalation procedures for high-severity incidents
- Integrate cybersecurity AI risks into broader cybersecurity risk management activities
- Regularly review security measures and update them in response to emerging risks, incidents, or major releases
- These responsibilities are delegated contractually to model or infrastructure providers
- Other(s) (specify): Two clarifications. Our AI-related security risks are not managed in a separate process , they are folded into our ISMS. The AI tool provider went through the same supplier evaluation as any other, the residual risk was accepted by the direction in writing, the account used for AI tools is covered by the quarterly access review, and AI-related events fall under our general incident procedure and register. On delegation , physical and environmental security and encryption at rest are delegated contractually to our infrastructure providers, whose agreements we reviewed and archived in 2026. No such delegation exists towards the provider of the models we use : we hold no data processing agreement with it, and we say so rather than let the checkbox imply otherwise. That gap is precisely why our primary control is minimisation of what may enter the tool at all.

11.B. Technical and operational controls

- Monitor system activity and log access to detect anomalies or misuse
- Conduct secure AI development and/or deployment practices
- These responsibilities are delegated contractually to model or infrastructure providers

- Other(s) (please specify): As a deployer, our technical controls sit at the boundary rather than inside the model. Before use, each AI tool's reachable surface is assessed which connectors and integrations it can access, and what it can do without human action. In operation, per-action permission gating applies, no AI tool performs an outbound action, and every output is validated by a named human before it leaves. Physical and environmental security of the infrastructure we use is delegated contractually to certified providers, no equivalent delegation exists towards the provider of the models themselves.

11.C. Protection of sensitive or high value AI assets

- Limit access to proprietary AI assets based on role or need-to-know
- Track and log access to sensitive AI resources
- These responsibilities are delegated contractually to model or infrastructure providers
- Other (please specify): On protection of AI assets, we hold no model weights, no training datasets and no proprietary model. What we protect is narrower, namely the source code of a deterministic rule engine held in private repositories with multi-factor authentication, and the credentials giving access to the third-party AI tools we use. The checkbox options in this subsection assume ownership of AI assets we do not have, and we select Other to say so rather than let the boxes imply otherwise.

11.D. Insider risk protection detections

- Monitor for unauthorized internal access or suspicious activity
- Implement insider-risk escalation or response procedures
- These responsibilities are delegated contractually to model or infrastructure providers

11.E. Please describe the main security measures your organization applies and indicate which are specific to particularly sensitive AI assets, where relevant.

Framework. An ISO/IEC 27001 information security management system, in force in version 1.0

Access and identity : Least privilege, need-to-know, attributability and revocation by default as stated principles. Role-based access control in the application with server-side enforcement, multi-factor authentication on supplier accounts, and a quarterly access review covering unknown users, unknown tokens and authorised third-party applications.

Application and data : Tenant isolation enforced on every data query and covered by dedicated test suites. Modern password hashing, anti-forgery protection, data access through an ORM rather than assembled queries, input validation, a full set of security response headers, and distributed rate limiting. Encryption in transit at TLS 1.2 or above.

Secrets and supply chain : No secret in source code, configuration by environment variables, secret files excluded from versioning, and a "zero secret in the repository" check at every internal audit. Dependencies pinned to exact versions with a traced vulnerability watch and defined correction deadlines.

On particularly sensitive AI assets. We hold none in the usual sense, no model weights, no training datasets, no proprietary model. What is sensitive is narrower, the credentials granting access to the third-party AI tools we use, and the source code of our deterministic rule engine. Both fall under the regime above · private repositories with multi-factor authentication, secrets never committed, and the AI-tool account included in the quarterly access review.

One limit: Encryption at rest is delegated to our infrastructure providers, application-level field encryption is under evaluation and not in place.

12. What is your organization's vulnerability management process? Does your organization take actions to address identified risks and vulnerabilities, including in collaboration with other stakeholders?

- Internal vulnerability disclosure receipt and tracking processes
- Investigation of reported or identified risks and vulnerabilities
- Development of patches or other mitigations for confirmed risks or vulnerabilities
- Testing of patches or other mitigations
- Scanning for and mitigation of similar risks or vulnerabilities
- Reporting to regulators / authorities when relevant
- Defined escalation criteria and follow-on actions for identified vulnerabilities (e.g., for external reporting or information sharing)
- Other(s) (please specify): Our vulnerability management runs through the same incident procedure and register as any other security event, with severity grades carrying their own deadlines. As a deployer we do not test the models we use · vulnerability management applies to our own software, its dependencies, and the credentials granting access to third-party AI tools.

13. What post-deployment practices does your organization use to monitor and respond to vulnerabilities, incidents, emerging risks and misuse of advanced AI systems?

- Conduct post-deployment monitoring
- Perform root-cause analysis and apply lessons learned to improve future development and/or deployment practices
- Monitoring and emergency controls specific to agentic AI systems (e.g., ability to interrupt autonomous action sequences, override tool-use permissions, or revoke delegated authority)
- Other(s) (please specify): On agentic controls our position is preventive rather than reactive. There is no autonomous action sequence to interrupt, every action requires approval before it runs, and no AI tool performs an outbound action. Tool-use permissions are overridden by removing the tool, an unused prospect-enrichment connector was disconnected on that basis, and the connector surface is reviewed rather than trusted to its defaults. No delegated authority is revoked because none is granted. Post-deployment monitoring is not AI-specific , AI-related events enter the same incident procedure and append-only register as any other security event, with root-cause analysis and lessons recorded per entry. In the interest of accuracy, that register holds no entry to date , the process is in force and has not yet been exercised.

14. Do you develop or deploy agentic AI systems (systems that can autonomously plan/act or use tools)?⁴

Yes

14.A. If yes, please indicate which categories of controls your organization applies. Please select all.

- Constraints on the agent's action space (e.g., tool-use permissions, environment sandboxing, allowlists/blocklists of resources, scope limits per task)
- Human oversight or interruptibility mechanisms (approval requirements for sensitive actions, escalation paths, interruption interfaces)

Please provide any further comments and implementation documentation.

Our controls are preventive rather than supervisory. Approval is required before every action, not only sensitive ones, so there is no autonomous sequence to interrupt and no unattended agent behaviour to monitor. The agents' reachable surface is assessed before use and reviewed rather than trusted to its defaults, an unused connector was disconnected on that basis.

We maintain no dedicated log of agent actions and no multi-agent interaction monitoring, our two agents do not interact with one another, and every output passes through a named human before it leaves the organisation.

Section 3 – Transparency Reporting on Advanced AI Systems

15. Does your organization publish public reports and/or technical documentation related to the capabilities, limitations, and domains of appropriate and inappropriate use of advanced AI systems? Are these clear and understandable by non-technical audiences?

Not applicable

16. How does your organization share information with a diverse set of stakeholders (other organizations, governments, civil society and academia, etc.) regarding the outcome of evaluations of risks and impacts related to an advanced AI system?

- Not applicable

Describe how evaluation results are shared and with which stakeholders (e.g. other organizations, governments, civil society and academia, etc). You may also indicate whether non-technical or plain-language summaries are provided for broader audiences.

This kind of report is our disclosure open to anyone, governments, peers, clients, civil society. For a deployer of our size that is the most complete and the most accountable channel available, and we chose it deliberately.

We evaluate suppliers, not models. Each AI tool is assessed before use and the residual risk is accepted in writing by the direction. That assessment is where our assurance comes from, and it is an internal record what belongs in public is the governance around it, which is what this report sets out.

Clients get more : A security posture document on request, and full answers to supplier security questionnaires.

Plain language is our default, not an add-on : We write for data protection officers and compliance leads, in French and English. We publish nothing that needs machine learning expertise to read,

because we produce nothing of the kind.

17. Does your organization publish public privacy policies addressing the use of personal data, user prompts, and/or the outputs of advanced AI systems?

Yes

If yes, provide links to privacy policies or notices covering personal data, user inputs, and AI-generated outputs.

<https://www.euroaiguard.com/confidentialite/>

Our public privacy policy addresses the processing of personal data. It does not address user prompts or the outputs of advanced AI systems, because our platform processes neither, it embeds no AI model, so there is no prompt and no model output involved in the service we provide.

18. How does your organization provide information about the sources of data used for the training of advanced AI systems, as appropriate, including information related to the sourcing of data annotation and enrichment?

- Not applicable

Please provide links to any publicly available information.

not applicable to us

19. Does your organization demonstrate transparency related to advanced AI systems through any other methods?

Not applicable

20. Does your organization publish statistical insights derived from anonymised, aggregated usage data to advance understanding of AI adoption and its economic impacts in areas including occupations and tasks by region, sector and country? To this end, do you contribute to international cooperation efforts?

No

Please estimate the total staff time and annual expenses associated with conducting compliance activities and include any further comments and implementation documentation (optional).

We do not publish figures, and the reason is substantive rather than commercial, in a founder-operated company the cost of compliance is almost entirely time.

What we can describe is its shape. Our management system was built over few months alongside commercial activity, and now costs a recurring cadence rather than a project, quarterly management and access reviews, a half-yearly restoration test, an annual policy review and awareness session. External spend is close to zero, we retain no external auditor, engage no compliance consultant and buy no dedicated compliance tooling, since we run our own governance on the platform we build.

Section 4 – Organizational Governance, Incident Management, and Transparency

21. How has AI risk management been embedded in your organization's governance framework?

- Executive oversight
- Other (e.g., distributed / multi-institutional governance arrangement) (please specify): Our methodology draws an explicit line between analysis, which may be tool-assisted, and decision, which belongs to the direction. Every decision on AI is a written, dated and signed act, including the acceptance of residual risk on the AI tools we use, recorded in an approval document. Role cumulation is stated openly in our scope document and compensated by systematic traceability, since separation of duties is not achievable at this size and we do not pretend otherwise. AI risk is not governed separately, it is embedded in the same management system as every other risk, through the same supplier evaluation, access review, incident procedure and quarterly management review.

Please describe how AI risk management is integrated into governance structures, including: which roles or bodies have advisory versus decision-making authority; which roles or bodies have final accountability for AI risk-related decisions; and how often policies are reviewed or updated (e.g., following major incidents, regulatory changes, or material updates to systems).

Advisory versus decision : Our methodology draws the line explicitly , analysis may be prepared with tool assistance, decision belongs to the direction. Written supplier evaluations, risk ratings and incident findings are advisory inputs. None of them decides. We have no advisory body and no committee, and we do not describe our own judgement as one.

Final accountability : For now, the direction owns every risk in the register and signs every acceptance of residual risk. The memo bringing the management system into force carries a name, a place and a date. Accountability here is a signature, not an org chart.

Cadence : Management review and access review quarterly. Backup restoration test half-yearly. Policy review and awareness session annually.

Between cycles, four things force a review. An incident whose root cause proves structural rather than one-off. Adoption of a new AI tool, which requires its own evaluation and its own decision. A supplier notifying a change of sub-processor. A regulatory change, the revision of the EU AI Act timetable was entered as an external-context input at our management review rather than noted and forgotten.

22. Are relevant staff trained on your organization's governance policies and risk management practices?

Yes

If yes, describe training methods (e.g., onboarding modules, regular sessions, and confidential channels to report concerns).

We wrote these governance policies and this risk methodology, and apply them daily, which is a stronger form of familiarity than attending a session on them. Competence is maintained deliberately rather than assumed, continuous professional development in AI governance and data protection,

including preparation for a recognised professional certification in AI governance, and an ongoing regulatory watch whose findings are entered as external-context inputs at our management review. Our management system also institutes an **annual awareness session** that applies to the founder and is attested like any other control. It falls due within the first annual cycle and has not yet been held (we report the state, not the intention).

23. Are steps taken to address reported incidents documented and maintained internally?

Yes

If yes, describe how incident records are kept and used for follow-up (e.g., incident logs, root-cause analysis reports, internal tracking systems). You may indicate whether and how you define escalation thresholds for when incidents must be surfaced to senior leadership or external parties.

Incidents are recorded in a single living register governed by a documented procedure. Entries are *append-only*, a correction is made by adding, never by erasing, and an event requalified as a false alarm keeps its trace because the requalification is itself part of the evidence.

Each entry carries the detection, a severity grade, a factual description, a timeline, preserved evidence, a personal-data assessment with the notification decision, a root-cause analysis, corrective actions with owner and deadline, the lessons drawn, and a closure validated by the direction. Where a root cause proves structural rather than one-off, the corrective action is carried into the risk treatment plan rather than closed locally.

The register is internal and stays internal. In the interest of accuracy, it holds no entry to date (the process is in force).

24. Does your organization share research and best practices on addressing or managing risk?

Yes

If yes, describe how research or good practices are shared (e.g., publications, conferences, collaborations with other sectors).

Yes, best practices rather than research. We publish practitioner guidance on managing AI and data protection risk, written analyses of regulatory developments, and a set of practical guides covering EU AI Act readiness for SMEs, ISO/IEC 42001, NIS2, DORA, the GDPR and AI, together with an impact assessment checklist and a register template.

Our published material also covers the risks of using AI in one's own work, including the effect of unexamined reliance on generative tools on the user's own understanding, with concrete counter-practices rather than warnings.

This report is part of the same intent, we joined the framework as a small deployer to document, at the smallest organisational scale, what the Code of Conduct looks like when actually applied.

We conduct and publish no academic research and take part in no research consortium.

25. Does your organization use international technical standards or best practices for AI risk management and governance policies?

Yes

If yes, please indicate which standards or frameworks apply:

- ISO/IEC standards (e.g., ISO/IEC 27001, ISO/IEC 42001)

If yes, describe which international technical standards or best practices your organization uses.

ISO/IEC 27001: our information security management system is the frame within which AI risk is handled. We run no separate process for AI supplier evaluation, access review, incident procedure and management review are the same for AI as for everything else.

The EU AI Act : we apply its risk categories to classify the AI systems we deploy, in our own inventory.

The Hiroshima Process International Code of Conduct : we commit to it by submitting this report, and the reporting framework is its concrete application.

One precision we think is owed , we track **ISO/IEC 42001** as a reference framework inside the tool we provide to clients. We operate no AI management system in the sense of that standard for ourselves, and we claim none.

Any further comments and implementation documentation.

No answer provided

Section 5 – Content Authentication & Provenance Mechanisms

26. What mechanisms, if any, does your organization put in place to allow end-users, where possible and appropriate, to know when they are interacting with an advanced AI system developed by your organization?

- Other (specify): We develop no advanced AI system, so no end-user ever interacts with one of ours. Our compliance platform embeds, calls and trains no AI model , it is a deterministic rule engine, and our public service pages state that classification is assisted by that rule engine with every decision remaining human and traced.

Provide any additional details or context on how or when users are informed they are interacting with advanced AI systems (e.g. if such mechanisms are planned, in testing, or applied only to certain products).

The two AI agents we deploy are internal. They interact with no one outside the organisation, take no outbound action, and produce nothing that reaches a third party without a named human validating it first.

One point we prefer to state rather than leave unsaid : Some of our published material is drafted with AI assistance, and we do not label it as such. That is a documented decision of the direction, not an omission, every published item carries named human editorial responsibility, and our production charter requires verification against primary sources and human validation before anything is published. We took the view that named accountability for the text is the substantive guarantee, and that a generic label would add less than it suggests.

27. Does your organization use content provenance detection, labeling or watermarking mechanisms that enable users to identify content generated by advanced AI systems?

Not applicable

Any further comments and implementation documentation.

No answer provided

Section 6 – Research & Investment to Advance AI Safety & Mitigate Societal Risks

28. Does your organization advance research and investment in the following areas?

- Other (please specify): not applicable to us

If yes, please describe key research or investment initiatives in the areas selected above (e.g., dedicated safety teams, funding of research programs, contributions to open-source safety tools, publishing research, partnerships and collaborative research initiatives). You may also indicate whether related outputs are publicly shared, shared with partners or within consortia or not publicly disclosed.

No answer provided

29. Does your organization pursue research or investment related to socio-economic or environmental considerations in the development and deployment of AI?

- Not applicable

Any further comments and implementation documentation.

On socio-economic considerations, our contribution is dissemination rather than research. We publish practitioner guides on AI Act readiness, ISO/IEC 42001 and impact assessment without charge, aimed at organisations too small to buy consultancy. We would not call that research, and we do not count our commercial product as an investment in access.

On environmental considerations, we pursue nothing. It is true that a deterministic rule engine consumes negligible compute compared with model inference, but that architecture was chosen for reproducibility and auditability, not for energy. We record the fact without claiming the intent.

Section 7 – Advancing Human and Global Interests

30. Does your organization support any digital literacy, education or training initiatives to improve user awareness and/or help people understand the nature, capabilities and limitations of advanced AI systems?

Yes

If yes, describe any digital literacy or training initiatives your organization supports (e.g., user education, public AI-education programs, school or university partnerships, community workshops). You may also indicate whether these efforts include or prioritise underserved or underrepresented communities.

We publish practitioner guides : EU AI Act readiness for smaller organisations, ISO/IEC 42001, NIS2, the GDPR and AI, plus an impact assessment checklist and a record template.

We also produce AI literacy training modules answering the AI Act's literacy duty, what AI systems are, what they can and cannot do, and where the obligations fall. We report no delivery figures because the material has not yet been delivered commercially.

Our articles address the limits of AI from the user's side, including the effect of unexamined reliance on generative tools on one's own understanding.

On underserved audiences we aim at small and founder-operated organisations, structurally underserved because AI governance advice is priced for large enterprises. We run no school or university partnership, no public education programme and no community workshop, and claim no outreach to communities we do not serve.

31. Does your organization pursue AI projects that aim to deliver societal or public benefits, including through partnerships with civil society and community groups?

No

Any further comments and implementation documentation.

No answer provided