Principles for Responsible AI Use in Trade Compliance

Prepared: September 14, 2026

LEGAL DISCLAIMER

These Principles are provided solely for general informational and educational purposes. They do not constitute, and should not be interpreted as legal, regulatory, technical, cybersecurity, compliance, or other legal advice. Publication or use of these Principles does not create an attorney-client, consultant-client, fiduciary, or other professional relationship between ISECS and any person or organization.

Export-control, sanctions, customs, data-protection, artificial-intelligence, and related requirements vary by jurisdiction and may change over time. Organizations remain solely responsible for identifying and complying with all laws, regulations, license conditions, contractual obligations, and internal requirements applicable to their activities. Organizations should obtain advice from qualified legal counsel and other appropriate professionals before implementing an AI system or relying on AI-assisted compliance determinations.

ISECS does not endorse or certify any AI system, vendor, product, compliance program, or organizational practice by publishing these Principles. No organization should rely exclusively on these Principles—or on the output of any AI system—to make legal or compliance determinations. References to controls or practices do not guarantee compliance, prevent violations, or establish a defense against regulatory or enforcement action.

To the fullest extent permitted by applicable law, ISECS and its officers, directors, employees, members, contributors, and affiliates make no representation or warranty regarding the accuracy, completeness, currency, suitability, or effectiveness of these Principles and disclaim responsibility for actions taken or not taken in reliance upon them, including any resulting loss, penalty, investigation, enforcement action, or other consequence.

RECOMMENDED PRINCIPLES

1. **<u>Principle 1:</u> Human Oversight and Accountability.** AI should support, rather than replace, qualified trade compliance professionals. Organizations remain responsible for decisions and communications based on AI-assisted outputs, and accountability should not be transferred to an automated system.

   Final decisions involving tariff or export control classification, valuation, licensing, sanctions screening, diversion risk, and regulatory filings should remain subject to appropriate human review, validation, and documented approval. Organizations should establish clear lines of responsibility identifying who may review, approve, override, escalate, or release AI-assisted determinations.

2. **<u>Principle 2:</u> Risk-Based Deployment.** AI tools should be deployed in a manner proportionate to the legal, regulatory, and operational risks associated with the relevant activity. Higher-risk applications should be subject to greater scrutiny, more conservative decision thresholds, and more extensive human involvement.

   Organizations should evaluate proposed AI use cases based on their products, services, customers, counterparties, destinations, transaction types, and geographic exposure. Routine and low-risk activities may be suitable for greater automation, while complex classifications, sensitive end-use assessments, high-risk counterparties, novel fact patterns, and transactions involving restricted destinations should receive enhanced review.

3. **<u>Principle 3:</u> Transparency and Explainability.** AI-assisted compliance decisions should be sufficiently transparent, explainable, and documented to permit meaningful internal review and, where necessary, scrutiny by regulators, auditors, customers, or other authorized parties. Compliance personnel should be able to understand and explain the material inputs, sources, assumptions, and processes that contribute to an AI-generated output. Explainability supports both internal governance and external accountability.

4. **<u>Principle 4:</u> Data Quality and Governance.** AI tools used in trade compliance should rely on accurate, authoritative, relevant, and current information. This may include sanctions and restricted-party lists, export control rules, tariff schedules, customs rulings, product specifications, customer information, and transaction data.

   Organizations should establish controls governing data quality, validation, lineage, access, storage, retention, and permitted use. Particular care should be taken when AI systems process export-controlled technical data, personal data, confidential business

information, or information subject to localization or cross-border transfer restrictions. AI risk assessments should also account for changes in sanctions, export controls, product lines, markets, and business operations.

Reliable outcomes depend on valid and reliable AI systems and on risk assessments that reflect changing sanctions and export-control environments.

5. **Principle 5: Continuous Monitoring, Testing, and Incident Response.** AI systems should be tested before deployment and monitored throughout their use. Testing should evaluate accuracy, reliability, false positives, false negatives, model drift, and performance against expert-reviewed historical decisions, particularly for high-risk functions such as tariff classification, export-control classification, restricted-party screening, and red-flag identification. Testing and monitoring should recognize that false positives and false negatives do not present equivalent risks. Although false positives may delay legitimate transactions and create operational burdens, false negatives may permit transactions involving sanctioned or restricted parties, prohibited end uses, or other regulatory violations, creating potentially significant legal and enforcement exposure.

   Enhanced scrutiny should not be limited to transactions flagged by the AI system. Organizations should conduct periodic, risk-based, and documented assurance sampling of transactions automatically cleared by the system to determine whether prohibited parties, red flags, classification errors, or other compliance concerns were missed. The sampling methodology, frequency, sample size, findings, corrective actions, and responsible reviewers should be documented. Sampling should be more frequent for higher-risk products, destinations, counterparties, end uses, or regulatory requirements.

   Organizations should establish feedback mechanisms that allow users to report errors, improve data quality, and support system recalibration. Controls should be reassessed when product lines change, new markets are entered, regulations are amended, assurance sampling identifies potential weaknesses, or audits reveal deficiencies.

   AI-related failures should be incorporated into the organization's compliance incident-response procedures. Organizations should establish thresholds for pausing or disabling an AI system and reverting to manual workflows when a systemic error, data-poisoning event, material model drift, unacceptable false-negative rate, or other significant failure is identified. Procedures should also address whether transactions processed or automatically cleared during the affected period require retrospective review.

6. **Principle 6: Legal and Regulatory Alignment.** AI use should comply with applicable sanctions, export controls, customs laws, privacy requirements, recordkeeping obligations, contractual commitments, and AI-specific laws and regulations.

   Because trade compliance is multi-jurisdictional, AI systems may need to account for overlapping or conflicting legal regimes, including unilateral sanctions, multilateral export control frameworks, blocking statutes, regional geographic privacy laws, and data transfer restrictions. Organizations should also evaluate whether AI models, algorithms, training data, software, semiconductors, technical data, or cloud access may independently trigger export control, sanctions, research security, or technology transfer obligations.

7. **Principle 7: Documentation and Auditability.** Organizations should maintain sufficient documentation to reconstruct how each material AI-assisted compliance determination was reached, reviewed, and approved. Documentation should support the governance and validation requirements established under Principle 4, as well as regulatory examinations, audits, investigations, litigation, and enforcement proceedings.

   For each material determination, organizations should retain the applicable model version, screening-list or control-list version, rule set, system configuration, data sources, prompts or queries, input data, outputs, confidence scores, timestamps, reviewer comments, overrides, approvals, and final determination. Where a screening list or other external data source changes over time, the organization should preserve the relevant version, dated snapshot, or other reliable record sufficient to establish the information against which the transaction was evaluated.

   These records should be retained for the full period required under applicable export-control, sanctions, and other recordkeeping laws, or any longer period required by organizational policy, contractual obligation, litigation hold, or government directive. Organizations should ensure that AI-assisted determinations are traceable and capable of being reconstructed. Because some AI systems are nondeterministic, reconstruction need not require generation of an identical output, but the organization must be able to identify precisely the system, model, lists, rules, inputs, configuration, and human-review process used at the time of the determination and demonstrate the due diligence performed.

8. **Principle 8: Bias and Fairness Mitigation.** Organizations should assess AI systems for unintended bias, unsupported assumptions, and systematically inaccurate outcomes. AI models trained on historical data may reproduce prior classification errors, screening deficiencies, or inconsistent risk judgments.

The determination of ownership and control along complex corporate chains — including aggregation of holdings and indirect control (for example, the "50% rule" and equivalent EU control tests) — is among the most delicate AI applications and the most exposed to systematic error.

Particular attention should be given to restricted-party screening, risk scoring, and anomaly detection, where inaccurate or unsupported patterns may adversely affect customers, counterparties, regions, or business units. Restricted-party screening presents specific technical risks because it frequently relies on fuzzy-matching logic and the transliteration of names from Chinese, Cyrillic, Arabic, and other scripts. Similarity thresholds that are improperly calibrated, incomplete transliterations, spelling variations, aliases, and differences in name order or structure can produce false positives or, more critically, false negatives. Organizations should establish procedures to evaluate and validate matching logic, similarity thresholds, transliteration methods, and other screening configurations, and to identify, investigate, correct, and document material errors, bias, or unfair outcomes.

9. **Principle 9: Security and System Integrity.** AI systems used in trade compliance should be protected from unauthorized access, manipulation, disclosure, and compromise throughout their lifecycle. Security controls should address sensitive products, technology, customer, routing, and transaction data.

   Organizations should assess the physical and cloud infrastructure used to host or operate AI systems. Processing export-controlled technical data, proprietary designs, personal data, or other restricted information through cross-border or multi-tenant environments may create unauthorized transfers or violate data localization requirements.

   Depending on the sensitivity of the information, organizations should consider secure deployment models, including private cloud environments or on-premises hosting. Controls should also address AI-specific vulnerabilities, including prompt injection, data poisoning, unauthorized model manipulation, compromised training data, and malicious inputs.

10. **Principle 10: Controlled Use of Generative AI.** Generative AI should be used cautiously in trade compliance and should not be relied upon as the sole basis for final legal or regulatory determinations. Organizations should clearly define approved and prohibited uses, when human review is mandatory, and which systems may be used for official compliance activities.

Policies should distinguish between public or open AI applications and enterprise-approved, secure environments. Compliance personnel should not use unapproved public AI tools, browser extensions, or personal accounts for official compliance work where entered information may be retained, disclosed, or used to train external models.

11. **Principle 11: External Communications, Reliance, and Legal Risk.** When AI tools are used to prepare classifications, compliance certifications, screening results, or other analyses that may be shared externally, organizations should assess the legal and commercial risks associated with third-party reliance.

    AI-assisted outputs that may influence regulatory compliance, contractual performance, customs declarations, shipping decisions, or other business activities should undergo appropriate review before release. Organizations should define which outputs may be communicated externally, the level of approval required, and the personnel authorized to release them.

    Organizations should also evaluate whether contractual provisions, disclaimers, limitations on reliance, or other legal safeguards appropriately address foreseeable use of the information. Disclosures that an output was AI-assisted may be appropriate, but disclaimers should not be treated as a substitute for accuracy, review, and responsible decision-making. The organization remains responsible for the information it communicates.

12. **Principle 12: Integration with Existing Compliance Programs.** AI tools should be incorporated into established compliance frameworks rather than operated as separate or informal systems. AI-generated red flags involving sanctioned parties, restricted end uses, military or intelligence end users, diversion patterns, controlled technology, embargoed destinations, or inconsistent transaction data should be routed through established escalation and investigation procedures.

    AI may assist in identifying anomalies, but compliance personnel should continue to evaluate unusual transaction circumstances and other indicators of prohibited end users, end uses, or destinations. Automation should strengthen, rather than weaken, know-your-customer, due diligence, escalation, licensing, and anti-diversion controls.

13. **Principle 13: Third-Party Vendor Management.** Organizations using third-party AI systems should conduct risk-based due diligence on the vendor and the relevant tool. The review should consider data use, model-training practices, security, hosting and processing locations, vendor personnel access, update cadence, change management,

explainability, audit rights, sanctions-screening logic, performance representations, and contractual commitments concerning confidentiality, accuracy, regulatory compliance, and liability.

Organizations should specifically evaluate cross-border hosting, data processing, technical support, and system-administration arrangements. For example, personnel of a non-EU vendor may be able to access controlled technical data while providing remote support, maintaining a multi-tenant cloud environment, troubleshooting the system, reviewing prompts or outputs, or accessing logs, backups, or training data.

Depending on the applicable law, the classification of the information, the personnel's location, and the manner in which access occurs, permitting such access may constitute an export or other controlled transfer of technology requiring authorization. Similar obligations may arise under non-EU export-control regimes. Organizations should therefore identify the vendor personnel and subcontractors who may access controlled data, determine where they are located, restrict access according to role and location, and obtain any required authorization before access is permitted.

Reliance on a third-party provider does not transfer regulatory responsibility away from the organization. Contracts should require the vendor to disclose relevant hosting locations, support arrangements, subcontractors, and material changes affecting access to controlled data. Vendor performance and access controls should be monitored throughout the relationship, and material changes to the system, service, personnel-access model, subcontracting arrangements, or hosting environment should be reassessed.

14. **Principle 14: Protection of Legal Privilege and Confidential Information.** Organizations should establish policies and technical safeguards to protect attorney-client privilege, attorney work-product protections, confidential business information, trade secrets, and other legally protected or sensitive information when using AI systems.

    Users should not enter privileged communications, legal advice, internal investigation materials, voluntary self-disclosures, litigation documents, regulatory correspondence, trade secrets, or other confidential information into systems that have not been expressly approved for that purpose or that lack appropriate security, confidentiality, retention, and data-use protections.

    Organizations should assess the risk of privilege waiver, unauthorized disclosure, breach of contractual confidentiality obligations, regulatory violations, and unintended

retention or reuse of information by AI providers. Legal counsel should be consulted before AI is used in matters involving privileged or highly sensitive information.

Approved AI environments should include controls governing access, retention, disclosure, and permitted use. Legal, compliance, and other personnel who handle sensitive information should receive appropriate training on these requirements.

15. **Principle 15: Proportional Benefits and Ongoing Governance.** Organizations should periodically evaluate whether an AI system provides measurable compliance benefits proportionate to its cost, complexity, and risk. Relevant considerations may include accuracy, efficiency, consistency, error reduction, escalation quality, and the system's effect on the overall effectiveness of the compliance program.

    AI governance should continue throughout the system's lifecycle, from procurement and development through deployment, monitoring, modification, and retirement. Organizations should maintain appropriate oversight, documentation, change management, and accountability to ensure that AI use remains aligned with legal requirements, organizational policies, operational needs, and the organization's risk tolerance.

INTERNATIONAL ALIGNMENT

- ISO/IEC 42001
- NIST AI Risk Management Framework
- EU AI Act

CONTRIBUTORS

Kevin Kurland, ISECS Trustee and Regional Ambassador for North America
Wendy Epley, ISECS Secretary and Ambassador for Arizona
Paul Lalonde, ISECS Ambassador for Toronto, Canada
Steve Wilcox, ISECS Chair
Harry Clark, ISECS Deputy Chair
Gianpaolo Porchiazzo, ISECS Board Member, Trustee for EMEA
Vahit Bicak, ISECS Ambassador for Turkiye