GEORGIA DOT RESEARCH PROJECT 25-24

Final Report

FOUNDATIONAL STUDY FOR TRANSPORTATION NETWORK CYBER THREAT DETECTION USING AI/ML

Office of Performance-based Management and Research
600 West Peachtree Street NW | Atlanta, GA 30308

September 2026

TECHNICAL REPORT DOCUMENTATION PAGE

1. Report No. FHWA-GA-01-2524	**2. Government Accession No.** N/A	**3. Recipient's Catalog No.** N/A
4. Title and Subtitle Foundational Study for Transportation Network Cyber Threat Detection Using AI/ML		**5. Report Date** July 2026 **6. Performing Organization Code** N/A
7. Author(s) Yatis Dodia; Saman Zonouz, PhD; Glen Chou, PhD		**8. Performing Organization Report No.** 25-24
9. Performing Organization Name and Address Georgia Tech Research Corp. 505 Tenth St. Atlanta, GA 30318		**10. Work Unit No.** N/A **11. Contract or Grant No.** PI-[0021139]
12. Sponsoring Agency Name and Address Georgia Department of Transportation Office of Performance-based Management and Research 600 West Peachtree Street NW Atlanta, GA 30308		**13. Type of Report and Period Covered** Final Report (Dec 2025 – Sep 2026) **14. Sponsoring Agency Code** N/A

15. Supplementary Notes

Conducted in cooperation with the U.S. Department of Transportation, Federal Highway Administration.

16. Abstract

This work lays a foundation toward developing anomaly detection techniques to identify suspicious network activity at machine speed. The focus of this effort was to identify, catalog, and assess the viability of data sources within the Georgia Department of Transportation's (GDOT) statewide network. From a cybersecurity perspective, the GDOT network can be viewed as a cyber-physical system (CPS), which is a set of internetworked nodes or sensors controlled by computer algorithms that affect changes in the physical world. Examples of such end nodes could include traffic control devices, cameras, or programmable logic controllers (PLCs). Georgia Institute of Technology (GT) investigators first identified and assessed data sources for use in artificial intelligence/machine learning (AI/ML) modeling. They obtained data from GDOT and packaged and ingested it into an initial threat-detection model to demonstrate proof of concept (PoC) capability. However, the identification, sourcing, and packaging of data amenable to ML-grounded CPS security modeling remained a primary focus of this initial phase. The researchers were successful in processing network capture data from traffic cabinets and investigating eight anomalous events using AI models to demonstrate PoC capabilities.

The final outcome of this effort is a technical assessment of available data and their applicability to anomalous cyber event detection as demonstrated by one PoC candidate ML model. Identifying data and assessing its value for anomaly detection provided insights into the available resources for ML-based modeling. Building this infrastructure will pave a path for efficient and effective use of at-scale transportation cyber network data to investigate increasingly complex cybersecurity scenarios in future work, allowing for a better understanding of the security posture of the State of Georgia's transportation network. Ultimately, efforts such as this will lead to improved cyber resiliency and security for all Georgians, resulting in safer travel, economic activity, and decreased time to resolution for inevitable cyber events.

17. Key Words Transportation Cybersecurity, Artificial Intelligence, Machine Learning, Programmable Logic Controllers, Traffic Cabinets, Anomaly Detection, Formal Logic, Formal Verification	**18. Distribution Statement** No Restrictions

19. Security Classification (of this report) Unclassified	**20. Security Classification (of this page)** Unclassified	**21. No. of Pages** 22	**22. Price** Free

Form DOT F 1700.7 (8-72) **Reproduction of completed page authorized.**

GDOT Research Project 25-24

Final Report

FOUNDATIONAL STUDY FOR TRANSPORTATION NETWORK CYBER THREAT DETECTION USING AI/ML

By

Yatis Dodia
School of Cybersecurity and Privacy (SCP)
ydodia3@gatech.edu

Saman Zonouz, PhD
School of Cybersecurity and Privacy (SCP)
School of Electrical and Computer Engineering (ECE)
Director, Cyber-Physical Systems Track in GTPE
szonouz6@gatech.edu

Glen Chou, PhD
School of Cybersecurity and Privacy (SCP)
School of Aerospace Engineering (AE)
School of Electrical and Computer Engineering (ECE)

chou@gatech.edu

Contract with

Georgia Department of Transportation

September 2026

The contents of this report reflect the views of the authors who are responsible for the facts and the accuracy of the data presented herein. The contents do not necessarily reflect the official views or policies of the Georgia Department of Transportation. This report does not constitute a standard, specification, or regulation.

METRIC CONVERSION CHART

SI* (MODERN METRIC) CONVERSION FACTORS

APPROXIMATE CONVERSIONS TO SI UNITS

Symbol	When You Know	Multiply By	To Find	Symbol
		LENGTH		
in	inches	25.4	millimeters	mm
ft	feet	0.305	meters	m
yd	yards	0.914	meters	m
mi	miles	1.61	kilometers	km
		AREA		
in^2	square inches	645.2	square millimeters	mm^2
ft^2	square feet	0.093	square meters	m^2
yd^2	square yard	0.836	square meters	m^2
ac	acres	0.405	hectares	ha
mi^2	square miles	2.59	square kilometers	km^2
		VOLUME		
fl oz	fluid ounces	29.57	milliliters	mL
gal	gallons	3.785	liters	L
ft^3	cubic feet	0.028	cubic meters	m^3
yd^3	cubic yards	0.765	cubic meters	m^3
		NOTE: volumes greater than 1000 L shall be shown in m^3		
		MASS		
oz	ounces	28.35	grams	g
lb	pounds	0.454	kilograms	kg
T	short tons (2000 lb)	0.907	megagrams (or "metric ton")	Mg (or "t")
		TEMPERATURE (exact degrees)		
°F	Fahrenheit	5 (F-32)/9 or (F-32)/1.8	Celsius	°C
		ILLUMINATION		
fc	foot-candles	10.76	lux	lx
fl	foot-Lamberts	3.426	$candela/m^2$	cd/m^2
		FORCE and PRESSURE or STRESS		
lbf	poundforce	4.45	newtons	N
lbf/in^2	poundforce per square inch	6.89	kilopascals	kPa

APPROXIMATE CONVERSIONS FROM SI UNITS

Symbol	When You Know	Multiply By	To Find	Symbol
		LENGTH		
mm	millimeters	0.039	inches	in
m	meters	3.28	feet	ft
m	meters	1.09	yards	yd
km	kilometers	0.621	miles	mi
		AREA		
mm^2	square millimeters	0.0016	square inches	in^2
m^2	square meters	10.764	square feet	ft^2
m^2	square meters	1.195	square yards	yd^2
ha	hectares	2.47	acres	ac
km^2	square kilometers	0.386	square miles	mi^2
		VOLUME		
mL	milliliters	0.034	fluid ounces	fl oz
L	liters	0.264	gallons	gal
m^3	cubic meters	35.314	cubic feet	ft^3
m^3	cubic meters	1.307	cubic yards	yd^3
		MASS		
g	grams	0.035	ounces	oz
kg	kilograms	2.202	pounds	lb
Mg (or "t")	megagrams (or "metric ton")	1.103	short tons (2000 lb)	T
		TEMPERATURE (exact degrees)		
°C	Celsius	1.8C+32	Fahrenheit	°F
		ILLUMINATION		
lx	lux	0.0929	foot-candles	fc
cd/m^2	$candela/m^2$	0.2919	foot-Lamberts	fl
		FORCE and PRESSURE or STRESS		
N	newtons	0.225	poundforce	lbf
kPa	kilopascals	0.145	poundforce per square inch	lbf/in^2

* SI is the symbol for the International System of Units. Appropriate rounding should be made to comply with Section 4 of ASTM E380. (Revised March 2003)

TABLE OF CONTENTS

EXECUTIVE SUMMARY 1

CHAPTER 1. INTRODUCTION 2

CHAPTER 2. NETWORK TOPOLOGY STUDY AND SENSOR/NODE/DEVICE CATALOG 4

CHAPTER 3. IDENTIFICATION AND IMPLEMENTATION OF ANOMALOUS EVENTS (NETWORK ANOMALY DETECTION) 5

CHAPTER 4. RISK ASSESSMENT OF TRAFFIC CABINETS AND PLC-ENABLED REVERSIBLE LANES CONTROL UNITS (CONTROLLER SECURITY) 7

CHAPTER 5. DATA PACKAGING AND ANALYSIS 9

CHAPTER 6. DESIGN, IMPLEMENTATION, AND CHARACTERIZATION OF ML MODELS 10

CHAPTER 7. CONCLUSIONS AND RECOMMENDATIONS 12

REFERENCES 15

LIST OF FIGURES

Figure 1. Diagram. The decode-to-mirror pipeline: The vendor .SwxCF project is exported to PME Extensible Markup Language (XML) (the ground-truth oracle), transpiled to International Electrotechnical Commission (IEC) Structured Text, and compiled through MATIEC to run on OpenPLC. There is 100 percent element coverage against the oracle, and all 15 sites are running on the lab fleet, so a finding shown on the mirror holds on the real device..... 7

Figure 2. Matrix. Flow level anomalous event classification confusion matrix in percentage. 11

LIST OF ABBREVIATIONS

ACL	Access Control List
ACS	Access Control System
AI	Artificial Intelligence
API	Application Computing Interface
Caltrans	California Department of Transportation
CISA	Cybersecurity and Infrastructure Security Agency
CMS	Changeable Message Sign
CMU	Cryptographic Management Unit
CPS	Cyber-Physical System
DHCP	Dynamic Host Configuration Protocol
DNS	Domain Name System
DPI	Deep Packet Inspection
ESBMC	Efficient SMT-Based Context-Bounded Model Checker
FIN	Finish
GDOT	Georgia Department of Transportation
GELU	Gaussian Error Linear Unit
GT	Georgia Institute of Technology
HMI	Human-Machine Interface
ICSA	Industrial Control Systems Advisory
IDS	Intrusion Detection System
IEC	International Electrotechnical Commission
IoT	Internet of Things
ITS	Intelligent Transportation System

L3	Layer 3
LD	Ladder Diagram
MAC	Media Access Control
ML	Machine Learning
MMU	Memory Management Unit
Nmap	Network Mapper
NTCIP	National Transportation Communications for ITS Protocol
NWC	Northwest Corridor
OOD	Out-Of-Distribution
OT	Operational Technology
PAC	Programmable Automation Controller
PCAP	Packet Capture
PLC	Programmable Logic Controllers
PME	PAC Machine Edition
PoC	Proof of Concept
RLCS	Reversible Lane Control System
SMT	Satisfiability Modulo Theories
SNMP	Simple Network Management Protocol
SNP	Series Ninety Protocol
SRTA	State Road and Tollway Authority
SRTP	Service Request Transport Protocol
SSH	Secure Shell
SWP	Southwest Corridor
TCP	Transmission Control Protocol
TMC	Transportation Management Center

TSC	Traffic Signal Controller
UDP	User Datagram Protocol
VINCE	Vulnerability Information and Coordination Environment
VLAN	Virtual Local Area Network
VM	Virtual Machine
VRF	Virtual Routing and Forwarding
XML	Extensible Markup Language

EXECUTIVE SUMMARY

This work lays a foundation toward developing anomaly detection techniques to identify suspicious network activity at machine speed. The focus of this effort was to identify, catalog, and assess the viability of data sources within the Georgia Department of Transportation's (GDOT) statewide network. From a cybersecurity perspective, the GDOT network can be viewed as a cyber-physical system (CPS), which is a set of internetworked nodes or sensors controlled by computer algorithms that affect changes in the physical world. Examples of such end nodes could include traffic control devices, cameras, or programmable logic controllers (PLCs). Georgia Institute of Technology (GT) investigators first identified and assessed data sources for use in artificial intelligence/machine learning (AI/ML) modeling. They obtained data from GDOT and packaged and ingested it into an initial threat-detection model to demonstrate proof of concept (PoC) capability. However, the identification, sourcing, and packaging of data amenable to ML-grounded CPS security modeling remained a primary focus of this initial phase. The researchers were successful in processing network capture data from traffic cabinets and investigating eight anomalous events using AI models to demonstrate PoC capabilities.

The final outcome of this effort is a technical assessment of available data and their applicability to anomalous cyber event detection as demonstrated by one PoC candidate ML model. Identifying data and assessing its value for anomaly detection provided insights into the available resources for ML-based modeling. Building this infrastructure paves a path for effective use of transportation network data to investigate complex cybersecurity scenarios. Ultimately, such efforts such lead to improved cyber resiliency and security for all Georgians, resulting in safer travel, economic activity, and decreased time to resolution for inevitable cyber events.

CHAPTER 1. INTRODUCTION

This report details a security assessment of the Georgia Department of Transportation (GDOT) statewide transportation network, focusing on reversible lane control systems, associated programmable logic controllers (PLCs), and network cybersecurity. GDOT's network is a complex cyber-physical system where internetworked, algorithm-controlled nodes directly impact the physical world. As transportation infrastructure becomes increasingly automated, securing these critical systems against cyber threats is essential. Reversible managed lanes concentrate significant safety consequences into a small, network-reachable control surface. A single cabinet controls gantry gates that enforce legal traffic direction. Unlike a typical office IT compromise, a logic or memory-integrity failure here has a direct physical, roadway-level consequence. Because these systems control public safety infrastructure, a thorough security analysis is vital to identify vulnerabilities before exploitation.

The Georgia Institute of Technology (GT) research team pursued a multifaceted approach across five key tasks. Task 1 (refer to chapter 2) surveyed critical subnetworks within GDOT’s infrastructure, mapping network topology and communication flows. Tasks 2 through 5 (refer to chapter 3, chapter 4, chapter 5, and

chapter 6, respectively) involved deep technical analysis. Using vendor project files, the team reconstructed complete PLC logic for every site, established a runtime lab reproducing the controllers, and validated security findings using live pass/fail attacks. The investigation yielded high-level insights into design-level safety properties and control logic. Regarding network security, the team assessed the applicability of available data to anomalous cyber event detection, demonstrated through a proof-of-concept machine learning model. This infrastructure establishes

a foundation for efficiently utilizing at-scale transportation cyber network data to investigate increasingly complex cybersecurity scenarios in future work.

CHAPTER 2. NETWORK TOPOLOGY STUDY AND SENSOR/NODE/DEVICE CATALOG

The GT team surveyed two aspects of GDOT's network: the subnetworks that contain the PLCs and the traffic controller cabinets. This task focused on characterizing GDOT's statewide transportation network and the traffic light cabinets deployed at intersections. Each cabinet houses sensors and communication devices that collectively manage signal operations. The GT team worked closely with GDOT to understand the network topology and its connected nodes and devices, and assess what data to capture and how. This phase required significant iteration to identify data sources that were both available and suitable for downstream analysis because the GDOT network represents a complex CPS consisting of internetworked nodes controlled by computer algorithms that affect changes in the physical world.

To accomplish this objective, the team focused on two critical subnetworks within GDOT's infrastructure: a representative traffic controller cabinet and PLC network. By systematically mapping the network topology, connected devices, and communication flows, the research established a baseline view of GDOT's cyber-physical infrastructure. Packet captures were obtained from both subnetworks to enable detailed analysis of network protocols, communication patterns, and device interactions. The PLC data were also collected to ensure they captured a lane reversal procedure and gate change. This data collection effort identified critical nodes and dependencies within the transportation control system, providing the foundation for subsequent threat modeling and anomaly detection activities. Additionally, the research identified and characterized a selection of artificial intelligence (AI)/ML models, including model performance metrics.

CHAPTER 3. IDENTIFICATION AND IMPLEMENTATION OF ANOMALOUS EVENTS (NETWORK ANOMALY DETECTION)

An anomaly detection framework was designed to identify security events within a transportation network. The team categorized potential anomalous events into groups based on their behavior and assigned them to primary detection strategies: supervised classification, unsupervised out-of-distribution (OOD) detection, and unsupervised OOD with additional metadata.

Supervised classification is a straightforward method that inputs network data features to output a specific class label. It is best suited for detecting anomalies that produce consistent signatures in network traffic. Unsupervised OOD detection offers more flexibility by learning representations of normal data to identify deviations. This method is better suited for traffic that may appear valid in isolation but differs from the network's characteristic baseline. It also allows for the detection of previously unknown attack types. The third strategy incorporates contextual metadata, such as device profiles or timing, to determine if activity is normal for a specific device but anomalous in a broader context.

The effort aimed to detect anomalous events by analyzing raw packet captures flowing through network hubs. Events were chosen carefully and aligned to the standard cyber taxonomy MITRE ATT&CK. They prioritized anomalies with consistent patterns for initial development to ensure reliable data labeling and performance evaluation, noting that this foundation could be expanded to detect more complex events.

Data collection involved capturing all network packets at hub levels during normal operations to establish a baseline of typical traffic. To test detection capabilities, anomalous events were emulated by inserting simulated traffic patterns into the recorded data or altering packet

timestamps to mimic activity occurring at unusual times. The data used for model training combined open-source security datasets with organization-specific network data. All information began as raw packet captures, which served as the standardized unit for processing.

Data processing converted raw captures into flows, or time-based groupings of packets between sources and destinations. One method used analysis tools to extract statistical features from these flows. This approach is useful for machine learning models requiring fixed inputs but omits payload information and requires a flow to complete before processing. The second method grouped raw packets without statistical conversion. This preserves timing relationships between packets and allows for real-time inference as time windows elapse. However, this method requires more storage and may be less effective if packet contents are encrypted.

CHAPTER 4. RISK ASSESSMENT OF TRAFFIC CABINETS AND PLC-ENABLED REVERSIBLE LANES CONTROL UNITS (CONTROLLER SECURITY)

The GT Cyber-Physical Security Laboratory conducted a comprehensive security assessment of the Access Control System governing reversible barrier gates on major express lanes. The assessment covered all field sites, each utilizing industrial controllers and human-machine interfaces to manage traffic flow direction. Working exclusively from vendor project files provided by the agency, the team reconstructed the complete logic of every site. They established a runtime laboratory that reproduces the controllers and validated security findings through live simulation and automated testing. This approach allowed for thorough analysis without requiring physical cabinet access or vendor software. The team developed a pipeline to turn opaque vendor files into ground truth logic, a faithful runnable mirror, and live, reproducible tests (see Figure 1).

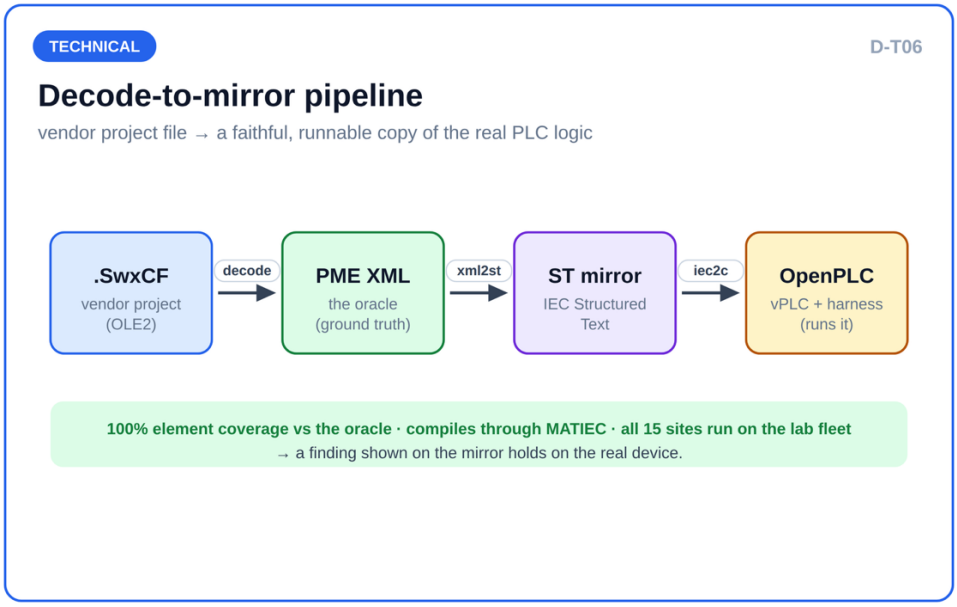

Figure 1. Diagram. The decode-to-mirror pipeline: The vendor .SwxCF project is exported to PME Extensible Markup Language (XML) (the ground-truth oracle), transpiled to International Electrotechnical Commission (IEC) Structured Text, and compiled through MATIEC to run on OpenPLC. There is 100 percent element coverage against the oracle, and all 15 sites are running on the lab fleet, so a finding shown on the mirror holds on the real device.

The investigation yielded significant safety and security findings regarding the fundamental design of the control systems. Primary results indicate discrepancies between safety enforcement logic and physical operation mechanisms. Furthermore, the systems lack independent hardware safeguards that would operate separately from the software control layer. Consequently, issues within the digital control logic cannot be mitigated by physical backups, leaving safety consequences solely dependent on software integrity.

Regarding network security, the assessment identified weaknesses in access control mechanisms governing communication between system components. The logic appears to be based on a reused commercial template, suggesting these structural weaknesses may affect other deployments using the same product line. Additionally, functionality exists within the system that allows operators or network entities to override safety checks or present false status information to monitoring systems.

Findings were validated through a combination of offline analysis and automated simulation. This confirmed that conditions violating core safety principles are reachable within the system architecture. The team demonstrated these issues through controlled testing, with data indicating that the identified risks are consistent across the entire fleet of installations rather than isolated to specific locations.

Recommendations include adopting automated testing as a standing safety oracle for future logic revisions. Override and simulation capabilities should be treated as safety-critical and write-protected. Crucially, because there is no independent hardware safety layer, digital detection is only a mitigation; durable fixes require physical safety engineering, such as external interlocks or

obstacle detection that the controller cannot override. The team emphasizes that findings are properties of the shared vendor template, warranting coordinated disclosure rather than an agency-specific fix.

CHAPTER 5. DATA PACKAGING AND ANALYSIS

The success of ML-based anomaly detection depends on data quality. Task 4 focused on sourcing and packaging network data from GDOT's infrastructure for AI/ML analysis, establishing a reusable pipeline.

Data came from two origins. First, GDOT-specific packet captures were collected at the hub level during normal operations from critical locations: Hub A (113 MB), Hub U (30 MB), and the reversible lanes PLC network (66 MB).[1] The reversible lanes data included complete lane reversal procedures. All GDOT data were raw PCAP files. Second, the publicly available CIC-IDS2017 intrusion detection dataset (50.15 GB) was incorporated, containing labeled attack patterns alongside benign traffic [1]. This hybrid approach allowed supervised training on a large corpus while reserving GDOT data for zero-shot evaluation.

Raw captures were processed through two pipelines. The first employed flow-based feature extraction using tools to compute statistical features per flow. The second preserved raw packet groupings, supporting deep learning architectures and streaming detection. Because collecting attack data from a live air-gapped network is unsafe, emulation strategies were developed. Port scanning and ping sweep attacks were performed on an isolated virtual network, with IP addresses replaced before insertion. Unusual time of activity scenarios were emulated by taking legitimate lane-change packets and modifying their timestamps [1]. The effort established three capabilities: a reusable data pipeline, validation that existing visibility contains sufficient signal

for threat detection, and a hybrid training methodology. This transforms archives into strategic assets for anomaly detection and resiliency.

CHAPTER 6. DESIGN, IMPLEMENTATION, AND CHARACTERIZATION OF ML MODELS

This report focuses on detecting anomalous network events, categorized by whether the traffic type itself indicates an attack (e.g., port scanning, ping sweeps) or depends on contextual observation. The primary focus is on the first category, though methods are described that could extend to context-dependent anomalies.

Machine learning approaches are divided into supervised methods, which require labeled training data, and unsupervised methods, which analyze unlabeled data through clustering or representation learning. While unlabeled traffic data is abundant, labeled datasets can be constructed by injecting attacks or using public repositories. Detection granularity is another key decision. Flow-based detection groups packets by their identifiers over a time window to compute statistics, naturally capturing client-server sessions but discarding payload information. Raw packet detection preserves payload and timing details but requires handling massive data throughput.

For flow-level detection, the team processed raw packet captures using feature extraction tools, retaining statistical features while excluding IP addresses to prevent overfitting. Values were normalized, and a classifier was trained on a large public intrusion detection dataset. When tested via zero-shot classification on operational data mixed with simulated attacks, the model achieved perfect recall for benign and port scan flows. However, recall for ping sweep flows was low, as illustrated in Figure 2. This indicates flow-level detection alone is insufficient.

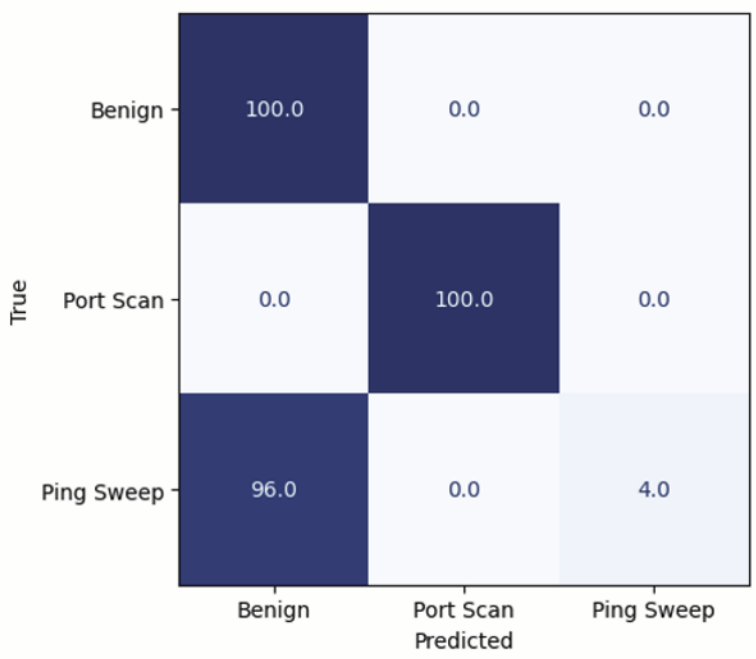

Figure 2. Matrix. Flow level anomalous event classification confusion matrix in percentage.

Consequently, the team explored packet-level detection using raw header and payload data. To address limited labeled data, they employed self-supervised learning via a masked byte reconstruction task. A neural network was trained to predict masked bytes within packets based on context, learning robust representations of packet structures. After pretraining on both public and operational packet data, a classification head was trained on labeled data. Zero-shot classification on operational data with simulated attacks achieved near-perfect accuracy. This suggests packet-level representation is sufficient for distinguishing malicious traffic.

Beyond supervised classification, the pretrained packet encoder supports unsupervised out-of-distribution (OOD) detection. Frozen packet embeddings can be aggregated to score traffic density against a benign baseline, potentially identifying novel or zero-day anomalies. This approach can be extended to metadata-augmented detection by appending contextual features, such as device profiles or timestamps, to the embeddings. This allows the system to identify

anomalies where traffic content is normal but context is not, such as operational changes occurring outside scheduled times. Evaluating these extensions on held-out data would confirm whether the self-supervised representations transfer effectively to open-ended detection tasks.

CHAPTER 7. CONCLUSIONS AND RECOMMENDATIONS

This assessment reconstructed the logic of all 15 express lane control cabinets using only vendor project files, creating faithful runnable mirrors. The analysis confirmed discrepancies between safety enforcement logic and physical operation mechanisms across all sites, with no independent hardware safety layer to compensate. Combined with insufficient access controls and available override features, unsafe physical outcomes are reachable through network commands. Because these design properties stem from a reused vendor template, the recommended path involves coordinated vendor disclosure alongside local hardening. The only durable fix for the core safety gap is an independent hardware safety layer, as digital detection cannot restore a physical backstop the cabinet never had.

This evaluation integrated infrastructure mapping, threat detection, and controller security analysis. Systematic characterization of the network topology established baseline communication flows. Dual-pipeline data processing, combining flow-based statistics with raw packet grouping, demonstrated proof-of-concept anomaly detection across multiple threat categories. Hub-level packet captures proved sufficient to train models achieving high accuracy on zero-shot classification, validating that existing network visibility can support machine-speed threat detection. A self-supervised learning architecture developed through masked reconstruction extends beyond supervised classification to unsupervised outlier detection, supporting multiple metadata-conditioned detectors without additional pretraining. Parallel

defensive tooling, including detection rules and deception assets, provides an interim detection layer while physical mitigations are engineered.

Recommended Actions

1. Pursue coordinated disclosure with vendors via a national coordination environment on a safety-adjusted embargo.
2. Install an independent hardware safety layer (e.g., contactor interlock, obstruction sensing), the only durable fix for the identified safety gaps.
3. Enable service authentication where firmware permits and restrict network reachability via access control lists.
4. Disable or lock down specific control override functionalities at the template level if unused; alert on any write to those ranges if they are used.
5. Remediate other critical operational technology assets surfaced by the asset inventory.
6. Deploy the defensive rule pack on the field network as an interim detection layer.

This assessment is deliberately conservative. The team never touched a production controller. Claims settled only on real hardware are permanent gaps. Proven claims are based on the vendor's authoritative serialization or reproduced live on a faithful mirror. Findings depending on unverified network assumptions are flagged as questions, not facts.

ACKNOWLEDGEMENTS

The GT team would like to thank GDOT for their contributions and assistance with technical knowledge sharing, data collection, and overall transportation guidance. Specifically, we thank Josuha Miller and Samantha Wagner and their team for their correspondence throughout the effort. We also thank Teague Buchanan for valuable feedback, helping us to understand the aspects most helpful to GDOT. Finally, we thank Sabbir Ahmed for his project management guidance.

REFERENCES

[1] Shafi, M., Lashkari, A.H., and Roudsari, A.H. (2025). "NTLFlowLyzer: Towards Generating an Intrusion Detection Dataset and Intruders Behavior Profiling Through Network and Transport Layers Traffic Analysis and Pattern Extraction." *Computers & Security*, *148*, p. 104160, Available online: https://doi.org/10.1016/j.cose.2024.104160.

[2] Sharafaldin, I., Habibi Lashkari, A., and Ghorbani, A. (2018). "Toward Generating a New Intrusion Detection Dataset and Intrusion Traffic Characterization." *Proceedings of the 4th International Conference on Information Systems Security and Privacy*, *1*, pp. 108–116. Available online: 10.5220/0006639801080116.

[3] Habibi Lashkari, A., Draper Gil, G., Mamun, M., and Ghorbani, A. (2017). "Characterization of Tor Traffic using Time based Features." *Proceedings of the 3rd International Conference on Information Systems Security and Privacy*, *1*, pp. 253–262. Available online: 10.5220/0006105602530262.

[4] Cybersecurity and Infrastructure Security Agency, "Emerson PACSystem and Fanuc," U.S. Dept. of Homeland Security, ICS Advisory ICSA-24-158-01, Jun. 6, 2024. [Online]. Available: https://www.cisa.gov/news-events/ics-advisories/icsa-24-158-01

[5] VersaMax PLC User's Manual, GFK-1503F, GE Fanuc Automation North America, Charlottesville, VA, USA, 2000.

[6] B. Sahin, D. Oygenblik, M. Yao, Y. Huang, B. Saltaformaggio, and S. Zonouz, "Fuzzing the physical space: Physics-aware testing of black-box industrial control systems," in Proc. 47th IEEE Symp. Security and Privacy (S&P), 2026.

[7] S. McLaughlin, S. Zonouz, D. Pohly, and P. McDaniel, "A trusted safety verifier for process controller code," in Proc. 21st Annu. Network and Distributed System Security Symp. (NDSS), San Diego, CA, USA, Feb. 2014.

[8] Functional Safety of Electrical, Electronic, and Programmable Electronic Safety-Related Systems, Part 1: General Requirements, IEC 61508-1:2010, 2010.

[9] K. Stouffer et al., "Guide to operational technology (OT) security," National Institute of Standards and Technology, Gaithersburg, MD, USA, NIST SP 800-82r3, Sep. 2023, doi: 10.6028/NIST.SP.800-82r3.

[10] OpenPLC, an open source programmable logic controller runtime (IEC 61131-3). [Online]. Available: https://openplcproject.com