KYG Trade™
Know Your Goods

Trust, But Verify: Guardrails for AI in Trade Compliance

February 6, 2026

Presented by **Todd R. Smith** | Founder CEO, Licensed Customs Broker

Today’s Topics

- The Application of AI in Trade Compliance Today
- Practical Guardrails
- Enterprise AI Adoption Checklist

Feeling liberated?

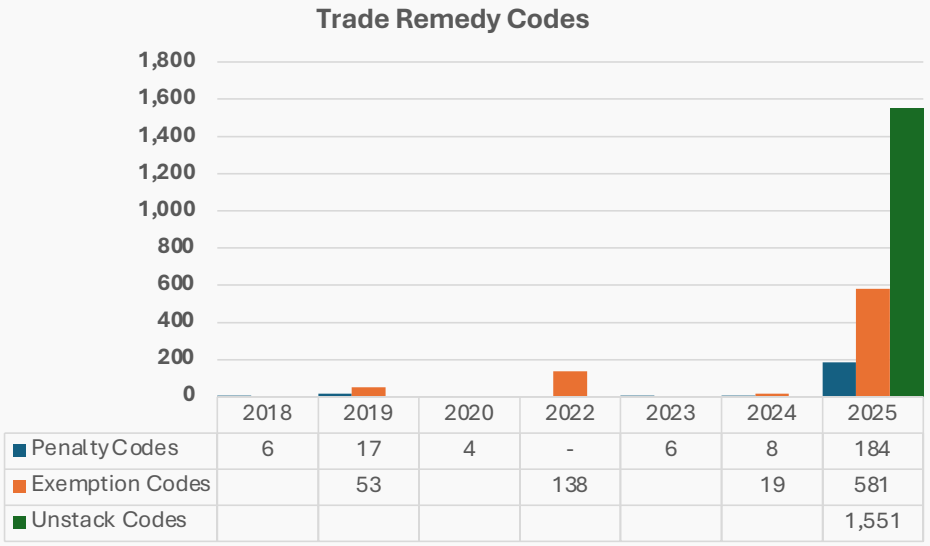

	2018	2019	2020	2022	2023	2024	2025
Penalty Codes	6	17	4	-	6	8	184
Exemption Codes		53		138		19	581
Unstack Codes							1,551

CBP Trade Stats (Fiscal year)

Trade Enforcement Activities	FY 2026[1]	FY 2025[2]	FY 2024
Number of Audits Completed	34	465	417
Total Collected as a result of Importer Audits	$17.67 million	$235.46 million	$117.67 million
Net Revenue Recovered due to Entry Summary Reviews (ESF)	$1.81 billion	$34.41 billion	$667.55 million
Total Trade Penalties Issued	423	2,432	2,204
Total Trade Liquidated Damages	15,655	53,052	22,399
Total Collected from Trade Penalties and Liquidated Damages	$5.58 million	$46.04 million	$26.21 million

Imports and Revenue Collection	FY 2026[1]	FY 2025[2]	FY 2024
Total Import Value for Goods	$569.25 billion	$3.61 trillion	$3.36 trillion
Total Entry Summaries	15.57 million	50.08 million	38.3 million
Total Informal Entry Summaries	9.78 million	15.4 million	5.16 million
Total Duty, Taxes, and Fees Collected*	$67.27 billion	$216.7 billion	$88.07 billion

I don’t have a budget for this. Really?

Fiscal Year 2025

Rank	Company	Penalty Amount	Key Violation
1	Ceratizit USA	**$54.4M**	Tariff evasion, misclassification, false COO
2	Sigma Corporation	**$26M**	False customs declarations to avoid tariffs
3	Allied Stone	**$12.4M**	Customs misclassification (quartz imports)
4 (tie)	Evolutions Flooring	**$8.1M**	False COO to evade AD/CVD duties
4 (tie)	Evolution Flooring	**$8.1M**	Customs misclassification (flooring imports)
6	Grosfillex	**$4.9M**	Customs-evasion scheme (extruded aluminum)

The Application of AI in Trade Compliance Today

The Era of Trade Intelligence: Capability Confirmed

Kay scores 93% on the CBLE …in 5 minutes.

So…how do we leverage AI efficiently and responsibly?

What are the new best practices?

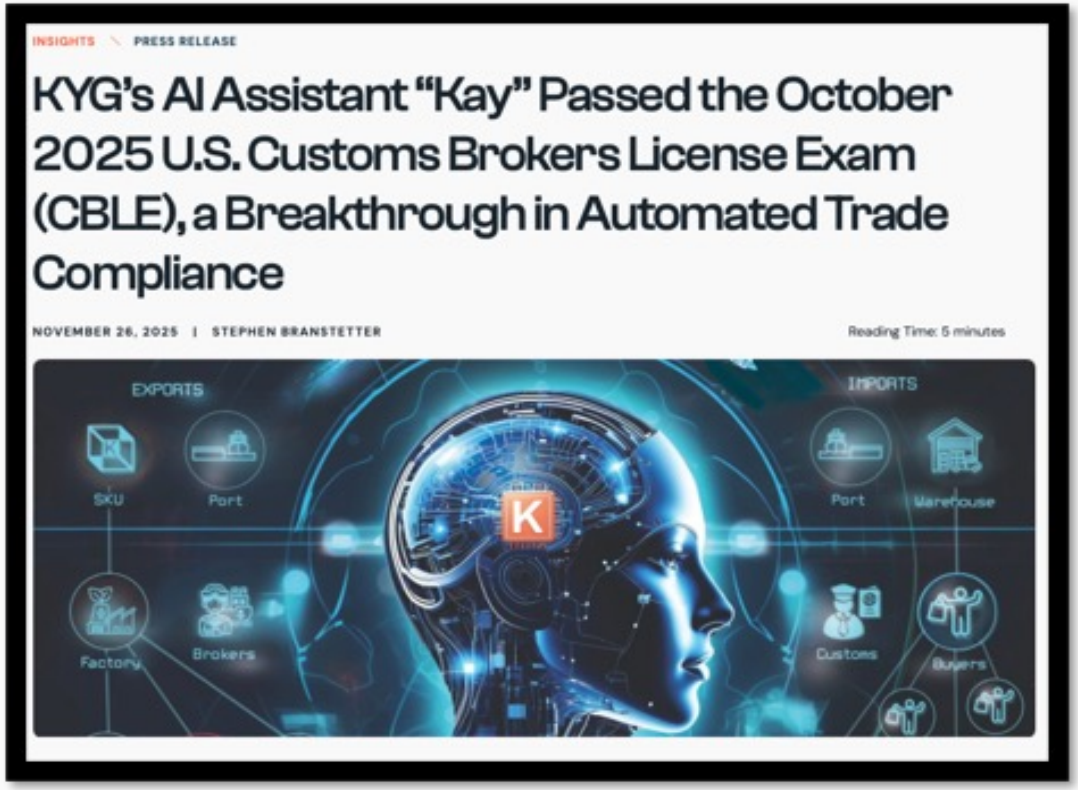

Reasonable Care: Leveraging AI is Now Expected

Consider the AI implications for malpractice in the medical field

Physician decision-making under the current medical liability legal framework, known as the **standard of care**.

Source: Johns Hopkins Carey Business research paper, “Artificial Intelligence on Call: The Physician’s Decision of Whether to Use AI in Clinical Practice.”

“Under the current medical liability system, **if physicians** take the step to consult an AI tool in a case of high uncertainty, then **decide to deviate from the AI recommendation, they are exposed to legal risks** in the event that something goes wrong with the patient”

If you're not using AI, are you exercising reasonable care?

Examples of AI in Trade Compliance Today

VAO, Customs Compliance AI

QuickCode, HS Classification

Caspian, Duty Drawback

Pax, Duty Drawback

Deep Cognition PaperEntry AI

- Current applications of AI in trade compliance today
  - HS Classification
  - Duty Drawback
  - Data extraction
- Increasing number of applications
  - Diverse technologies: beyond OCR
- Challenges
  - Potential for less integration, less efficiency
    - Need for solutions that are open and extensible
  - Potential for risk exposure
  - How to exercise reasonable care?
    - **By establishing guardrails.**

Enterprise AI Adoption Considerations

Discussion Points for Cross-Functional Teams

The Core Problem: AI Semantics Fail in Regulated Industries

- LLMs produce plausible but imprecise answers that may violate statutory boundaries.
- Vector search retrieves similar-sounding content, not necessarily the correct regulation or time period.
- Fragmented enterprise data leads to inconsistent, incomplete compliance answers.
- Lack of traceability creates audit and regulatory exposure.

Why This Matters for Global Trade Compliance

- Trade laws require literal, not semantic, interpretation.
- Regulatory decisions must be defensible and source-grounded.
- Standard RAG pipelines cannot enforce business rules or date/authority filters.
- Fragmented systems increase the risk of divergent compliance outcomes.

What a Governance Framework Must Contain

- Risk-based oversight aligned to established AI governance practices.
- Deterministic filters and rules applied before semantic ranking.
- Full traceability and audit logs for all retrieval events and sources.
- Model governance including approvals, versioning, and human-in-the-loop review.
- Strong security and data-handling controls throughout the lifecycle.

Practical Guidance for Trade Compliance Teams

- Require source-grounded answers only; avoid free-form LLM output.
- Use governed hybrid retrieval with deterministic rules to prevent near-match errors.
- Establish a model governance committee to review AI use and risks.
- Build evaluation scenarios measuring groundedness, accuracy, and leakage.
- Embed compliance subject-matter experts in the AI lifecycle.

Practical Guardrails

Guardrail: Updated Regulations & Resources

Ensure that classifiers *and* AI Tools have access to the necessary resources.

- CBP Title 19 of the Code of Federal Regulations
- Harmonized Tariff Schedule (HTS)
- Customs Bulletin and Decisions
- CBP website
- Customs Rulings Online Search Service
- Explanatory Notes
- Customs Valuation Encyclopedia
- Informed compliance publications
- Court cases
- …or other research service to permit you to establish reliable procedures and facilitate compliance with customs laws and regulations

PLUS

- Your company's policies and procedures related to classification, etc.

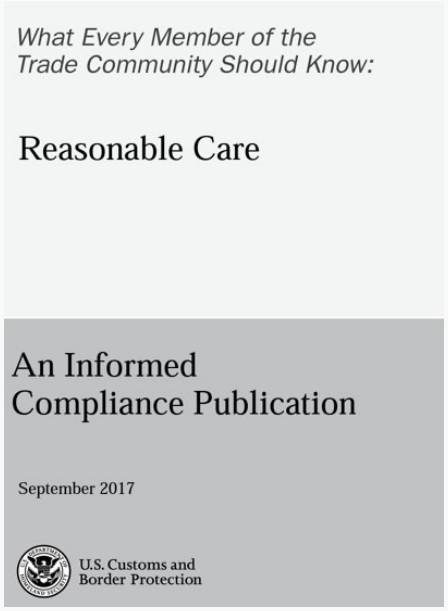

Guardrail: Controlled Web Sources for AI Reference

Reduce Inaccuracies by Limiting Web Sources

- **Provide as much product data as possible.**

- **Configure/restrict web sources.**
  - If the AI tool pulls information from the internet to fill gaps in product details, clearly define which source websites are acceptable.

Guardrail: Confidence Scores & Automated Workflows

Set Risk Thresholds

- Set specific scores to trigger manual human review based on risk tolerance.
  - Check for configuration capability: If the default settings are not adequate, are they configurable? If so, to what degree?

Human-In-The-Loop: Automated Workflows

- Rapidly adjust confidence parameters to respond to shifting global trade regulations.
- Keep a human in-the-loop and tools for workflow triggers.
  - E.g., Supervisor workflow triggers for complex/updated classifications.

Confidence Scores: Example

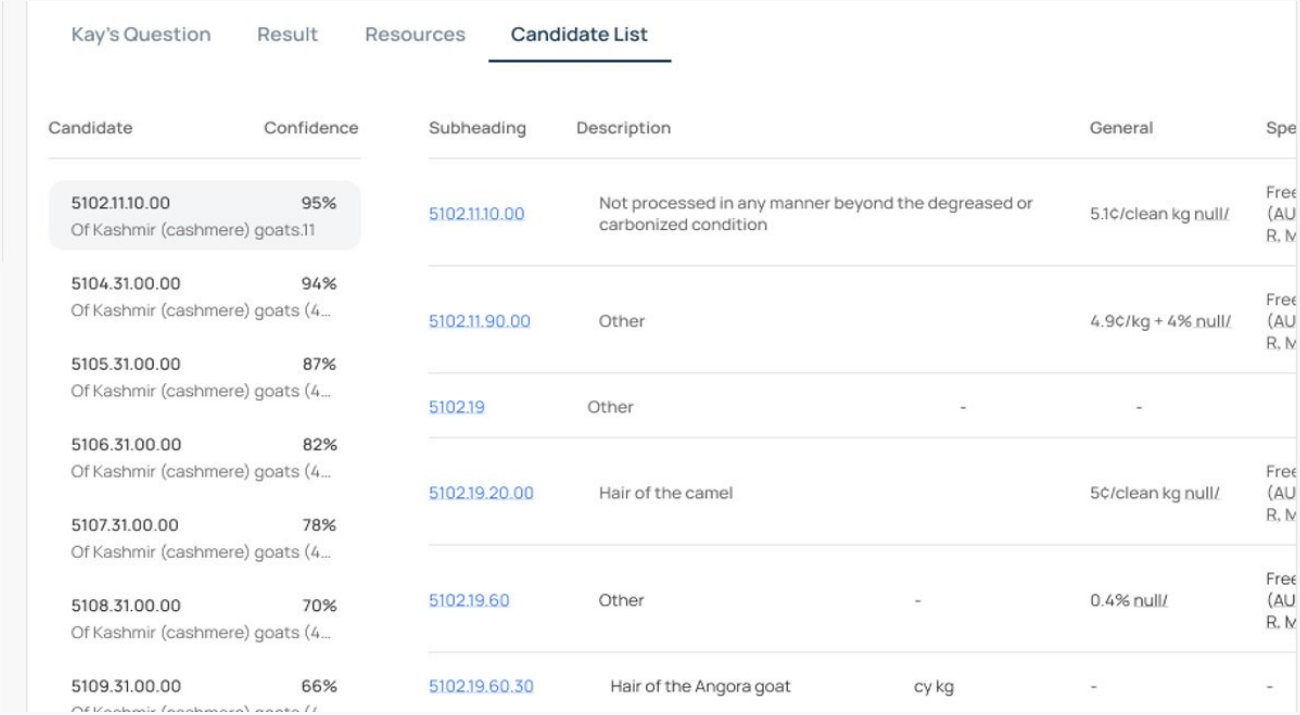

AI + Human in the loop

Product data, attributes, and attestations are optionally reviewed and approved by a human

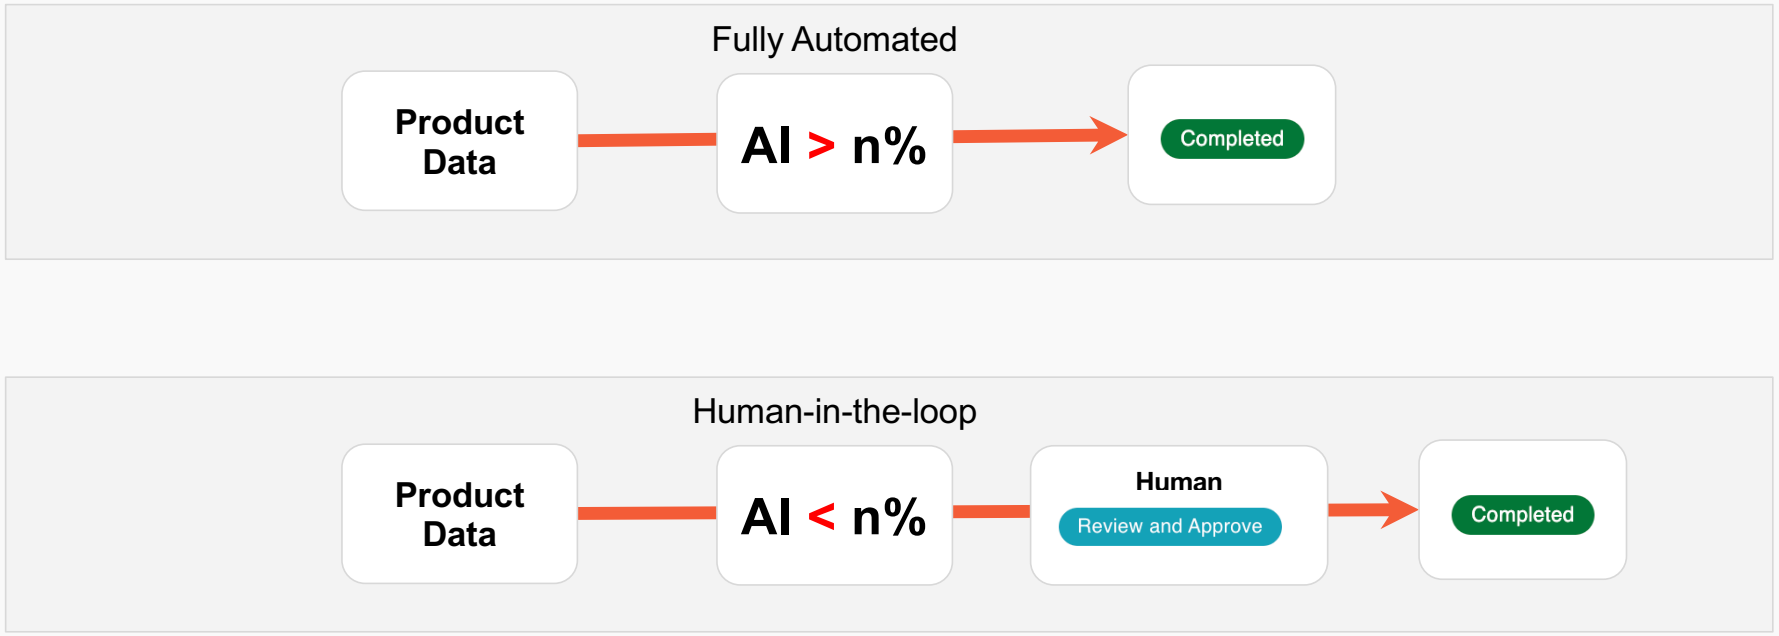

n% = user specific confidence score

Guardrail: Multi-Layered Audit Architecture

Precision at Scale: Enhancing Audit Depth and Breadth through AI Restructuring

Internal Audits

- Conduct audits of AI reasoning at least annually, recommended quarterly.
- Same sampling approach as a broker audit

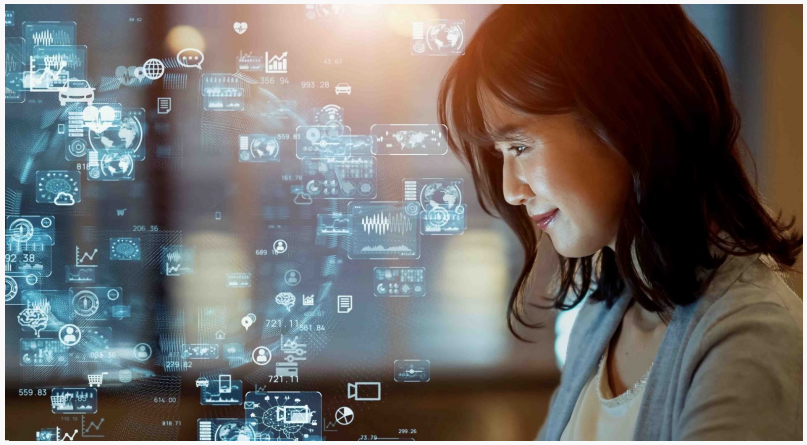

AI <> Human: Multi-Layered Audit

- Humans can audit AI output
  - E.g. Classifications and confidence scoring
- AI can audit human output
  - E.g. Pre- Post- Entry Reviews
  - Post Summary Corrections
  - Leverage AI to find discrepancies

Summary: Guardrail Suggestions

Practical Guardrails

- Informed Compliance: Ensure that the tool has access to updated rules and regulations.
- Web Sourcing: Limit and define internet source sites
- Confidence Scoring: AI output is rated by confidence score
- Customizable Workflows: Configure triggers that prompt a human-in-the-loop
- Internal Audit Processes: Maintain a regular internal audit process. Leverage AI to audit human output.

Additional Considerations for AI Tools

- Is there a click-ready audit trail?
- Is the tool designed by industry professionals?

Enterprise AI Adoption Checklist

A Discussion Guide for Cross-Functional Teams

Enterprise AI Adoption Considerations

Discussion Topics for Cross-Functional Teams

Outbound Data Governance & Human Review

Consider human signoff for data signoff; conclusion reasoning

Inbound Data Scope & Enablement Risk

Data acquisition alternatives to integration; data gap management

Structured Outputs & Decision Decomposition

AI output structure; rules vs. inference

Control Architecture & Monitoring

Application-level controls; AI response logs

Source of Regulatory Truth & Sustainability

Regulatory data citations; updated regulatory materials

Internal Audits

Audit the AI agents' results similar to a customs broker humanoid

Enterprise AI Adoption Checklist

Outbound Data Governance & Human Review

(Different standards for downstream system data vs. audit-facing outputs)

- Is there a bright-line rule that no AI-generated data enters enterprise systems without explicit human signoff (or a documented exception)?

  Why this matters: Automation convenience often erodes governance quietly. Enterprises must consciously decide which data, if any, can bypass review, and why.

- Are all AI-generated outputs that flow into downstream systems fully visible to users before ingestion?

  Why this matters: Hidden or opaque write-backs are the fastest way AI shifts from “assistive” to “decision-making” without accountability. Visibility is the prerequisite for control.

- Is human review focused on validating conclusions and evidence, not re, performing the underlying analysis?

  Why this matters: If humans must redo the work, the AI has failed operationally. Review should confirm defensibility, not replicate effort.

- Are audit-facing AI outputs structured into discrete compliance elements (assumptions, sources, determinations) rather than long-form narrative?

  Why this matters: Requiring reviewers to read 100 pages defeats the purpose of AI. Auditable AI should reduce review effort by isolating what actually needs validation.

Enterprise AI Adoption Checklist

Inbound Data Scope & Enablement Risk

(The risk of starving the AI of necessary inputs)

Is the minimum required inbound data scope clearly defined for reliable AI operation?

Why this matters: AI accuracy is bounded by what it can see. Vague input requirements obscure responsibility when outputs degrade.

If full ERP or system integration is not possible, does the vendor offer alternative data acquisition mechanisms (e.g. domain-constrained web search)?

Why this matters: AI will always "sound smart", but if it doesn't have the right inputs, it won't have the right outputs. Implementing without API integrations is a risk decision that must be discussed and addressed.

Does the AI explicitly flag determinations made with incomplete or inferred data?

Why this matters: Silent inference creates false confidence. Users must know when outputs rest on assumptions rather than verified inputs.

Are data gaps captured as part of the output (e.g., "known unknowns")?

Why this matters: Good AI documents uncertainty so risk can be managed consciously.

Enterprise AI Adoption Checklist

Structured Outputs & Decision Decomposition

(Understanding how the AI arrived at the result)

Are AI outputs decomposed into inputs, assumptions, logic, and conclusions?

Why this matters: If a decision cannot be broken apart, it cannot be explained or defended to regulators or auditors.

Can reviewers distinguish between deterministic rules and probabilistic inference in the output?

Why this matters: Treating probabilistic outputs as facts is a common and dangerous misconception in AI-assisted compliance.

Are confidence indicators or validation flags exposed, not hidden?

Why this matters: Confidence is part of the data. Suppressing it encourages automation-bias and over-reliance.

Can AI-generated decisions be audited programmatically over time?

Why this matters: Point-in-time explainability is insufficient. Enterprises need trend, level visibility to detect drift, degradation, or systemic error.

Enterprise AI Adoption Checklist

Control Architecture & Monitoring

(Where governance moves from promises to proof)

Is the AI constrained behind application-level controls rather than open-ended user prompting?

Why this matters: Control architecture determines behavior more than policy statements. Guardrails must be enforced technically, not culturally.

Are inputs validated and outputs schema-constrained before use?

Why this matters: Most AI failures are integration failures. Schema enforcement prevents "creative" outputs from becoming operational defects.

Is AI activity logged at a level sufficient for incident reconstruction and audit defense?

Why this matters: If you can't reconstruct what happened, you can't credibly respond to regulators, or customers.

Enterprise AI Adoption Checklist

Source of Regulatory Truth & Sustainability

(Will this still work, and be defensible, next year?)

Are the sources of regulatory data explicitly identified and contractually supported?

Why this matters: “The model knows” is not a data strategy. Enterprises need assurance that sources are authoritative, licensed, and durable.

Is there a documented process for regulatory updates and version control?

Why this matters: Static compliance data creates delayed risk. Update cadence is as important as accuracy.

Can historical AI outputs be tied back to the regulatory version in effect at the time?

Why this matters: Retroactive inconsistency is a litigation and audit nightmare. Temporal traceability is non-negotiable.

Is the vendor economically and operationally incentivized to maintain regulatory data long-term?

Why this matters: Regulatory drift is inevitable if the vendor’s business model can’t support ongoing source-of-truth maintenance for regulatory data.

Thank You

CONTACT

Todd R. Smith | Founder CEO, Licensed Customs Broker
todd@kyg.ai 949-903-0338

Anand Raghavendran | Chief Technology & Product Officer (CTPO), Licensed Customs Broker
anand@kyg.ai 949-923-1357

Leslie Levy August | Chief Marketing Officer
leslie@kyg.ai 650-391-7759