OECD AI Transparency Report

Organization: Odysai (FRA) **Reporting Period:** Q4 2026 **Published:** November 2, 2026

Section 0 – Your role in the AI system lifecycle

Responding as (please select all that apply):

- Application developer (/provider)

> **If any of the fields above are selected, please indicate your organization size:**
>
> Small enterprise: 10-49 employees

Section 1 – Risk Identification and Evaluation

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

- 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): ISO/IEC 27001:2022 (certified, certificate IC-IS-2506288, INTERCERT); Regulation (EU) 2024/1689 (AI Act), under which Odysai SAS has qualified itself as the provider of an AI system that is not a high-risk AI system; France Digitale Legaltech Generative AI Code of Conduct (signatory).

> **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.**
>
> As an application developer, Odysai SAS (Jimini) integrates third-party general-purpose models, accessed for inference only under zero data retention terms, into a legal assistant for legal professionals, and does not train foundation models. It therefore does not define capability thresholds of its own. AI-related risks are assessed within its ISO/IEC 27001:2022-certified information security management system (risk assessment and Statement of Applicability), and the company has qualified itself under Regulation (EU) 2024/1689 as the provider of an AI system that is not a high-risk AI system.

> Given its clientele (notably lawyers bound by professional secrecy), the risk treated as most severe is a breach of the confidentiality of client data. Mitigations include hosting of customer data in France, inference exclusively within the EU, contractual zero data retention and no use of customer data for model training by model providers, and strict logical isolation of data per organisation.

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?

- Adversarial testing to detect technical vulnerabilities
- Internal benchmarks, standard metrics (e.g., frequency, severity scores)
- Qualitative assessments (e.g., expert judgment, narrative scoring)
- Impact assessments such as personal data protection assessment and fundamental rights impact assessments
- Vulnerability reports
- Incident reports, including reports shared by other organizations
- Automated monitoring / remote status reporting (e.g., telemetry pipelines, automated alerts when risk thresholds are breached)

> **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.**
>
> Every change to the service follows Jimini's Secure Software Development Lifecycle. Changes move through four segregated environments (local, development, pre-production, production). A pull request must be approved by a code owner other than its author, with a successful continuous-integration run (unit, integration and end-to-end tests), and functional acceptance is carried out on staging and pre-production before promotion. Each production deployment must be approved by someone other than the change author, the approval being recorded through Git tags; emergency fixes are reviewed after the fact within 24 working hours. Production systems, container images and third-party dependencies are scanned continuously, and an external penetration test is commissioned annually from an independent third party. Internal evaluation campaigns of model answers are also run by legal experts against an internal question grid scoring legal reliability and answer quality, including hallucinations.

> **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.**
>
> Security events are reported through a defined channel, with an alternate point of contact, and are qualified by the Security Officer and the Chief Technology Officer, who assign a severity (low, medium, high or critical). Incidents are recorded in an internal Security Incident Register. Affected customers are notified within 24 working hours of qualification and in any event within 72 hours, and the CNIL within 72 hours where Odysai SAS acts as controller. A blame-free retrospective is held within one week of resolution and corrective actions are recorded in a corrective action register; this part of the process is being formalised within the ISMS. Security advisories from CERT-FR, guidance from the CNIL and

security advisories from hosting and model providers are monitored and fed into vulnerability management.

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)	□		
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”

> Customers and users can report issues through Jimini's support channel (support@jimini.ai). Under the Security Incident Management Policy, any person may report a suspected security event through a defined channel with an alternate point of contact, and no adverse measure may be taken against a person reporting in good faith. Reporting obligations towards customers and authorities are set by law and contract: affected customers are notified within 24 working hours of qualification and in any event within 72 hours (Data Processing Agreement), and the CNIL is notified within 72 hours under Article 33 GDPR where Odysai SAS acts as controller. Odysai also receives and monitors security advisories from CERT-FR and from its hosting and model providers.

> **How does your organization share relevant information as appropriate?**

- Reporting to regulatory authorities as relevant

> **Please describe what type of information is shared.**
>
> In the event of a personal data breach, the notification to affected customers includes at minimum the nature and description of the breach, the categories and approximate number of data subjects and of records concerned, and the measures taken or envisaged. Odysai assists customers with their own notification to the CNIL under Article 33 GDPR and notifies the CNIL directly for processing for which it is controller.

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

- Research grants, contracts, and hiring red teamers
- Operational incentives for customers and partners
- Internal incentives and culture

> **Please provide public URLs and/or details on your incentive program(s). If none, write “Not applicable.”**
>
> 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)
> - Industry-led best practices or codes of conduct

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	☒		
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.): *French cloud infrastructure provider (Scaleway), sovereign cloud provider for public-sector deployments, LLM providers, legal content and data providers, and document management system integration partners.*		☒	
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.

Jimini AI is an application developer serving lawyers, in-house counsel and notaries in France, where AI errors carry professional secrecy and liability consequences. We joined the G7 Hiroshima AI Process in May 2026 at the initiative of the French DGE, sit on the France Legaltech strategic committee, and are one of ten vendors selected for the French State’s legal AI experiment run by the DAE and the DITP with around 150 civil servants. We signed a national partnership with the Conseil national des barreaux, the French national bar council, after meeting its documented requirements on professional secrecy and independence, and we are assessed against its AI tool selection grid and against bar evaluation questionnaires. User and client reports are centralised and drive our mitigation roadmap: reports of

unsupported outputs in 2026 led to a capability that verifies every case law and statutory citation against its source. Our ISMS is ISO/IEC 27001 certified with annual external audits and an annual independent penetration test. We train trainee lawyers through the bar schools with a module on critical distance towards AI outputs. We have no structured collaboration with civil society organisations today.

Any further comments and implementation documentation.

No answer provided

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).

Risk management is run through Jimini's ISO/IEC 27001:2022-certified information security management system (certificate IC-IS-2506288, INTERCERT), whose scope covers infrastructure, applications and delivery systems, customer and business data management, HR, engineering and product, support, sales and marketing, and finance. Plan and design: secure-by-design and privacy-by-design principles are applied from the design stage of any change. Data: customer data is hosted in France (Scaleway by default; S3NS for public administrations), backups are kept in the EU, inference is performed exclusively in the EU, and model providers apply zero data retention and may not use customer data to train or improve models. Build, verify and validate: secure development lifecycle with peer review by a code owner, continuous-integration testing, acceptance on staging and pre-production, continuous vulnerability scanning and an annual independent penetration test. Deploy: each production deployment is approved by someone other than the change author. Operate and monitor: logging of access and operations (audit logs retained one year) with alerting, quarterly access reviews, daily

encrypted backups with restoration tests, and an incident management process with customer notification within 24 working hours of qualification. Governance: an ISMS Governance Council meets quarterly, a management review is held in Q4 2026 and a first internal audit is scheduled for Q1 2027. The policy set is reviewed at least annually.

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).

Jimini AI does not train or fine-tune foundation models and does not use client data for training. Data quality work therefore concentrates on the legal corpus we curate and on the documents our users bring.
Our legal corpus is built exclusively from official and authoritative French public sources, primarily Légifrance for legislation and regulation and Judilibre for Court of Cassation case law, together with administrative and constitutional case law. Using official sources means provenance is traceable by construction and no harmful or personal content enters the corpus beyond what the public authorities themselves publish. The corpus is refreshed so that answers reflect the law in force, and every answer surfaces the underlying source so users can check it against the official text.
Client documents are processed strictly to serve the user's own request, are logically segregated by organisation, and are never pooled, reused across clients or used to improve models. An anonymisation capability is available to users who need to remove identifying information before processing. Known limitations, including the fact that coverage focuses on French law and that outputs must be verified by a qualified professional, are documented in our code of conduct and covered in the training we deliver to users.

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

- Isolation of staging environments (from production / operational environments)
- Controlled security vulnerability testing
- Access controls, including roles-based permissions
- Audit logging and monitoring
- Isolated and secure test environments
- Network isolation from external access
- Use of "AI-safe" test-data, as appropriate
- Use of synthetic, anonymized, or dummy data, as appropriate
- Controlled environments for red-teaming and adversarial testing
- Rate limits and resource caps
- Human-in-the-loop controls for higher-risk test cases

8.A. Please indicate, if relevant, the team or role responsible for overseeing secure testing and whether this responsibility is shared across multiple teams and/or functions.

Responsibility for secure testing is shared, under a single owner. Overall accountability sits with the co-founder responsible for technology and security, who owns the information security management

system, commissions the annual independent penetration test and maintains the daily vulnerability watch.
Day-to-day execution sits with the engineering team, through a rotating daily on-duty role that owns incoming bug and security reports, monitors the alerting channels and escalates according to a documented severity scale. Infrastructure and platform engineers handle environment isolation, access control and remediation. The legal and compliance function, headed by our Data Protection Officer, covers data protection impact assessments, client security questionnaires, audit rights and incident notification obligations. External assurance is provided by an independent certification body for our ISO/IEC 27001 management system, an independent third party for the annual penetration test, and a compliance partner supporting internal audit.

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:

- Other(s) (please specify): Jimini AI does not train or fine-tune foundation models, so training-data safeguards do not apply to us. Our safeguards operate at the sourcing and output stages. Our legal corpus is built exclusively from official French public sources, principally Légifrance for legislation and regulation and Judilibre for case law, which are published by the public authorities under open licensing. We do not scrape third-party publishers and we do not source content from repositories associated with piracy. Client documents remain the property of the client, are processed solely to serve that client's request and are never reused to train or improve models. Every generated answer cites the source it relies on, so users can trace content back to the official text, and our terms of use set out the allocation of intellectual property rights between Jimini and its clients.
- Consult and/or retain an intellectual property attorney to protect, enforce, and comply with, intellectual property rights.

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

- Designating clear points of contact so creators or rights holders can raise concerns about infringing outputs
- Other(s) (specify): Jimini AI is a text-only professional tool and does not generate images, audio or video, which removes the categories of infringing output that this question is principally concerned with. Outputs are legal analyses and drafts produced for a named professional user in the context of their own matter, not content published to the public. Our main safeguard is grounding. Answers are built from official public legal sources and from the documents the user supplies, and every answer carries a citation to the source it relies on, so users can verify the provenance of the content they receive rather than relying on unattributed generation. Users retain and control their own outputs, and our terms of use require them to review and validate any output before use, consistent with the professional responsibility of lawyers and in-house counsel. Contact points are published on our website and in our terms of use, including a dedicated address for data protection matters, so that any third party, including a rights holder, can raise a concern and obtain a response.

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
- Apply Privacy-Enhancing Technologies (PETs) where applicable (e.g. differential privacy, federated learning, synthetic data generation, homomorphic encryption)

> **Describe key privacy measures and data-protection standards applied.**
>
> Odysai SAS acts as processor for customer content (uploaded documents, prompts and outputs) and as controller for account, authentication, billing and analytics data, under a Data Processing Agreement whose Annex A lists all sub-processors. Customer data is stored in France (Scaleway, ISO/IEC 27001 and HDS; S3NS, SecNumCloud-qualified, for public administrations), backups are kept in the EU and inference is performed exclusively in the EU, with no transfer of customer data outside the EU. All model providers apply zero data retention and are contractually prohibited from using customer data to train or improve their models. Data is encrypted at rest (AES-256) and in transit (TLS 1.2 minimum), logically isolated per organisation, and accessed on a least-privilege basis with multi-factor authentication and quarterly access reviews; access is logged. Customer content is kept for the term of the contract and deleted within 30 days of account closure, and users can delete it at any time; backups are retained 30 days and audit logs one year. Product analytics are pseudonymised. A dedicated clause addresses lawyers' professional secrecy. Odysai is ISO/IEC 27001:2022 certified, has appointed a Data Protection Officer, publishes a privacy policy and maintains a public Trust Center (trust.jimini.ai).

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
> - Regularly review security measures and update them in response to emerging risks, incidents, or major releases

> **11.B. Technical and operational controls**
>
> - Apply multi-factor authentication (MFA) and role-based access controls for advanced AI systems
> - Encrypt AI assets (datasets, models, configurations) in storage and in transit as appropriate
> - Monitor system activity and log access to detect anomalies or misuse
> - Conduct secure AI development and/or deployment practices

> **11.C. Protection of sensitive or high value AI assets**
>
> - Limit access to proprietary AI assets based on role or need-to-know

- Store model weights, datasets, and algorithms in secure, access-controlled environments
- Track and log access to sensitive AI resources
- Require NDAs or confidentiality agreements for staff and contractors with access

11.D. Insider risk protection detections

- Conduct pre-employment screening or role-sensitive vetting
- Monitor for unauthorized internal access or suspicious activity

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

Main measures, operated within an ISO/IEC 27001:2022-certified ISMS: multi-factor authentication on all internal access and adaptive MFA for users; single sign-on via OpenID Connect and OAuth 2.0; role-based, least-privilege access reviewed quarterly, with separate privileged accounts; AES-256 encryption at rest and TLS 1.2 minimum in transit; logical isolation of customer data per organisation; logging of logins, data operations and administrative access to customer data, with alerting (security and audit logs retained one year); managed, encrypted and remotely wipeable endpoints; managed secrets; infrastructure as code; network segregation; perimeter protection (WAF, anti-DDoS); continuous vulnerability scanning and an annual independent penetration test; daily encrypted backups in the EU with restoration tests; a formal incident response plan; security responsibilities and confidentiality undertakings in employment and service contracts, background checks on new hires where permitted by law, and security awareness training. Specific to the most sensitive assets (customer documents, prompts and outputs): hosting in France, inference exclusively in the EU, zero data retention at model providers, and no permanent Odysai access to the content of customer files (any exceptional access requires the customer's agreement and is logged). Physical security of data centres is the responsibility of the hosting providers (Scaleway, S3NS); Odysai operates no site of its own and stores no customer data on its premises.

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?

- 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
- Reporting to regulators / authorities when relevant
- Defined escalation criteria and follow-on actions for identified vulnerabilities (e.g., for external reporting or information sharing)

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

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)
> - Monitoring (anomaly detection, multi-agent interaction monitoring)
> - Logging of agent actions
> - Other (specify): Output-level verification: a dedicated capability checks the case law and statutory references produced by the agent against the official sources, so that citations the agent cannot ground are surfaced rather than passed on to the user. Agent tool accuracy is also measured against fixed benchmarks of realistic professional requests, scored deterministically rather than by another model.

Please provide any further comments and implementation documentation.

Our platform includes an agent that plans a sequence of steps and calls tools to carry out a legal task: searching official legal sources, retrieving documents from the client’s own library, and reading or editing documents, spreadsheets and presentations. It is built on third-party foundation models that we integrate; we do not train or fine-tune models.
Autonomy is deliberately bounded. The agent runs only within a session a professional user has started, on that user’s own documents and within that user’s permissions, and it produces a draft or an analysis for that user rather than taking actions outside the platform. It cannot send communications, make payments, file documents with a court or act on third-party systems. Every answer cites the source it relies on, and the output is a draft that the qualified professional reviews and validates before use. Agent tool accuracy is measured against fixed benchmarks of realistic professional requests, scored deterministically rather than by another model, and a dedicated verification capability checks that the case law and statutory references produced by the agent correspond to existing sources.

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?

No

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?

- Public documentation
- Protected information sharing
- Direct engagement with AI safety institutes, civil society or academic partners

> **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.**
>
> As an application developer, we do not evaluate foundation models ourselves. What we share are the results of the evaluations we do run: security testing, conformity assessments and the performance and reliability of our platform in professional legal use.
> Publicly, we publish a code of conduct describing how the platform works, what it relies on and its limitations, together with our data protection policy and a trust page giving our ISO/IEC 27001 certification and the list of controls covered. These are written in plain language for legal professionals rather than for a technical audience.
> Under protection, we share detailed results with clients and their auditors: the summary of our annual independent penetration test and the status of each remediation, delivered annually under contract, responses to client security questionnaires and risk assessments, and access for a client's auditor to the underlying reports in the context of a contractual audit. Several clients have used these exchanges to co-design contractual risk annexes covering inaccuracy, bias, confidentiality and automation bias.
> With public bodies, we report usage and findings to the authorities running the French State's legal AI experiment, and we submit our practices to the professional bodies that assess legal AI tools, including the national bar council and bar digital commissions, through their evaluation grids and questionnaires.
> With academia, we present how the tool behaves and where it fails to trainee lawyers at the bar schools, at a university legal technology diploma and at panels hosted by the French National School for the Judiciary. These sessions are deliberately non-technical and include the limitations users need to know.

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.**
>
> Privacy policy: https://www.jimini.ai/politique-de-protection-des-donnees ; Trust Center: https://trust.jimini.ai

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: Odysai SAS does not train AI models. Related public information: privacy policy https://www.jimini.ai/politique-de-protection-des-donnees and Trust Center https://trust.jimini.ai

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

Yes

> **If yes, describe additional method(s), including where relevant transparency measures related to AI agents**
>
> Jimini maintains a public Trust Center (https://trust.jimini.ai) showing, without authentication, the status of the ISO/IEC 27001:2022 Annex A controls. Its Data Processing Agreement lists, in Annex A, every sub-processor, including each model provider used for inference, with its role, location and guarantees, and customers are informed of changes to this list with 30 days' notice. Jimini is also a signatory of the France Digitale Legaltech Generative AI Code of Conduct.

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 statistical analyses of AI adoption. We do produce aggregated usage statistics on request for the public bodies running the French State's legal AI experiment, covering active accounts, volume of requests and main use cases, as an input to their own assessment of AI adoption in the public legal function. We also contribute to international cooperation through our participation in the G7 Hiroshima AI Process, which we joined in May 2026 at the initiative of the French Directorate General for Enterprise, and through this report.

On compliance effort, as a company of around forty people we do not run a separate compliance department. Compliance is carried out by our legal and data protection function, which is led by our Data Protection Officer and represents the largest share of this effort, with support from a lawyer on the legal team, and by the co-founder responsible for technology and security, who owns the information security management system. External expenditure covers our ISO/IEC 27001 certification and surveillance audits, a compliance management platform with internal audit support, and an annual independent penetration test. For an organisation of our size this is a significant and deliberate investment, driven by the confidentiality requirements of the legal profession and by the public sector clients we serve.

Section 4 – Organizational Governance, Incident Management, and Transparency

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

- Governance council or committee

> **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).**
>
> AI risk management is embedded in Jimini's ISO/IEC 27001:2022-certified information security management system. Roles are held by named individuals: the Chief Technology Officer (Information Security Management Leader), the Head of Engineering (Security Officer), the Data Protection Officer and Head of Legal, and the Chief of Staff (HR functions). An ISMS Governance Council meets quarterly; its fourth-quarter session serves as the management review. The policy set is reviewed at least annually, and the ISMS provides for an annual independent internal audit, the first being scheduled for Q1 2027.

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).**
>
> All personnel complete security awareness training at onboarding and annually thereafter, and phishing simulation results are measured against the ISMS information security objectives. Each person signs an attestation of adherence to the security policies as a condition of engagement. Security training is mandatory for new employees and included in the tech team onboarding; developers receive secure development training annually (including an annual review of the OWASP Top 10). Any person may report a suspected security event through a defined channel with an alternate point of contact, and no adverse measure may be taken against a person reporting in good faith.

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.**
>
> Security incidents are recorded in an internal Security Incident Register. Each reported event is qualified by the Security Officer and the Chief Technology Officer and assigned a severity (low, medium, high or critical). Evidence relating to an incident is preserved before normal log retention expires, stored separately with restricted and logged access, and retained for five years. A blame-free retrospective is held within one week of resolution and corrective actions are recorded in a corrective action register;

this follow-up process is being formalised within the ISMS. External escalation follows legal and contractual thresholds: affected customers are notified within 24 working hours of qualification and in any event within 72 hours, and the CNIL within 72 hours where Odysai SAS acts as controller.

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

No

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)
- OECD AI Principles
- Industry-led best practices or codes of conduct
- Other international or regional standards (please specify): EU General Data Protection Regulation, with a Data Protection Officer, data protection impact assessments and processing records; the EU AI Act, which we track and apply to our transparency and user information obligations; and the OWASP Top 10 methodology, used as the reference for our annual independent penetration test.

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

Our information security management system is certified ISO/IEC 27001:2022 across the whole company, with annual surveillance audits by an accredited certification body, and we maintain the associated documentation set, including the statement of applicability, the risk treatment plan, the incident response plan and the business continuity policy.
Our annual independent penetration test follows the OWASP Top 10 methodology and covers the authenticated application, the API and cross-tenant isolation, with auditor re-verification of each remediation.
On data protection, we apply the EU General Data Protection Regulation under a designated Data Protection Officer, with data protection impact assessments, processing records, documented sub-processor governance and contractual incident notification deadlines. We track the EU AI Act and apply it to our transparency and user information obligations.
On sector best practice, we follow the guidance and the AI tool evaluation grid published by the Conseil national des barreaux, the French national bar council, which cover professional secrecy, professional independence and the verifiability of outputs, and we are assessed against bar evaluation questionnaires. We publish our own code of conduct setting out how the platform works, what it relies on and its limitations.
We joined the G7 Hiroshima AI Process in May 2026 and align our governance with the OECD AI Principles. We are not certified ISO/IEC 42001 and do not claim to be, though we use it as a reference as our AI governance matures.

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?

- User disclosure (e.g., text notice, verbal statement, or system label)
- Persistent visual or audio indicators (e.g., watermark, icon, chime, banner) during AI interaction
- Documentation or FAQs explaining when and how AI is used in the system
- Disclosure to end-users when an AI system is acting autonomously or executing actions on their behalf (e.g., making purchases, sending communications, modifying files)

> **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).**
>
> *No answer provided*

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

No

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?

- Security
- Data quality and representativeness
- Information integrity
- System performance and outcomes (e.g. accuracy, consistency)
- Transparency
- Robustness or reliability

> **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.

As an application developer, our investment goes into making legal AI outputs verifiable in professional practice. Answers are grounded in official French legal sources and in the user's own documents, each one carries a citation, and a dedicated capability verifies that the case law and statutory references produced by the system actually exist in those sources. Agent accuracy is measured against fixed benchmarks scored deterministically rather than by another model. Robustness and security rely on an automated quality assurance agent, continuous dependency and container scanning, an annual independent penetration test and a daily national CERT vulnerability watch.
We publish technical and analytical articles on our blog and a public code of conduct setting out how the platform works and its limitations. Under our partnership with the national bar council, we train every trainee lawyer in France on the ethical and professional conduct dimensions of AI, including critical distance towards AI outputs and the duty to verify them.

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

- Socio-economic considerations (e.g., AI adoption and productivity, workforce transitions, skills development, access to AI tools)

If yes, please provide examples of relevant research or investment initiatives for the areas selected above.

Our socio-economic work is about skills and access. Under our national partnership with the Conseil national des barreaux, we train every trainee lawyer in France on the professional and ethical use of legal AI. We trained around 150 civil servants from more than twenty ministries and public bodies during the French State's legal AI experiment, and we run recurring training and public webinars for practising lawyers. Our partnerships with local bars are designed to give small practices access to legal AI on the same terms as large firms, which is the main access gap in the French legal market.

Any further comments and implementation documentation.

No answer provided

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?

No

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.

Under our national partnership with the Conseil national des barreaux, the French national bar council, we train every trainee lawyer in France on the professional and ethical use of legal AI, covering what these systems can and cannot do, their tendency to produce plausible but unsupported answers, and the duty to verify outputs. We trained around 150 civil servants from more than twenty ministries during the French State's legal AI experiment, run recurring public webinars with open replays for practising lawyers, teach on a university legal technology diploma and take part in panels hosted by the French National School for the Judiciary. Our code of conduct, setting out how the platform works and its limitations, is public.