

OECD AI Transparency Report

Organization: OpenText (CAN) Reporting Period: Q2 2026 Published: October 2, 2026

Section 0 – Your role in the AI system lifecycle

Responding as (please select all that apply):

- Deployer

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

Large enterprise: 250 or more 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): OpenText’s AI risk classification and governance approach draws on established international frameworks, standards and regulatory guidance, including the OECD AI Principles and associated definitions, the Hiroshima AI Process International Code of Conduct, the EU AI Act’s definitions and risk-based approach, ISO/IEC 27001, and relevant NIST cybersecurity practices. OpenText also takes account of industry-led responsible AI practices, including Canada’s Voluntary Code of Conduct on the Responsible Development and Management of Advanced Generative AI Systems and the EU General-Purpose AI Code of Practice. These external reference points are operationalised through OpenText’s developing enterprise AI Governance Framework and existing enterprise security, privacy, legal, compliance and risk-management processes. The AI Governance Framework uses the OECD and EU AI Act definition of an AI system and applies a proportionate, use-case-based classification methodology. AI use cases are classified according to complexity and potential business impact, considering factors including the type of AI capability, degree of autonomy, system integrations, data sensitivity, external exposure, potential effects on individuals, and financial or operational impact. OpenText’s governance model applies progressively enhanced review, approval, documentation, monitoring and escalation requirements as risk increases. Generative and agentic capabilities receive particular consideration, including the scope and reversibility of actions, human oversight and approval points, identity and permission controls, and whether the system can write to or trigger actions in integrated systems. Relevant use cases may also be subject to targeted security, privacy, legal, data-

governance and third-party risk assessments. OpenText's Information Security Management System is aligned with ISO/IEC 27001, and its security practices are informed by recognised NIST cybersecurity practices. OpenText is also evaluating the NIST AI Risk Management Framework and ISO/IEC 42001 as its AI-specific governance model continues to mature. This reference should not be understood as a claim that OpenText is currently certified against, or has fully implemented, ISO/IEC 42001 or the NIST AI Risk Management Framework. OpenText periodically reviews its taxonomy, governance controls and risk-classification approach to reflect developments in applicable law, international standards, authoritative guidance, emerging AI capabilities and lessons from operational use cases. OpenText also contributes to international and industry initiatives, including the OECD HAIP Reporting Framework and the EU General-Purpose AI Code of Practice, to support greater international alignment, interoperability and transparency in AI governance.

- Other (please specify): AI use case are risk tiered by complexity and business impact in OpenText to enable proportionate governance oversight across the lifecycle.

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.

Deployment decisions require evidence from security testing, privacy reviews (DPIA where applicable), legal/compliance checks, and product approvals under OpenText governance (e.g., PSDL/SoP). Incident reports (internal and shared by others) inform updates to risk assessments and corrective actions.

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

- Internal red teaming³
- Adversarial testing to detect technical vulnerabilities
- Internal benchmarks, standard metrics (e.g., frequency, severity scores)
- Qualitative assessments (e.g., expert judgment, narrative scoring)
- Impact assessments such as personal data protection assessment and fundamental rights impact assessments
- Vulnerability reports
- Incident reports, including reports shared by other organizations
- Specific risk identification practices for agentic AI capabilities, e.g., autonomous action-taking, tool use, multi-step planning (please specify): Specific considerations are applied when assessing agentic AI capabilities.
- Assessment of risks arising from interactions among multiple AI agents
- Other(s) (please specify): Post-deployment performance monitoring.

How does your organization monitor the interactions of AI agents under your control with other agents / assess risk from multiple agents interacting?

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.

OpenText uses a risk-based, stage-gated process to approve AI deployment, where solutions must pass defined checkpoints for advancing AI use cases through their lifecycle stages—from Proof of Concept to Pilot and ultimately to Production with evidence on performance, risk, and compliance, decision and consultation authorities scale from a centralized AI Office (low risk) to a governance committee (higher risk) with participations from relevant SME groups.

Progression to deployment requires validated results (e.g., performance and value realization against KPIs), completion of applicable security, privacy, legal, and vendor assessments, and confirmed risk mitigations.

Final production approval requires formal sign-off - where appropriate, at management level-, defined ownership, and safeguards such as monitoring, controls, and a rollout plan, ensuring ongoing oversight and compliance.

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, OpenText reviews relevant internal incident records and information from external sources to detect new or recurring risk signals. Where relevant to an AI-enabled capability, these insights are incorporated into updated risk assessments and may result in investigation, remediation, strengthened safeguards, changes to technical or governance controls, and escalation through established incident-management processes.

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

Not applicable

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

If contributes or uses, please select all that apply:

- Industry-led best practices or codes of conduct
- General Purpose AI Code of Practice

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.

OpenText collaborates across internal and external stakeholders to support responsible AI deployment and to identify, assess, and mitigate risks associated with AI systems in enterprise environments. From an external engagement perspective, OpenText participates in multi-stakeholder and industry initiatives focused on AI governance, data protection, information security, and responsible innovation, including policy discussions, consultations, and working groups involving public authorities, international organizations, and industry forums. These engagements allow OpenText to share operational insights from an enterprise deployer perspective and to stay aligned with evolving regulatory and best-practice expectations.

OpenText also engages with public sector bodies, regulators, and standards communities on an as-needed basis, particularly in relation to transparency, incident management, cybersecurity, privacy, and compliance with emerging AI frameworks. Where appropriate, OpenText contributes feedback to consultations and participates in workshops or knowledge-sharing sessions addressing AI risk management, governance, and deployment challenges.

Internally, OpenText maintains cross-functional governance structures involving product engineering, security, privacy, legal, procurement, and compliance teams to ensure coordinated oversight of AI deployments. These structures enable internal knowledge sharing, escalation of emerging risks, alignment with legal and regulatory developments, and consistent application of safeguards across AI use cases.

In addition, OpenText collaborates with external subject-matter experts, vendors, and assurance providers as part of due-diligence, risk assessment, and security testing activities relating to AI-enabled products. This includes vendor engagement on security, privacy, and contractual safeguards, as well as reliance on external intelligence such as public incident reports, government advisories, and industry guidance to inform ongoing monitoring and risk mitigation.

Through these combined internal and external collaborations, OpenText seeks to promote responsible AI deployment, share knowledge, and address systemic and operational risks in a way that is proportionate to its role as an enterprise application deployer.

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).

OpenText manages risks and vulnerabilities across the lifecycle of AI-enabled capabilities through its established enterprise information security, privacy, and risk management framework, complemented by AI-specific application and platform controls where appropriate with the prime objective of maintaining its systems and data confidentiality, integrity and availability.

Security and risk management practices are implemented through OpenText's Information Security Management System (ISMS), aligned to ISO/IEC 27001 and ISO27034 and supported by a defense-in-depth control model across identity and access management, infrastructure security, application security, monitoring, and incident response.

AI-related risks are incorporated into product governance and lifecycle reviews across design, development, release, deployment, and operation. Reviews consider factors such as system functionality, data sensitivity, intended use, deployment context, and applicable legal or regulatory requirements. Where appropriate, application-level controls include authentication and authorization, encryption, retrieval constraints, content filtering, evaluation and observability, audit logging, human review, and business-specific policy enforcement.

Where OpenText incorporates third-party models, AI services, or infrastructure, OpenText governs their selection, integration, and use through due diligence, security and privacy review, contractual requirements, and ongoing third-party risk management. This allows OpenText to maintain governance over the AI capabilities it delivers while assigning technical responsibilities to the party that controls each underlying layer.

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).

OpenText addresses data quality and fairness through a layered governance approach covering both the AI technologies selected for use and the controls implemented within OpenText products and services.

Where OpenText incorporates third-party foundation models or AI services, OpenText may not control the large-scale datasets used by the provider to train the underlying model. OpenText nevertheless remains responsible for how those technologies are selected, integrated, governed, and used within OpenText solutions. Provider practices and safety capabilities are considered through vendor due diligence, security and privacy review, contractual requirements, and ongoing third-party risk management.

Within OpenText products and services, controls may include appropriate use-case and data scoping, data classification and access controls, privacy review, content filtering, retrieval constraints, evaluation pipelines, human review, documented limitations, and operational monitoring. These measures are designed to help mitigate risks such as bias, inappropriate content, data leakage, hallucinations, and other unintended outcomes, while maintaining the needed transparency.

This layered approach enables OpenText to address data quality and fairness through both direct application-level controls and governance of the external AI technologies incorporated into the solutions delivered to customers.

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

- Not applicable

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

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

- Not applicable

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

- Not applicable

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

- Make information about data privacy, practices, and controls accessible to a diverse set of stakeholders
- Use security controls, including encryption at rest and in transit

Describe key privacy measures and data-protection standards applied.

OpenText protects privacy and helps prevent unauthorized disclosure of confidential or sensitive data through its established enterprise security and privacy framework, complemented by controls applied to AI-enabled applications and services.

Controls include data classification, least-privilege access management, encryption in transit and at rest, secure development and deployment practices, logging and monitoring, privacy impact assessments, security risk assessments, and product governance reviews.

At the AI application layer, OpenText applies controls where appropriate such as authentication and authorization, retrieval constraints, content filtering, audit logging, human review, evaluation, and business-specific policy enforcement. Where applicable, customers are provided with controls to configure access, retention, permissions, and data-handling settings within the relevant product or service. When designing and applying controls, OpenText takes into account data protection and AI related applicable laws and standards.

Third-party AI services incorporated into OpenText solutions are also subject to appropriate security and privacy review, contractual data-protection requirements, and ongoing vendor governance.

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
- 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
- These responsibilities are delegated contractually to model or infrastructure providers
- Other(s) (specify): OpenText governs AI-enabled capabilities through an integrated enterprise governance, policy, and risk management framework supported by its Information Security Management System aligned to ISO/IEC 27001 and informed by industry leading recognized cybersecurity practices. AI-related risks are incorporated into established product, security, privacy, and risk-governance processes rather than managed in isolation. Relevant assessments consider the capabilities being deployed, data sensitivity, intended use, security and privacy risk, deployment context, and applicable legal and regulatory requirements. Cross-functional teams spanning engineering, legal, compliance, security, privacy, and AI governance work together to identify, assess, and mitigate risks associated with AI-enabled capabilities. Where third-party models or infrastructure are incorporated, OpenText manages associated risks through due diligence, contractual controls, provider governance, and ongoing oversight. OpenText continues to evaluate emerging AI-specific frameworks, including the NIST AI Risk Management Framework and ISO/IEC 42001. These frameworks should not be represented as formally implemented or certified unless that status is established.

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
- Other(s) (please specify): AI-enabled capabilities are implemented within OpenText's established technical and operational security environment, including role-based access management, strong authentication, encryption, secure configuration standards, monitoring and logging, secure

development and deployment practices, vulnerability management, and patching. These foundational controls are supplemented at the AI application layer, where appropriate, by authentication and authorization enforcement, content filtering, retrieval constraints, audit logging, evaluation pipelines, human review, and business-specific policy enforcement. OpenText also uses approved foundation models with provider safety capabilities as part of its model-selection approach. Where AI capabilities incorporate third-party models or infrastructure, OpenText governs the integration and associated risks through architectural controls, vendor due diligence, contractual requirements, and ongoing third-party risk management.

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

- Limit access to proprietary AI assets based on role or need-to-know
- Track and log access to sensitive AI resources
- Require NDAs or confidentiality agreements for staff and contractors with access
- These responsibilities are delegated contractually to model or infrastructure providers
- Other (please specify): OpenText protects sensitive and high-value AI assets according to the nature of the asset, its sensitivity, and OpenText's role in controlling or operating it. Deployment-layer AI assets - including sensitive customer data, prompts, configurations, integrations, logs, traces, credentials, access permissions, embeddings, and application components - are protected using established controls such as data classification, least-privilege access, encryption, monitoring and logging, confidentiality requirements, and secure development and deployment processes. Where models or AI components are packaged or otherwise controlled as part of OpenText products, applicable OpenText product and supply-chain security requirements are applied. Where third-party foundation model weights, training datasets, or infrastructure remain under the provider's control, OpenText governs those dependencies through provider selection, due diligence, contractual requirements, and third-party risk management.
- Use hardware security modules (HSMs) or trusted execution environments (TEEs) where feasible

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
- These responsibilities are delegated contractually to model or infrastructure providers
- Other (please specify): Provide staff training on AI-specific threats. OpenText manages insider risk through established enterprise security and workforce governance processes, and those controls apply to systems, services, and sensitive assets supporting AI-enabled capabilities. Controls include pre-employment screening or role-sensitive vetting, role-based and least-privilege access controls, monitoring for unauthorized access or anomalous activity, cybersecurity and information-protection training, and established escalation and 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.

OpenText applies layered security measures to AI-enabled capabilities as part of its mature enterprise information security program. The program is aligned to ISO/IEC 27001 and ISO27034 and supported by a defense-in-depth approach across identity and access management, infrastructure security, application security, monitoring, vulnerability management, and incident response.

These foundational controls are complemented, where appropriate, by AI application-level safeguards such as authentication and authorization, retrieval constraints, content controls, evaluation and monitoring, audit logging, human review, and policy enforcement.

Where AI systems involve particularly sensitive data, configurations, integrations, or use cases, additional controls are applied proportionately, including enhanced access restrictions, monitoring, confidentiality requirements, and change governance.

Where third-party AI technologies are incorporated, OpenText governs their selection and integration through security and privacy review, contractual requirements, and ongoing vendor risk management.

This integrated approach combines OpenText's established enterprise security foundation with controls appropriate to AI-enabled products and services.

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?

- Not applicable

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
- Other(s) (please specify): OpenText manages post-deployment AI risks through established enterprise security and operational processes, complemented by AI- and agent-specific observability and controls where appropriate. These processes include continuous monitoring, vulnerability management, incident response, structured root-cause analysis, evaluation, and audit logging. For agentic capabilities, agents are configured to operate with explicitly scoped tools and permissions, and tool-use and execution events are logged and available for inspection where supported. Tools and integrations are subject to security review and operate within defined user, role, and application access boundaries. These controls support detection and containment of misuse, anomalous activity, or unexpected behavior. Identified issues are investigated and addressed through formal incident management and remediation processes, with lessons learned incorporated into product, security, evaluation, and governance improvements. As part of OpenText's AI Center of Excellence (CoE), emerging AI threats and AI technologies are constantly monitored across the industry. Relevant ones (such as the likes of Mythos/MDASH/etc), are analyzed, evaluated and are addressed by leveraging new tools/techniques and process, as well as updating relevant policies and procedures.

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

Yes

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

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

Please provide any further comments and implementation documentation.

OpenText deploys agentic AI capabilities through its Aviator portfolio and Aviator Studio. These capabilities can reason, plan, orchestrate, and use tools to perform tasks, while operating within defined governance and security boundaries.

OpenText's approach is designed for bounded enterprise autonomy rather than unrestricted self-directed operation. Planning and execution are scoped by human-provided instructions, defined tools and sub-agents, user or role permissions, and application and security policies. Human oversight is maintained for actions with significant operational, legal, or customer impact.

OpenText supports these capabilities with enterprise security and privacy controls together with AI-specific guardrails, evaluation, observability, auditability, and policy enforcement. Where third-party models or AI services are incorporated, they are selected and governed through due diligence, contractual requirements, and ongoing third-party risk management.

This approach enables OpenText to deliver practical agentic AI that can automate and orchestrate enterprise work while preserving customer control, accountability, security, and human oversight.

Section 3 – Transparency Reporting on Advanced AI Systems

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

Not applicable

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

- Not applicable

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

Not applicable as above.

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.

OpenText publishes a global Privacy Policy and maintains a public Privacy Center describing how personal data is processed across its products and services, including AI-enabled functionality. In addition, OpenText provides product-specific privacy and data-handling disclosures for AI-enabled offerings such as the Aviator portfolio, reflecting their deployment within customer-controlled, enterprise environments.

<https://www.opentext.com/about/privacy>

<https://www.opentext.com/about/opentext-privacy-center>

<https://www.opentext.com/aviator-ai/content-aviator>

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.

No applicable as noted above.

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

Not applicable

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

No

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

No answer provided

Section 4 – Organizational Governance, Incident Management, and Transparency

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

- Executive oversight
- Senior management oversight
- 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).

Governance structure and roles

AI-related risks are overseen through a cross-functional governance framework involving executive leadership, senior management, and dedicated governance forums.

Governance councils and cross-functional committees (involving Legal, Privacy, Security, Product, Engineering, and Compliance) provide advisory oversight, review AI-related risks, assess proposed use cases, and recommend safeguards, mitigations, or escalation where appropriate.

Senior management (including product

and functional leaders) holds decision-making authority for day-to-day AI risk matters, including approval of deployments, remediation actions, and operational controls.

Executive leadership has final accountability for AI risk-related decisions and ensures that AI risk management is aligned with OpenText's enterprise risk management, security, and compliance obligations.

This structure ensures that AI risk decisions are informed by multidisciplinary expertise while remaining accountable within established management hierarchies.

Integration into operational processes

AI risk management is embedded into existing workflows, including:

- product intake and design reviews
- security and privacy assessments (including DPIAs where applicable)
- vendor due-diligence and third-party risk management
- release approval processes under the Product Security Development Lifecycle

Policy review and update cadence

AI-relevant policies, standards, and controls are reviewed on a regular basis as part of broader governance and compliance cycles and are updated on an event-driven basis, including:

- following material system changes or new AI-enabled functionality
- in response to security or privacy incidents or near-misses
- when significant regulatory, legal, or authoritative guidance changes occur
- where post-incident reviews identify the need for enhanced safeguards

This approach ensures that AI risk management remains dynamic, proportionate, and responsive to technical, regulatory, and operational developments, while leveraging OpenText's existing enterprise governance structures.

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).

Relevant OpenText staff are trained on governance policies and risk management practices through a combination of mandatory enterprise-wide training and role-specific training tailored to their involvement in AI-enabled systems.

Training methods include:

- Onboarding and recurring mandatory training covering information security, data protection, code of conduct, and compliance obligations, which apply to all employees and form the baseline for responsible AI use.
- Role-based training and briefings for employees involved in the design, deployment, operation, or oversight of AI-enabled products, including product, engineering, security, privacy, legal, and compliance functions. These sessions address AI-specific risks, governance expectations, privacy-by-design, security-by-design, and human-in-the-loop requirements.
- Targeted sessions and communications led by Privacy, Security, Legal, or Compliance teams to address emerging AI risks, regulatory developments, or material updates to policies and systems.
- Internal guidance and documentation, including policies, standards, playbooks, and recorded materials, made available through internal knowledge platforms.
- Confidential reporting and escalation channels, including established ethics, compliance, and security reporting mechanisms, which allow employees to raise concerns related to safety, ethics, privacy, or policy compliance without fear of retaliation.

Training content and delivery are reviewed and updated periodically and refreshed on an event-driven basis, including following significant regulatory developments, policy updates, or identified risk patterns.

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.

OpenText maintains internal records of reported incidents and the steps taken to address them through established enterprise incident management, security, privacy, and compliance processes.

How incident records are kept:

- Reported incidents are logged in internal incident tracking systems appropriate to the nature of the issue (e.g., security, privacy, operational).
- Records typically include incident details, assessment notes, actions taken, remediation status, and closure outcomes.
- Where relevant, root-cause analysis and post-incident review documentation are prepared to identify underlying causes and control improvements.

Use of incident records for follow-up:

- Incident records are used to track remediation actions, verify resolution, and confirm that appropriate safeguards have been implemented.
- Lessons learned from incidents or near-misses are fed back into governance, risk assessments, technical controls, training, or documentation updates, as appropriate.
- Patterns or recurring issues inform continuous improvement of controls and risk management practices.

Escalation thresholds:

- Escalation thresholds are defined based on severity, impact, and risk, including potential effects on security, privacy, compliance, customers, or service availability.
- Higher-severity incidents are escalated to senior management and executive leadership in line with established enterprise incident response and governance procedures.
- External notification or engagement is assessed and handled in accordance with contractual requirements and applicable legal or regulatory obligations.

This approach ensures that incident handling is consistent, auditable, and proportionate, and that incident records actively support governance oversight and ongoing AI risk management.

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).

OpenText does share research and best practices on managing risk, including AI-related risk, through:

- participation in industry and policy initiatives (e.g., OECD, G7, voluntary codes)
- public thought leadership, white papers, blogs, and conference contributions on AI governance, privacy, security, and responsible AI
- customer-facing guidance and implementation materials on secure and responsible deployment of AI capabilities, including Aviators
- engagement with standards-aligned practices, including ISO and NIST-aligned security and privacy approaches.

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

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

OpenText aligns its AI risk management and governance practices with internationally recognised frameworks, including ISO/IEC information security standards, NIST-aligned security and privacy frameworks, and the OECD AI Principles, applying them in a risk-based and proportionate manner appropriate to its role as an enterprise AI deployer.

Any further comments and implementation documentation.

These frameworks are evaluated / operationalised through OpenText's existing enterprise governance processes, including product security and privacy reviews, vendor due-diligence, incident management, and ongoing policy review. Alignment is reviewed periodically and refreshed in response to material system changes, incidents, or regulatory developments.

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

OpenText informs users when they are interacting with AI-enabled functionality through clear in-product disclosures, feature naming (e.g., Aviator branding), and supporting documentation and FAQs that explain when and how AI is used. These mechanisms are applied selectively and proportionately, depending on the product and use case, and are designed to provide transparency without overstating system autonomy.

Aviator capabilities operate with human-in-the-loop controls and within user-initiated workflows, and are not meant to act autonomously or execute actions on a user's behalf. As such, disclosures focus on

informing users of AI assistance rather than continuous provenance or automated-action indicators.

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.

As of the date of this submission, OpenText does not currently deploy technical content provenance detection, labelling, or watermarking mechanisms. Transparency is instead provided through third party vendor own mechanisms (where appropriate), or otherwise clear user disclosures, feature labeling, and documentation that explain when AI-enabled functionality is used.

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

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

- Other (please specify): Not applicable to deployers.

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?

- 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.

OpenText invests in mitigating socio-economic risks associated with AI through a combination of product design choices, workforce initiatives, and governance programs.

Examples include the design of AI-enabled capabilities (such as the Aviator portfolio) with human-in-the-loop controls, role-based access and safeguards, and clear limitations on autonomous behaviour, ensuring AI augments rather than replaces human judgment. OpenText also invests in AI literacy, workforce training, and internal awareness programs to support responsible adoption, reduce skills gaps, and promote fair and appropriate use of AI across roles and functions.

These efforts are complemented by cross-functional AI governance and review processes that address fairness, bias, access, and appropriate use, and by customer-facing guidance that supports responsible

deployment of AI in enterprise contexts. Collectively, these initiatives focus on improving productivity, access, and inclusion while avoiding socio-economic harms associated with unchecked automation. As an enterprise AI deployer, OpenText does not conduct or control large-scale AI model training and therefore does not undertake AI-specific environmental impact research or measurement, which primarily resides with infrastructure and model providers.

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.

OpenText supports AI literacy, education, and training initiatives for both employees and customers. This includes mandatory baseline training on privacy, security, and ethical use, role-specific AI governance and risk awareness training. We provide AI foundation trainings for all employees and targeted briefings on the capabilities and limitations of AI-enabled systems.

In addition, OpenText provides customer-facing documentation, FAQs, and enablement materials that explain when and how AI is used, the role of human-in-the-loop controls, and appropriate use of AI-enabled features in enterprise environments.

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.

OpenText advances the UN Sustainable Development Goals primarily through enterprise-wide ESG programs, technology enablement for customers, and responsible AI governance. While OpenText does not currently pursue AI projects explicitly framed as SDG-specific initiatives or delivered through partnerships with civil society or community groups, its AI-enabled offerings are designed to support inclusive growth, productivity, and responsible digital transformation consistent with broader sustainability objectives.

