



ARTIFICIAL INTELLIGENCE IN TRADE COMPLIANCE

AN EDUCATIONAL GUIDE OF THE TRADE

PRESENTED BY

AAEI American Association
of Exporters & Importers

TABLE OF CONTENTS

OVERVIEW	1
THE RACE FOR EFFICIENCY.....	2
BALANCING EFFICIENCY WITH REGULATORY RISK	4
USE CASES	5
EVALUATING AI OPTIONS AND OPPORTUNITIES	8
OPERATING AN AI SYSTEM	15
INTERNAL CONTROLS	16
THE FUTURE OF AI IN INTERNATIONAL TRADE	18
CLOSING	20

Overview

AI's impact on trade-based regulatory compliance is undeniable. Companies have implemented technology that automates routine tasks associated with filings and declarations, and are increasingly leveraging systems that can make more complex decisions and take independent action. Likewise, governments, especially U.S. federal agencies, are using AI to ensure compliance, catch mistakes, and detect fraud.

Less than four years after the public release of ChatGPT, artificial intelligence (AI) is reshaping the practice of international trade, just as it has reshaped virtually every aspect of business. AI-related automation and machine learning advances are already incorporated into a vast array of technologies, often seamlessly and opaquely. They are leveraged everyday all across the planet – whether the user knows it or not. While some companies may simply use chat bots to perform internet research or write emails, others have already substantially integrated AI into their decision-making and automation processes. Even as the technology evolves at a breakneck pace, cries for guardrails and speed bumps will undoubtedly affect the way import and export companies use AI going forward. At the very least, AI is fundamentally changing the way products are designed, produced, and distributed, the way services are delivered, and the way business operations function. Even those who do not interact directly with it every day are nevertheless impacted by it continually. This is especially true for industries and businesses participating in international trade.

Due to international trade's complexity, expense, and the risks associated with logistics, governmental and regulatory policies, and legal compliance, businesses that import and export products and services must perpetually find ways to achieve greater efficiencies – particularly in regulatory matters – even as they promote innovation, protect quality, and deliver the greater value that customers demand as inflationary pressures remain stubborn. It is clear that importers and exporters must embrace AI or risk competitive disadvantage, but less clear is how to leverage the technology responsibly without incurring penalties due to improper usage or reliance on untrustworthy information. In this new frontier, a roadmap is needed.

This paper explores the integration of AI into import and export operations and practices, particularly regulatory compliance, including:

- Technological advances
- Human supervision
- Compliance implications and risks
- Best practices and uses

Even as AI promises revolutionary advances in machine-centered efficiencies, the uncertainty about the road ahead cannot be ignored. Its impact heretofore cannot be undone, but its future usage has raised concerns in every segment of society, from taxpayers who are balking at the building of data centers to CEOs and top researchers at leading AI tech companies from inside the AI industry, who have sounded loud alarms about the damage unchecked AI decision making could cause. AI implementation must be pursued in balance with humanity's ability to evolve with it in a way that most benefits the safety, economic, and personal interests of real people, and that can only happen with constant, controlled and effective human intervention. Boldly crossing into the new technological frontier, trade companies and governments are both racing to use AI to improve regulatory effectiveness and compliance, and corporations are seeking more profit and scaled international business operations.

The Race For Efficiency

Rapid regulatory change and shrinking revenue margins have driven AI adoption in a wide range of export and import industries, and automation initiatives have quickly moved from pilot to production. The technology helps companies speed up regulatory compliance, clearing products faster so they can make their way to customers more efficiently – increasing profits, reducing time-based opportunity costs, and streamlining supply chains. Early adopters have already blazed trails by partnering closely with AI technology providers to work through a myriad of typical challenges associated with new technology – especially regulatory technology. Today, the path to adopting AI tech is much shorter and swifter, even when a system requires human intervention.

Business

Companies use AI passively and actively. Passive use, the most common function, often employs tools like “bots,” which, like a robot, can perform automated tasks like research or organization faster than a human. Tools like “agents,” which are experiencing growth in usage, can execute more active operations, such as analysis or a series of tasks in a workflow. When combined and integrated with other technologies, the passive and active functions of AI can provide exporters and importers a virtually unlimited menu of operational options. For example, AI can quickly comb through millions of lines of data to extract relevant information, analyze the data without input from a human, and make a filing decision.

Many exporters and importers have used AI to increase efficiency in a number of trade activities, including:

- Classifications in accordance with Harmonized Schedules (HS), a common classification code system for imports and exports
- Free Trade Agreement (FTA) qualifications, which can reduce import fees
- Post entry review and documentation
- Forced labor compliance processes to verify that supply chains do not include forced labor
- Automated freight pricing adjusted to market demands

AI also processes massive amounts of data that even subject matter experts cannot realistically sort through, glean insights which can be used to manage risks. The technology can be used to crunch hundreds of thousands of origin, destination, or commodity sets to identify if a supply chain is at risk, for example, or detect business changes and reprocess an HS, origins, or FTA trade declaration in real time, which eliminates the need for costly amendments. Firms using these tools may find themselves clearing goods faster, spend less per entry, and catch errors before declaration – avoiding detention and penalties.

From logistics providers to clothing companies to global entertainment powerhouses, businesses are using AI to scale trade. They can serve customers in a more timely and transparent manner, without the need for analysts and spreadsheets in the background. Attorneys can expedite their work, and trade compliance professionals can reduce risk, repurpose time to more meaningful tasks, and focus on complex, exceptional problems where judgment matters most.

Governments

The massive resources of the U.S. government are being leveraged to harness the power of AI. Key regulatory agencies are also incorporating it into their operations, though uneven adoption continues to foster some uncertainty and challenges for businesses. While important agencies like the Department of Commerce's Bureau of Industry and Security (BIS) have few, if any, publicly documented uses^[1], others, such as the Directorate of Defense Trade Controls (DDTC)^[2], and the Department of Treasury's Office of Foreign Assets Control (OFAC)^[3], assertively use AI to assess risk at scale.

The U.S. Customs and Border Protection (CBP) agency's import risk detection process is an example of how the federal government is changing the way it operates through AI^[4]. A June 2026 funding law added roughly \$3.5 billion for technology and enforcement to the agency's IT budget, part of which has been used to upgrade its AI-enabled systems. For instance, CBP's Advanced Trade Analytics Platform collects trade data from across government and industry to perform a pattern analysis review, and its Automated Targeting System uses the data to score inbound and outbound cargo with regard to risk. As a practical matter, this data is read by AI-enabled machines before a human at CBP sees it. This allows the review process to leverage more historical information than ever, as the systems comb through years of past data entries for questionable anomalies – values that drift from market prices, classifications or origins that change after a tariff increase, or suppliers that are replaced or appear and then quickly vanish. An audit officer no longer has to rely on a subjective hunch when deciding whether to launch a fraud referral or penalty case. The system now synthesizes the data into lists of ranked leads from which the officer can simply select.

While importers answer chiefly to CBP, the agencies that regulate exports are also ramping up AI use. Exporters answer to the BIS for dual-use items under the Export Administration Regulations (EAR), the Directorate of Defense Trade Controls (DDTC) for defense articles under the International Traffic in Arms Regulations (ITAR), and the Office of Foreign Assets Controls (OFAC) for sanctions and restricted-party issues. Though the regulations and filings are different, the AI-related issues for exporters are very similar to those facing importers.

^[1] At the time of publishing, AAEI contributors and researchers could find no BIS AI use cases. What exists is documentation of the gap: a CSIS study found "BIS analysts generally only have access to outdated versions of Microsoft Excel, a far cry from the more modern knowledge graph and data analytic platforms," with enforcement officers spending 80% of their time finding data and 20% analyzing it. GAO's 2025 export controls report cites BIS asking for "investments in modern data analytics systems" The DOJ and Commerce mention use of real-time intelligence and 21st century data analytics in the 2023 announcement of the Disruptive Technology Strike Force inferring that these agencies are employing the most advanced technologies to enforce export control regulations.

^[2] DDTC launched an AI-Enabled Virtual Agent on December 19, 2025, providing "24/7 access to IT assistance" and searching the DDTC IT Knowledge Base. Some trade press framed this as licensing help; DDTC's own text scopes it strictly to IT support, not ITAR adjudication or commodity jurisdiction.

^[3] TFI-1 "Public Chatbot," a GenAI interface for searching public OFAC FAQs, executive orders, and general licenses. Status: pre-deployment, described in Treasury's own words as "Aspirational, not yet funded."

^[4] <https://www.dhs.gov/ai/use-case-inventory/cbp>

Balancing Efficiency with Regulatory Risk

AI is augmenting productivity for businesses around the globe, and changing the way they interact with governments, especially the U.S. federal government. The technology provides opportunities to reduce expenses, operate faster, and mitigate errors— all of which can enhance a company's ability to comply with governmental regulations. Importers/Exporters (I/Es) and their service providers are using AI to prepare declarations and filings with efficiency, for example. Embracing these technological advances can enable businesses to scale their global operations, seamlessly adjusting to trade requirements which may have previously been prohibitive due to their complexity and the resources needed to comply.

Trade compliance has relied on software for decades, and AI represents the natural evolution of that technological application. First came rules engines, which did exactly what they were told, such as flag an entry when a value crossed a predetermined threshold or when a country appeared on a special watch list. Machine learning followed, in which computers followed patterns inferred from historical data rather than written rules – ranking a shipment's overall risk without someone specifying every risk factor, for example. Next, advances in natural language processing enabled software to read unstructured text, such as an invoice description, a ruling, or a contract clause.

Now, generative models can produce new text, propose a classification and explain its rationale, and hold a conversation with a human about a tariff question. The newest step is agentic — systems that do not merely answer, but act from start to finish with limited human prompting — pulling a document, filing a form, sending a reply, and other mission-critical tasks. Each step forward has widened the breadth and scope of software's capacity for independent decision making and narrowed the technological need for human intervention.

Using AI to satisfy regulatory requirements presents its own set of risks and challenges, however, starting with the potential for errors. AI is only as accurate as the data it uses to learn and make decisions. An error made by an AI tool is ultimately the responsibility of business owners, who have a duty to exercise the reasonable care principle mandating that due diligence is used to ensure accuracy and transparency. That principle may make allowances for human error, but, as government policies evolve, those allowances may not hold for errors made by machine automation. A mistake made by a trained employee may yet demonstrate reasonable care, while the same mistake made by an untrained and unsupervised algorithm might suggest negligence. As regulations change midstream and AI systems receive an influx of potentially conflicting information supplied by importers and exporters, careful human judgment is still paramount.

Software is not a substitute for human decision making, as confirmed by CBP Ruling H350722^[1]. Approaches to using AI in the scope of customs business should also be viewed through a reasonable care lens, both in filing an entry directly and performing work on behalf of another party under a principal or agent relationship, such as a broker acting for an importer or exporter. An unlicensed entity's optical character recognition (OCR) tool that extracts data destined for an entry can independently constitute customs business, for example, because CBP has held that an unlicensed entity cannot decide what data should appear on an entry — whether that data is pulled manually or through an automated tool. By this logic, AI tools that pre-populate entry forms may extract and organize document data, but a licensed customs brokerage must review that output and make the final determination before it is used on an entry. The tool's role is to inform the decision, not to decide. Industry interpretation of H350722 has varied since its January 2026 release, and open questions remain about how this standard applies to newer AI-assisted platforms. AAEI is actively monitoring CBP developments in this area, including the possibility that additional or revoked rulings could further clarify — or shift — these boundaries.

The reasonable care requirement may be a holdover from old laws, but it still applies to new tools. In this respect, AI is no different from other technologies used to support trade compliance or operations. When electronic filing became the norm, for example, the Customs Modernization Act still required importers and exporters to ensure the accuracy and precision of their submissions. The advent of AI has not changed that obligation. The AI tool may do much of the work, but the liability for decision making still rests with a person the law can hold accountable. It helps to view today's AI tools as a continuum, especially as it relates to compliance questions. The further a system evolves from simple mechanical functions towards independent judgement and action, the more deliberately a company must create controls that keep human beings in charge.

^{1]} <https://rulings.cbp.gov/ruling/H350722>

Use Cases

The global buzz around AI can lead trade business leaders to conclusions on opposite ends of the spectrum: either AI is a tool that will solve all of a company's problems or it is an uncontrollable, existential menace. In reality, AI's utility depends on the individual needs of each company and that company's ability to effectively implement the technology. This extends beyond initial procurement and activation and extends to a company's ability to train and supervise the AI system. Value for a company comes through embracing platforms that can show validated results using a company's own products, keep a record of evidence and review, and govern where data goes. Though AI technology providers use marketing tactics to sell systems, the best use cases are anchored in deployments that are publicly documented by major importers and exporters or regulatory agencies.

The sensible places for exporters and importers to begin AI exploration are areas where AI can enhance operations on a scale that the company has the bandwidth and resources to manage. Even starting with basic AI tools can lead to significant efficiency. Prior to purchasing trade-specific software, many companies are leveraging chatbot tools such as ChatGPT or Claude to audit import and export files, look for data anomalies, and summarize the impact of regulatory changes.

Whether a company is already employing AI for a myriad of functions or is only now preparing to make significant investments for the first time, it helps to review key use cases that are driving the industry today. It is important that the company draws the lines between mechanical and judgement functions before the tool is in production rather than after a determination is questioned. Sorting the use cases into tiers is helpful.

- *Automated end to end (mechanical)*. Document intake, entry-against-source reconciliation, and restricted-party screening of known counterparties.
- *AI proposes, a person determines*. HS classification for new SKUs in established families, FTA qualification where the bill of materials is complete, post-entry review findings.
- *Never automated (judgment)*: First-of-kind classification, valuation involving royalties, assists, or related-party pricing, antidumping and countervailing duty scope calls, forced-labor detention responses, and voluntary self-disclosures.

Classification. Modern AI tools propose Harmonized System (HS) classifications and provisions, accompanied by an explanation of the reasoning, including the relevant General Rules of Interpretation, applicable rulings, and tariff notes. Used well, they can turn what was previously an hour of research into minutes and leave a written rationale behind. For exporters, the same tools support ECCN and USML classification, with the controlling entry, the technical parameters that drove it, and the license-exception analysis attached to the determination.

Origin and Trade Agreements. Qualification analysis depends on complete supplier documentation and stays accurate only as long as sourcing holds still. AI helps on both fronts, flagging what's missing and keeping the work current as supply chains change. Document intake. Document intake is mechanical work, and a strong candidate for automation. AI-powered extraction software reads invoices, packing lists, and turns bills of lading into structured data, which is the raw material for everything downstream – including the supply chain traceability files that forced labor enforcement now demands. In some cases, the data is used to populate import and export declarations.

Reviews and Audits. A machine can review every entry against its source documents, catch errors while correction windows are still open, and assemble the record an auditor may eventually request. Additionally, AI has the capability to perform continuous review, highlighting needed corrections in real-time. But an AI audit tool only catches errors it knows to look for. It cannot flag the dubious items an experienced human would notice simply by asking the right questions, such as an undeclared value that doesn't appear correct given the type of commodity, conflicting countries of origin listed across separate documents in the same package, or an undisclosed relationship between a buyer and a seller that changes the analysis. Those findings come from judgment and curiosity, not pattern-matching, and no amount of continuous review replaces a human who knows when to ask the right questions.

Drafting Responses to Government Inquiries. Requests for information, notices of action, audit questionnaires, detention responses, and subpoenas all start with assembling the record — the entry, its source documents, the determination at issue, and the evidence behind it. AI assembles that file, drafts a first response, and checks it against a company's actual filing. Assertions, however, remain a judgment for counsel and licensed personnel. The response becomes part of the record, so nothing should go out unreviewed.

Sanctions Compliance. Screening platforms utilize AI to streamline restricted party lists and mitigate third party risk. These systems apply machine learning and fuzzy logic to account for transliterations, aliases, and complex ownership tiers, dramatically reducing false positives and capturing subtle red flags that manual reviews can miss. This includes screening against the Entity List, the Denied Persons List, and OFAC's SDN list. Know-your-customer screening matters more now than ever, and the content behind the screen determines its value. A tool run against stale list data, or applied across full ownership chains without tuning, will generate false positives as readily as it catches true matches unless it is trained via user request to answer a question or enter information.

License Assignment Validations. AI supports license determination workflows run through BIS's SNAP-R and DDTC's DECCS systems, testing a proposed transaction against the applicable control entry, end use, end user, and destination to identify whether an authorization is required and which exceptions may apply. Some firms are also piloting AI-assisted drafting of technology control plans and voluntary self-disclosures, though that use case carries its own risk given how much those documents depend on precise, defensible language.

Regulatory Change Monitoring and Horizon Scanning. AI is used to monitor dynamic international trade regulations across multiple jurisdictions at the product, part, and determination level. Organizations can automatically track new policies, tariffs, or sanctions updates, receiving product-level alerts and ensuring rapid adjustments to compliance programs.

Supply Chain Visibility. Supply chain visibility particularly benefits from AI use because thousands of disparate data sources, multiple language documents, and unstructured data sets must be read, understood, and mapped. Visibility is the map: who is actually in the supply chain beyond the supplier named on the invoice. The tier one is a matter of record. Everything above it — the mills, smelters, growers, subcontractors, and intermediate manufacturers a product passes through — sits in other parties' records, and no single importer, exporter, or broker holds it. AI illuminates that structure by reasoning across shared, networked data that no one party contributes alone: declarations, shipment records, corporate registries, and facility relationships. The output is not determination. It is a picture of the chain, refreshed as sourcing shifts, and it is often where importers and exporters start to understand where risk may exist. Note: some supply chain visibility tools infer supplier relationships, and others are evidence-based requiring documentary evidence of supplier relationships.

Supplier Data Traceability. Companies are using AI in tracing exercises, especially when they must prove supplier relationships. Traceability includes the evidence and its relationship to each supplier that supports the map. A visible chain becomes a defensible one only when each link is supported by documents tying a specific input to a specific product: purchase orders, invoices, packing lists, mill certificates, bills of lading, affidavits, and supplier declarations – dated and linked at the SKU or part level. This is mechanical work at scale, and AI does it well — soliciting the documents, parsing them into structured records, matching them to the right SKU and period, scoring what is missing or internally inconsistent, and maintaining the file as suppliers and sourcing change. What it produces is a record rather than a conclusion – the chain of custody a regulator asks for when an origin, forced-labor, tax, or emissions position is questioned, and the point at which recordkeeping stops being retained paperwork and becomes evidentiary proof. Note: there is a distinction between a product/SKU level trace and tracing supplier relationships in the chain to identify risk overall. Regulators are primarily interested in the product/SKU level trace.

Supplier Risk Assessment. Supplier risk assessments, created by AI tools, are being generated as part of the tracing output. Supply chain assessment is where regulation meets the traced supply chain map — noting where exposure actually sits across a mapped and documented supply chain. UFLPA reaches entities and inputs several tiers upstream, putting the burden on the importer to prove the negative. Screening has to run against every entity the map surfaces rather than only the counterparty on the invoice, and against ownership as well as identity. For example, OFAC's 50 percent rule treats an entity owned in aggregate 50 percent or more by blocked persons as blocked itself, and BIS's Affiliates Rule extends Entity List, MEU, and certain SDN restrictions to entities 50 percent or more owned by listed parties. Other regimes layer their own tests onto the same chain — the EU Forced Labour Regulation, EUDR, CBAM, and Russia sanctions. AI is well suited to applying all of these continuously across thousands of mapped entities and re-testing when a designation, ownership change, or a sourcing shift lands. It shows where risk lingers and the evidence that supports each position. Whether that exposure is acceptable, and what to do about it, remains a judgment — and the determination stays with a person. Users of these emerging tools will need to decide what output they need, is it an indication of where risk lies overall or is it risk specific to a product or shipment.

When evaluating the three supply chain use cases above, it is important to note that the tools and methodologies currently being used and developed are in their infancy, and no single tool can provide a comprehensive solution today. Currently in use, however, are “tech stacks” that are mapping product-level supply chains, identifying missing evidence, producing supplier risk assessments, and generating responses to detentions.

Evaluating AI Option And Opportunities

Today's AI is superb at mechanical work — extraction, comparison, calculation, and screening at a scale no human team can match. It is weaker at judgment work, such as the weighing of ambiguous facts against legal standards, which valuation, complex classification cases, and disclosure decisions regularly demand. Between the two sits the determination — the classification, origin, or license conclusion — that binds the company and has to be attributable to a person using reasonable care.

Automating mechanical work is usually good business. Automation that results in determinations requires careful validation to ensure the tool is properly trained and outcomes are accurate. It needs to reference the proper legal, policy, and guidance records, use the proper order of review, enrich facts, and provide supporting evidence.

There are several issues importers and exporters should consider as AI technology continues to evolve. While the pace of refinement has been fast and fierce, AI is not perfect. Good judgement, accountability, compassion and common sense remain the province of human beings. As with any sound business decision, importers and exporters must carefully assess both the benefits and risks associated with implementing this promising, but in some ways untested technology.

Though a thorough analysis of AI opportunities requires examination of a range of issues, key areas of common concern include: fabrications, server location and access, security, comprehensive inputs, human oversight, governance, and testing and risk assessment.

Fabrications and Hallucinations

Generative AI can dramatically accelerate research and analysis, but its value depends on how it is built and deployed. One issue to understand is fabrication, commonly called hallucination. When asked something its sources do not directly answer, an open-ended generative tool may assert confidence in an invented answer rather than report it found none – sometimes even citing rulings, cases, or journals that do not exist. A related concern is that the tool may not know where to find all the information an analysis requires, producing output that looks comprehensive and well-supported but is incomplete or wrong.

These risks are real, but they are largely a function of design, not an inherent limit of the technology. Fabrications can be reduced — though not eliminated — if the tool answers only from a defined corpus, cites the source document behind each statement, and returns nothing when no source supports an answer. That makes procurement the key control point. Before purchase, buyers should specifically ask vendors what their AI products do when they cannot find an answer, rather than assuming the output will be transparent and reliable. As with any tool, results should be verified before they are relied upon.

Server Location and Access

U.S. government regulations place guardrails around access to servers in terms of location and user. AI software is run by computer servers, which must be physically located somewhere. Cloud-based computing allows software to run instantaneously, receive inputs, and deliver output or action anywhere in the world. Export and import rules put important parameters around this practice, which can create risk for AI-enabled companies. It is important that those companies know exactly where the servers are located, control who accesses the servers, and analyze the related implications on their regulatory responsibilities.

1. Nationality of person accessing the server.

If a foreign individual accesses technical data residing on a U.S. server – even to perform routine maintenance – their view, download, or use of that data is considered an export, according to BIS and DDTC. An export license may be needed to perform that work.

2. Server location.

Data used to report imports and exports to the U.S. government have to reside on servers physically located in the United States. This applies to Electronic Export Information (EEI) filings, which are used to declare goods leaving the U.S. and destined for a foreign country under certain conditions, including value, intended use, destination, and licensing requirements. They are generally filed through the Automated Commercial Environment (ACE) module, part of the Automated Export System (AES) operated by CBP to process trade reporting. The regulation also applies to the import declarations made by most importers using the ACE Automated Broker Interface (ABI) to provide information to CBP and handle customs and fee requirements.

3. *Decryption.*

Using a commercial AI service may pose risks related to decryption. By definition, these systems have to decrypt prompts in order to answer questions and may not be able to meet security standards. EAR and ITAR both assert that sending or storing unclassified technology or technical data is an export, unless an exception applies. For example, when data is secured from end to end using cryptographic modules compliant with FIPS 140-2 or its successors, the means of decryption is never provided to a third party, and the data is not intentionally stored in a prohibited location, such as an arms-embargoed country or Russia.

The carve-out also only ever covered the act of storing or transmitting — it never authorized giving a foreign person access to the decrypted data, so a foreign-national engineer at the provider who can read the prompt raises a deemed export question of its own. For an exporter, the most important questions for a vendor selling an AI product relate to not only where servers are located, but probe who can read the data, who controls the encryption keys, and what authorization — if any — covers access by foreign persons. These are export licensing and compliance questions, not simply a preference for data storage locations.

4. *Security.* As most AI tools use cloud computing, accessible from anywhere, it is important that AI-enabled companies incorporate security management into their compliance strategy, so that rules governing who can access data and from where the data is accessed are not compromised.

5. *Specifically Designed Articles.* Certain commodities are classified as “specially designed articles” when they are intended to be part of a specifically identified end product, such as one to be used for military purposes. The engineering intent of these goods is critical to their classification in adherence to governing rules, and AI systems may have difficulty discerning that intent.

Comprehensive Inputs

Human intelligence is used to make decisions by synthesizing a lifetime of comprehensive information from many different sources, modulating the right balance between competing interests and priorities. Likewise, AI systems associated with trade are most effective when they are trained using comprehensive inputs based on both regulatory requirements and related business operations. For example, an AI classification tool that isn't trained on agency rulings, court decisions, and the current state of that body of law (including which rulings have been modified or revoked) will only make correct legal determinations through luck. But even a tool with full regulatory coverage will have gaps in decision making if it isn't also fed complete and accurate product information — specifications, materials, function, engineering changes — because classification and valuation decisions turn on the facts of the product as much as on the law. An AI tool may make assumptions when the proper input is missing, confidently producing potentially inaccurate output that adds to a company's risks.

System Types

When evaluating AI systems, import and export companies often choose between purpose-built systems and general-purpose Large Language Models (LLMs), which can include enterprise versions. These systems are based on architectural codes that use a skeletal framework incorporating “model weights” to help its “intelligence” determine which input or accumulated knowledge should most influence the result it delivers. These inputs may be proprietary files or export controls, for example.

Purpose-Built systems are meant to handle a specific workflow. Systems designed for importers and exporters include those created for 1) regulatory/legal use and 2) the specific application of global trade compliance regulations. They are trained using very particular information and make relatively few errors when carrying out the function for which they are built.

LLM models are trained to understand the nuances of human language. They can handle a large variety of tasks, but are more generally trained and may make mistakes when carrying out a highly specialized function. They can be released openly, so their information is published for anyone to download, run on their own hardware, and modify. They can also be closed, reached only through a vendor's interface so the model never leaves the vendor's control. An open model can run entirely inside a company's own environment, so no data ever leaves — attractive when the inputs are confidential bills of materials or sourcing records — but the company then owns everything: security, updates, and the discipline to keep it current. A closed model hands those burdens to the vendor and usually offers stronger capability, at the cost of sending data externally and depending on a provider, who can change or negate the model.

One system is not necessarily categorically safer than another. It is important to know what kind of system is powering the AI tool in use, so that the user knows where the data is going and, just as importantly, who can stop the information flow if necessary for security purposes. In any case, any AI system will still need to undergo standard company technical security and privacy requirements.

The Need for Human Oversight

AI systems sometimes deliver results which users may believe are infallibly accurate, when in fact, AI systems are still evolving and making errors. Although they are built to mirror human interaction, at their core, they are still making apathetic judgements based on the mathematical likelihood of correctness. This is especially true when the AI is called upon not just to perform a routine mechanical task, but rather to make a judgement.

Three important technical considerations deserve review when considering utilizing an AI system for trade activities:

1. *AI decisions are based on probabilities.* AI systems are probabilistic and calculate the most likely correct outcomes based on patterns or odds. The same input used twice, with slight variations in facts, can produce two different answers. Users must take care to verify and measure accuracy, especially when using multiple systems.

2. *Oversight is critical. AI systems make decisions using their training and coding, but human beings are ultimately responsible for the results. Proper regulatory compliance requires that human supervisors oversee the AI system's decision making at every stage. An audit trail should provide a record logging both the AI model's internal reasoning steps (reasoning traces) and evidence of supervision and control, including human inputs, reviews, and overrides. A trail that shows only the machine thinking may, in effect, document the absence of oversight, and the security it offers is not sufficient.*
3. *Confidence can be inflated. AI models routinely produce confidence scores, which show the degree to which the system is sure its results are correct. Those scores, however, are akin to the model grading its own homework. Whatever blind spot produced a wrong answer usually survives into the model's assessment of that answer, so a high confidence score offered by an AI system is no substitute for human review. Well-designed systems will give the user the ability to set their own parameters for confidence scores. This would be customization in legacy systems but now can be easily built into newer systems to align with a user's needs.*

The old “garbage in, garbage out” problem mattered most when companies trained their own models. Few do now, because frontier models like OpenAI’s ChatGPT and Anthropic’s Claude have improved faster than any in-house training set could. The modern version of the problem occurs when a capable model is fed inaccurate or incomplete data, such as partial product information, and produces a polished but inaccurate answer. While a truly capable tool can also enrich what's missing, much as an analyst can turn unstructured data into usable information by pulling a specification from a supplier's website or reconciling a description against a prior entry, an enriched fact needs an accompanying, reliable source and should be reviewed as an assertion, not quietly absorbed into a determination.

AI cannot replace the human experience required to use it responsibly. Overreliance on AI falls hardest on companies which have not yet built the capacity to notice when an answer provided by AI is wrong. A firm that replaces all of its junior-level work with AI will eventually discover it has stopped producing senior-level expertise, because a large part of that expertise is the result of judgment developed through the hands-on, junior-level work. That judgment is built, not installed. It comes from education and experience, such as years spent personally observing government declarations go right and wrong, not merely from seeing the process unfold through the eyes of a tool.

Governance

AI usage carries numerous risks, so the decision to employ AI systems should balance that risk against the company's priorities. The company has to ensure that its productivity expectations are based on the company's realistic capacity, including its employees' expertise in supervising and managing the system. A company's corporate governance protocols are a major defense against these risks, establishing internal protocols for how AI is supervised, evaluated, implemented, adopted, and used. The protocols can help the organization answer important questions associated not only with risk, but also with return on investment. These questions include:

- Does the organization have sufficient clean data to make the tool effective?
- Was adoption of the system unduly influenced by the vendor's sales practices?
- Has the tool actually been trained to understand relevant regulations, how they're applied in practice, and the reasoning behind them?
- How will the AI system function if the models underlying its programming change?
- Can the system's output withstand the hard questions a regulator may ask?

These answers to these questions won't stay the same over time, so it's good practice to reassess governance protocols periodically as the technology and the business environment change.

Testing and Risk Assessment

AI system usage requires ongoing testing to mitigate risk, related both to security and performance.

Security

Important security tests, such as a penetration test, can assess whether or not an AI tool will not leak information.

Evaluation by Subject Matter Experts

Beyond security, substantial risk remains if the tool does not deliver the correct output. Before workflows are established which incorporate AI, subject matter experts should test the tool, using several hundred of the organization's own products to score results, evaluate accuracy and consistency, and keep records.

Pilot Testing

Pilot testing is often a recommended part of the RFP/RFI process. After soliciting proposals based on the basic technical perimeters the AI must meet as outlined in the RFP/RFI, a company may then participate in demonstrations of the prospective AI tools. The selection should only occur, however, after a company performs a pilot of several AI tools. During the pilot, the company runs its own data through the AI tool and then analyzes results, allowing potential stakeholders to get a true feel of the GUI ("Graphical User Interface") and whether or not the AI functions as advertised by the vendor. A pilot also allows a company to more accurately compare several AI vendors against each other in a real-world environment, verifying the accuracy of their sales pitches.

Recognizing Data Insufficiency

A production-grade trade AI system should recognize when the available facts are insufficient to support a defensible determination. Organizations should therefore analyze three factors: 1) how strongly the system favors an answer, 2) whether the necessary business facts are present when an answer is generated or a decision is made, and 3) the legal or financial consequences that would follow if the answer were wrong. The model's confidence in its own answer is neither a substitute for data sufficiency nor a company's assessment of compliance risk.

Continuous Monitoring

Continuous monitoring deserves as much attention as the initial validation. Vendors frequently publish model updates, so tools validated in March may behave differently in September, for example. Throughout the testing process, it is important to establish a company's risk tolerance in advance – written down in plain, actionable terms with a clear understanding of decisions and determinations that should not be automated.

National Security Implications

Adopters should also weigh a newer kind of risk: the models themselves are now instruments of national security policy.

On June 12, 2026, three days after Anthropic released its two newest AI models, the U.S. Department of Commerce ordered the company to stop making them accessible to any foreign national – outside or inside the United States, including the company's own foreign national employees. This directive came after researchers reported discovery of a method to bypass the model's guardrails and reach restricted cyber capabilities. Because the platform could not verify the nationality of users in real time, the practical effect was a worldwide shutdown of those models.

The Commerce Department lifted the controls on June 30, and access resumed on July 1 – nearly three weeks later – though the more capable of the two models was authorized to be used only for a set of approved U.S. organizations that defend critical infrastructure. It was the first time export controls were enforced against access to an AI model rather than the hardware on which it was housed and, given the dual-use character of frontier models, it is unlikely to be the last. This raises continuity questions. What is the impact on a company's operations if a regulator switches it off for three weeks and then turns it back on with conditions attached? For export-controlled work, this makes tool selection a licensing decision rather than a productivity choice, and it should be made by someone with intricate knowledge of the technical data the company holds.

Recordkeeping and Review

While AI streamlines compliance workflows, it does not relieve companies of their legal obligations. Importers and exporters remain strictly liable for maintaining complete, accurate records, regardless of the automation tools used. For example, CBP's own AI systems will review filings, so any tool an importer adopts must leave a trail showing the company acted responsibly even when decisions are made by its AI systems. These decisions have to be reconstructable as of the entry date, include the version of the model in use, and stored for a full five-year recordkeeping period. The EAR and ITAR carry their own multi-year retention requirements, so the same requirement to reconstruct any decision on demand applies to export classifications and license determinations made through AI, and to the prompts, outputs, and human review behind them.

Context is important. For example, a portfolio heavy in partner-agency imports or antidumping exposure calls for tighter human review than a catalog of lower-risk goods. Conversely, regulators should endeavor to strike the right balance between leveraging AI for detection while reserving matters of judgment for its public servant workforce.

Operating An AI System

Once the decision is made to move forward with AI tools to support trade compliance, importers and exporters should create and adhere to defined rules by which to operate. Governance is the foundation that any new technology must fit into, not the other way around. Evaluation is less a single test than an ongoing discipline. AI is already in the trade compliance workflow. The practical questions now are which tools are in use, what work they handle, and what controls keep them honest? In regular work, AI arrives through two kinds of tools and a growing set of connections between them and everything else. Chatbots (think ChatGPT or Claude) answer questions and draft text; agents act rather than merely respond, carrying a task through on their own – pulling the document, generating the statement, sending the reply.

In trade compliance, the more consequential AI integrations connect regulatory information with product, supplier, engineering, and transaction data. For example, a change in a product's material composition may call an existing classification or origin determination into question. An integrated system can flag the affected determination, request the missing facts, prepare a revised analysis, route material changes for expert approval, and preserve the resulting decision record. The value lies not merely in producing an answer, but in maintaining the connection between the answer, the facts supporting it, and the business events that may cause it to change.

The work administered through AI systems is the familiar practice of compliance: classification, documentation, recordkeeping, traceability, vendor due diligence, government filings, licensing, and license/permit management included — and, on the export side, classification and licensing under the EAR and ITAR. The emerging pattern is that mechanical work is converted to AI operations first, determinations are delegated to AI systems on exception, and judgment is never ceded to the technology. This makes exception management and judgment central to the human job description.

An exception should not be defined by model confidence alone. Relevant triggers may include missing or contradictory product facts, competing conclusions with materially different consequences, a change in governing authority, a change in product or sourcing data, disagreement with an approved system-of-record determination, substantial financial exposure, or a pattern of repeated human overrides. Each organization should define and document the conditions that require review and who has authority to resolve them. The government has machines of its own; OCR lets government agencies read documents that once consumed countless auditor hours, so audits can be expected more often and to proceed at a faster pace.

AI may change how compliance work is performed, but it does not change who is legally responsible. An importer or exporter still has to uphold requirements related to reasonable care and reason to believe standards, and a broker is still responsible for supervision and control. Protocols have to change to accommodate new processes. For example, someone has to be designated to review filings before they go to the broker, before they are submitted for governance, and again before transmission to government customs – with decision records kept in a format that would survive an audit. Guardrails built into the models are helpful but vary considerably from one provider to another and are subject to misuse. A company's compliance program should treat an AI model's guardrails as the vendor's controls, not a substitute for the company's human oversight.

Ongoing monitoring should test the entire operation, not merely the model. Organizations should set their own thresholds based on risk and should retain enough information to explain both automated recommendations and human decisions. Useful testing measures to incorporate include testing the share of records with sufficient input data, the frequency and cause of human escalation, the rate and reason for overrides, the completeness of supporting authority and evidence, the stability of results over time, and the number of determinations reopened after a business or regulatory change.

Data controls are a final critical piece of AI operations. Employees with open access to commercial models can leak confidential information without meaning to, so control systems should monitor traffic between the company and consumer AI services. On the export side the exposure is sharper than inadvertent disclosure. For example, pasting ITAR technical data or EAR-controlled technology into a consumer AI service can itself be an unauthorized export, whether as a deemed export to a foreign-national engineer at the AI vendor or as a transmission routed through a server outside the United States. Vendors warrant substantive vetting, including inquiries into software vendors locations and data security adequacy. After all, the cybersecurity chain is only as strong as its weakest link. Furthermore, confidentiality should be explicit and understood by every vendor and employee who touches an importer's or exporter's data.

An AI operation run this way is nothing exotic. It is ordinary compliance carried out with newer instruments, and organizations that treat it as such will be better positioned than their slow-to-adopt peers.

Operational resilience and recordkeeping also require portability. A company should retain control of its product facts, approved determinations, supporting evidence, review history, and applicable effective dates independently of any particular model or vendor. If a model changes, a provider is replaced, or access is interrupted, the company should be able to reconstruct its prior decisions and continue operating without losing the compliance knowledge accumulated through the system. When the AI system is updated or replaced, records of the decision-making rationale from the prior system must be retained to preserve auditability.

Internal Controls

Introducing AI into trade compliance does not change the underlying control obligation — it simply changes where controls must be applied. The same components that govern any internal control structure apply with equal force.

Control Environment. Ownership must be assigned before adoption — who owns the tool, the compliance review, the security review, and staff adoption. Duties without named owners lapse quietly, and monitoring is usually the first to go. For ITAR-registered exporters, the Empowered Official who signs export authorizations should appear by name in that ownership map, since that person carries personal responsibility for what an AI-assisted classification or license application ultimately says.

Risk Assessment. The primary risk is a confident wrong answer. Staff needs training not just on what these tools can do, but on how they fail — recognizing fabricated or unsupported outputs, prohibiting the tool from filling gaps with guesswork, and knowing what information should never be uploaded.

Control Activities: Control activities often involve access restrictions, but they should also verify that the tool performs as intended. This means testing before deployment, comparing the outputs of the AI system against results known to be correct. It also includes spot-checking a sample of the tool's work on an ongoing basis, and requiring a second-level review for any output that feeds directly into a filing or disclosure. Data controls should also follow the same discipline as prudent personnel access — an AI tool sees only what its job requires. Narrow access protects confidential information and keeps the model's working material clean, reducing the risk of degraded output.

Information and Communication. Every determination must be traceable to a person who approved it. An accountable person must be documented along with recordkeeping of inputs, outputs, and the rationale behind each decision. It allows a transaction to be reconstructed after the fact – something regulators and auditors will expect to see when a decision is questioned. On the export side, that reconstruction standard reaches the prompts and outputs behind an ECCN or license determination, not only the authorization that was filed.

Monitoring: Monitoring should combine periodic audits with active human oversight of the tool's day-to-day use. Audits confirm the controls in practice by testing a sample of outputs against expected results and review overrides — along with the reasoning behind them. Between audits, a human must remain in the loop on consequential outputs, so that drift or failure is caught in real time rather than at the next scheduled review.

Recordkeeping: The use of software tools, including AI-enabled systems, does not alter, transfer, or diminish an importer's or exporter's responsibility for regulatory compliance related to recordkeeping, data integrity, or the ability to produce supporting documentation upon request. All records generated or reliant upon AI must remain complete, retrievable, compliant with applicable retention requirements, and available for government examination. Importers are responsible for maintaining and producing records in accordance with the recordkeeping requirements of 19 CFR Part 163 and related CBP regulations.^[1] Similarly, exporters remain responsible for compliance with the EAR and FTR recording keeping requirements set forth in 15 CFR Part 762^[2] and section 30.10,^[3] respectively, as well as reporting and recordkeeping requirements for any other government agency that regulates trade transactions.

Compliance with Standards. BIS has its Export Compliance Program (ECP) standards that all exporters must implement. Two of the key eight elements relate to internal auditing and risk management.^[4] When an exporter utilizes AI it must consider these two elements and ensure that it is considered in the exporters overall ECP.

These internal control obligations are not new obligations imposed by AI; they are the same controls that have always governed compliance operations, now extended to a new class of tool. Organizations that fail to formalize them will find that the absence of structure, not the technology itself, is what ultimately creates exposure.

The Future Of AI In International Trade

AI is a new, disruptive technology already remaking the workflows, professional requirements, and expectations of users in trade industries in much the same way as it has impacted other fields. As in the legal profession, AI drafts and reviews documents for trade companies. In logistics, it routes shipments and forecasts demand. As AI becomes more prevalent in trade, companies in the industry will need to navigate both its benefits and its pitfalls.

As its ubiquity grows, governments will also increasingly integrate AI into their processes, enhancing efficiency but also potentially adding bureaucratic hurdles faster than smaller businesses can adapt. The practical implication for exporters and importers is that the regulatory agencies will catch AI-assisted mistakes before the firms that made them, creating compliance gaps.

The next stage of trade technology will move beyond transaction-by-transaction automation toward persistent compliance. Classifications, origin analyses, licensing conclusions, and other determinations will increasingly function as living digital records connected to the facts and authorities that support them and capable of being instantaneously reconsidered when regulations, products, suppliers, sourcing, or transactions change. In that environment, the compliance function will not only process a larger volume of work more quickly. It will also manage a continuously changing portfolio of determinations and exceptions. Although AI systems are capable of making many determinations independently, the most valuable systems will be those that identify what has changed, determine which prior determinations may be affected, assemble the relevant evidence, and direct consequential questions to the appropriate professional.

As the regulator, the U.S. government is engaging with the same challenges and opportunities presented by AI technology as companies are encountering – both as a referee and as a user. For example, the CBP customs business ruling discussed earlier suggests the government is increasingly tolerant of AI involvement in customs compliance, within limits. Software will steadily absorb mechanical work – drafting, extraction, and first-pass analysis – while the

^[1] <https://www.ecfr.gov/current/title-19/chapter-I/part-163>

^[2] <https://www.ecfr.gov/current/title-15/subtitle-B/chapter-VII/subchapter-C/part-762>

^[3] <https://www.ecfr.gov/current/title-15/subtitle-B/chapter-I/part-30/subpart-A/section-30.10>

^[4] <https://www.bis.gov/developing-an-export-compliance-program>

determinations that bind an importer will continue to be the province of licensed and trained personnel. More rulings in this vein should be expected as vendors test existing boundaries and as the tools grow capacity in ways the rules could not anticipate. Regulators are exploring the use of AI to detect forced labor, validate countries of origin, review classification and valuation, and improve audit efficiency. This trend is likely to continue into the next decade, as the current federal administration directs substantial investment toward those emerging innovations.

Export agencies are moving in the same direction. BIS and DDTC are building AI-assisted technology to screen lists and are testing automated classification review. OFAC's sanctions screening already leans heavily on the same list-matching technology. Though AI-related rulings resemble customs business regulations that have not yet been issued by these agencies, the need for human intervention in the review loop will undoubtedly remain. For companies who are not abreast, the most common surprise is the speed with which the government is embracing the new technologies. An importer or exporter that has not recognized the shift underway will experience customs agencies' growing capability as a sudden shock rather than witnessing the steady build as it grows over time.

Although AI systems will continue to grow in ability, their greater technical capability should not, by itself, justify greater autonomy. The appropriate level of autonomy will differ by use case and risk. The governing principle is that control should scale with consequence. As rote processes are automated, professional value will shift toward defining policy, resolving exceptions, evaluating evidence, and exercising accountable judgment in the most complex matters. Governing the systems that propose determinations will be paramount – validating them against correct answers, reading monitoring data for drift, and being able to explain a determination at least a year after it was made.

Expect AI governance to appear in trade compliance job descriptions as a specific function rather than as a duty absorbed between other activities. Organizations will also need deliberate ways to develop junior professionals. If automation removes the work through which practitioners traditionally learned, training must increasingly come from supervised review of AI recommendations, rationales, exceptions, and overrides. Instead of treating expertise as a resource the technology can simply replace, the best positioned firms will combine automation with professional development.

AI technology will be an asset to firms that balance efficiency with accountability and a liability to those that ignore its potential and risk. The tools in use today are early versions and the improvements that matter most will be shaped by people with deep domain expertise. The chips and the software are the delivery system, but the trade knowledge is what makes them useful. Increasing AI capabilities should be accompanied by increasing transparency about performance, and claims made by the systems must be evaluated against reproducible tasks using known outcomes, with the scope of the test, the quality of the input data, the handling of exceptions, and the role of human review clearly stated. The industry will benefit from standards that reward demonstrated accuracy and defensibility rather than those that only demonstrate speed or fluency.

Closing

The AI frontier is vast and expanding, but as with all unexplored territories, there is risk ahead. The technology presents revolutionary opportunities for import and export companies to improve efficiency, enhance regulatory compliance, and scale operations. At the same time, U.S. agencies and other governments are leveraging their enormous procurement power to develop their own AI systems and automate processes which previously relied on the expertise and judgement of human regulators.

An importer or exporter that rejects automation/AI and remains fully manual, has chosen, whether they know it or not, to fall behind both their competitors and their regulators. At the same time, an importer or exporter that fails to implement proper AI selection and governance has converted a productivity tool into a potential penalty generator. The sensible path runs between them, and it asks something of each side. Importers and exporters should adopt deliberately - relying on honest testing, meaningful reviews, and verifiably accurate determinations.

Regulatory compliance will be affected as much by AI as will commercial applications, and it is critical that trade businesses and the government act as good faith partners. As trade businesses navigate the complexity of the emerging AI era, it will be important for the government to draw clear lines in advance so firms know what responsible adoption looks like in advance instead of after an audit. At the same time, the government should exercise the same governance controls over its own AI usage that it expects of those it regulates.

Neither wholesale automation nor reflexive restriction serves the trade. The firms and agencies that will do best over the next five years are the ones that are prepared to show their work and keep a human firmly in the loop. Even as AI systems become inconceivably powerful, they are still best viewed as tools which humans can use to achieve organizational priorities, not as accountable decision makers in themselves. Careful evaluation, monitoring, and supervision remain crucial as AI tools remain a means to an end, not the drivers of the business.

Evaluating AI for Trade Compliance

A companion to the AAEI white paper on artificial intelligence in international trade, for compliance professionals evaluating AI tools.

Main theme. Companies not using these tools are falling behind their regulators, which already use them. Companies using them without controls increase their penalty exposure. An AI tool is only as reliable as the data it reasons over and enriches and the humans who supervise them.

The threshold question for any task is, "does automation support mechanical, determination or judgment work?" Automating mechanical work is usually sound. Automating determination requires validation. Automating judgment should not happen.

QUESTIONS TO ANSWER BEFORE PURCHASING AN AI TOOL

ASK THE VENDOR

Which regulations, lists, rulings, and revocations does it reason over, and how current is each?

When a required fact is missing, does the tool flag it or supply it? If it supplies it, is the source recorded?

Does the audit trail record human review and overrides, or only the model's reasoning?

Does the same input produce the same answer? Which outputs are deterministic and which are generative?

How is a confidence score calculated, and has it been validated against known outcomes?

Does the platform allow collaboration internally and externally to enrich/validate facts?

When a designation, an ownership change, or a regulation changes, which prior determinations are reopened automatically?

For supply chain tools: how does it map the chain, document it, and/or test it against regulatory requirements?

Can we export our determinations, evidence, and review history if we leave?

Which model is underneath the tool, who hosts it, and what happens if a regulator withdraws or conditions access?

Is the model open or closed? Where does our data go, and who can disable it?

Who can decrypt our data, and which of your personnel can read our inputs? Server location is not the controlling test; end-to-end encryption with keys we hold is.

ASK YOUR OWN ORGANIZATION

Our business data complete enough to support the tool?

Who is the named owner of AI governance: tool inventory, validation records, monitoring thresholds, override logs?

Which determinations will we not automate? Record them before the pilot.

For each determination the tool touches, where does AI propose or decide?

Who is the named approver for each? For ITAR registrants, is the Empowered Official included?

What is our risk tolerance, stated in writing?

Where are trained staff performing mechanical work today?

Does our portfolio — partner-agency goods, antidumping exposure — require tighter review?

Can we reconstruct any determination as of its date, including the model version used?

What is our contingency if the tool is unavailable for three weeks?

PILOT TESTING

RFP or RFI sets technical parameters and a demonstration shows the interface. Neither establishes accuracy.

Test on company data: several hundred products with known correct answers, scored and documented.

Apply the same test to every vendor, and require a stated scope, input data quality, exception handling, and role of human review.

WARNING SIGNS

The tool was not trained in the proper regulations, guidance, or policies applicable to the intended use.

Determinations that are not reopened when the underlying rule changes.

No means of exporting determinations and evidence.

Confidence scores offered in place of review.

Server location offered as the answer to an export-control question.

Claims supported only by unnamed references.

Vendor guardrails presented as the customer's compliance controls.

A confident determination built on incomplete facts, or a supplied fact with no recorded source.

An audit trail that records only model reasoning, without human review or overrides.

WHAT DOES NOT CHANGE

Reasonable care, reason to believe, and responsible supervision and control. A tool can propose; a person must determine.

Recordkeeping: five years for customs, multi-year under the EAR and ITAR, reconstructable as of the determination date.

Controlled technical data entered into a consumer AI service can constitute an unauthorized export, including to a foreign-national employee of the provider.

Liability cannot be transferred to software. A tool cannot sign a filing or be deposed.

WHERE TO START

Target existing manual effort. Work where trained staff perform mechanical tasks: rekeying, document comparison, repeating a determination after a minor change

Automate mechanical work: document intake, entry-to-source reconciliation, restricted-party screening, regulatory change monitoring, post-entry review.

Identify the determinations. List each determination the tool will touch and decide whether AI proposes or decides. Any output supporting a filing requires a named approver.

Do not automate judgment: ambiguous classification, valuation, first-of-kind origin, disclosures.

Assign a governor. A new role responsible for the curation and understanding of the AI and the compliance results.

IN PRODUCTION

Match the mechanical work to the tool and determinations to a person on exception. Exception handling becomes the primary human task.

Track three measures separately: model confidence, data sufficiency, consequence of error.

Revalidate after every model update, and review the full evaluation at least annually.

Document the basis for each determination, and keep human review on consequential outputs.

About AAEI White Paper Group

The AAEI White Paper Working Group brings together trade compliance professionals and subject-matter experts to identify emerging issues, develop informed perspectives, and produce actionable white papers for the international trade community. The group draws on the expertise of AAEI members and industry leaders to provide clear analysis and recommendations that support members in navigating complex trade and customs challenges.

About AAEI

The American Association of Exporters and Importers (AAEI) is a leading trade association representing companies and professionals engaged in international trade, including importers, exporters, manufacturers, distributors, retailers, and service providers. Since 1921, AAEI has advocated for effective trade policy and customs practices while providing its members with education, expertise, and a forum to address the evolving challenges of global commerce.