Cyber-Physical Convergence in Critical Logistics Infrastructure: Threat Patterns and Detection Methods

Michael A. Torres

PLIANT Institute, Palo Alto, CA

International Journal of Critical Infrastructure Protection, Vol. 19, No. 3 (September 2025), pp. 44–79

Abstract

The convergence of cyber and physical security domains in critical logistics infrastructure has created a threat landscape that neither traditional cybersecurity frameworks nor physical security protocols are individually adequate to address. This paper examines the patterns of cyber-physical attack targeting logistics infrastructure, with particular attention to the detection challenges that arise when digital intrusions are designed to produce physical disruption outcomes. Drawing on a dataset of 67 confirmed or suspected cyber-physical incidents affecting logistics facilities between 2018 and 2024, we identify four primary attack patterns, characterize their detection signatures, and assess the adequacy of current detection methodologies across each pattern type. We find that detection rates vary substantially across attack patterns, with intrusions designed to manipulate rather than disable terminal operating systems presenting the lowest detection rates and the longest mean dwell times. Recommendations for improving detection capability across pattern types are provided.

Keywords: cyber-physical attacks, logistics infrastructure, terminal operating systems, threat detection, supply chain security, critical infrastructure

---

1. Introduction

Critical logistics infrastructure — ports, terminals, intermodal facilities, and the digital systems that connect them — occupies a position of increasing strategic significance in the contemporary threat landscape. The accelerating digitization of terminal operations, cargo tracking, and customs processing has created an expanding attack surface that sophisticated threat actors have demonstrated both the capability and the willingness to exploit. The consequences of successful attacks extend well beyond the directly affected facility, propagating through interconnected logistics networks in ways that can impose supply chain disruption costs orders of magnitude larger than the direct impact at the point of attack.

The analytical challenge this creates is compounded by the convergence of cyber and physical security domains that characterizes modern logistics infrastructure. Terminal operating systems control physical processes — crane movements, berth assignments, gate operations — through digital interfaces. Cargo management platforms integrate with customs systems, shipping line databases, and freight forwarder portals through APIs that cross organizational boundaries. Port community systems aggregate data from dozens of actors across a single port's operational ecosystem. In this environment, the distinction between a cyber incident and a physical disruption event is increasingly artificial — digital intrusions produce physical consequences, and physical disruption events generate digital artifacts that can support or complicate cyber incident response.

Despite this convergence, the institutional and analytical frameworks governing cybersecurity and physical security in logistics facilities have remained largely separate. Cybersecurity teams focus on network intrusion

detection, malware analysis, and system hardening; physical security teams focus on access control, CCTV monitoring, and incident response to physical disruption events. The gap between these functions creates detection blind spots that sophisticated adversaries have learned to exploit — designing intrusions that fall below the detection threshold of either framework individually while producing measurable disruption outcomes.

2. Dataset and Methodology

The incident dataset underlying this analysis was compiled from four primary sources: open-source reporting in trade and technical press; public regulatory filings and government disclosures; PLIANT Atlas database records incorporating partner-shared intelligence; and direct engagement with affected facility operators conducted under confidentiality arrangements. The resulting dataset of 67 incidents covers the period January 2018 through December 2024 and spans maritime terminals, intermodal facilities, and airport cargo operations across 23 countries.

Incident classification followed a structured protocol based on the MITRE ATT&CK; for ICS framework, adapted for logistics-specific operational contexts. Each incident was classified along three dimensions: attack vector (the technical method of initial access and payload delivery), intended effect (the physical or operational outcome the attacker sought to produce), and detection method (how the incident came to the attention of the affected organization or external authorities). Attribution confidence was assessed using the framework developed in Torres (2024), with High and Medium confidence attributions disaggregated by threat actor category in the analysis.

2.1 Dataset Limitations

The dataset is subject to several limitations that constrain the generalizability of findings. First, confirmed cyber-physical incidents are substantially underreported relative to their actual frequency — affected organizations face strong commercial and reputational incentives not to disclose incidents, and regulatory disclosure requirements are inconsistent across jurisdictions. The dataset therefore overrepresents high-severity incidents that attracted external attention and underrepresents lower-severity incidents that were contained without public disclosure. Second, attribution confidence varies substantially across incidents, and findings disaggregated by threat actor category should be interpreted with corresponding caution.

3. Four Primary Attack Patterns

3.1 Pattern A: Disruptive Ransomware

Disruptive ransomware operations targeting terminal operating systems represent the most frequently documented attack pattern in the dataset (n=24, 36%). These operations deploy ransomware payloads that encrypt TOS data and demand payment for decryption keys, producing immediate operational disruption as terminal managers revert to manual processes. The 2017 NotPetya attack on Maersk, which disrupted 17 terminal facilities globally and is estimated to have cost $300 million in direct losses, remains the most consequential documented example.

Detection rates for Pattern A are relatively high (71% detected within 24 hours) because the operational impact — sudden TOS unavailability — is immediately apparent to terminal managers. Mean dwell time prior to payload execution is 14 days across the dataset, during which the attacker establishes persistence and maps the target environment without triggering anomaly detection. This pre-execution dwell period represents the primary detection opportunity that current frameworks consistently miss.

3.2 Pattern B: Data Exfiltration with Disruption Cover

Pattern B incidents involve cyber intrusions primarily designed for intelligence collection — exfiltrating cargo manifests, shipping schedules, customs data, or operator credentials — that use a disruptive payload as cover for the primary operation or to complicate forensic investigation (n=18, 27%). These operations are more difficult to attribute than Pattern A incidents because the disruptive element attracts investigative attention that may be deliberately misdirected away from the exfiltration activity.

3.3 Pattern C: Manipulation without Disruption

Pattern C incidents — intrusions designed to manipulate operational data or processes without producing immediately detectable disruption — are the most analytically significant category in the dataset despite representing a smaller share of documented incidents (n=14, 21%). These operations modify cargo manifests, berth assignments, or gate release authorizations in ways that produce operational consequences — sanctioned cargo reaching its destination, priority shipments delayed, specific vessels receiving preferential treatment — without triggering the disruption-level alerts that would initiate incident response.

Mean dwell time for Pattern C incidents is 47 days, nearly three times the Pattern A figure, reflecting both the difficulty of detection and the operational value of sustained access for manipulation purposes. Detection rate is the lowest of any pattern (38%), and in most documented cases detection was achieved through external notification rather than internal monitoring — suggesting that current detection frameworks are not calibrated to identify low-and-slow manipulation operations in logistics environments.

3.4 Pattern D: Physical Access Exploitation

Pattern D incidents involve the exploitation of physical access to logistics facilities to install hardware implants, access air-gapped systems, or conduct reconnaissance that supports subsequent cyber operations (n=11, 16%). This pattern represents the most direct expression of cyber-physical convergence — using physical presence to overcome digital security controls — and is the most difficult to detect through purely technical monitoring.

4. Detection Adequacy Assessment

Across all four pattern types, detection capability is most adequate for high-impact, immediately disruptive incidents (Pattern A) and least adequate for low-and-slow manipulation operations (Pattern C). This asymmetry has strategic implications: sophisticated threat actors with the capability to conduct Pattern C operations have strong incentives to do so, as the lower attribution risk and extended dwell time provide substantially more operational value than the higher-impact but more detectable Pattern A operations.

Current detection frameworks in most logistics facilities are oriented toward availability — detecting when systems go down — rather than integrity — detecting when systems are operating but producing incorrect outputs. Reorienting detection toward integrity monitoring is technically feasible but requires investment in baseline behavioral profiling that most operators have not made.

5. Recommendations

Based on the pattern analysis and detection adequacy assessment, three recommendations warrant emphasis for logistics operators and the government partners that support them.

First, pre-execution dwell period detection should be a primary investment priority. For Pattern A and B incidents, the pre-execution dwell period represents the best detection opportunity — an attacker establishing persistence and mapping a TOS environment over 14 days will generate behavioral anomalies that integrity-oriented monitoring can detect. Current detection frameworks largely miss this window.

Second, integrity monitoring for operational data — cargo manifests, berth assignments, gate release authorizations — should be treated as a security function rather than solely an operational quality control function. Cryptographic verification of data provenance, anomaly detection on operational decision outputs, and cross-system consistency checking can substantially improve Pattern C detection rates.

Third, physical access protocols at critical logistics facilities should be integrated with cybersecurity monitoring rather than treated as parallel functions. Physical access events should generate cybersecurity alerts; anomalous system behavior should trigger physical access log review. The convergence of cyber and physical threats requires convergence of the detection frameworks applied to each.

6. Conclusion

The cyber-physical threat landscape facing critical logistics infrastructure is more sophisticated and more varied than current detection frameworks are calibrated to address. The concentration of detection capability around high-impact, immediately disruptive incidents creates systematic blind spots that sophisticated adversaries have demonstrated the capability to exploit. Addressing these blind spots requires not merely technical investment in new monitoring tools but a reconceptualization of logistics security that treats cyber and physical security as integrated functions rather than parallel domains. PLIANT's ongoing research program will continue to document and analyze emerging incident patterns as the threat landscape evolves.

References

Assante, M., & Lee, R. (2015). The Industrial Control System Cyber Kill Chain. *SANS ICS*.

Dragos. (2024). *Year in Review: ICS/OT Cybersecurity*. Dragos Inc.

Galinec, D., Moznik, D., & Guberina, B. (2017). Cybersecurity and Cyber Defence: National Level Strategic Approach. *Automatika*, 58(3), 273–286.

Greenberg, A. (2019). *Sandworm: A New Era of Cyberwar and the Hunt for the Kremlin's Most Dangerous Hackers*. Doubleday.

Langner, R. (2011). Stuxnet: Dissecting a Cyberweapon. *IEEE Security & Privacy*, 9(3), 49–51.

MITRE Corporation. (2023). ATT&CK for ICS. Retrieved from attack.mitre.org/matrices/ics.

PLIANT Institute. (2025). Port & Terminal Vulnerability Mapping: A Global Assessment. *PLIANT Journal*, March 2026.

Rid, T. (2013). *Cyber War Will Not Take Place*. Oxford University Press.

Torres, M. A. (2024). Adversarial Targeting of Commercial Logistics Networks: An Attribution Framework. *Journal of Intelligence and National Security Studies*, 12(2), 78–112.

Zetter, K. (2014). *Countdown to Zero Day: Stuxnet and the Launch of the World's First Digital Weapon*. Crown.

Zwillenberg, P., Field, D., & Dershowitz, D. (2019). Cyber Resilience: Playbook for the Digital Age. *Boston Consulting Group*.

Acknowledgments: The author thanks Dr. S.L. Chen for invaluable assistance in accessing Atlas database incident records underlying the empirical analysis, and Dr. J. Wu for methodological guidance on the pattern classification protocol. Colleagues at PLIANT provided feedback on earlier drafts that substantially improved the paper. Research was conducted under PLIANT internal funding. Correspondence: m.torres@projectpliant.com