

Artificial Intelligence Readiness for Compliance Risk Management: A Qualitative Case Study of a Vietnamese Maritime Logistics Firm

Võ Thị Thu Hồng¹; Vũ Kiều Sa²

^{1,2} Saigon International University, Ho Chi Minh City, Vietnam

Publication Date: 2026/08/25

Abstract: **Purpose** – This study investigated how a maritime logistics firm in Vietnam could develop organizational readiness for the responsible use of artificial intelligence (AI) in compliance risk management. It addressed three connected questions: which compliance risks should be prioritized, what organizational conditions constrain AI adoption, and which implementation sequence is feasible for a resource-conscious firm. **Design/methodology/approach** – An applied qualitative single-case design was used at TRA-SAS. Evidence comprised three semi-structured interview sessions involving eight participants from leadership and key functional areas, internal process and business documents, public corporate reports for 2021–2025, and a 5 × 5 risk-assessment matrix. Thematic analysis was combined with a purpose-specific AI-readiness assessment covering data, technology, people, processes, finance, leadership, compliance culture, and security/legal governance. **Findings** – Eight material compliance risks were identified. Customs compliance (CR04) received the highest priority within the portfolio (likelihood 3, impact 5, score 15; High). The firm displayed an uneven M2–M3 digital-maturity profile: software and selected processes were near M3, while data standardization and system integration remained closer to M2. High-value AI opportunities were document checking, contract and obligation analysis, regulatory change monitoring, early-warning analytics, and compliance reporting. However, these use cases depended on data governance, standardized workflows, role clarity, human review, and model monitoring. **Originality/value** – The study connects compliance-risk materiality with purpose-specific AI readiness in an under-researched maritime logistics setting. It proposes a risk-first, readiness-gated, human-in-the-loop pathway that avoids treating AI adoption as a technology procurement decision.

Keywords: Artificial Intelligence; AI Readiness; Compliance Risk; RegTech; Maritime Logistics; Responsible AI; Vietnam.

How to Cite: Võ Thị Thu Hồng; Vũ Kiều Sa (2026) Artificial Intelligence Readiness for Compliance Risk Management: A Qualitative Case Study of a Vietnamese Maritime Logistics Firm. *International Journal of Innovative Science and Research Technology*, 11(8), 1718-1725. <https://doi.org/10.38124/ijisrt/26aug789>

I. INTRODUCTION

Maritime logistics firms operate at the intersection of transport, customs, trade, contracting, taxation, safety, environmental requirements, and data exchange. A single shipment may require the coordinated treatment of customer instructions, transport documents, tariff classifications, specialized inspections, contractual obligations, payment records, and communications with public authorities. This density of rules and handoffs makes compliance risk both operational and strategic. Errors may delay clearance, increase storage and demurrage costs, trigger penalties or tax reassessments, weaken customer trust, and consume managerial attention. The problem is intensified when legal requirements change faster than internal guidance, when information remains dispersed across systems, and when quality assurance depends heavily on experienced employees.

AI and regulatory technology (RegTech) offer a plausible response because many compliance tasks involve repeated comparison, classification, retrieval, anomaly detection, and deadline monitoring. Natural-language processing can support the review of contracts and regulatory documents; rule-based and machine-learning tools can flag inconsistencies in operational records; and dashboards can consolidate indicators for managerial oversight (Davenport & Ronanki, 2018). Yet the value of these tools is frequently overstated. Compliance decisions are context-sensitive, data may be incomplete, and model outputs can be inaccurate or difficult to explain. In high-consequence settings, automation without governance may scale errors instead of reducing them (Micheler & Whaley, 2020; National Institute of Standards and Technology [NIST], 2023).

The central managerial issue is therefore not whether AI is technically available, but whether an organization is ready to use it for a defined compliance purpose. Organizational AI readiness concerns the assets, capabilities, knowledge, commitment, and governance arrangements that enable a firm to adopt AI responsibly. Prior research shows that readiness is context- and purpose-specific rather than a one-time threshold (Jöhnk et al., 2021). This observation is especially important for mid-sized firms that have accumulated operational software and digital records but have not yet achieved integrated data architectures or formal AI governance.

This study examined TRA-SAS, a Vietnamese maritime transport and services company, during a period of business expansion and regulatory change. Between 2021 and 2025, the company's net revenue increased from approximately VND 519.96 billion to VND 1,134.43 billion (TRA-SAS, 2022–2026). Growth increased the scale of transactions and the volume of compliance-relevant data, but internal evidence also indicated fragmented data, continuing manual checks, and uneven analytical capability. The case is analytically valuable because it represents a common organizational condition: digital tools are present, practical interest in AI exists, and management supports digital transformation, yet foundational readiness is not uniform.

The study pursued three research questions: (1) Which compliance risks are most material in the case organization? (2) What is the organization's current readiness to apply AI to these risks? (3) Which use cases and implementation sequence are appropriate given its capabilities and governance obligations? The article contributes by integrating risk materiality, digital maturity, and AI readiness into a single decision logic. Instead of beginning with a vendor or model, the approach begins with the compliance-risk portfolio, identifies control gaps, evaluates whether the organization can support a specific use case, and imposes governance gates before deployment. This risk-first pathway extends AI-readiness research into maritime logistics and offers a practical alternative to technology-led adoption.

II. LITERATURE REVIEW AND ANALYTICAL FRAMEWORK

➤ *Compliance Risk in Maritime Logistics*

Compliance risk is the possibility that an organization will incur legal, regulatory, financial, operational, or reputational harm because it fails to meet applicable obligations. ISO 37301:2021 treats compliance as a management system requiring leadership, planning, support, operational control, performance evaluation, and improvement. ISO 31000:2018 complements this view by emphasizing the iterative identification, analysis, evaluation, treatment, monitoring, and communication of risk. COSO's enterprise risk management framework further connects risk with strategy and performance, while the Three Lines Model clarifies the complementary responsibilities of operations, specialist oversight, and independent assurance (Committee of Sponsoring Organizations of the Treadway Commission [COSO], 2017; Institute of Internal Auditors [IIA], 2020;

International Organization for Standardization [ISO], 2018, 2021; Kaplan & Mikes, 2012). In Vietnam, the case also operates under the Law on Customs and the Vietnam Maritime Code (National Assembly of Vietnam, 2014, 2015).

In maritime logistics, compliance does not sit in a single department. Front-line units create and transform compliance-relevant information, documentation and customs teams translate transaction facts into declarations, contract personnel define obligations, finance units process payments and taxes, and information-technology staff protect the digital environment. This interdependence means that a formal policy can coexist with control gaps. A procedure may be documented, for example, while its execution still depends on spreadsheets, unstructured email, or tacit knowledge. Risk assessment must therefore distinguish the existence of a control from its consistency, traceability, and capacity to generate early warning.

➤ *AI, RegTech, and Compliance Management*

RegTech refers to the use of digital technologies to interpret, operationalize, monitor, and demonstrate compliance. Its most ambitious form translates legal or policy requirements into machine-executable logic, but real-world applications often combine rule engines, information retrieval, workflow automation, and human judgment (Arner et al., 2017; Micheler & Whaley, 2020). AI expands this toolkit by extracting text, identifying patterns, ranking anomalies, summarizing regulatory changes, and generating drafts. The relevant unit of analysis, however, should be the business use case rather than "AI" in the abstract.

A document-checking use case can compare mandatory fields across invoices, bills of lading, declarations, and internal records. Contract analytics can identify clauses, duties, renewal dates, liability limits, and deviations from approved templates. Regulatory intelligence can retrieve new instruments, compare versions, and route alerts to accountable owners. Early-warning models can combine operational and historical indicators to flag unusual cases, while compliance dashboards can consolidate exceptions and response times. These applications differ in data requirements, error consequences, explainability needs, and acceptable levels of automation. A low-risk summarization aid may require different controls from a tool that recommends a customs classification.

AI also creates new risks. Training or prompt data may contain confidential or personal information; models may produce plausible but incorrect outputs; third-party systems can change without notice; and excessive reliance can erode professional judgment. Responsible deployment therefore requires traceability, access control, data minimization, testing, escalation, human review, monitoring, and incident response (Organisation for Economic Co-operation and Development [OECD], 2019). NIST AI RMF 1.0 organizes these responsibilities through Govern, Map, Measure, and Manage. ISO/IEC 42001:2023 embeds similar concerns in an AI management system, while ISO/IEC 23894:2023 provides guidance for AI risk management. These frameworks support a central principle for compliance applications: the

organization remains accountable even when a model assists the task.

➤ *Organizational AI Readiness*

Organizational readiness is not equivalent to having software or a budget. Jöhnk et al. (2021) identify strategic alignment, resources, knowledge, culture, and data as interacting categories and argue that readiness must be assessed continuously against the intended adoption purpose. This study translated that insight into eight case-relevant dimensions: data, technology, people, processes, finance, leadership, compliance culture, and security/legal governance. The dimensions also reflect the technology–organization–environment logic of innovation adoption (Tornatzky & Fleischer, 1990), while adding use-case governance demanded by contemporary AI-risk frameworks.

A maturity label should not conceal unevenness. An organization may have operational applications near a defined or managed level while its data integration remains at an earlier stage. The lowest enabling capability can become the binding constraint for a specific use case. For example, a capable contract model cannot reliably track obligations if contract versions, templates, responsible owners, and completion evidence are not standardized. This study therefore treats readiness as a profile, not a single score, and uses maturity levels descriptively rather than as a statistically validated scale.

➤ *Integrated Analytical Framework*

The analytical framework follows five linked decisions. First, identify and rate material compliance risks. Second, diagnose gaps in current controls. Third, assess purpose-specific readiness across the eight dimensions. Fourth, match risks and control gaps to bounded AI use cases. Fifth, apply governance and implementation gates before scaling. A use case proceeds only when the data source, process owner, human reviewer, performance metric, error threshold, access rule, and fallback procedure are defined. This ordering separates three concepts that are often conflated: risk level (business exposure), AI suitability (the extent to which AI can assist), and AI readiness (the organization's capacity to deploy and govern the tool).

III. METHODOLOGY

➤ *Research Design and Case Selection*

The research used an applied qualitative single-case design. A case-study strategy was suitable because the phenomenon—the readiness to deploy AI for compliance risk management—could not be separated from the organization's processes, data environment, regulatory setting, and managerial practices (Yin, 2018). TRA-SAS was selected as an information-rich case with diverse logistics and maritime-service activities, an established operational history, available documentary evidence, and an explicit interest in digital transformation. The unit of analysis was the organization's compliance-risk management system and its readiness for bounded AI-assisted use cases, not the attitudes of individual employees in isolation.

➤ *Data Sources and Participants*

Evidence was assembled from four sources. First, three semi-structured interview sessions involved eight participants representing leadership, operations, documentation/customs, finance-accounting-human resources, and IT/digital-transformation perspectives. The sample was purposive and sought functional coverage of the workflows in which compliance risks emerge or are controlled. Second, internal process materials and compliance-related records were reviewed where available. Third, public corporate and financial reports covering 2021–2025 were used to establish organizational context. Fourth, the analysis considered applicable legal and professional frameworks.

Interview questions addressed major compliance risks, existing controls, recurrent process weaknesses, data availability, digital maturity, perceptions of feasible AI use cases, implementation barriers, and governance conditions. Findings are therefore presented as cross-source thematic syntheses. The study does not claim theoretical saturation; it claims role-based coverage and triangulation within the focal case.

➤ *Data Analysis*

The analysis proceeded in four stages, following a thematic logic (Braun & Clarke, 2006). First, interview notes and documentary evidence were organized into meaning units concerning risk sources, consequences, current controls, process weaknesses, data conditions, technology conditions, human capabilities, and implementation requirements. Second, these units were coded into eight compliance-risk categories (CR01–CR08) and eight AI-readiness dimensions. Third, the risks were reviewed through a 5 × 5 likelihood–impact matrix. Risk exposure was interpreted separately from AI suitability: a risk could be highly material but poorly suited to automation, or less material but useful as a safe pilot. Fourth, themes and risk ratings were triangulated against process evidence, business context, and international governance frameworks.

The maturity assessment used a five-level descriptive logic ranging from largely ad hoc/manual (M1) to optimized and continuously improved (M5). It was not treated as a psychometric instrument. Evidence was examined for consistency across sources, and negative or cautionary observations—such as data fragmentation, uneven skills, and security concerns—were retained rather than averaged away. This approach improved confirmability while preserving the contextual nature of the case.

IV. FINDINGS

➤ *Compliance-Risk Portfolio*

The case evidence identified eight interrelated compliance risks. Seven were rated High and one Medium under the harmonized 5 × 5 matrix. Customs compliance (CR04) received likelihood 3 and impact 5, producing a score of 15 and the highest implementation priority within the portfolio. Under the adopted thresholds, this is High—not Very High. This correction matters because an inflated

category can distort resource allocation and weaken methodological credibility.

Table 1. Compliance-Risk Portfolio and Candidate AI Support

Code	Risk	L	I	Score / level	Candidate AI support
CR01	Complex legal and sectoral requirements	4	4	16 – High	Regulatory search and synthesis
CR02	Regulatory change	4	4	16 – High	Change detection and routed alerts
CR03	Interpretation and application difficulty	4	4	16 – High	Retrieval, comparison, expert support
CR04	Customs compliance	3	5	15 – High	Cross-document checks and flags
CR05	Specialized inspection requirements	3	4	12 – High	Completeness and requirement matching
CR06	Service-contract compliance	3	4	12 – High	Clause extraction and obligation tracking
CR07	Information security	3	3	9 – Medium	Anomaly support; strict data governance
CR08	Operational process compliance	4	4	16 – High	Workflow monitoring and anomaly detection

Source: Authors' synthesis from interview and documentary evidence.

The scores do not imply that CR01, CR02, CR03, or CR08 should automatically be automated before CR04. Priority was determined by combining materiality, transaction frequency, data availability, rule clarity, controllability of model output, and pilot feasibility. CR04 was considered especially consequential because errors can delay clearance and create direct cost or sanction exposure. Yet the recommended AI role remains assistive: pre-checking, comparison, and exception flagging before an accountable employee validates the declaration.

➤ Existing Controls and Persistent Gaps

TRA-SAS had several favorable control foundations. Responsibilities were allocated across relevant functions; employees possessed substantial operational experience; regulations were monitored; and documents were reviewed before external release. Managers also adjusted procedures when requirements changed and used multiple levels of checking for selected tasks. These practices explain why the case should not be characterized as uncontrolled or digitally immature in absolute terms.

Four persistent gaps limited proactive compliance management. First, risk identification and interpretation relied heavily on individual experience, creating variation when experienced staff were unavailable. Second, data were digitized but distributed among operational systems, files, and communication channels, which reduced a complete and timely view of risk. Third, several checks remained manual, making consistency and scalability difficult as transaction volumes grew. Fourth, incident, exception, and corrective-action information was not yet consolidated into a sufficiently structured dataset for systematic trend analysis and early warning. The organization therefore possessed controls, but not a fully integrated evidence loop from risk event to corrective learning.

➤ Digital Maturity and AI-Readiness Profile

The overall digital-maturity finding was an uneven M2–M3 profile. Software availability, selected digitized workflows, and initial AI experimentation were near M3. Data standardization and integration were closer to M2. Reporting and automated early warning also lagged behind the existence of operational software.

Table 2. Purpose-Specific AI-Readiness Profile

Dimension	Observed condition	Profile	Immediate requirement
Data	Digitized but fragmented; inconsistent structures and ownership	M2	Data dictionary, master data, quality rules, lineage
Technology	Operational software; partial integration; early AI exposure	M3-	Secure integration, logging, test environment
People	Strong domain expertise; uneven AI/data literacy	M2–M3	Role-based training and cross-functional team
Processes	Defined and partly software-supported; manual handoffs	M3-	SOPs, decision and escalation points

Dimension	Observed condition	Profile	Immediate requirement
Finance	Capacity for selective investment; benefits untested	M2–M3	Stage-gated pilots and total-cost review
Leadership	Visible interest in digital transformation	M3	Sponsor, risk appetite, decision rights
Compliance culture	Control awareness and multi-level checking	M3	Consistent issue logging and learning loop
Security/legal	Awareness exists; AI governance immature	M2	Data classification, approved tools, impact assessment

Source: Authors' synthesis from interview and documentary evidence.

Leadership commitment and domain expertise were the strongest enablers. Data fragmentation and AI-specific governance were the principal gating constraints. These findings show why procurement alone would be insufficient: the organization could acquire a capable model while remaining unable to supply controlled data, document the basis of outputs, or assign responsibility when the tool is wrong.

➤ Priority AI Use Cases

Table 3. Prioritized AI-Assisted Compliance Use Cases

Rank	Use case	Why suitable	Minimum viable output	Human-in-the-loop control
1	Document pre-checking	High-volume, rules explicit, outputs reviewable	Missing fields and cross-document inconsistency	Human validates exceptions and final filing
2	Contract and obligation analysis	Material duties and deadlines; text-intensive	Clause extraction, deviations, obligation calendar	Legal/contract owner confirms interpretation
3	Regulatory change intelligence	High update burden and dispersed sources	Version comparison, topic tags, owner alerts	Specialist validates applicability
4	Compliance early warning	Consolidates recurring exceptions and deadlines	Indicators, thresholds, anomaly flags	Manager investigates; no automatic sanction
5	Compliance reporting/dashboard	Improves visibility and audit trail	Risk status, overdue actions, trends	Owners certify data and corrective actions

Source: Authors' synthesis from interview and documentary evidence.

Document checking was the preferred first pilot because it combines high frequency, comparatively clear validation rules, measurable processing time, and direct human review. Contract analytics and regulatory intelligence could follow once source repositories and ownership rules are stable. Early-warning analytics should not precede disciplined issue logging: without a structured history of errors, near misses, deadlines, causes, and corrective actions, a model would have little reliable signal. Dashboards should visualize governed data rather than create a polished interface over inconsistent records.

➤ Implementation Roadmap

Table 4. Phased Roadmap for 2026–2030

Period	Stage	Core actions	Decision evidence
2026	Foundation	Standardize data and risk taxonomy; approve AI policy; assign owners	Data dictionary; risk register; approved-tool list; baseline KPIs
2027	Controlled pilot	Pilot document checks and selected contract review	Pilot report; error taxonomy; reviewer logs; go/no-go decision

Period	Stage	Core actions	Decision evidence
2028	Integration and expansion	Connect approved sources; extend alerts and obligation tracking	Integrated workflow; access controls; monitored use cases
2029	Optimization	Tune thresholds; assess costs, quality, adoption, residual risk	Validated KPIs; model-change process; assurance review
2030	Institutionalize	Embed continuous monitoring and governance	Management review; audit trail; portfolio refresh

Source: Authors' synthesis from interview and documentary evidence.

The roadmap uses gates rather than calendar dates alone. Progress to the next stage requires evidence that the preceding controls work. A pilot should therefore have a defined owner, bounded dataset, test cases, reviewer instructions, false-positive and false-negative measures, response-time targets, incident procedures, and a manual fallback. The 2026–2030 dates express strategic sequencing; they should be updated against actual implementation status before managerial adoption.

V. DISCUSSION

➤ *A Risk-First and Readiness-Gated Adoption Logic*

The findings support a risk-first interpretation of AI readiness. Conventional adoption discussions often begin with technological potential and then ask where the tool can be used. The present case reverses that sequence. Material compliance risks and control gaps define the problem; purpose-specific readiness determines whether AI is an appropriate response; governance defines the boundary of acceptable use. This sequence reduces solutionism and clarifies that some gaps—such as ambiguous process ownership—must be solved managerially before they can be assisted technologically.

The case also confirms that readiness is uneven and dynamic. TRA-SAS did not fit a binary ready/not-ready category. It had leadership support, domain expertise, and operational software, yet it lacked uniformly governed data and mature AI controls. This is consistent with Jöhnk et al.'s (2021) argument that readiness depends on context and adoption purpose and must be revisited through the adoption process. A firm may be ready for a narrow internal document-checking pilot but not ready for autonomous customer-facing advice or automated customs decisions.

➤ *The Weakest-Enabling-Layer Proposition*

A second insight is that the lowest critical enabling layer constrains use-case readiness. Aggregate maturity scores can hide this bottleneck. For document checking, the relevant layers include document quality, identifier consistency, process rules, exception ownership, and reviewer capacity. If any indispensable layer remains below the minimum threshold, the use case should be redesigned or delayed. This suggests a weakest-enabling-layer proposition: for a bounded compliance use case, deployment readiness is determined less

by the average maturity of the organization than by the least mature capability that is necessary for safe operation.

This proposition has practical implications for investment. Raising an already-adequate technology layer may yield less value than improving data definitions, issue logging, or reviewer training. It also explains why small pilots can be useful when carefully bounded: a pilot reveals which enabling layer is actually binding and can build experience without exposing the entire compliance system.

➤ *Human Oversight as an Operating Design*

Participants consistently viewed AI as a support mechanism rather than a replacement for accountable expertise. Human-in-the-loop control should not be reduced to a generic statement that “a person remains responsible.” It must be designed into the workflow: who reviews which output, what evidence is displayed, when escalation is mandatory, how disagreement is recorded, and how model errors improve future controls. The Three Lines Model offers a practical allocation. Operations own the decision and first-line validation; compliance or legal specialists establish standards and monitor use; internal assurance independently evaluates design and effectiveness (IIA, 2020).

This design aligns with NIST's Govern–Map–Measure–Manage cycle and ISO/IEC 42001. It is also consistent with the direction of contemporary AI regulation, which emphasizes documented responsibilities, risk management, human oversight, and data governance (European Parliament & Council of the European Union, 2024). For a Vietnamese firm, the governance baseline must additionally reflect the Law on Artificial Intelligence No. 134/2025/QH15, Decree No. 142/2026/ND-CP, the Law on Personal Data Protection No. 91/2025/QH15, and Decree No. 356/2025/ND-CP (Government of Vietnam, 2025, 2026; National Assembly of Vietnam, 2025a, 2025b). Legal applicability should be assessed for each system and data flow rather than assumed from the general use of AI.

➤ *Theoretical Contribution*

The article contributes to three streams. First, it adds a maritime-logistics case to the organizational AI-readiness literature, which has largely developed through cross-industry expert studies or technology-adoption research. Second, it distinguishes risk materiality, AI suitability, and AI readiness, thereby preventing a high risk score from being

interpreted automatically as a mandate for automation. Third, it links readiness to compliance-management architecture. Data, processes, people, and technology are necessary, but for high-consequence use cases they become operationally meaningful only when embedded in decision rights, traceability, and assurance.

VI. MANAGERIAL AND POLICY IMPLICATIONS

➤ *Managerial Implications*

Managers should establish a compliance AI portfolio rather than launch disconnected experiments. Each candidate should have a one-page use-case charter describing the risk addressed, user, data, decision supported, prohibited uses, human reviewer, success measures, cost, and residual risk. Portfolio review should favor cases with material value, bounded scope, available data, reviewable outputs, and reversible implementation. Document pre-checking meets these conditions more readily than autonomous legal interpretation.

Data work should be treated as a control program, not an IT cleanup. TRA-SAS should define authoritative sources, common identifiers, document versions, retention rules, access permissions, and data-quality responsibilities. An exception register should record the event, process, date, cause, consequence, owner, corrective action, and closure evidence. This creates the audit trail needed for compliance learning and, later, for responsible analytics.

Pilot evaluation should use operational and risk metrics together. Relevant indicators include document-processing time, reviewer effort, precision of flagged exceptions, false-negative rate on a controlled test set, overdue obligations, repeated error rates, user overrides, security incidents, and cost per processed case. Accuracy alone is insufficient because a model can appear accurate while missing rare but severe cases. Management should define tolerances based on consequence and maintain a manual fallback until the system has demonstrated stable performance.

Workforce development should combine compliance expertise with AI literacy. Training should cover appropriate system use, verification of sources, confidentiality, bias and error, escalation, and documentation. Cross-functional design teams should include operations, compliance/legal, IT/security, and intended users. Participation reduces resistance and improves process fit because the people closest to the work can identify exceptions that a generic technical specification may overlook.

➤ *Governance and Policy Implications*

An internal AI policy should define approved tools, data classifications, prohibited inputs, access and logging requirements, procurement due diligence, testing, human oversight, incident reporting, third-party responsibilities, model changes, and retirement. High-consequence compliance uses should undergo an impact assessment before deployment. Vendor contracts should address confidentiality, data location and reuse, security controls, subcontractors, service changes, audit rights, and exit arrangements. These

requirements translate general responsible-AI principles into enforceable operating controls.

For policymakers and industry associations, the case highlights a need for sector-specific implementation guidance. General AI and personal-data rules establish essential obligations, but logistics firms also need practical mappings to customs, transport documents, commercial confidentiality, cross-border data flows, and specialized inspection processes. Shared taxonomies, anonymized test cases, and guidance for human review could lower the cost of responsible adoption without weakening accountability.

VII. CONCLUSION

This study examined how a Vietnamese maritime logistics firm could prepare for AI-assisted compliance risk management. Eight material risks were identified, with customs compliance treated as the highest priority within the harmonized matrix (score 15, High). TRA-SAS showed an uneven M2–M3 digital-maturity profile: operational software and selected processes were relatively developed, while data standardization, integration, and AI-specific governance remained less mature. The most feasible initial use cases were document pre-checking, contract and obligation analysis, regulatory change intelligence, early warning, and compliance reporting.

The core conclusion is that AI readiness is neither a procurement decision nor a single maturity score. It is a purpose-specific configuration of data, technology, people, process, resources, leadership, compliance culture, and security/legal governance. For compliance work, this configuration must be joined to explicit human oversight and auditable decision rights. A phased 2026–2030 roadmap is therefore appropriate, beginning with data and governance foundations, moving to bounded internal pilots, and expanding only when evidence supports the next gate.

The study has limitations. It is a single organizational case with three interview sessions and eight participants; it does not provide statistical generalization, saturation claims, independent inter-coder reliability, or realized outcomes from a deployed AI system. Some evidence was available only in synthesized form rather than as a complete publication-ready transcript audit trail. Future research should test the framework across logistics firms, validate the readiness dimensions quantitatively, compare use cases by error consequence, and evaluate pilot performance longitudinally. Research should also examine how emerging Vietnamese AI and personal-data regulation changes the governance costs and adoption choices of medium-sized firms.

REFERENCES

- [1]. Arner, D. W., Barberis, J., & Buckley, R. P. (2017). FinTech, RegTech, and the role of compliance in 2020. *Northwestern Journal of International Law & Business*, 37(3), 371–413.

- [2]. Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77–101. <https://doi.org/10.1191/1478088706qp063oa>
- [3]. Committee of Sponsoring Organizations of the Treadway Commission. (2017). *Enterprise risk management: Integrating with strategy and performance*.
- [4]. Davenport, T. H., & Ronanki, R. (2018). Artificial intelligence for the real world. *Harvard Business Review*, 96(1), 108–116.
- [5]. European Parliament & Council of the European Union. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence. *Official Journal of the European Union*.
- [6]. International Organization for Standardization. (2018). *ISO 31000:2018 risk management—Guidelines*.
- [7]. International Organization for Standardization. (2021). *ISO 37301:2021 compliance management systems—Requirements with guidance for use*.
- [8]. International Organization for Standardization & International Electrotechnical Commission. (2023a). *ISO/IEC 23894:2023 information technology—Artificial intelligence—Guidance on risk management*.
- [9]. International Organization for Standardization & International Electrotechnical Commission. (2023b). *ISO/IEC 42001:2023 information technology—Artificial intelligence—Management system*.
- [10]. Institute of Internal Auditors. (2020). *The IIA's Three Lines Model: An update of the Three Lines of Defense*.
- [11]. Jöhnik, J., Weißert, M., & Wyrski, K. (2021). Ready or not, AI comes—An interview study of organizational AI readiness factors. *Business & Information Systems Engineering*, 63(1), 5–20. <https://doi.org/10.1007/s12599-020-00676-7>
- [12]. Kaplan, R. S., & Mikes, A. (2012). Managing risks: A new framework. *Harvard Business Review*, 90(6), 48–60.
- [13]. Micheler, E., & Whaley, A. (2020). Regulatory technology: Replacing law with computer code. *European Business Organization Law Review*, 21(2), 349–377. <https://doi.org/10.1007/s40804-019-00151-1>
- [14]. National Assembly of Vietnam. (2014). *Law on Customs No. 54/2014/QH13*.
- [15]. National Assembly of Vietnam. (2015). *Vietnam Maritime Code No. 95/2015/QH13*.
- [16]. National Assembly of Vietnam. (2025a). *Law on Artificial Intelligence No. 134/2025/QH15*.
- [17]. National Assembly of Vietnam. (2025b). *Law on Personal Data Protection No. 91/2025/QH15*.
- [18]. National Institute of Standards and Technology. (2023). *Artificial intelligence risk management framework (AI RMF 1.0) (NIST AI 100-1)*. <https://doi.org/10.6028/NIST.AI.100-1>
- [19]. Organisation for Economic Co-operation and Development. (2019). *Recommendation of the Council on artificial intelligence*.
- [20]. Government of Vietnam. (2025). *Decree No. 356/2025/ND-CP detailing the Law on Personal Data Protection*.
- [21]. Government of Vietnam. (2026). *Decree No. 142/2026/ND-CP detailing and implementing the Law on Artificial Intelligence*.
- [22]. Tornatzky, L. G., & Fleischer, M. (1990). *The processes of technological innovation*. Lexington Books.
- [23]. TRA-SAS. (2022–2026). *Financial statements and annual reports, 2021–2025*.
- [24]. Yin, R. K. (2018). *Case study research and applications: Design and methods (6th ed.)*. SAGE.