OECD AI Transparency Report

Organization: Cognisant LLC (JPN) **Reporting Period:** Q4 2026 **Published:** November 10, 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]

- 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): Japan's AI Guidelines for Business Ver1.2 (Ministry of Internal Affairs and Communications; Ministry of Economy, Trade and Industry); ISO/IEC 27001 as a reference for information-security risk
- Maintain an internal taxonomy, classification system, and/or set of risk area definitions

> **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.**
>
> We do not develop foundation models and do not use numeric capability thresholds. We decide case by case. We do not build or use AI in ways that create discrimination or unfair bias, could harm life, body or property, or break the law (AI Use Policy, Art. 6), and we decline such uses. For systems that support important decisions, a person makes the final decision by design (Art. 1). For generative AI, the main risks we track are inaccurate output reaching customer work without review, and customer data being

> sent beyond the agreed scope. For AI agents, we treat actions that reach outside the company or cannot be undone as the main risk.

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)
- 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)
- 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): we review what each agent may do, which actions need a person's approval, and what is logged

> **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.**
>
> For customer systems, the customer decides on production use, based on criteria and procedures agreed with us (AI Use Policy, Art. 3). On our side, the Representative Member (the managing member and chief executive of the LLC) is accountable. Evidence: accuracy measured on evaluation questions drawn from the customer's actual work, including paraphrased questions; checks of which data fields are sent to a model; access control, logging and security settings agreed at the start of the project. Caveats: results from synthetic data or small samples are treated as an initial diagnosis and measured again on real use. Safeguards: uncertain results go to a person; we release in stages where possible and keep measuring after release. We do not currently use independent external testing.

> **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.**
>
> We record incidents and security events in an internal log with the cause, the handling and the prevention measure. Staff report to the Representative Member within one hour, and prevention measures are decided within two weeks. We follow external information through a regular internal digest and public security advisories. When an issue applies to us, we update our procedures and report to any affected customer.

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.): *cloud and AI model providers, through their support channels*		☐	
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”

Anyone can write to info@cognisant.io (listed in our privacy policy, https://cognisant.io/privacy). Customers also use the contact points set in each contract. We do not run a dedicated AI-incident form. Reports to authorities follow legal requirements, including Japan's Act on the Protection of Personal Information.

How does your organization share relevant information as appropriate?

- Information-sharing partnerships (with confidentiality agreements as appropriate)

- Reporting to regulatory authorities as relevant

Please describe what type of information is shared.

To affected customers: the facts of an incident, its scope, the handling and the prevention measure. To Japan's Personal Information Protection Commission: breach reports where required.

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

- Not applicable

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

Uses

If contributes or uses, please select all that apply:

- ISO/IEC standards (e.g., ISO/IEC 27001, ISO/IEC 42001, ISO/IEC 23894, ISO/IEC 42005)
- Other international or regional standards (please specify): Japan's AI Guidelines for Business Ver1.2

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.): *cloud and AI model providers*		□	
Academia / researchers			□

	Always-on	Ad hoc / as relevant	No current collaboration
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.**
>
> On each project we work with the customer on evaluation criteria, on which data may be sent to a model, and on which results must go to a person. We do not take part in multi-stakeholder initiatives at present.

Any further comments and implementation documentation.

Notes on the selections above.
Q2: Adversarial testing: prompt-injection test cases on some projects.
Q2: Structured techniques: measuring each processing stage separately; paraphrased versions of evaluation questions.
Q2: Automated monitoring: cloud security monitoring and error alerts on some projects.
Q3: Media / general public: public contact address.
Q5: ISO/IEC standards: ISO/IEC 27001 as a reference; we are not certified.

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).

The Representative Member owns AI risk and decides on escalations. We follow our published AI Use Policy and internal security rules. As part of our ISO/IEC 27001 preparation, we keep a risk register with AI entries, and in November 2026 we started a monthly internal check of devices, account permissions and incident records. We review policies and rules every April and when laws, customer requirements or the services we use change. In November 2026 we completed a self-assessment against Appendix 7 of Japan's AI Guidelines for Business. For generative AI, the focus is on accuracy and data scope. For AI agents, it is on permitted actions and approval.

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).

We do not train foundation models. Source documents and evaluation data come from the customer's work, and we agree with the customer how they are kept current. Personal data that is not needed is removed or masked before it is sent to a model. Formal bias testing is not yet a standard step; we are adding a check for bias factors to our project-start checklist, and we will apply it to uses that affect individuals or groups.

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

- Isolation of staging environments (from production / operational environments)
- 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.

The Representative Member.

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

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

- Not applicable

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

- Other(s) (specify): staff review output before it is used in customer work (AI Use Policy, Arts. 1 and 4). Concerns can be sent to our general contact address (info@cognisant.io); a notice for this is being added to the AI Use Policy

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)
- 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.

We follow Japan's Act on the Protection of Personal Information and publish a privacy policy. Customer information is entered into generative AI only within the scope the contract allows. We use only services whose terms or settings keep inputs out of training, and we send only the fields the work needs. On some projects, the system itself limits which fields can be sent. Staff must not put work information into personal AI accounts.

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

11.A. Governance, policy and assessments

- Maintain or update security policies aligned with recognized standards (e.g., NIST CSF, ISO/IEC 27001, SOC 2, CIS Critical Security Controls), and review them regularly in response to emerging risks or regulatory guidance
- Maintain vulnerability management processes (scanning, channels for external reporting, patching, remediation)
- Implement incident response, disaster recovery, or business continuity plans, including escalation procedures for high-severity incidents
- Integrate cybersecurity AI risks into broader cybersecurity risk management activities
- Regularly review security measures and update them in response to emerging risks, incidents, or major releases
- Other(s) (specify): policies and rules are explained to staff on joining, with acknowledgement recorded since November 2026; training focused on AI-specific threats is in preparation

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
- Conduct security testing focused on AI-specific risks (e.g., red-teaming, adversarial testing, stress tests)
- Other(s) (please specify): physical security of the servers that host models is the responsibility of our cloud and AI model 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
- Require NDAs or confidentiality agreements for staff and contractors with access

11.D. Insider risk protection detections

- Monitor for unauthorized internal access or suspicious activity
- Other (please specify): see Q11A on staff training

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

We have declared two stars under SECURITY ACTION, the self-declaration program of Japan's Information-technology Promotion Agency (a self-declaration, not a certification). Work devices must have storage encryption, a screen lock within five minutes, automatic updates and malware protection, and the Representative Member checks the device list. MFA is required for administrator accounts and is being extended to all accounts. Our rules require a password vault. A check for secrets before commit is in place in the Representative Member's development environment and is being extended to all staff. We are preparing an ISO/IEC 27001 information security management system; we are not certified.

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

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.**
>
> - 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)
> - Logging of agent actions
> - Other (specify): documents sent outside the company go through machine checks and a separate AI review before a person approves them

Please provide any further comments and implementation documentation.

Q7 practices by lifecycle stage.
Plan and design: agree the use, the data scope and the role of people with the customer; design so that people make important final decisions; use deterministic rules where a rule can decide, with the model choosing among fixed options where possible
Collect and process data: limit the fields sent to a model; remove or mask personal data that is not needed; keep confidential data out of source repositories; check for secrets before code is committed (in place in the Representative Member's development environment and being extended to all staff)
Build, integrate and interpret model: use third-party foundation models under terms or settings where inputs are not used for training; keep code and prompts under version control
Verify and validate: measure each processing stage; test with paraphrased questions; prompt-injection test cases on some projects; review by a person
Deploy: use criteria agreed with the customer; separate staging from production; release in stages where possible
Operate and monitor: keep measuring accuracy after release and investigate drops; cloud audit logs and security monitoring on projects we run
Notes on the selections above.
Q9A: we do not train foundation models.
Q10: Controls for organizational customers: the scope of data sent to a model is agreed for each project.
Q10: Privacy-Enhancing Technologies: masking, k-anonymization and synthetic test data on some projects.
Q11A: Incident response, disaster recovery or business continuity plans: incident response in place; business continuity plan in preparation.
Q11B: MFA and role-based access controls: see Q11E.
Q11B: Security testing focused on AI-specific risks: on some projects.
Q13: Post-deployment monitoring: for systems we operate.

Q13: Monitoring and emergency controls specific to agentic AI systems: agents can be interrupted and their access revoked.
Q14: We use AI agents for software development and back-office work, and we build retrieval and answering systems with agent-like steps for customers.
Q14A: Constraints on the agent's action space: tool-use permissions, separate accounts, scope limits.
Q14A: Human oversight or interruptibility mechanisms: outbound messages to customers and authorities, payments and signatures need a person's approval, except routine actions the Representative Member has approved in advance.
Q14A: Logging of agent actions: including a log of actions completed automatically.

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?

- Protected information sharing

> **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.**
>
> We share evaluation results and their limits with the customer concerned, under the confidentiality terms of the contract. We do not publish project results that would identify a customer.

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

Yes

> **If yes, provide links to privacy policies or notices covering personal data, user inputs, and AI-generated outputs.**
>
> https://cognisant.io/privacy covers personal data; https://cognisant.io/ai-policy covers inputs to generative AI (Art. 2). Outputs are covered only by our rule that a person reviews them (Art. 1). Both pages are in Japanese.

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.

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**
>
> We publish our AI Use Policy (https://cognisant.io/ai-policy), information security policy (https://cognisant.io/security-policy) and a page on how we use AI and AI agents in our own work (https://cognisant.io/dx). All are in Japanese.

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).

Notes on the selections above.
Q18: we do not train foundation models.

Section 4 – Organizational Governance, Incident Management, and Transparency

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

- Executive oversight

> **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).**
>
> We are a company of five people, including the Representative Member, and have no board. The Representative Member decides on AI risk and is finally accountable, and serves as the person responsible for information security and personal data. There is no separate advisory body. Policies and rules are reviewed every April and after incidents or changes in laws, customer requirements or the

services we use. We have not yet assigned an independent reviewer; we will decide this as part of our information security management system preparation.

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).**
>
> On joining, or at the start of a contract, staff are told about our internal rules and the three public policies. We started recording acknowledgements in November 2026. Staff can raise concerns directly with the Representative Member. Material focused on the reliability and risks of generative AI is in preparation.

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

Yes

> **If yes, describe how incident records are kept and used for follow-up (e.g., incident logs, root-cause analysis reports, internal tracking systems). You may indicate whether and how you define escalation thresholds for when incidents must be surfaced to senior leadership or external parties.**
>
> Incidents and events go into a log kept on an access-controlled shared drive, with what happened, the handling and what will change. Staff report to the Representative Member within one hour. Affected customers are told promptly, and breaches of personal data are reported to the regulator within the legal time limits. The procedure does not yet list AI-specific incidents such as wrong output or prompt injection; we are adding them.

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).**
>
> We publish articles in Japanese on evaluating and improving AI systems, including how to build evaluation sets and how to find where accuracy is lost (https://cognisant.io/insights).

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): Japan's AI Guidelines for Business Ver1.2

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

We used Appendix 7 of Japan's AI Guidelines for Business Ver1.2 for a self-assessment in November 2026 and will repeat it every year. We are preparing an information security management system based on ISO/IEC 27001, and we plan to consider ISO/IEC 42001 after that. These are self-assessments and preparation, not certifications.

Any further comments and implementation documentation.

Notes on the selections above.
Q25: ISO/IEC standards: ISO/IEC 27001 as a reference; not certified.

Section 5 – Content Authentication & Provenance Mechanisms

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

- Other (specify): we explain to each customer which work uses which AI, with its limits and risks (AI Use Policy, Art. 5)

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).

On customer systems, how users are told is agreed with the customer for each project. A company-wide rule for on-screen AI labels is not yet in place; we are adding it to our project-start checklist. Our own public tools (an AI diagnosis tool and a chat assistant) are being reviewed for on-screen AI labels.

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?

- System performance and outcomes (e.g. accuracy, consistency)
- 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.

This is practice-based work rather than formal research. We develop evaluation methods in our own work, such as measuring each processing stage separately and testing with paraphrased questions. Parts of this are shared in our public articles. We have no dedicated research team.

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.

We publish plain-language articles in Japanese on what AI systems can and cannot do in business use (https://cognisant.io/insights). We do not run programs aimed at underserved communities.

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

No

Any further comments and implementation documentation.

No answer provided