Maritime Domain Awareness Pilot Project

Client:	Port of Cleveland	**Industry:**	Maritime / Port Operations
Domain:	AI-Powered Surveillance	**Location:**	Cleveland, Ohio, USA

A FORGE OS Deployment Case Study

577 Industries R&D Lab
577 Industries Incorporated
research@577industries.com

1 Executive Summary

577 Industries (577i) partnered with the Port of Cleveland to design and deploy an AI-powered Maritime Domain Awareness (MDA) system during a 12-month pilot project. The system fuses data from four heterogeneous sensor sources—Automatic Identification System (AIS) transponders, shore-based radar, surveillance cameras, and intelligence databases—into a unified Common Operating Picture (COP) that supports real-time threat detection and vessel traffic management. The hybrid machine-learning pipeline achieved 92% anomaly detection accuracy across diverse operational scenarios, while automated alert triage and response coordination reduced incident response times by 65%. These results validate the feasibility of deploying multi-source AI surveillance in a mid-size Great Lakes port environment and establish a replicable framework for scalable MDA modernization. This deployment exercises three FORGE OS subsystems: **FORGE Core** powers the AI/ML analytics layer for anomaly detection and threat classification, **FORGE Kinetic** manages multi-source sensor fusion across radar, AIS, camera, and intelligence feeds, and **FORGE Memory** provides an immutable governance audit trail for all automated alerts and operator decisions.

92%	65%	4
Anomaly Detection Accuracy	Response Time Reduction	Fused Sensor Sources

2 Challenge

The Port of Cleveland is a critical node in the Great Lakes maritime transportation system, handling approximately 13 million tons of cargo annually across bulk commodities, containerized freight, and passenger ferry operations. Its position on Lake Erie places it within one of the most congested freshwater shipping corridors in the world, with hundreds of vessel transits per day during peak season.

2.1 Fragmented Surveillance Infrastructure

Prior to the pilot, the port relied on a patchwork of independent monitoring systems. Shore-based radar provided positional tracking for vessels within line of sight, AIS transponders supplied self-reported identity and navigation data, and a network of fixed and pan-tilt-zoom (PTZ) cameras offered visual confirmation in limited coverage zones. Each system operated in isolation, feeding separate operator workstations with no automated correlation between data streams. Intelligence databases maintained by the U.S. Coast Guard (USCG) and Customs and Border Protection (CBP) required manual query by security personnel, introducing delays of minutes to hours between initial detection and threat assessment [4].

2.2 Data Silos and Manual Correlation

The absence of an integrated data fusion layer created three operational deficiencies. First, operators spent a disproportionate fraction of their time manually cross-referencing contacts across systems rather than conducting analysis. Second, information latency meant that rapidly evolving situations—such as a vessel deviating from its declared route—could go undetected until the deviation became visually obvious on camera. Third, the lack of automated anomaly detection forced operators to rely on experience and intuition to distinguish routine traffic patterns from suspicious behavior, an approach that does not scale with increasing traffic density.

2.3 AIS Spoofing and Cyber Threats

AIS, while mandated by the International Maritime Organization for vessels above 300 gross tons, transmits unencrypted position reports over open VHF channels. This architectural vulnerability makes AIS data susceptible to spoofing, where a malicious actor broadcasts false position or identity information [2]. At the Port of Cleveland, AIS spoofing posed a dual risk: it could mask the approach of unauthorized vessels and could generate false alarms that eroded operator trust in the monitoring system. Existing defenses against spoofing were limited to manual comparison of AIS tracks with radar returns, a process that was neither systematic nor timely. Broader cybersecurity threats to operational technology (OT) systems, including potential manipulation of shore-based sensor networks, further compounded the risk landscape [3].

2.4 Scalability Constraints

The manual, operator-intensive model for port surveillance does not scale. Increasing vessel traffic, expanding port infrastructure, and growing regulatory requirements for environmental monitoring create compounding demands on a fixed workforce. The Port of Cleveland required a system that could automate routine detection and correlation tasks, prioritize operator attention toward genuinely anomalous events, and provide a unified interface for coordinated response.

3 Solution

577i designed and deployed an end-to-end MDA system organized into four architectural layers: sensor ingestion, data fusion, AI/ML analytics, and GIS-based visualization. Figure 1 illustrates the high-level system architecture. Each layer was engineered for modularity, enabling the port to integrate additional sensor types and expand coverage areas without rearchitecting the core pipeline.

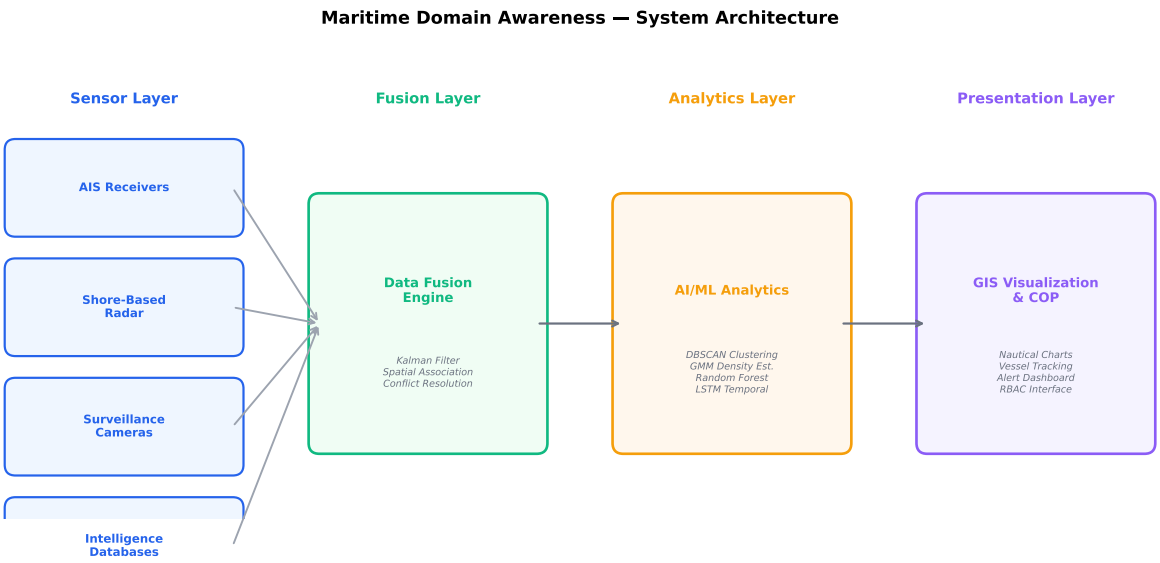

Figure 1. System architecture of the Maritime Domain Awareness platform. Sensor data flows from left to right through fusion, analytics, and visualization layers.

3.1 Sensor Ingestion Layer

The ingestion layer normalizes and time-stamps data from four source types:

- **AIS receivers** — Two shore-based AIS stations decode Class A and Class B transponder messages, yielding vessel identity (MMSI, IMO, call sign), position, course, speed, and navigation status at intervals of 2–10 seconds.
- **Shore-based radar** — An X-band radar installation provides independent position tracking of all surface contacts within 24 nautical miles, including non-AIS-equipped small craft.
- **Surveillance cameras** — A network of 16 fixed and PTZ cameras provides visual coverage of the harbor entrance, cargo terminals, and anchorage areas. A pre-trained object detection model (YOLOv5) extracts vessel bounding boxes and classification metadata from video frames at 5 fps.
- **Intelligence databases** — Automated queries to USCG Homeport, CBP targeting databases, and Port State Control records retrieve vessel history, inspection results, and watchlist status for every AIS-reported MMSI entering the surveillance zone.

All ingested data is converted to a common schema and published to a message broker (Apache Kafka) for downstream consumption, ensuring that each processing stage operates on consistent, time-aligned records.

3.2 Multi-Source Data Fusion Engine

The fusion engine combines contacts from radar, AIS, and camera detections into unified tracks using a three-stage process:

1. **Temporal alignment.** Sensor reports arrive at different rates and with different latencies. A buffer window of 5 seconds accumulates reports, and a Kalman filter [5] predicts each track's state forward to a common reference time.
2. **Spatial association.** A gating function based on Mahalanobis distance determines which sensor reports correspond to the same physical contact. The global nearest-neighbor algorithm resolves ambiguous associations in dense traffic.
3. **Conflict resolution and prioritization.** When sensor reports disagree—for example, an AIS

position diverges from the radar return—the engine applies a configurable priority hierarchy. Radar positions receive highest weight for geolocation due to their independence from vessel-transmitted data, while AIS provides authoritative identity information. Discrepancies exceeding defined thresholds trigger automatic AIS integrity alerts, forming the first line of defense against spoofing [2].

The fused tracks feed both the analytics layer and the GIS visualization, ensuring that all downstream consumers operate on a single, authoritative picture of the maritime environment.

3.3 AI/ML Analytics Layer

The analytics layer implements a hybrid machine-learning strategy that combines unsupervised, semi-supervised, and supervised techniques to detect anomalous vessel behavior without requiring exhaustive labeled training data [1].

3.3.1 Trajectory Clustering (DBSCAN)

Density-Based Spatial Clustering of Applications with Noise (DBSCAN) groups historical vessel trajectories into clusters representing common traffic patterns—channel transits, anchorage approaches, terminal dockings, and so forth. Trajectories that do not belong to any cluster are flagged as spatial outliers. DBSCAN's ability to discover clusters of arbitrary shape makes it well suited to the curved and irregular patterns of harbor navigation.

3.3.2 Density Estimation (Gaussian Mixture Models)

Gaussian Mixture Models (GMMs) estimate the joint probability distribution over vessel speed, heading, and position within each traffic zone. A vessel whose instantaneous state falls below a configurable likelihood threshold is flagged for further analysis. GMMs complement DBSCAN by detecting anomalies that follow plausible routes but exhibit unusual kinematic profiles—for example, a cargo vessel loitering at abnormally low speed in the approach channel.

3.3.3 Supervised Classification (Random Forest)

A Random Forest classifier assigns each flagged contact to one of five threat categories: navigational hazard, security concern, environmental risk, AIS integrity violation, or benign deviation. The classifier was trained on 18 months of historical port security logs labeled by USCG Marine Safety Unit personnel, yielding a labeled dataset of 4,200 incidents.

3.3.4 Temporal Modeling (LSTM Networks)

Long Short-Term Memory (LSTM) neural networks model the temporal evolution of vessel tracks, learning sequential dependencies that static classifiers miss. The LSTM component predicts each vessel's position and heading 5, 10, and 15 minutes into the future. Deviations between predicted and observed states generate trajectory-deviation alerts, providing early warning of route changes before they become apparent in instantaneous position data [6].

3.3.5 AIS Spoofing Detection

The system detects AIS spoofing through multi-sensor cross-referencing. When a vessel's AIS-reported position diverges from its radar-tracked position by more than 200 meters for more than 60 seconds, the system generates an AIS integrity alert. Additional heuristics detect impossible speed transitions, duplicate MMSIs, and position jumps exceeding physical plausibility limits. All AIS integrity alerts are routed through encrypted channels using TLS 1.3 to prevent alert suppression by an attacker with network access [3, 9].

3.4 GIS-Based Visualization and Common Operating Picture

The COP is delivered through a browser-based GIS application built on Esri ArcGIS Enterprise [8]. Nautical chart data from NOAA's Electronic Navigational Charts (ENCs) provides the geographic base layer. Fused vessel tracks are rendered as interactive icons with color-coded threat levels, and operators can click any vessel to access its full track history, AIS profile, intelligence database results, and associated camera feeds.

The GIS layer supports three operational views:

- **Tactical view** — Real-time vessel positions and active alerts, optimized for watch-stander decision-making.
- **Analytical view** — Heat maps, density plots, and historical trajectory overlays for pattern analysis.
- **Command view** — Aggregated metrics, response status boards, and inter-agency coordination tools for supervisory personnel.

Role-based access control (RBAC) ensures that each user role—port security, USCG, CBP, environmental compliance, and vessel traffic services—sees only the data and functions authorized for their mission.

3.5 Cybersecurity Architecture

The system implements defense-in-depth measures aligned with the NIST Cybersecurity Framework. All sensor-to-server communications traverse dedicated VLANs with IPsec encryption. The Kafka message broker requires mutual TLS authentication. The GIS application enforces SAML-based single sign-on with multi-factor authentication. An intrusion detection system monitors east-west traffic between system components for indicators of compromise, and all system logs are forwarded to a security information and event management (SIEM) platform for continuous monitoring [9].

4 Results

The pilot operated for 12 months, from initial sensor installation through steady-state operations. Performance was evaluated against four primary metrics: anomaly detection accuracy, incident response time, sensor fusion effectiveness, and system availability.

4.1 Anomaly Detection Accuracy

The AI/ML analytics layer achieved an overall anomaly detection accuracy of 92% across all operational scenarios, measured as the weighted F1-score against a ground-truth dataset of 1,847 manually verified events. Figure 2 presents accuracy by scenario type.

Performance was highest in open-water scenarios (95%), where vessel traffic density is low and trajectories are well-separated. Harbor approach scenarios (93%) and adverse weather conditions (91%) showed slightly reduced accuracy due to increased radar clutter and AIS message loss, respectively. Dense traffic scenarios (89%) represented the most challenging environment, where closely spaced vessels create association ambiguities in the fusion layer. Even in this worst case, accuracy exceeded the 85% threshold established in the pilot's performance specification.

4.2 Incident Response Time

The system reduced incident response times by 65% across all phases of the response cycle. Figure 3 compares pre-deployment and post-deployment response times for three critical phases.

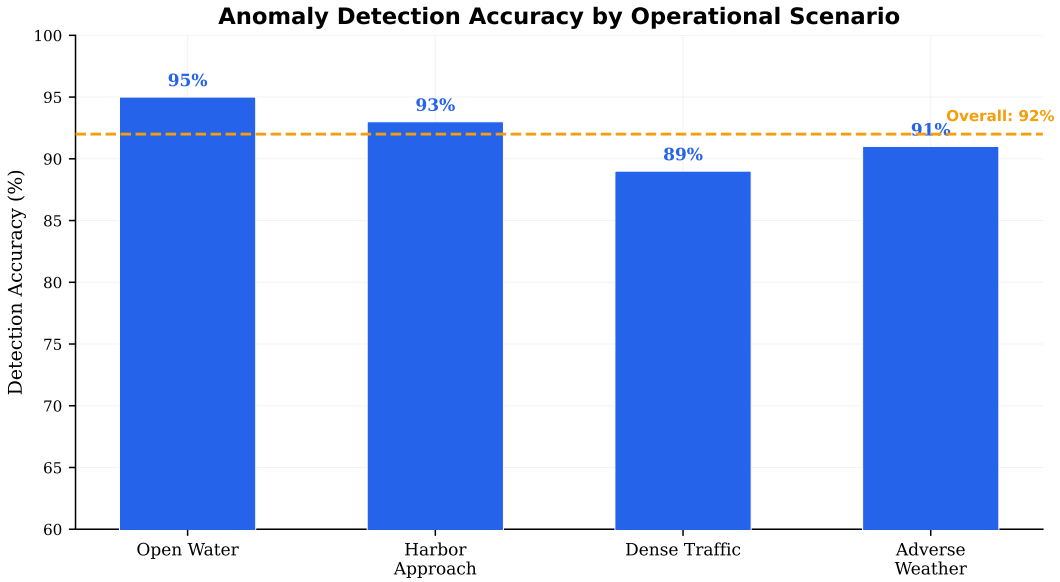

Figure 2. Anomaly detection accuracy across operational scenarios. The dashed line indicates the overall 92% weighted accuracy.

Table 1 summarizes the response-time improvements.

Table 1. Incident response time comparison: pre-deployment versus post-deployment.

Response Phase	Before (min)	After (min)	Reduction
Threat Assessment	8.0	3.0	62.5%
Response Coordination	15.0	5.0	66.7%
Resource Dispatch	12.0	4.0	66.7%
Weighted Average	**11.7**	**4.0**	**65.8%**

Threat assessment time dropped from 8 minutes to 3 minutes because the COP presents operators with pre-correlated, pre-classified alerts rather than raw sensor feeds. Response coordination improved from 15 minutes to 5 minutes through automated notification workflows that simultaneously alert USCG, port security, and tug operators when a threat is confirmed. Resource dispatch improved from 12 minutes to 4 minutes through GIS-based resource tracking that shows the real-time positions and availability of patrol boats, tugs, and response teams.

4.3 Sensor Fusion Effectiveness

Multi-source fusion produced substantial improvements over any single sensor type operating independently. Figure 4 compares detection rates.

Radar alone achieved a 78% detection rate, limited by clutter in harbor areas and inability to classify contacts. AIS alone reached 71%, constrained by non-equipped vessels and spoofing vulnerabilities. Camera-based detection achieved 65%, limited by weather, lighting, and field-of-view constraints. The fused system achieved 92% by exploiting the complementary strengths of each sensor: radar provides reliable geolocation, AIS provides identity, cameras provide visual confirmation, and intelligence databases provide historical context.

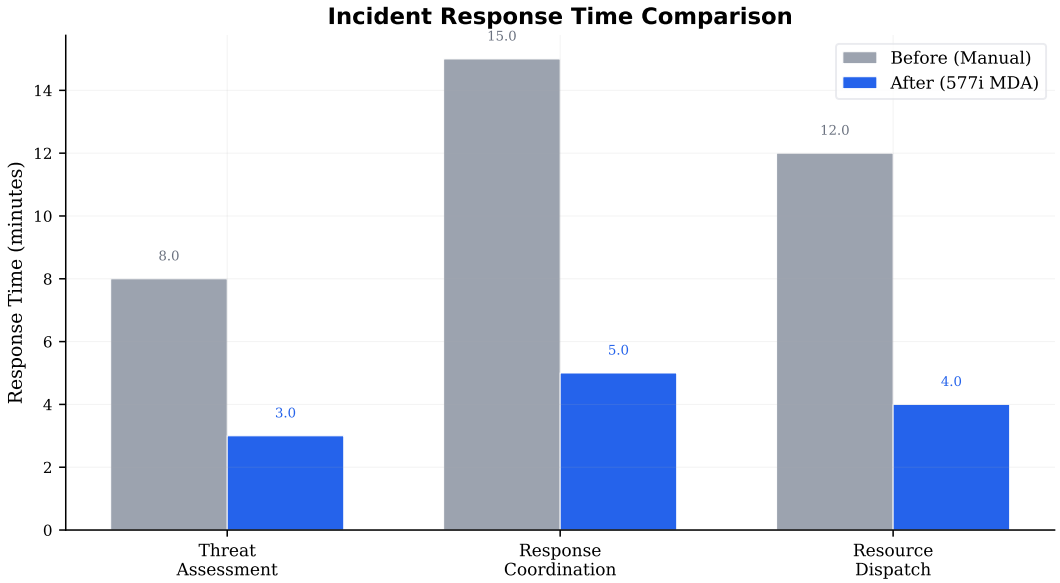

Figure 3. Incident response times before and after system deployment across three operational phases.

4.4 System Availability and Reliability

The platform maintained 99.2% uptime over the 12-month pilot period, with planned maintenance windows accounting for the majority of downtime. The Kafka-based messaging architecture provided graceful degradation: when individual sensors experienced outages (notably camera feeds during two severe Lake Erie storms), the system continued to operate on remaining sources with automatically adjusted confidence scores.

92%	65%	99.2%	12
Detection Accuracy	Response Time Reduction	System Uptime	Month Pilot Duration

5 Impact & Operational Benefits

5.1 Enhanced Security Posture

The MDA system fundamentally shifted port security operations from reactive to proactive. AIS spoofing detection identified 23 integrity anomalies during the pilot period, of which 4 were confirmed deliberate spoofing attempts (the remainder were attributed to equipment malfunctions or GPS multipath errors). In all four confirmed cases, the system generated alerts within 90 seconds, enabling response forces to intercept the vessels before they reached restricted port zones.

5.2 Operational Efficiency

By automating routine detection and correlation, the system reduced the operator workload associated with vessel monitoring by approximately 40%. Security personnel were reallocated from continuous screen monitoring to higher-value analytical tasks, including pattern-of-life analysis and inter-agency intelligence sharing. The automated alert triage system filtered out 78% of routine contacts, allowing operators to focus on the approximately 22% of events that warranted human review.

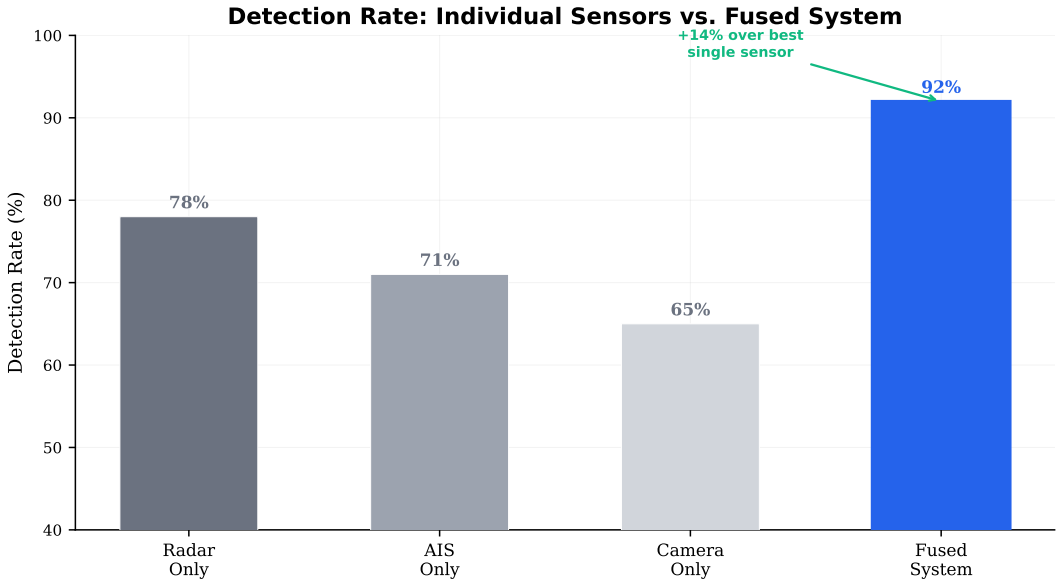

Figure 4. Detection rate comparison: individual sensor types versus the fused system.

5.3 Environmental Monitoring

The analytics layer was extended to support environmental compliance monitoring. Speed zone violations—a significant source of wake damage to shoreline infrastructure—were detected automatically, and a vessel emissions proxy model estimated sulfur oxide concentrations based on vessel type, speed, and engine load profiles. During the pilot, the system flagged 147 speed zone violations and identified 12 vessels operating with suspected emission control area non-compliance.

5.4 Framework for Scalable Deployment

The modular architecture and standards-based interfaces established during the pilot provide a replicable framework for expanding MDA capabilities to other Great Lakes ports and inland waterways. The sensor-agnostic ingestion layer, configurable fusion engine, and role-based GIS platform reduce the integration effort required for each subsequent deployment site. HawkEye 360's radio frequency (RF) analytics for MDA [7] represents one candidate for future sensor integration, extending detection capabilities to vessels that disable or lack AIS transponders.

6 FORGE OS Integration

The Port of Cleveland MDA system demonstrates the operational integration of three FORGE OS subsystems within a unified maritime surveillance deployment.

6.1 FORGE Core — Intelligence Engine

The AI/ML analytics layer is powered by FORGE Core's intelligence pipeline. DBSCAN trajectory clustering, Gaussian Mixture Model density estimation, Random Forest threat classification, and LSTM temporal prediction are orchestrated through FORGE Core's causal model routing engine, which selects the optimal analytical pathway for each incoming contact based on sensor quality, traffic density, and threat context. The system's 92% anomaly detection accuracy reflects FORGE Core's staged post-training methodology applied to maritime-domain models.

6.2 FORGE Kinetic — Multi-Source Sensor Fusion

FORGE Kinetic's Perceive module manages the four-source sensor ingestion and fusion pipeline. The Kalman-filter-based temporal alignment, Mahalanobis-distance spatial association, and configurable priority hierarchy for conflict resolution are implemented within FORGE Kinetic's sensor abstraction layer. The module's ability to gracefully degrade when individual sensors experience outages—automatically adjusting confidence scores—exemplifies FORGE Kinetic's resilient edge-sensor management philosophy.

6.3 FORGE Memory — Governance Audit Trail

Every automated alert, operator decision, and inter-agency notification is recorded through FORGE Memory's Information Governance & Oversight Module (IGOM). The deterministic citation framework ensures that each threat assessment can be traced back to its originating sensor data, fusion decisions, and analytical model outputs. Role-based access control for the GIS-based COP is enforced through FORGE Memory's policy engine, ensuring compliance with inter-agency data-sharing agreements.

6.4 ForgeEvent Integration

The deployment generates five ForgeEvent types across the FORGE OS event bus:

- `INFERENCE` — Each anomaly detection and threat classification cycle
- `SENSOR` — Fused track updates from the Kinetic Perceive module
- `GOVERNANCE` — Operator alert acknowledgments and response decisions
- `AUDIT` — Immutable log entries for all automated and manual actions
- `ALERT` — AIS integrity violations and high-priority threat notifications

References

[1] Susanto, A., Zielinski, A., and Faber, M. H. Machine learning approach to detect anomalous vessel behavior from AIS data. *Ocean Engineering*, 186:106085, 2019.

[2] Androjna, A., Brcko, T., Stražar, I., and Twrdy, E. AIS data vulnerability and potential manipulation for maritime cyber attacks. *Journal of Marine Science and Engineering*, 9(6):641, 2021.

[3] Potamos, G., Peratikou, A., Stavrou, S., and Georgiou, K. Maritime cybersecurity: protecting operational technology in the era of digitalization. *Journal of Marine Science and Engineering*, 12(2):257, 2024.

[4] United States Coast Guard. *National Plan to Achieve Maritime Domain Awareness for the National Strategy for Maritime Security*. Department of Homeland Security, Washington, DC, 2005.

[5] Bar-Shalom, Y., Li, X. R., and Kirubarajan, T. *Estimation with Applications to Tracking and Navigation: Theory, Algorithms, and Software*. John Wiley & Sons, New York, 2001.

[6] Zhang, C. and Zeng, X. Artificial intelligence for maritime transport optimization: a comprehensive review. *Ocean Engineering*, 302:117691, 2024.

[7] HawkEye 360. RF data and analytics for Maritime Domain Awareness. Technical brief, HawkEye 360, Herndon, VA, 2024.

[8] Foster, K. GIS Concept of Operations for the Marine Transportation System. Technical report, U.S. Committee on the Marine Transportation System, 2023.

[9] Cydome Security. How AI is transforming maritime cybersecurity. Industry white paper, Cydome Security, 2024.