International Journal of

Supply Chain and Logistics

(IJSCL)

Autonomous Multi-Jurisdictional Logistics under 21st Century Trade Fragmentation

CARI Journals

International Journal of Supply Chain and Logistics
ISSN 2520-3983 (Online)
Vol. 10, Issue No. 4, pp 1 - 16, 2026
www.carijournals.org

Autonomous Multi-Jurisdictional Logistics under 21st Century Trade Fragmentation

Ravindra Kumar Patro[1*], Susruta Satapathy[2], Subhash Chandra[3], Nitin Rakheja[4]

[1]Zum Services Inc.

https://orcid.org/0009-0003-9256-9263

[2]Tata Consultancy Services

[3]Amazon.com, Inc.

[4]Zoox

Accepted: 5th May, 2026, Received in Revised Form: 12th May, 2026, Published: 24th May, 2026

ABSTRACT

Purpose: This study develops an operating model for autonomous, multi-jurisdictional freight flows that must simultaneously meet requirements for customs risk management, cybersecurity resilience, AI governance, and cross-border data stewardship, addressing the compliance and governance challenges created by the convergence of trade fragmentation and autonomous logistics deployment.

Methodology: A multi-case, document-based comparative analysis was applied, anchored in recent EU border-risk reforms (ICS2 Release 3), security baselines (WCO SAFE), and horizontal regulatory regimes (NIS2, the EU AI Act, and eIDAS 2.0). Documents were coded using a structured scheme covering data requirements, trust and identity expectations, and operational governance mechanisms.

Findings: The study synthesizes an evidence-first compliance architecture mapping: (i) data elements required for advance cargo security filings; (ii) cryptographic identity and audit trails for automated decisions; and (iii) operational escalation protocols to preserve safety and service continuity under disruption. Corridors implementing decision-level audit trails exhibit lower documentation defect rates and shorter exception cycle times.

Contribution to Theory, Policy and Practice: The paper contributes a practical governance blueprint that is role-segregated, auditable, and interoperable across regulators, alongside a framework for measuring compliance latency and exception handling in autonomous logistics. It enables managers to quantify risk, reduce clearance variability, and maintain uptime in safety-critical delivery networks, while extending compliance theory from episodic documentation to continuous, architecture-embedded assurance.

Keywords: *Autonomous Logistics, Customs Compliance, AI Governance, Cybersecurity Resilience, Trade Fragmentation*

JEL Codes: *F1, L87, K24, O3, D73*

1.0 INTRODUCTION

Trade policy volatility and 'security-first' border regimes are reshaping how logistics networks are designed, instrumented, and governed. While tariffs and export controls are highly visible, the more operationally consequential shift is the expansion of advance cargo information requirements and the tightening of digital risk management at borders (Hesketh, 2010; European Commission, n.d.). For logistics operators, particularly those integrating autonomy (e.g., automated dispatch, sensor-driven routing, robotic yard moves, and decision support for exceptions) (Hofmann & Rüsch, 2017), the compliance burden increasingly depends on the integrity, timeliness, and provenance of data rather than on paper documentation (Bharosa et al., 2013).

Autonomous logistics increases the volume and velocity of operational telemetry (location, condition monitoring, safety events, device identity signals, and operator interventions). At the same time, regulators and security frameworks demand structured, high-quality data for risk scoring and auditability. The EU's Import Control System 2 (ICS2) is emblematic: it expands pre-arrival security filings to additional transport modes and introduces multi-party filing arrangements (e.g., information provided by carriers, freight forwarders, and other actors at 'house' and 'master' levels) to support customs risk analysis (European Commission, n.d.). Operationally, this creates a distributed responsibility model in which data completeness and quality must be coordinated across actors and systems well before border arrival.

Concurrently, horizontal regulatory regimes shape the governance environment for automated systems that touch critical infrastructure. The NIS2 Directive raises baseline cybersecurity obligations for important and essential entities and their supply-chain security practices (European Parliament and Council, 2022). The EU AI Act establishes requirements for high-risk AI systems and mandates risk management, documentation, transparency, and human oversight where relevant (European Parliament and Council, 2024a). eIDAS 2.0 reinforces a European digital identity framework intended to enable interoperable trust services and identity credentials (European Parliament and Council, 2024b). In parallel, cross-border data transfer mechanisms—such as the EU-U.S. Data Privacy Framework adequacy decision—shape where data may flow and under what safeguards (European Commission, 2023; EDPB, 2023).

These developments challenge a common managerial assumption: that compliance can be and led 'after the fact' by a documentation team, while operations and technology optimize for cost and speed. In autonomous logistics, compliance is upstream and embedded. Systems must be engineered to produce trustworthy evidence continuously; otherwise, automation increases failure modes (e.g., data drift, untraceable overrides, cross-border transfer violations) that manifest as border holds, forced manual intervention, or-worst case-safety incidents.

This paper addresses the managerial problem created by the convergence of these forces: how can autonomous logistics systems operate across jurisdictions while maintaining (a) customs-grade data integrity, (b) cybersecurity and AI governance controls, and (c) operational safety and service

International Journal of Supply Chain and Logistics
ISSN 2520-3983 (Online)
Vol. 10, Issue No. 4, pp 1 - 16, 2026
www.carijournals.org

continuity? The objective is not to provide a jurisdiction-by-jurisdiction legal treatise, but to propose a reusable operating model that logistics leaders can implement as they scale autonomy across corridors.

The remainder of the paper is structured as follows. Section 2 synthesizes the literature and policy context. Section 3 explains the methodology. Section 4 presents results as layered compliance architecture with templates, workflows, and measures. Section 5 discusses theoretical contributions and limitations. Section 6 offers actionable managerial implications, including a maturity model and implementation roadmap. Section 7 concludes and proposes future research.

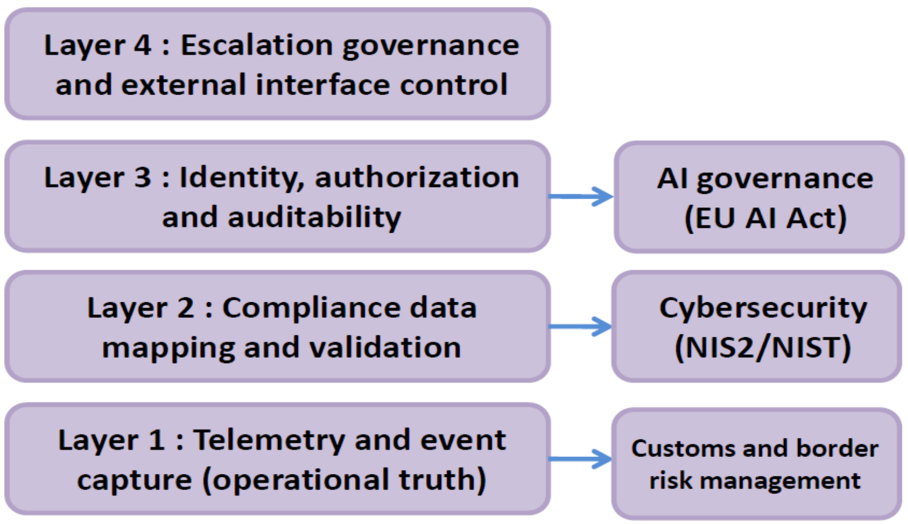

Figure I-B. Conceptual model: layered compliance architecture interacting with external regimes.

2.0 LITERATURE AND POLICY CONTEXT

2.1 Border risk management, advance cargo information, and operational variability

Border agencies increasingly rely on pre-arrival risk screening driven by standardized electronic data (Hesketh, 2010; WCO, n.d.). This shift changes the locus of control from border-adjacent document review to upstream data generation and coordination. Practically, the same shipment can experience very different outcomes depending on whether required data fields are complete, consistent across actors, and submitted within prescribed time windows.

ICS2 operationalizes this approach in the EU through phased releases and progressive expansion of scope (European Commission, n.d.). Release 3 extends to additional transport modes and introduces multiple filing arrangements. The managerial consequence is that compliance becomes a distributed process: different parties (carrier, forwarder, warehouse operator, platform provider) become responsible for specific data subsets, while overall risk screening depends on the assembled record. Distributed filing introduces classic coordination problems—missing data, inconsistent versions, and unclear authority to amend—now with real-time service consequences (Bharosa et al., 2013).

2.2 Supply-chain security standards and facilitation incentives

Frameworks such as the WCO SAFE Framework of Standards create an internationally recognized baseline for secure and facilitated trade (WCO, n.d.). SAFE emphasizes advance electronic information, risk management, and partnerships such as Authorized Economic Operator (AEO) programs (WCO, n.d.). From an organizational perspective, SAFE encourages investment in internal controls (security management, traceability, and documentation discipline) by coupling them to facilitation benefits (reduced inspections, prioritized treatment).

In autonomous logistics, SAFE's core idea-'trust is engineered'-is particularly relevant. Autonomy produces high-volume data and frequent micro-decisions (reroutes, substitutions, schedule changes). Without governance architecture, these micro-decisions can undermine traceability and increase regulatory scrutiny. SAFE implies that facilitation depends on the ability to demonstrate control and evidence, not on the sophistication of automation per se.

2.3 Cybersecurity governance as an operational requirement

Cybersecurity is a central determinant of physical continuity in digitized logistics (Christopher & Peck, 2004). Ransomware incidents and supply-chain compromises can immobilize fleets, break visibility systems, or corrupt data needed for filings and audits. NIS2 strengthens requirements for governance, incident reporting, and supply-chain security across a broader set of entities (European Parliament and Council, 2022).

For autonomous logistics, cybersecurity governance intersects with operational design in several ways: (i) asset identity and secure provisioning for devices and software; (ii) monitored integration surfaces with partners and authorities; (iii) segmentation between corporate IT and operational environments; and (iv) incident response playbooks that prioritize safety and service continuity. Importantly, autonomous operations increase the attack surface by introducing additional sensors, connectivity, and remote management interfaces.

2.4 AI governance and the need for auditable decisioning

AI systems in logistics are frequently deployed as optimization tools (routing, slotting, demand prediction). However, as AI becomes embedded into exception handling and safety decisioning-especially in semi-autonomous operations-governance requirements increase. The EU AI Act requires risk management, documentation, and human oversight for high-risk systems (European Parliament and Council, 2024a).

The central management challenge is traceability: when a model's output influences an action (reroute, hold, escalation, or override), operators must be able to reconstruct what happened. This is required not only for external audits but also for internal continuous improvement. Without auditable decisioning, organizations cannot systematically reduce exception rates or demonstrate that safeguards prevented unacceptable outcomes.

2.5 Digital identity, cross-border data, and trade fragmentation

Autonomous logistics depends on coordination across actors who do not share a corporate boundary. Digital identity frameworks and trust services support non-repudiation and interoperability for filings, approvals, and attestations; eIDAS 2.0 (Regulation (EU) 2024/1183) strengthens the EU's framework for digital identity and trust services (European Parliament and Council, 2024b). Cross-border transfer regimes shape where data may be processed and stored: the EU-U.S. Data Privacy Framework adequacy decision illustrates how transfers can be enabled under defined safeguards while still requiring organizations to operationalize privacy and security controls (European Commission, 2023; EDPB, 2023).

Trade fragmentation compounds these requirements. Beyond visible tariff changes, divergence in technical and administrative requirements—MFN duties, preferential arrangements, origin rules—interacts with customs procedures to increase documentation burden (WTO, 2025). World Bank LPI evidence reinforces that border management and reliability are core determinants of logistics competitiveness (World Bank, 2023). Together, these sources point to a research gap: strong policy descriptions and emerging autonomy capabilities exist, but fewer implementable governance architectures connect regulatory primitives to operational design.

2.6 From episodic compliance to continuous assurance: KPI implications

Most autonomy roadmaps assume operational KPIs (on-time performance, cost per mile, utilization) can be optimized first and compliance added on later. In multi-jurisdictional corridors, this ordering fails: compliance exceptions drive a disproportionate share of cycle-time variance, and exception handling becomes a dominant source of dispatch and broker workload. A two-layer KPI model is therefore useful. Layer 1 captures standard network metrics (OTIF, damage, claims rate, utilization, empty miles, cost-to-serve) (Pajić et al., 2024). Layer 2 captures assurance quality (documentation defect rate, inspection hold rate, identity/authorization defect rate, exception cycle time, audit completeness). Layer 2 indicators are leading: when documentation defects rise, OTIF degradation often follows; when audit completeness drops, dispute remediation weakens.

Traditional cross-border compliance is episodic: filings are prepared before crossing and post-entry audits may occur months later. Autonomy introduces continuous, data-dense decision-making (route changes, load rebalancing, re-tendering, autonomy mode changes, custody transfers), each with potential compliance implications. Continuous operations therefore demand continuous assurance. This paper positions "audit-by-design" (OECD, 2014; Bharosa et al., 2013) as that mechanism: each materially relevant decision is recorded at the moment it is made, with justification, authority, and evidence pointers, producing replayable narratives rather than mere event logs. The posture shifts from individual heroics to a distributed accountability model that withstands multi-agency scrutiny—an approach more analogous to safety-critical engineering than to legacy trade compliance.

2.7 Research gap and propositions

Across the reviewed domains, three gaps remain salient. First, the compliance and customs literature offers limited guidance on how to translate real-time operational telemetry into audit-grade records that can withstand multi-agency scrutiny. Second, autonomy discussions often treat cross-border constraints as fixed "externalities" rather than variables that can be engineered through system design (data architecture, identity, exception workflows). Third, empirical work rarely links governance design choices to measurable operational outcomes such as inspection holds, documentation defect rates, or exception cycle time.

In response, this paper advances three propositions for future empirical testing. (P1) Corridors that implement decision-level audit trails will exhibit lower documentation defect rates and shorter exception cycle time than corridors that rely on post-hoc reconciliation. (P2) Identity first operating models (role-bound permissions, device trust, and authority logging) will reduce inspection holds linked to authorization and documentation issues. (P3) Pairing standard service KPIs (e.g., OTIF) with assurance indicators (documentation defects, audit completeness) will improve operational learning by increasing the organization's ability to attribute service failures to operational versus compliance root causes.

These propositions are designed to guide both field studies and design-science evaluations where the "intervention" is governance and data architecture change rather than a pure optimization model.

3.0 METHODOLOGY

This study uses a multi-case, document-based comparative methodology. The unit of analysis is the 'autonomous cross-border shipment' as a socio-technical system: a sequence of events and decisions spanning origin operations, line-haul transport, border risk filings, inspection outcomes, and last-mile distribution.

Case materials include: (1) primary regulatory and standards documents (ICS2 Release 3 guidance; NIS2; EU AI Act; eIDAS 2.0; EU-U.S. Data Privacy Framework materials), (2) security frameworks (WCO SAFE), and (3) institutional facilitation indicators (World Bank LPI; WTO tariff profile analysis). Documents were selected because they define stable governance primitives—advance data, risk screening, identity, auditability, and incident response-that can be translated into operational requirements.

Documents were coded using a structured scheme with three categories: (i) data requirements and timing constraints; (ii) trust, identity, and audit expectations; and (iii) operational governance mechanisms (roles, accountability, incident handling, escalation, continuity). Within each category, the analysis extracted (a) explicit requirements and (b) implicit design implications (e.g., if multiple actors file, then responsibility must be allocated; if cyber incidents must be reported, then monitoring and classification must exist).

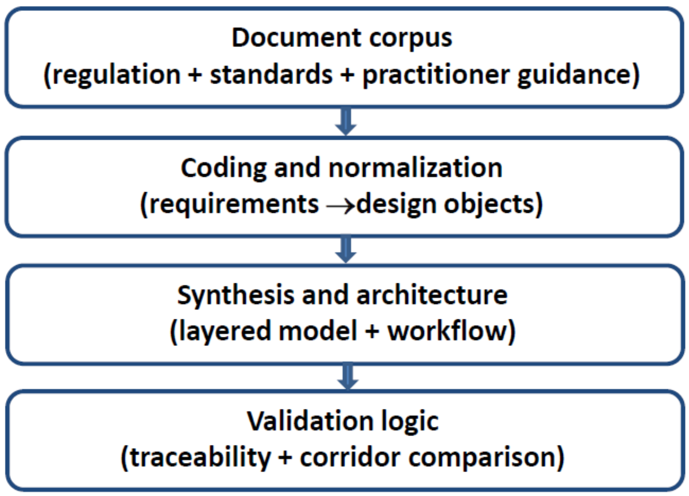

Figure M-A. Research design and synthesis pipeline.

Analytical synthesis followed a two-step process. First, the study normalized recurring requirements into reusable design objects: a compliance data dictionary, audit trail templates, an escalation matrix, and performance indicators. Second, these objects were assembled into a four-layer operating model designed to be implementable in logistics control towers and fleet/yard operations. The resulting architecture is normative and intended to guide design and measurement; empirical validation is proposed as future research.

4.0 RESULTS: A COMPLIANCE ARCHITECTURE FOR AUTONOMOUS MULTI-JURISDICTIONAL LOGISTICS

4.1 Overview

The proposed operating model is built on four layers: (L1) operational telemetry and event capture; (L2) compliance data mapping and validation; (L3) identity, authorization, and audit trails for automated decisions; and (L4) escalation governance and external interface controls. The model assumes autonomy is deployed within a safety-critical environment and therefore treats continuity and auditability as core requirements.

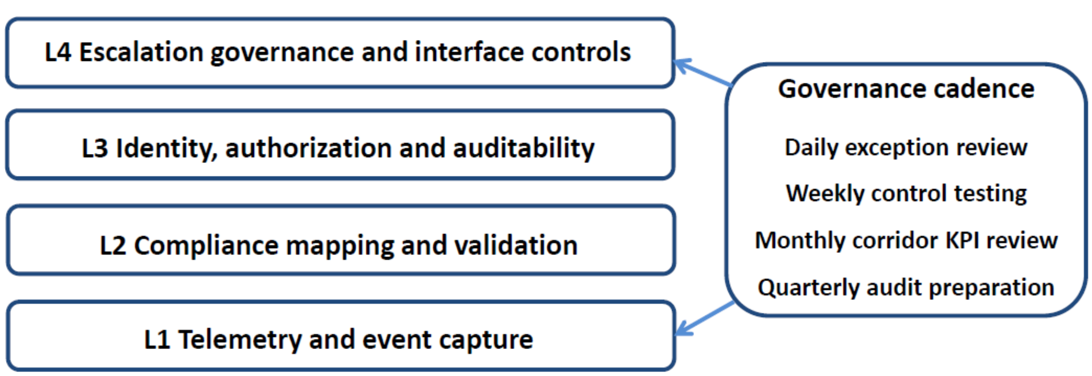

Figure F-A. Four-layer AMJL compliance architecture (overview).

International Journal of Supply Chain and Logistics
ISSN 2520-3983 (Online)
Vol. 10, Issue No. 4, pp 1 - 16, 2026 www.carijournals.org

4.2 Layer 1: Telemetry and event capture (operational truth)

Layer 1 defines the operational 'truth set' required to reconstruct shipment history and support both safety and compliance outcomes. Autonomous logistics relies on high frequency data (position traces, geofence transitions, condition monitoring, equipment diagnostics) (Hofmann & Rüsch, 2017; Bartlett et al., 2007) and event data (handoffs, loading/unloading, seal changes, damage reports, operator interventions). Governance requirements include time synchronization, stable identity binding, and tamper-evident retention. If events cannot be reliably attributed to an asset, shipment, and authorized operator, then downstream filings and audits become contestable.

Implementation guidance: an event ledger should (a) normalize events into a canonical schema, (b) attach cryptographic hashes for tamper evidence (Wu et al., 2024), (c) maintain references to raw data when necessary (e.g., images or sensor logs), and (d) enforce retention/access controls aligned to risk and legal requirements.

Table I. Regulatory-to-Data-Element Mapping Template

Regulatory instrument	Filing object	Required data element	Source system of record	Responsible party	Validation rule	Timing trigger	Amendment rule	Evidence link
—	—	—	—	—	—	—	—	—

Table II. Audit Trail Template for Automated Decisions

Decision ID	Timestamp	Decision type	Triggering event(s)	Data inputs (ledger links)	Model/version	Output/confidence	Human oversight action	Authority level	Rationale	Downstream actions	Exception flags	Post-event review outcome	Retention/access classification
—	—	—	—	—	—	—	—	—	—	—	—	—	—

Table III. Escalation Matrix for Autonomous Cross-Border Exceptions

Exception category	Example triggers	Primary owner	Backup owner	Decision authority	Required evidence	Customer notification	Regulatory notification	Fallback option	Time-to-resolve target
—	—	—	—	—	—	—	—	—	—

4.3 End-to-end workflow walkthrough

The architecture activates at six checkpoints along a shipment's life. At order creation and booking, Layer 2 validation flags missing or ambiguous fields that will later be required for advance filings (the "shift-left" principle: resolve data ambiguity before physical movement begins). During origin consolidation and packing, operational events (pack, seal, weigh, label) are logged in the event ledger (Layer 1) and linked to the shipment's identity; re-packing or consolidation triggers an amendment workflow with version control, and audit objects are created for any automated decisions affecting compliance-relevant attributes (Layer 3). During line-haul execution, when autonomous or semi-autonomous routing systems propose changes that could affect the border crossing, the system assesses whether advance filings need updates (Layer 2), validates and submits them with evidence links, captures the decision via the audit template (Layer 3), and triggers human oversight at high-risk thresholds.

At border pre-arrival, the compliance stack must ensure filings are complete, consistent, and submitted; Layer 4 governance ensures a designated "compliance single owner" coordinates with partners to close gaps, and integration alerts route filing errors to the escalation matrix. During inspection or hold handling, Layer 4 governs evidence collection, customer and authority interaction, and route re-planning; the event ledger records the inspection timeline while decision objects record any automated or human overrides made under pressure, and post-event review feeds the weekly governance cadence. Finally, at destination delivery and post-release audit, the system completes audit packages automatically (filings, amendments, decision logs, supporting evidence), measures audit completeness, and categorizes exceptions for continuous improvement.

4.4 Performance measures and governance cadence

To operationalize governance, managers need measurable indicators. The architecture proposes: compliance latency (event-to-filing time), exception resolution time, audit completeness rate, clearance variability, and interface health (integration uptime and error rates). These measures support a daily triage (exceptions, holds, integration failures) and a weekly governance review (root causes, recurring defects, partner performance).

Importantly, these indicators are corridor specific. A mature organization benchmarks performance by corridor and by partner, enabling targeted interventions (e.g., fix a specific data feed, renegotiate a partner data responsibility, adjust internal training).

4.5 Interoperability under cross-border data constraints

Autonomous logistics telemetry can include sensitive data. Cross-border transfer regimes shape where data can be processed and stored. The EU-U.S. Data Privacy Framework illustrates how adequacy decisions can enable transfers under defined safeguards while requiring organizations to operationalize privacy and security controls (European Commission, 2023).

A practical design principle is modularization: separate compliance-critical datasets (needed for filings and audit packages) from high-volume telemetry. Implement policy-driven routing and retention and support localized processing for sensitive data when required. This allows organizations to adapt to divergent regimes without losing end-to-end auditability.

4.6 Illustrative corridor case comparisons

To demonstrate how the operating model adapts across regimes, Table IV contrasts two illustrative corridors: U.S.-Canada movements under USMCA/CBP/CBSA operating realities and EU-UK movements where data requirements and border controls have evolved post-Brexit. Across both, the common requirement is replayability—the ability to reconstruct what was known, who decided, what control checks were performed, and what changed after a decision. What varies is the content of the record (data fields, filings, evidence artifacts) and the shape of exceptions (inspection holds, documentation mismatches, port congestion, data-sharing restrictions).

Table IV. Illustrative corridor comparison for AMJL design

Design element	U.S.-Canada corridor (illustrative)	EU-UK corridor (illustrative)
Primary assurance focus	Customs/security filings; carrier and equipment eligibility; admissibility and inspection holds	Data and filings heterogeneity; evolving border processes; safety/security and data-sharing constraints
Identity controls	Role-bound permissions for dispatch, broker, and documentation actions; device trust for telemetry sources	Role-bound permissions plus stronger data minimization and provenance for shared datasets across agencies and partners
Exception pattern	Inspection holds, paperwork defects, equipment/driver eligibility issues; weather disruptions	Documentation mismatches, rule/format divergence, port congestion; safety/security inspections and queue volatility
Audit trail emphasis	Replayable decision narratives for route/custody changes to defend admissibility and duty/tax posture	Replayable narratives plus provenance to prove integrity across multiple systems and changing filing formats

5.0 DISCUSSION

5.1 Theoretical contribution: compliance as operational architecture

The central contribution is a shift in how compliance is conceptualized. Rather than treating compliance as external constraint managed by specialized teams, the paper treats compliance as an operational architecture composed of primitives: data, identity, audit trails, and authority. These primitives can be designed, instrumented, and measured. In autonomous logistics, where decisions occur continuously, this architecture becomes essential to prevent automation from amplifying failure modes.

ICS2 multiple filing and advance data requirements illustrate why these reframing matters: distributed filing makes compliance inseparable from coordination and data engineering (European Commission, n.d.; Bharosa et al., 2013).

5.2 Linking cybersecurity, AI governance, and continuity

NIS2 and the EU AI Act imply that organizations must provide governance evidence—not only of what systems do, but of how risks are managed (European Parliament and Council, 2022, 2024a). In safety-critical logistics, continuity is itself a governance outcome: the ability to keep operating safely under disruption. This paper proposes that audit trails and escalation matrices are the organizational mechanisms that connect policy requirements to continuity-by-design.

WCO SAFE reinforces this interpretation by positioning secure supply chains and trusted operator relationships as foundations for facilitation (WCO, n.d.).

5.3 Propositions for future empirical testing

To support future research, the architecture yields testable propositions. Proposition 1: Lower compliance latency (faster validated amendments) reduces clearance variability and downstream delivery variance. Proposition 2: Higher audit completeness rates correlate with lower inspection intensity over time for similar shipment profiles (conditional on jurisdiction and commodity). Proposition 3: A defined escalation matrix reduces exception resolution time and decreases the share of exceptions that require ad hoc executive intervention. Proposition 4: Interface health (lower integration error rates) is a leading indicator of border friction and customer service incidents.

These propositions can be evaluated using shipment-level event data, customs outcome data (e.g., hold/inspect/release), and operational performance data (on-time delivery, rework rates).

5.4 Limitations

The study is document-based and therefore cannot claim quantified causal effects. Additionally, regulatory landscapes evolve; however, the architecture emphasizes stable primitives that remain relevant across regimes even as specific filing formats change. Finally, autonomy varies by context; the architecture assumes safety-critical operations where human oversight remains available and required for certain decisions.

6.0 MANAGERIAL IMPLICATIONS

6.1 A maturity model for autonomous cross-border compliance

Managers often need to stage investments. A four-stage maturity model can guide adoption.

Stage 1 (Reactive): Compliance handled by manual document teams; limited telemetry integration; audits assembled after incidents. Stage 2 (Instrumented): Event ledger exists; basic validation gates; ad hoc decision logs. Stage 3 (Governed): Compliance data dictionary in place; role-segregated authority; standardized audit templates; escalation matrix; routine governance cadence. Stage 4 (Adaptive): Modularized data architecture supporting localization; automated audit package generation; corridor benchmarking; continuous reduction of compliance defects; partners integrated through standardized trust mechanisms.

This maturity framing allows leadership to connect compliance investment to operational outcomes (predictability, reduced rework, faster exception resolution).

A practical roadmap follows six steps: (1) stand up a compliance data product team with clear ownership and KPIs; (2) build the compliance data dictionary and map it to operational processes (Table I); (3) implement validation gates and amendment workflows and measure compliance latency; (4) deploy identity and audit templates (Table II) and enforce them through tooling; (5) create the escalation matrix and governance cadence (Table III) and measure resolution time; (6) modularize data routing and retention to handle cross-border constraints and validate transfer safeguards where applicable.

6.2 Connecting governance to performance narratives

Executives and customers frequently evaluate logistics partners on reliability. Facilitation and performance indicators such as the LPI signal that border capability and reliability are strategic assets (World Bank, 2023). Managers can use the architecture's metrics to tell a performance story: fewer late amendments, fewer preventable holds, faster exception closure, and higher audit readiness. In fragmented trade environments, this story is itself a competitive differentiator.

7.0 CONCLUSION

Autonomous logistics will increasingly operate in a world where border risk management, cybersecurity governance, and AI oversight are intertwined. The managerial challenge is to maintain safety and service continuity while meeting distributed data-intensive compliance obligations. This paper contributes a layered operating model that maps operational telemetry to customs-grade data, binds automated decisions to identity and audit trails, and defines escalation governance for exceptions. By reframing compliance as an operational architecture with measurable indicators, the model enables logistics leaders to reduce clearance variability, improve audit readiness, and scale autonomy across jurisdictions.

Future research should test the proposed propositions empirically using corridor-level event data, evaluate how audit completeness influences inspection outcomes over time, and study how interoperable digital identity ecosystems change compliance costs and partner coordination dynamics.

8.0 SUPPLEMENTARY IMPLEMENTATION DETAIL

8.1 Organizational design, roles, and partner interface contracts

Autonomous cross-border logistics often fails not because the technology cannot optimize, but because accountability is unclear at the moment of disruption. The architecture uses three lines of responsibility—Operations (service continuity), Compliance (regulatory interface and filing quality), and Technology (platform integrity and interface health)—with five named roles: Corridor Compliance Owner, Control Tower Duty Manager, Integration Steward, Security & Risk Liaison, and Audit Readiness Lead. A RACI mapping makes ownership explicit at each handoff (for example, a route change that affects a border crossing is Responsible: Control Tower Duty Manager; Accountable: Corridor Compliance Owner; Consulted: carrier/forwarder partner; Informed: customer and yard operations).

In distributed filing contexts, partner interfaces should be governed like production systems. Interface contracts should specify message formats, required fields, acknowledgement semantics, retry logic, dead-letter handling, maximum acceptable latency, and escalation thresholds; data SLAs should specify field completeness, timeliness, and amendment turnaround targets. Governance cadences review partner SLA adherence and address systemic defects through root-cause analysis rather than repeated firefighting.

8.2 Engineering disciplines: quality, drills, and stress testing

Compliance data should be subjected to quality engineering practices similar to safety engineering: unit tests for data transformations, synthetic transaction testing for integrations, and "canary shipments" to validate end-to-end workflows after system changes. A "golden dataset" of representative shipments enables nightly validation to detect schema drift, mapping regressions, or partner feed changes. Training should be operationally framed—how validation gates prevent holds, how audit trails prevent blame games, how escalation ladders protect safety commitments—and reinforced through periodic exception drills analogous to safety drills.

Scenario stress testing complements drills by evaluating whether the architecture holds under high-stress conditions: peak-season surges, mass weather reroutes, cyber incidents that degrade an integration endpoint, partner outages affecting multiple filings, and regulatory changes introducing new mandatory fields. Even without simulation tools, tabletop exercises using real shipment histories with synthetic disruptions yield measurable resilience indicators: time-to-detect, time-to-amend, time-to-notify, and time-to-close.

8.3 KPI dashboard and compliance latency as a leading indicator

Traditional logistics KPIs (on-time delivery, cost per mile, utilization) are lagging indicators: they reveal outcomes after disruption has propagated. Compliance latency—the "event-to-filing clock" for triggers such as route change, carrier substitution, consolidation, or seal change—is a leading indicator. For each trigger, set a target SLA (for example, 30 minutes for route change amendments; 2 hours for consolidation amendments), track median and tail latency, and root-cause tail latency by attributing delays to missing data, partner dependency, approval bottlenecks, or integration failures. A governance dashboard should combine Layer 1 service KPIs (OTIF, cost-to-serve, utilization) with Layer 2 assurance KPIs (documentation defect rate, inspection hold rate, identity/authorization defect rate, exception cycle time, audit completeness, interface health) and corridor-segmented "defect heatmaps" that show where holds originate by partner, filing type, and defect category.

8.4 Decision logging, AI lifecycle governance, and escalation triggers

Not all decisions warrant the same logging rigor. Tier A covers safety and regulatory decisions (full audit template, mandatory human oversight where required); Tier B covers customer commitment decisions with potential compliance impact (reduced fields); Tier C covers routine optimization decisions (aggregate metrics only). The taxonomy reduces operator fatigue while preserving defensibility where it matters. An illustrative Tier A record for a compliance-affecting reroute would capture: timestamp, trigger (weather closure risk), inputs (telemetry, weather feed, filing status), model and version, output and confidence, human oversight action and authority level, rationale, downstream actions (dispatch update, partner notification, filing amendment, customer notification), exception flags, post-event review outcome, and evidence links to the event ledger.

Even when the EU AI Act does not classify a specific logistics use case as high-risk, the disciplines it codifies are useful for safety-critical operations (European Parliament and Council, 2024a): training data provenance, versioning and change approval, pre-deployment validation, post-deployment drift monitoring, retraining triggers, and decommissioning procedures. The audit trail template should capture model version and confidence wherever a model influences an action. Explicit triggers must require human oversight or controlled shutdown, including repeated identity-credential integrity failures, sensor anomalies suggesting tampering, persistent mismatch between declared routing and telemetry, and repeated integration acknowledgement failures—encoded into the escalation matrix so that autonomy does not "fail silently".

8.5 Partner onboarding and adapting to regulatory change

Because multiple filing distributes responsibility, partner onboarding becomes a compliance control: interface certification tests, data dictionary alignment, identity and authorization verification, and escalation contact validation. Contracts should embed governance clauses (data SLAs, notification duties, remediation timelines). To absorb regulatory change without platform rewrites, the compliance data dictionary should be treated as a configuration layer rather than hard-coded logic, and interface contracts should version schemas with backward compatibility windows. A "regulatory change pipeline" analogous to software release management monitors upcoming changes, updates the dictionary, validates on the golden dataset, certifies partner interfaces, and deploys with canary shipments.

8.6 Economics of compliance-by-design and data minimization

Compliance engineering is best analyzed as a variability-reduction investment (Christopher & Peck, 2004). Border holds, preventable inspections, and manual rework impose costs that can be quantified as "avoidable friction hours"—delay hours from data defects multiplied by downstream cost coefficients—providing a finance-ready justification for validation, integration, and audit automation. Data minimization is a parallel resilience strategy: classifying telemetry into compliance-critical, operationally useful, and optional tiers supports policy-driven routing and retention without losing compliance capability.

REFERENCES

Bartlett, P. A., Julien, D. M., & Baines, T. S. (2007). Improving supply chain performance through improved visibility. *The International Journal of Logistics Management*, *18*(2), 294–313. https://doi.org/10.1108/09574090710816986

Bharosa, N., Janssen, M., van Wijk, R., de Winne, N., van der Voort, H., Hulstijn, J., & Tan, Y.-H. (2013). Tapping into existing information flows: The transformation to compliance by design in business-to-government information exchange. *Government Information Quarterly, 30*(Suppl. 1), S9–S18. https://doi.org/10.1016/j.giq.2012.08.006

Christopher, M., & Peck, H. (2004). Building the resilient supply chain. *The International Journal of Logistics Management, 15*(2), 1–14. https://doi.org/10.1108/09574090410700275

European Commission. (n.d.). Import Control System 2 (ICS2): Release 3 (including multiple filing and extension to additional transport modes). European Commission, Taxation and Customs Union.

European Commission. (2023). EU-U.S. Data Privacy Framework: adequacy decision documentation and related materials.

European Data Protection Board [EDPB]. (2023). Materials related to the EU-U.S. Data Privacy Framework.

European Parliament and Council. (2022). Directive (EU) 2022/2555 on measures for a high common level of cybersecurity across the Union (NIS2). Official Journal of the European Union.

European Parliament and Council. (2024a). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union.

European Parliament and Council. (2024b). Regulation (EU) 2024/1183 on a European Digital Identity Framework (amending Regulation (EU) No 910/2014 — eIDAS). Official Journal of the European Union.

Hesketh, D. (2010). Weaknesses in the supply chain: Who packed the box? *World Customs Journal*, *4*(2), 3–20.

Hofmann, E., & Rüsch, M. (2017). Industry 4.0 and the current status as well as future prospects on logistics. *Computers in Industry, 89*, 23–34. https://doi.org/10.1016/j.compind.2017.04.002

OECD. (2014). *Tax compliance by design: Achieving improved SME tax compliance by adopting a system perspective*. OECD Publishing. https://doi.org/10.1787/9789264223219-en

Pajić, V., Kilibarda, M., & Andrejić, M. (2024). Enhancing distribution efficiency through OTIF performance evaluation. *Mathematics, 12*(21), 3372. https://doi.org/10.3390/math12213372

World Bank. (2023, April 21). World Bank releases Logistics Performance Index 2023 (press release).

World Customs Organization [WCO]. (n.d.). SAFE Framework of Standards to Secure and Facilitate Global Trade. https://www.wcoomd.org/

World Trade Organization [WTO]. (2025). World tariff profiles 2025. WTO Publications.

Wu, L., Lu, W., & Chen, C. (2024). Compliance checking for cross-border construction logistics clearance using blockchain smart contracts and oracles. *International Journal of Logistics Research and Applications, 27*(12), 2778–2812. https://doi.org/10.1080/13675567.2024.2328718