

OECD AI Transparency Report

Organization: Felaris Global (USA) Reporting Period: Q2 2026 Published: October 2, 2026

Section 0 – Your role in the AI system lifecycle

Responding as (please select all that apply):

- Application developer (/provider)
- 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): EU AI Act risk tiers; NIST AI Risk Management Framework; ISO/IEC 42001; OECD AI Principles.

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.

Felaris does not develop foundation models and does not set frontier capability thresholds. As an application developer and deployer, and as a micro enterprise, it applies a risk-based approach built on the EU AI Act risk tiers, the NIST AI Risk Management Framework, ISO/IEC 42001, and the OECD AI Principles. Severe or unacceptable risk is treated as any use that a client's data classification, its regulatory obligations, or the EU AI Act would place off limits, and such uses are excluded in the governance framework Felaris installs and applies to itself. For its own applications, the primary risk is

data quality, assessed on ingest before anything is modelled, and each application is scoped to decisions where sufficient history exists; where a situation has no precedent, the system is designed to say so. Where third-party foundation models are used for extraction, outputs pass human confirmation before use. Felaris does not currently build agentic AI.

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?

- Qualitative assessments (e.g., expert judgment, narrative scoring)
- Impact assessments such as personal data protection assessment and fundamental rights impact assessments

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.

Felaris applications are pre-production, so deployment decisions are at the design stage. Decision authority sits with the CEO, who holds internal AI accountability. The criteria to move toward deployment are: data quality and readiness assessed on ingest; a governance framework and data rule in place; human confirmation of any AI-extracted input; data residency and per-tenant isolation; and separation of development and production environments. First release is a scoped pilot on one decision and one data domain, tested against a recent event the organisation already lived through, before any wider rollout. Sign-off is against the Felaris AI Governance framework, and monitoring for model and data drift is a precondition of production.

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.

Felaris has no deployed systems, so there is no live incident-handling record. The design provides for a use register and an incident register with named owners and a quarterly review, the same discipline Felaris installs for clients. Felaris follows publicly reported AI incidents and ecosystem guidance to inform its governance advisory and its own design decisions. At release, incident reports, including those shared by other organisations, will feed the register, update the risk assessment, and trigger corrective action through the governance framework.

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

No

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?

Contributes and 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)
- NIST AI Risk Management Framework (AI RMF)
- Industry-led best practices or codes of conduct
- Other international or regional standards (please specify): EU AI Act; OECD AI Principles; the Hiroshima International Code of Conduct; and the CARICOM AI and Data Governance Readiness Index, a Felaris-published benchmark of CARICOM jurisdictions against EU, US, and UK frameworks.

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.)		☐	

	Always-on	Ad hoc / as relevant	No current collaboration
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.

Felaris represented the private sector at the United Nations' inaugural Global Dialogue on AI Governance in Geneva in July 2026. It is co-founding the Caribbean Dialogue on AI, a standing non-commercial forum for the region's public and private sectors, and engages regional bodies including CANTO and the CTU Caribbean AI Forum, alongside academia through the University of the West Indies. It publishes open governance references (the CARICOM AI and Data Governance Readiness Index, the Obligation Clock, and a rules self-check) so leaders across the region can see which obligations apply to them, and it delivers governance and training in partnership with a regional ICT partner. This HAIP submission is itself a contribution to a multi-stakeholder transparency mechanism.

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).

Felaris runs a governance-by-design approach, applied to its own AI work the way it advises clients. At the plan and design stage, a use case is scoped against the EU AI Act risk tiers and the Felaris five-dimension readiness assessment (data, process, people, governance, technology). At the data stage, a data-quality report runs on ingest so gaps are visible before anything is modelled, and per-tenant isolation keeps client data separate. Build and integration use third-party foundation models for extraction, with a human confirmation step on every extracted input. Verification uses a scoped pilot on one decision and one data domain, tested against a recent event the organisation already lived through. Deployment is gated by dev/production separation, data residency, and sign-off against the Felaris AI Governance framework. Operation is designed for model and data drift monitoring, a use register and an incident register with named owners, and a quarterly review. Practices are reviewed at each release and on any incident. Felaris does not build agentic AI, and its applications are pre-production.

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).

Felaris does not train foundation models, so it does not curate large-scale training corpora. Its data-quality work sits at ingest for its own applications: uploaded data is validated against a defined template schema, and a data-quality report flags missing fields, stale dates, and gaps before anything is modelled, so the customer sees what is usable first. Any figure extracted from unstructured documents by a third-party model passes human review and confirmation before it enters the system. Per-tenant isolation keeps one client's data from mixing with another's. Fairness is handled by scoping each application to decisions with sufficient history, by stating limits where a situation has no precedent, and by human oversight of outputs. In its published CARICOM Readiness Index, every score is sourced, dated, and human-reviewed before publication.

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

- Isolated and secure test environments
- Isolation of staging environments (from production / operational environments)
- Human-in-the-loop controls for higher-risk test cases
- Access controls, including roles-based permissions
- Audit logging and monitoring
- 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.

Secure testing is overseen by the CEO, who holds internal AI accountability, working with a senior Felaris AI engineering lead. As a micro enterprise, this responsibility sits with a small senior team rather than being split across many functions.

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:

- 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:

- Not applicable

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

- 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.

Felaris applies per-tenant data isolation (database-per-tenant for the first clients and for anyone with residency needs) so no client's data mixes with another's, and region-specific hosting to support data residency and sovereignty. Access is role-based and least-privilege through a managed identity provider with multi-factor authentication, and data is encrypted in transit and at rest. Personal and confidential data is handled under a data-classification overlay and a one-page data rule, the same instruments Felaris installs for clients, with GDPR-aware practice. Any data extracted by a third-party model passes human confirmation before use. Felaris carries records and audit-traceability discipline from regulated-industry work under FDA, EMA, and EU MDR conditions. In line with this framework's guidance, no confidential, proprietary, or personal data is included in this report.

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
- 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
- These responsibilities are delegated contractually to model or infrastructure providers

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
- These responsibilities are delegated contractually to model or infrastructure providers

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
- These responsibilities are delegated contractually to model or infrastructure providers

11.D. Insider risk protection detections

- Monitor for unauthorized internal access or suspicious activity
- Provide staff training on cybersecurity best practices and AI-specific threats
- 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.

Felaris runs a proportionate, standards-aligned security posture for a micro enterprise. Core measures are a managed identity provider with multi-factor authentication, role-based and least-privilege access, encryption in transit and at rest, separation of development and production environments, and audit logging. Infrastructure and identity security is handled through specialist managed providers rather than built in-house. Staff receive cybersecurity and AI-threat awareness, and security policies are aligned to recognised standards and reviewed on emerging risks and releases. Measures specific to sensitive AI assets: per-tenant data isolation (database-per-tenant where residency requires it), access to proprietary data and design assets limited by role and need-to-know and logged, and NDAs or confidentiality agreements for staff and contractors with access. Felaris does not hold third-party model weights; where foundation models are used, that risk sits with the provider under contract.

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

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

- Other(s) (please specify): Felaris applications are pre-production, so no post-deployment monitoring is running yet. Post-deployment monitoring, model and data drift detection, root-cause analysis, and an incident register with named owners and a quarterly review are designed and will be in place before any production release.

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

No

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?

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
- 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.

Because Felaris applications are pre-production, there are no product evaluation results to share yet. Felaris shares governance knowledge publicly through its maintained instruments at ai.felarisglobal.com: the CARICOM AI and Data Governance Readiness Index, the Obligation Clock, and a rules self-check, all written in plain language for non-technical leaders, with sources and dates shown. It engages directly with the public sector, civil society, and academia through the UN Global Dialogue on AI Governance, the Caribbean Dialogue on AI, CANTO, the CTU Caribbean AI Forum, and

the University of the West Indies. When its applications reach evaluation, results will be shared with clients and, where appropriate, summarised publicly for broader audiences.

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

No

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. Felaris does not develop or train foundation models. It uses third-party foundation models under contract for narrow extraction tasks, so it does not source or annotate training data.

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

Felaris demonstrates transparency through its published governance instruments (the CARICOM Readiness Index, the Obligation Clock, and a rules self-check), through this HAIP report, and through public participation in the UN Global Dialogue on AI Governance and the Caribbean Dialogue on AI. Felaris does not build AI agents.

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).

No answer provided

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

- Executive oversight
- Senior management oversight
- Independent advisory body

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).

Felaris is a micro enterprise, so AI risk management sits with its senior leadership rather than a large committee structure. The CEO holds final accountability for AI risk-related decisions and internal AI governance. An independent Advisory Board provides external input and challenge on strategy and direction, in an advisory rather than decision-making role. Felaris runs its own AI activity inside the same AI Governance framework it installs for clients, with a named AI owner, a data rule, and use and incident registers. Policies are reviewed on any incident, on regulatory change, on material updates to systems, and at least annually.

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).

As a micro enterprise, training is delivered directly by senior leadership. New team members and contractors are briefed on the internal AI-use policy and data rule at onboarding, and receive regular updates as the governance framework and approved tools change. The Felaris AI Academy curriculum the firm delivers to clients also informs its own staff capability. Concerns can be raised directly to the CEO, and AI tool failures and issues are logged internally.

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.

Felaris has no live production incidents, since its applications are pre-production. It maintains an internal log of AI tool failures and issues in its own operations, recording the issue and the corrective step taken. For client governance engagements, Felaris stands up use and incident registers with named owners and a scheduled quarterly review. The design for its own products includes an incident register, root-cause analysis, and escalation to the CEO, with external escalation to a client or regulator where an obligation requires it.

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).

Felaris shares governance research and good practice publicly. It publishes and maintains the CARICOM AI and Data Governance Readiness Index, the Obligation Clock, and a rules self-check. It participates in the UN Global Dialogue on AI Governance and the CTU Caribbean AI Forum, is co-founding the Caribbean Dialogue on AI, and collaborates with the University of the West Indies. It also shares practice through articles and posts and through this HAIP report.

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)
- NIST AI Risk Management Framework (AI RMF)
- OECD AI Principles
- Industry-led best practices or codes of conduct
- Other international or regional standards (please specify): EU AI Act; the Hiroshima International Code of Conduct; and the CARICOM AI and Data Governance Readiness Index, which benchmarks CARICOM jurisdictions against EU, US, and UK frameworks.

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

Felaris aligns its AI governance and risk-management work to ISO/IEC 42001 and related ISO/IEC standards, the NIST AI Risk Management Framework, the OECD AI Principles, the EU AI Act, and the Hiroshima International Code of Conduct. It also works under the regulated-industry standards its clients operate within, including FDA, EMA, EU MDR, SOX, and IT general controls, and it benchmarks against EU, US, and UK frameworks through the CARICOM Readiness Index.

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)
- Documentation or FAQs explaining when and how AI is used in the system

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).

These mechanisms are planned, since Felaris applications are pre-production and have no end-users yet. Veritas Predict is a business-to-business decision tool whose users know they are using an AI decision system; at release it is designed to disclose AI use through a system label and documentation, and its outputs carry confidence bands and an "as of" timestamp. It does not act autonomously on a user's behalf.

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?

- Data quality and representativeness
- Explainability or interpretability
- Transparency

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 a micro enterprise, Felaris invests in governance, transparency, and data quality rather than frontier safety research. It builds and maintains public governance instruments (the CARICOM AI and Data Governance Readiness Index, the Obligation Clock, and a rules self-check), all publicly shared with sources and dates. In its products it treats data quality as the primary risk, with a data-quality report on ingest, and it invests in explainability through confidence bands and clearly stated limits in Veritas Predict.

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.

Felaris's socio-economic focus is AI readiness and workforce capability in an under-served region. The Felaris AI Academy delivers tiered, role-based AI training to enterprise and government audiences (skills development and access to AI capability). Its AI readiness assessments help organisations adopt AI productively, and the CARICOM Readiness Index and the Caribbean Dialogue on AI support regional capability and data sovereignty.

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.

The Felaris AI Academy is the firm's training arm: a tiered, role-based curriculum (AI Essentials, AI Fluency, AI Foundations, and AI Practitioner) delivered live to enterprise and government audiences, teaching safe and effective AI use and the limits of these tools. It is delivered in Trinidad and Tobago through a regional partner. Felaris also publishes free public references (the rules self-check, the Obligation Clock, and the CARICOM Readiness Index) that help non-technical leaders understand which obligations apply to them. This work prioritises the Caribbean, an under-served region for AI literacy and readiness.

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.

Felaris publishes free public governance goods for the Caribbean: the CARICOM AI and Data Governance Readiness Index, the Obligation Clock, and a rules self-check. It is co-founding the

Caribbean Dialogue on AI, a standing non-commercial forum for the region's public and private sectors, and it represented the private sector at the UN Global Dialogue on AI Governance in Geneva. It engages CANTO, the CTU Caribbean AI Forum, and the University of the West Indies. The through-line is regional AI readiness, data sovereignty, and helping a small region take part in AI governance on its own terms.

Any further comments and implementation documentation.

No answer provided