

OECD AI Transparency Report

Organization: Jus Mundi (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):

- Model developer (/provider)
- Application developer (/provider)
- Deployer

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

Medium-sized enterprise: 50-249 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 ISO/IEC 42001
- Consult with external stakeholders (e.g., from academia, industry groups, civil society, or standard development organisations) when developing or updating risk taxonomies

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.

Jus Mundi does not currently rely on a single quantitative threshold to determine when an AI-related risk becomes “unreasonable.” Instead, severe AI risks are assessed through a risk-based governance process aligned with the organization’s role as a legal intelligence and technology provider, and with its broader security, privacy, compliance and AI governance frameworks. The most relevant references are

ISO/IEC 42001, ISO 27001, the EU AI Act, OECD AI Principles, data protection requirements, and internal information security and risk management policies.

A risk may be considered unreasonable or intolerable when, after considering available mitigations, it could create a severe impact on individuals, clients, the legal process, confidentiality, security, regulatory compliance, or the integrity of Jus Mundi's services. Examples include: unauthorized disclosure or processing of confidential legal or personal data; AI outputs that could materially mislead users in legal research or decision-making if presented without appropriate caveats or human review; system behavior that could compromise the availability, integrity or security of the platform; discriminatory or unfair outputs in sensitive contexts; or autonomous actions performed without adequate authorization, traceability or human oversight.

In practice, Jus Mundi evaluates these risks using decision criteria such as severity of potential harm, likelihood, scale of exposure, sensitivity of the data involved, reversibility of harm, detectability, regulatory impact, customer impact, and the effectiveness of controls. Risks that remain high or critical after mitigation require escalation to the relevant governance bodies and accountable functions, including Security, Legal/DPO, Product, Engineering and AI governance stakeholders. Deployment may be blocked, delayed, limited to a controlled environment, or subject to additional safeguards where residual risk is not acceptable.

Mitigations are graduated according to the capability and risk profile of the AI system. Lower-risk use cases may require standard controls such as documentation, user guidance, access control, logging, human review and monitoring. Higher-risk use cases may require enhanced review, security and privacy assessments, vendor due diligence, model or application testing, stricter access restrictions, output validation, contractual safeguards, phased deployment, incident response procedures and formal approval before release.

For generative AI, Jus Mundi pays particular attention to hallucination, explainability, source attribution, confidentiality, prompt/input handling, copyright and data protection risks, and user reliance on AI-generated outputs. For agentic AI or tool-using systems, additional scrutiny is applied to autonomy, permissions, ability to take actions, access to sensitive systems or data, auditability, rollback capacity, and human-in-the-loop controls. Agentic capabilities would be restricted or prohibited where they could perform high-impact actions without appropriate authorization, monitoring, and intervention mechanisms.

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?

- Independent third-party red teaming
- Adversarial testing to detect technical vulnerabilities
- Publicly available benchmarks and measurement tools
- Internal benchmarks, standard metrics (e.g., frequency, severity scores)
- Structured techniques to explore the behavior or performance of a model and/or application, such as model elicitation or bottleneck assessments
- Qualitative assessments (e.g., expert judgment, narrative scoring)
- Impact assessments such as personal data protection assessment and fundamental rights impact assessments

- Internal model, system, or application-independent research initiatives
- 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)
- Specific risk identification practices for agentic AI capabilities, e.g., autonomous action-taking, tool use, multi-step planning (please specify): /
- Multilingual and cross-cultural evaluation, including collaboration with local ecosystems (e.g., national AI institutes, language-specific evaluation consortia) for under-represented languages and cultural contexts (please specify): /

Are any national AI safety / security institutes of the AISI network involved?

No answer provided

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.

Jus Mundi uses a risk-based approval process before deploying AI systems. Decision authority depends on the risk level and involves the Product and Engineering owners, with Security, Legal/DPO, and AI governance review where confidentiality, privacy, compliance, customer impact, or service integrity may be affected.

Before release, the team reviews the intended use, data involved, vendor or model assessment, security and privacy risks, testing results, output quality, robustness, misuse scenarios, access controls, monitoring, and compliance requirements. Deployment is allowed only when performance is acceptable for the intended use, critical risks are mitigated, and the residual risk is within the organization's risk appetite.

Higher-risk or customer-facing systems require additional sign-off and may be escalated to governance bodies. Required safeguards may include human oversight, access controls, data minimization, logging, user guidance, source attribution, monitoring, incident response, rollback mechanisms, and phased deployment. For agentic AI, Jus Mundi applies stricter controls on autonomy, permissions, approval gates, audit trails, and the ability to suspend unsafe behavior.

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.

After deployment, Jus Mundi reviews internal incidents, vulnerability reports, customer or user feedback, monitoring alerts, audit findings, and relevant external incident reports to identify emerging AI-related risks. These inputs are assessed for severity, likelihood, affected systems or users, data sensitivity, regulatory impact, and relevance to Jus Mundi's AI use cases. When an incident or external report indicates a new or increased risk, the relevant AI system risk assessment is updated and the issue is escalated to the appropriate Product, Engineering, Security, Legal/DPO, or AI governance stakeholders. Corrective actions may include fixing vulnerabilities, changing prompts or model behavior,

improving monitoring, strengthening access controls, adding user guidance or human review, restricting or pausing a feature, notifying affected parties where required, and updating policies, procedures, or controls. Lessons learned are incorporated into future testing, deployment criteria, incident response, and AI governance reviews.

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.

No answer provided

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”

Not applicable

How does your organization share relevant information as appropriate?

No answer provided

Please describe what type of information is shared.

/

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

- Paid rewards (“bug bounty”)
- 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:

No answer provided

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.)		☐	
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.

Jus Mundi collaborates with internal and external stakeholders through a risk-based governance model. Internally, Security, Legal/DPO, Product, Engineering, Compliance and business teams work together to assess AI-related risks, define mitigations, review incidents, and decide whether AI systems or features can be deployed. Externally, Jus Mundi collaborates with enterprise customers, users, suppliers, cloud and AI service providers, auditors, penetration testers, and legal or regulatory advisors. These

collaborations support vendor due diligence, security and privacy assessments, independent testing, audits, incident handling, contractual safeguards, and improvement of risk controls.

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).

Jus Mundi manages AI-related risks through an integrated security, privacy and AI governance process covering the full lifecycle: design, data handling, development, validation, deployment and monitoring. Risks are identified through AI/security assessments, DPIAs where relevant, threat modeling, architecture reviews, vendor due diligence and secure development controls. Findings are recorded, assigned owners, prioritized by severity and tracked through remediation or formal risk acceptance. Cross-cutting governance is coordinated by Security, Engineering, Legal/DPO and relevant business owners. Higher-risk issues, incidents, major vulnerabilities, privacy concerns or material product changes are escalated through the security incident/risk management process and, where appropriate, management review. Generative AI systems receive additional controls for prompt injection, data leakage, output reliability and human review. Agentic AI systems, where applicable, require stronger safeguards around tool permissions, autonomy limits, logging, interruption/override mechanisms and monitoring.

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).

Jus Mundi applies data quality, confidentiality and fairness controls during data collection and curation. Data sources are reviewed before use to confirm relevance, legitimacy, quality and suitability for the intended AI use case. Personal, confidential or unnecessary data is minimized, filtered or excluded where possible. Access to datasets is restricted on a need-to-know basis, and sensitive data is handled under privacy, security and contractual controls. For AI systems relying on third-party models or providers, Jus Mundi also assesses provider practices for data handling, training, retention and security. Where internal datasets are used, teams review data representativeness, known limitations and potential sources of bias, and apply human review or testing where outputs may affect users or legal workflows. Limitations are documented so that AI outputs remain contextualized and subject to appropriate human oversight.

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

- Isolated and secure test environments
- Network isolation from external access
- Isolation of staging environments (from production / operational environments)
- Controlled environments for red-teaming and adversarial testing
- Containment measures for prompt-injection and jailbreak attempts
- Controlled security vulnerability testing
- Rate limits and resource caps
- Human-in-the-loop controls for higher-risk test cases
- Access controls, including roles-based permissions
- Disabling, restricting or simulating high-impact capabilities during testing
- Audit logging and monitoring
- Use of “AI-safe” test-data, as appropriate
- Use of synthetic, anonymized, or dummy data, as appropriate

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.

No answer provided

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:

- Publishing “model cards” or contextual information on training data and model limitations
- Apply provenance controls to guard against training data being sourced from websites or repositories that have been determined to engage in large-scale piracy
- Refrain from the unlawful circumvention of technological protection measures in data sourcing

- Maintain policies or practices to respect copyright holders' preferences through, where appropriate/applicable, mechanisms like Robots.txt "do not train" tags or other machine-readable protocols developed through voluntary, industry-led standards processes
- 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

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)

Describe key privacy measures and data-protection standards applied.

Jus Mundi implements comprehensive technical and organizational measures to protect personal data across the AI lifecycle, in compliance with the GDPR and applicable privacy laws.

Security program: Jus Mundi maintains an information security program aligned with **ISO/IEC 27001:2013**, reviewed at least annually.

Encryption: All subscriber data is **encrypted at rest and in transit** using industry-standard cryptographic techniques with secure key management.

Access control: Access to personal data is granted on a strict **need-to-know basis**, enforced through **multi-factor authentication**.

Incident management: Formal processes are in place to detect, investigate, and respond to security incidents, including timely notification to subscribers and relevant supervisory authorities.

Sub-processors & international transfers: All sub-processors undergo due diligence assessments. International transfers are governed by appropriate safeguards, including the **EU-U.S. Data Privacy Framework** and **Standard Contractual Clauses (SCCs)**.

Staff training: Personnel with access to personal data receive regular privacy and security training.

Business continuity: A disaster recovery plan is maintained and tested annually to ensure operational resilience.

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

- Conduct regular cybersecurity risk and vulnerability assessments tailored to AI models or systems
- Perform threat modeling that includes AI-specific risks (e.g., data poisoning, prompt injection)

- 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
- Train staff on cybersecurity best practices and AI-specific threats, including how to recognize and respond to misuse or unintended outcomes of advanced AI systems
- 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

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
- Restrict physical access to servers, GPUs, or devices hosting AI models, systems, and/or applications
- Monitor system activity and log access to detect anomalies or misuse
- Conduct secure AI development and/or deployment practices
- Conduct security testing focused on AI-specific risks (e.g., red-teaming, adversarial testing, stress tests)

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
- Provide staff training on cybersecurity best practices and AI-specific threats
- Implement insider-risk escalation or response procedures

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

Jus Mundi applies security measures including MFA and role-based access controls, encryption at rest and in transit, secure development and deployment practices, vulnerability management, security testing, audit logging, monitoring, incident response, disaster recovery, and regular security reviews. For particularly sensitive AI assets, such as proprietary datasets, model configurations, prompts, logs, and AI system credentials, access is limited to authorized personnel on a need-to-know basis. These assets are stored in secure, access-controlled environments, monitored through logging, and protected

by confidentiality obligations for employees and contractors. Additional controls may include enhanced review, restricted permissions, segregation of environments, and human oversight for higher-risk AI use cases.

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
- Bug bounty or incentive programs for external reporters
- 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
- Collaborative vulnerability sharing (e.g., with partners, industry groups)
- Reporting to regulators / authorities when relevant
- Recognition of external finders of confirmed risks or vulnerabilities
- 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
- 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)

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.

- Other (specify): Jus AI uses an agentic architecture: specialized components plan, decompose and execute legal-research tasks over multiple loops, and select search filters and research skills autonomously. Its autonomy is deliberately bounded to research within the Jus Mundi corpus and the user's own documents. It does not browse the open internet, take actions in external systems, or execute transactions, and every output is returned to the user for their judgment. Human oversight, interruption and override are retained at all times (see Principle 1 of our AI Manifesto).

Please provide any further comments and implementation documentation.

No answer provided

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?

Yes

15.A. If yes, which of the following information is included in your organization's publicly available documentation for any of your AI products?

- Other (specify): AI Manifesto, Trust Center, Privacy Charter

Please provide links or references showing where this information is made public.

No answer provided

15.B. If yes, how often are such reports usually updated?

No answer provided

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

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.

Jus Mundi shares information about its AI systems primarily through its publicly available **AI Manifesto**, which provides plain-language summaries of AI capabilities, limitations, safeguards, and intended use cases. This document is accessible to all stakeholders without requiring technical expertise.

At the contractual level, clients (organizations and institutions) receive additional detail through the **Data Processing Agreement (DPA)** and Terms of Service, which document data handling practices, security measures, and risk mitigation commitments.

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.

No answer provided

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?

- Public documentation

Please provide links to any publicly available information.

<https://trust.jusmundi.com/>

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

No answer provided

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).

Jus Mundi has a dedicated DPO/legal compliance function, with an estimated 1 FTE per year allocated to AI and data protection compliance activities (GDPR, EU AI Act, DPA management, vendor due diligence, staff training). Product and engineering teams contribute additional time on an ad hoc basis for privacy-by-design implementation.

Estimated annual compliance expenses: €50,000–€100,000, covering internal staff costs, external legal counsel, security tools, and training.

Section 4 – Organizational Governance, Incident Management, and Transparency

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

- Senior management oversight
- Independent advisory body
- 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 at Jus Mundi is integrated into governance through a dedicated cross-functional body, the AI Concilium, which brings together relevant teams to supervise AI risk management and ensure alignment across legal, product, and engineering functions.

Advisory roles: The AI Concilium provides ongoing oversight, risk assessment, and recommendations on AI-related matters.

Final accountability: Ultimate responsibility for AI risk-related decisions rests with the Head of Legal and the Head of AI, who have authority to approve policies, respond to incidents, and validate material changes to AI systems.

Policy review: All AI and data protection policies are reviewed at least annually, and updated as needed following regulatory changes, significant product updates, or material incidents.

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).

/

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.

/

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).

Through our AI Manifesto, the AI Hub, the AI Core newsletter and our training for the legal community.

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)
- Other international or regional standards (please specify): SOC 2 Type II

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

ISO/IEC 42001 (certified), ISO/IEC 27001, SOC 2 Type II

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
- Distinct branding or labeling for AI-generated content (e.g., “AI-generated” tag, metadata label)

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).

Users of Jus Mundi's AI features (Jus AI) are informed of their interaction with an AI system through multiple touchpoints: a clear disclosure in the user interface at the point of interaction, labeling of AI-generated content within the product (“AI-generated output”), and explicit mention in the Terms of Service and user documentation.

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

Yes

27.A. If yes, how?

No answer provided

27.B. If yes, does your organization use international technical standards or best practices when developing or implementing content provenance?

No answer provided

27.C. Does your organization collaborate on or invest in research to advance content authentication and provenance?

No answer provided

Describe how provenance mechanisms are implemented and which standards or tools are used, and any collaborations or investments to improve content authentication (e.g., C2PA credentials, joint R&D projects, open-source initiatives, standard-setting groups). You may also indicate whether these mechanisms are tested for robustness against removal, tampering, or circumvention.

No answer provided

Any further comments and implementation documentation.

Jus Mundi labels AI-generated content directly within the user interface, allowing users to clearly distinguish between AI-generated outputs and source legal documents. AI-generated summaries, translations, and analytical responses are systematically identified as such at the point of display, ensuring full transparency on content provenance.

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

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

- Data quality and representativeness
- Information integrity
- System performance and outcomes (e.g. accuracy, consistency)

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.

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?

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.

Jus Mundi delivers regular AI literacy programmes for academic and professional audiences, including university seminars, onboarding sessions and role-specific bootcamps for foundation, growth and leadership levels. The programme combines core AI concepts with practical instruction on capabilities and limitations, confidentiality and security, ethical use and the continued importance of human judgment. Training is adapted to participants' responsibilities and is supported by documented curricula, presentation materials, practical case studies and a centralised training calendar and resource hub.

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

Yes

If yes, please provide examples, including any collaborations with civil society, community groups, or efforts to address broader societal challenges.

No answer provided

Any further comments and implementation documentation.

No answer provided