Aperia

ICM
Powered by Tapestrii

- AI IN LOGISTICS & TRANSPORTATION SECURITY -

The *Execution* Gap

AI meets reality: The gap isn't detection. It's response.

tapestriisupport@aperia.com
www.aperia.com/platforms/tapestrii-icm

The Wrong Starting Point

Most AI pitches begin with what the model can do. The real question comes earlier.

A suspicious cargo manifest was flagged. The alert was logged. The investigator on shift had seven other open items. The cross-reference with a prior incident from three months ago was never made. The shipment cleared.

No one failed. The process failed. Or more precisely: there was no process that could have succeeded, because the information that would have connected those two events lived in a different system, managed by a different team, documented in a different format. The analyst would need to know to look. There was nothing in the workflow to tell them to.

It takes a particular kind of environment for this pattern to repeat: a port with fragmented systems and a security team under pressure. That describes most ports in Latin America.

The more common failure is not detection. It is process execution. Alerts logged but not escalated. Incidents documented but not linked. Investigations stalled because no one had visibility into their status. Evidence never attached because attaching it was a manual step at the end of a long shift.

The more common failure is not detection. It is process execution. Alerts logged but not escalated. Incidents documented but not linked. Investigations stalled because no one had visibility into their status. Evidence never attached because attaching it was a manual step at the end of a long shift.

In port security, the question AI should help answer is not: what did we miss? It is: why did the process that should have caught this not run the way it was designed?

Those are different questions. They lead to different answers. The second one is the one worth asking.

CENTRAL ARGUMENT: ***Port security doesn't fail because security teams can't see the threat. It fails because the process that should have connected the signals didn't run. AI deployed into that environment doesn't fix the process. It adds a layer on top of a structure that was already broken.***

The Systems Problem

When the foundation is fragmented, the failures become predictable.

Port security operations often run on systems that were not designed to work together. Cargo management. Access control. CCTV. Incident logs. HR records. Each one holds a piece of the picture. None of them share it. The result is an environment where meaningful signals exist but have

no reliable path to the people or processes that would act on them. Integration efforts can close the gaps between systems. They do not, on their own, reveal how information is shared, where it slows down, or what gets lost between hand-offs.

The end-of-shift handover is the clearest example. What gets preserved at the end of a shift is whatever someone had time to write down. Go back two months and ask what the team knew about a specific vessel or a specific individual at the time of a flagged incident; the honest answer, in most operations, is that you cannot reconstruct it. The information existed. It was never captured in a form that survives the handover.

It is a systems problem. When there is not a structured environment to capture operational knowledge, the knowledge does not persist. The next shift starts from scratch. The next investigation starts without context. The next incident follows patterns that were documented somewhere but are not findable.

The failure modes that result are predictable:

- **Alert fatigue.** High volume, no structured triage. Analysts prioritize based on current workload. Some signals wait. The ones that wait are not always the low-risk ones.
- **Inconsistent triage.** The same signal, handled by two different analysts on two different shifts, produces different outcomes. Not because either made a mistake but because there was no enforced process to produce consistent ones.
- **Context fragmentation.** An analyst handling an incident today has no efficient path to what a colleague documented about the same entity six weeks ago. Finding it requires knowing where to look and having time to look.
- **Escalation gaps.** Time-sensitive signals surface when an analyst notices them, not when protocol requires action. The difference between those two moments is where exposure accumulates.
- **Audit gaps.** Evidence is attached when analysts attach it. Decision trails are reconstructed after the fact. That reconstruction is expensive, incomplete, and usually prompted by an incident that has already cost someone money.

Add AI to this environment and you have added a capability on top of a fragmentation problem. Automated alerts can fire. Escalation rules can trigger. But the underlying workflow is still disconnected. The gaps do not close, they move.

CENTRAL ARGUMENT: ***The workflow is the intervention. AI deployed into a fragmented environment inherits its fragmentation. Before you can automate, you must unify.***

The Workflow Layer

Standardization comes before enforcement.

The workflow is where technology and people meet every day. Before enforcement is possible, it must be understood; not assumed and not inferred from the systems that currently hold the data. That requires mapping how information moves through the operation: where it stalls, transfers, or drops. These are questions that determine whether AI can create value. Without answers, the platform still runs, but much of its enforcement value is left unrealized.

The platform is built around that understanding. AI is woven into the workflow, not bolted onto the data. Bring investigations, alerts, evidence, entities, and case history into a single operational environment. Then enforce the process within that structure. Once that is in place, AI does something it cannot do in a fragmented environment: it runs the process consistently, at scale, regardless of who is on shift.

Four failure points, and what changes when the workflow is in place:

Failure Point	Without Workflow Unity	What The Platform Enforces
Alert Triage	Analyst prioritizes ad hoc based on workload. High-volume environments mean some signals wait unreviewed.	Every alert follows the same triage sequence, cross-referenced against existing profiles before human review.
Context Retrieval	Analyst searches prior cases if they know where to look and have time to do it. Institutional memory depends on individual effort.	Incoming signals automatically enriched with platform context: related entities, prior incidents, and open investigations, before the analyst opens the record.
Escalation Timing	Escalation happens when an analyst notices it is necessary. Time-sensitive items surface when noticed, not when protocol requires action.	Defined escalation rules trigger automatically. Decision-makers receive notifications on schedule.
Audit Trail	Evidence is attached when analysts attach it. Case documentation reflects what was logged, not necessarily what occurred.	Workflow steps are recorded as they execute. The audit trail is a byproduct of the process, not a separate task.

The budget argument follows from this directly. Not that AI will catch what humans miss, but that AI ensures the process runs every time. Those are different claims. The second one holds up under scrutiny.

Where Friction Exists

The theory is straightforward. Knowing where to address it first is the harder question.

Unify the workflow, then enforce the process within it. That logic is not complicated. What is harder is knowing where to address it first, because in a functioning port security operation, everything feels like a priority, and nothing feels like it can be paused long enough to fix.
The right place to start is wherever the friction is highest. Where is work piling up? Where are experienced analysts spending time on coordination and follow-up instead of judgment? Where are approvals sitting because nobody has visibility into their status? Where are handoffs dropping because the next person in the chain has no automatic way to know what the previous one knew?

Unresolved friction is not just inefficiency. It is exposure. It determines where signals are delayed, where decisions stall, and where risk accumulates unnoticed.

AI deployed into operational reality behaves differently from AI deployed into a clean environment. It does not just run the process. It absorbs the operational friction that currently falls on the people carrying the most, and it moves that friction to the platform.

The analyst who spent time reconstructing context before opening an investigation now opens it already oriented. The supervisor who got notified when someone remembered to notify them now gets notified when the process requires it. The coordination overhead that used to live in senior staff moves to a workflow step. They do not disappear from the process. They step up a level.

At scale, something else follows. Every investigation, every alert, every entity connection gets captured consistently, in one environment, over time. That is institutional memory a fragmented operation cannot build. Not reports. Not logs. A continuously evolving intelligence layer where patterns, entities, and behaviors compound across cases, across time, and across the team. This is where advantage forms, not in individual investigations, but in the system's ability to learn across them.

In an environment where organized criminal networks time their entries to the gaps in process execution, that compounding intelligence is not a secondary benefit. It is the long argument for addressing the workflow layer in the first place.

Three Questions Worth Asking

Not AI questions. Operational questions. The answers tell you where your actual exposure is.

If you run port security operations and you are evaluating AI, these are the questions that matter. They are not questions about technology. They are questions about process maturity, and the answers tell you whether the foundational work has been done, or where it has not.

One. Do you have a reliable way to confirm that every alert is triaged according to the same process, regardless of who is on shift? If the answer is no, or 'we think so', process inconsistency is your most exploitable gap. It is also the one most directly accessible to platform enforcement.

Two. When an incident occurs, can you reconstruct within 24 hours what was known, by whom, and what action was taken? If not, you have an audit gap. That is a governance risk and an intelligence loss. The next incident that follows the same pattern will likely follow it again.

Three. Do investigators working on related cases have automatic visibility into each other's work, or does cross-case connection depend on individual initiative? If the latter, institutional memory is leaking out of your operation every time a case closes.

A well-designed platform addresses all three. But only after the workflow is unified enough to give it something stable to enforce.

Operational intelligence platforms purpose-built for port security, such as ICM by Tapestrii, are built around this sequence. The platform does not add AI to the existing process. It builds the unified environment first, then enforces the process within it. Investigation management, automated triage, AI-assisted search, and audit trail generation are not features layered on top. They are the workflow.

About the AI in Port Security Series

This is the first paper in a three-part series on artificial intelligence in logistics and transportation security. Published May 2026. The series runs in parallel with Ports Under Pressure, a four-paper analysis of structural security challenges facing Latin American container ports. If you would like to discuss how these issues are affecting your port or terminal, please contact us at: tapestriisupport@aperia.com

About Aperia

Aperia delivers ICM by Tapestrii, an intelligent case management platform that helps terminal operators and port authorities connect fragmented incidents, entities, and operational data to detect emerging risks earlier and strengthen decision-making in complex, high-stakes environments. We welcome dialogue with port operators, terminal managers, and security leaders committed to staying ahead of these dynamics. Learn more at https://www.aperia.com/platforms/tapestrii-icm

Bill Biddy is the Senior Vice President of Technology, with a background in enterprise technology and cloud architecture, leading the company's technology strategy and systems.

Aaron Coats is a Senior Product Manager, with a background in enterprise technology and national security, leading go-to-market strategy for ICM by Tapestrii.