

Article

Towards Capability Development: A Layered Analysis of AIS Replay and Ghost Ship Attacks

Georgios Potamos , Savvas Theodoulou, Eliana Stavrou and Stavros Stavrou 

Faculty of Pure and Applied Sciences, Open University of Cyprus, Latsia, 2220 Nicosia, Cyprus; savvas.theodoulou1@ouc.ac.cy (S.T.); eliana.stavrou@ouc.ac.cy (E.S.); stavros.stavrou@ouc.ac.cy (S.S.)

* Correspondence: georgios.potamos@st.ouc.ac.cy

Abstract

The Automatic Identification System (AIS) is a fundamental component of modern Maritime Situational Awareness (MSA) systems, supporting navigation safety, vessel traffic management, and maritime surveillance. However, the absence of authentication, integrity protection, and message verification makes AIS vulnerable to a range of spoofing attacks capable of manipulating the maritime operational picture. Although existing research has proposed various techniques for detecting AIS spoofing, comparatively little attention has been given to understanding how different spoofing mechanisms influence operational behaviour and the capabilities required to detect and respond to them. To address this gap, the paper introduces a focused taxonomy that distinguishes replay and ghost ship attacks, while accounting for hybrid scenarios that combine characteristics of both mechanisms. A layered threat analysis characterises their observable manifestations across the physical, network infrastructure, AIS protocol, AIS message validation, and operational behaviour layers and identifies the distinct operational capabilities required to detect, distinguish, and respond to these attacks. Building upon this analysis, the paper proposes a structured capability-driven cyber range exercise design approach that systematically translates these capabilities into learning objectives and differentiated training scenarios. Furthermore, the work contributes guidelines for evaluating participants' performance in alignment with the proposed operational capabilities and learning objectives. Overall, the work provides a structured analytical framework for understanding AIS replay and ghost ship attacks, and their combination, and a practical approach for developing capability-driven maritime cybersecurity training aligned with their operational characteristics.

Keywords: Automatic Identification System (AIS); AIS spoofing; replay attacks; ghost ship attacks; maritime cyber range training



Academic Editor: Lei Wang

Received: 18 August 2026

Revised: 14 September 2026

Accepted: 17 September 2026

Published: 26 September 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and

conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

The Automatic Identification System (AIS) has evolved into one of the most important technologies supporting Maritime Situational Awareness (MSA) [1]. Originally introduced as a collision avoidance and navigational safety mechanism [2], AIS enables vessels to automatically exchange dynamic and static information, including vessel identity, position, speed, heading, navigational status, and voyage-related information [3]. Through this capability, AIS supports a wide range of maritime operations, including collision avoidance, vessel traffic management, maritime surveillance, search and rescue (SAR), border protection, fisheries monitoring, and port operations [4].

AIS, mandated for most commercial vessels, plays a critical role in ensuring maritime safety. Despite its operational importance, the AIS protocol presents significant cybersecurity vulnerabilities due to its lack of authentication, integrity protection, and encryption mechanisms [4,5]. These weaknesses expose AIS to numerous cyber threats, including message manipulation, identity spoofing [4,6], replay attacks [7,8], and the generation of fabricated vessels [9,10]. Such attacks compromise the integrity of maritime operations by manipulating the maritime operational picture, reducing trust in AIS information, and potentially affecting navigational safety, maritime surveillance, and operational decision support [11].

While AIS was intended to enhance navigational safety, cybersecurity was not prioritized during its development. AIS messages are transmitted in plaintext and are accepted by receiving systems under a trust-based architecture that assumes transmitted information is legitimate. Consequently, receiving systems cannot verify either the origin or the integrity of AIS messages, allowing adversaries to manipulate, replay, or fabricate maritime information without being readily detected. These architectural limitations have been extensively documented and remain one of the principal cybersecurity challenges affecting modern maritime systems [9].

Existing research has primarily concentrated on AIS anomaly detection [12], trajectory prediction [6,13], protocol enhancements [5], and machine learning techniques for identifying suspicious vessel behaviour [6]. While these studies have significantly advanced AIS security, spoofing attacks are frequently considered as a single threat category despite important differences in their technical implementation and operational effects. In particular, replay attacks and ghost ship attacks represent fundamentally different forms of AIS deception. Replay attacks reuse previously legitimate AIS transmissions to introduce outdated but authentic information into the maritime picture, whereas ghost ship attacks generate entirely fabricated vessel information to create non-existent or impersonated maritime targets. As a result, replay attacks primarily challenge information update and temporal consistency [8], while ghost ship attacks challenge vessel validation [14], multi-sensor correlation [15], and target authenticity [16,17]. Hybrid attacks may combine these mechanisms. Other forms of AIS manipulation, such as the falsification of live vessel data through GNSS manipulation, are out of the scope of this work. Therefore, the categories examined in this paper are not intended to constitute an exhaustive classification of all possible AIS manipulation.

Despite these differences, replay and ghost ship attacks are frequently treated as variations in the broader AIS spoofing problem within both the research literature and cyber range training activities. Existing studies primarily focus on improving detection algorithms or proposing defensive mechanisms [18–20]. Comparatively little attention has been given to systematically examining how replay and ghost ship attacks differ and how these differences influence operational detection, response, and capability development. Consequently, training approaches often generalise AIS spoofing scenarios [14,19,21], without considering the distinct technical and operational challenges associated with each attack type. Furthermore, existing AIS attack classifications are generally organised according to attack vectors, affected components, security properties, or observable anomalies, but provide limited guidance on how the distinguishing characteristics of replay and ghost ship attacks can be translated into operational capability requirements that can ultimately drive the design of appropriate maritime cyber range exercises.

To address this gap, the paper makes two principal contributions. First, it develops a focused taxonomy, informed by the existing literature and a layered threat analysis that differentiates replay and ghost ship attacks based on their technical characteristics, operational manifestations, and defensive implications across the Physical, Network Infrastructure, AIS

Protocol, AIS Message Validation, and Operational Behaviour layers of the AIS ecosystem. The taxonomy also clarifies the relationship of hybrid attacks to the two attack categories forming the focus of the paper. The analysis identifies the observable characteristics that influence operational reasoning and the capabilities required to recognise, differentiate, and respond to each attack type during maritime operations. Second, building upon this analysis, the paper proposes a structured approach to systematically translate the derived operational capabilities into learning objectives, differentiated scenario designs, and operational evidence to support the development of the intended operational analytical capabilities. Furthermore, this work contributes guidelines for evaluating participants' performance in alignment with the proposed operational capabilities and learning objectives. Overall, this work aims to provide insights to trainers on how the distinct characteristics of replay and ghost ship attacks can be represented in capability-driven maritime cybersecurity exercises. A cyber range exercise design is provided as an example to illustrate how threat analysis can be transformed into structured capability development for maritime cybersecurity training.

The remainder of the paper is organised as follows. Section 2 presents the research methodology, describing the overall study design and the process followed to develop the proposed taxonomy and layered threat analysis. Section 3 reviews the existing literature on AIS spoofing, compares the proposed taxonomy with existing AIS attack classifications, and identifies the limitations of current approaches. Section 4 introduces the proposed taxonomy and presents the layered threat analysis of replay and ghost ship attacks. Section 5 discusses the operational implications of the analysis, examining the distinct detection and response competency requirements associated with each attack category. Section 6 presents how the layered threat analysis and operational capability requirements can be systematically translated into learning objectives, scenario design, and operational evidence, alongside guidelines for capability-driven assessments. Section 7 provides an example of a maritime cyber range exercise implementation, illustrating how the study's insights can inform the design of exercises. Section 8 discusses limitations and outlines directions for future work. Finally, Section 9 concludes the paper.

2. Methodology

This study adopted a qualitative analytical research approach to examine the differences between AIS replay and ghost ship attacks and translate the findings into operational capability requirements to drive the design of maritime cyber range training scenarios. The research was conducted in four sequential phases, as illustrated in Figure 1.

First, a structured review of the literature was undertaken to examine existing research on AIS spoofing attacks, detection approaches, and maritime cyber operations. Relevant publications were identified using combinations of terms including "AIS spoofing," "AIS replay attack," "ghost ship," "fabricated AIS vessel," "AIS taxonomy," and "maritime cyber range." Studies were included when they addressed attack mechanisms, observable indicators, detection requirements, operational implications, or related training activities. Information concerning attack characteristics, execution paths, detection methods, and operational implications was extracted and thematically synthesised.

Second, a focused taxonomy of AIS spoofing attacks was developed by synthesising evidence from the literature. The taxonomy classifies replay, ghost ship, and hybrid attacks according to the mechanism through which deceptive maritime information is introduced into the AIS ecosystem. Replay and ghost ship attacks formed the principal categories examined, while hybrid attacks and related forms of AIS manipulation were considered to clarify the scope and boundaries of the taxonomy.

Third, the identified attack categories were examined using a layered threat analysis. Each attack type was systematically analysed across the Physical, Network Infrastructure, AIS Protocol, AIS Message Validation, and Operational Behaviour layers to identify the observable characteristics that distinguish replay and ghost ship attacks. Rather than focusing solely on technical implementation, the layered analysis examined how different attack mechanisms influence operational reasoning, detection processes, and response activities. These observations were subsequently synthesised to derive the operational capability requirements for effective detection and response.

Finally, the derived operational capability requirements informed the design of maritime cyber range exercises through a structured design approach. Specifically, the adopted design approach demonstrates how operational capabilities can be systematically transformed into learning objectives, scenario design, operational evidence, implementation, and assessment. The cyber range exercise presented in this paper serves as an illustrative example, demonstrating how the analytical findings can be operationalised within a realistic maritime training environment. The exercise was implemented using a specialised Maritime cyber range which is developed by the Cyprus Centre of Cybersecurity, Telecommunications & Critical Infrastructure Protection (CTRL) at the Open University of Cyprus.



Figure 1. Study Methodology.

3. Related Work

Research on AIS security has expanded considerably as maritime operations have become increasingly dependent on AIS-derived Maritime Situational Awareness (MSA) for navigation, vessel traffic management, surveillance, and decision support. Early security studies demonstrated that AIS lacks authentication, integrity protection, and message freshness mechanisms, making it vulnerable to message manipulation, vessel impersonation, fabricated vessel generation, and replay attacks [4,22–25]. Although recent developments in maritime communication standards, including the IEC 61162-450/460 series, introduce mechanisms that improve the security of certain maritime network infrastructures [26], these mechanisms do not authenticate open AIS radio broadcasts and are not universally deployed [27]. Consequently, replay attacks, fabricated vessel generation, and other AIS spoofing techniques remain feasible in operational environments [28].

Cybersecurity surveys consistently identify AIS as one of the most vulnerable components of modern maritime systems due to its trust-based architecture and its central role in operational decision-making [29–31]. As a result, existing research has primarily focused on detecting AIS spoofing, validating transmitted information, and improving protocol security.

3.1. AIS Spoofing Detection Approaches

A substantial body of research investigates AIS anomaly detection using vessel trajectories, behavioural modelling, statistical analysis, machine learning, and deep learning techniques [12,32–36]. These approaches identify abnormal vessel movements by analysing route deviations, unexpected speed changes, anomalous headings, AIS silence, loitering behaviour, and physically implausible trajectories. Although highly effective for recognising suspicious vessel behaviour, these methods generally identify anomalies without determining the underlying cyber attack responsible for them.

Complementary research focuses on validating AIS messages and detecting manipulated information. Existing techniques include kinematic validation, protocol compliance verification, TDMA timing analysis, trajectory consistency, behavioural plausibility assessment, and integrity verification of static and dynamic vessel information [19,37,38]. Other studies have also demonstrated that multisensor fusion, combining AIS with radar, RF observations, satellite imagery, or other surveillance sources, significantly improves the detection of manipulated maritime targets [39,40].

Collectively, these approaches have substantially advanced AIS spoofing detection and contributed to improving maritime situational awareness. However, they remain primarily focused on identifying anomalous behaviour or validating AIS information rather than explaining how different spoofing mechanisms influence operational interpretation, response activities, or capability development.

3.2. Replay and Ghost Ship Attacks

Among AIS spoofing techniques, replay attacks and ghost ship attacks represent two fundamentally different attack mechanisms. Replay attacks reuse previously legitimate AIS transmissions or trajectories, introducing temporally incorrect but authentic vessel information into the maritime picture. Existing detection approaches primarily rely on timestamp verification, freshness validation [8,41], behavioural consistency, SIEM-based monitoring [42], protocol analysis, and cryptographic mechanisms designed to detect the reuse of historical information [17].

Ghost ship attacks, in contrast, generate fabricated AIS messages describing vessels that do not physically exist or deliberately impersonate legitimate vessels. Research on ghost ship detection focuses on validating vessel identities, analysing trajectory plausibility, detecting inconsistencies in static vessel information, and correlating AIS data with independent surveillance sources such as radar or electro-optical (EO) systems [19,37,43,44]. Because fabricated vessels frequently lack corresponding physical observations, multisensor correlation has become one of the most effective approaches for identifying this class of attack.

These categories describe the nature of the deceptive information rather than the technical point at which manipulation occurs. Both replayed and fabricated information may be introduced through RF transmission, network manipulation, compromised AIS feeds, or application and database manipulation. Conversely, GNSS manipulation may cause a genuine AIS transponder to broadcast falsified live position or navigational information without replaying historical data or generating a separate fictitious vessel [45,46]. It is therefore treated as a related form of live-data falsification rather than being forced into either the replay or ghost ship category. Hybrid attacks combine two or more mechanisms, such as the reuse of a legitimate historical trajectory together with fabricated identity, timing, or voyage information.

Although detection techniques, e.g., [18,25,38], address different aspects of AIS spoofing, they are generally developed independently. Consequently, existing studies do not consistently distinguish replay, fabricated-vessel, and combined attacks according to the provenance and operational validity of the deceptive information. This limits the guidance available to operators and trainers regarding the different evidence and reasoning processes required to distinguish among them.

3.3. Existing AIS Attack Classifications

Existing AIS security classifications have been developed for different purposes. Early studies classify attacks according to the AIS function targeted or the type of manipulation performed, including false information generation, vessel impersonation, and disruption of

availability [4]. Broader maritime cybersecurity taxonomies support adversarial modelling, risk analysis, and resilience planning [16], while other approaches classify suspicious activity according to data-integrity violations [18,22,29], message falsification [18,38], identity manipulation [30,33,37,39], behavioural anomalies [12,33,39], or the detection technique employed [12,18,22,30,38].

The taxonomy developed in this study complements these approaches by focusing on the nature and provenance of the deceptive information utilized. Replay attacks reuse previously legitimate AIS information, whereas ghost ship attacks introduce fabricated vessel information. Hybrid attacks combine these mechanisms, while live-data falsification is recognised as a related but analytically distinct threat. Therefore, the contribution does not lie in introducing previously unidentified attack types, but in systematically comparing their manifestations across five analytical layers and translating findings into operational capability requirements to drive the design of cyber-range training activities.

3.4. Operational Perspective and Research Gap

The operational impact of AIS spoofing extends beyond the technical detection of manipulated messages. Maritime operators must assess the reliability of AIS information, correlate multiple surveillance sources, distinguish cyber attacks from navigation anomalies, and determine appropriate response actions while maintaining situational awareness [47,48]. Replay attacks primarily challenge temporal consistency and information freshness, whereas ghost ship attacks require vessel validation, multisensor correlation, and verification of target authenticity. These differences influence not only attack detection but also operator decision-making, incident classification, escalation procedures, and response strategies [49].

Overall, the reviewed literature has significantly advanced the detection and mitigation of AIS spoofing attacks. Nevertheless, existing studies and classifications predominantly investigate replay and ghost ship attacks as isolated technical problems and devote comparatively little attention to understanding how their differing characteristics influence operational behaviour or the capabilities required to detect and respond to them. As a result, the literature provides limited analytical guidance for systematically distinguishing these attack categories from an operational perspective or for translating their characteristics into capability development and training requirements.

To address this gap, this paper systematically characterises replay and ghost ship attack behaviour across the AIS ecosystem. Building upon this analysis, the paper derives the operational implications associated with each attack category and demonstrates, through an illustrative cyber range use case, how the findings can inform the design of differentiated, capability-informed maritime cybersecurity training.

4. Taxonomy and Layered Threat Analysis of AIS Spoofing

4.1. AIS Spoofing Taxonomy

AIS spoofing is commonly discussed as a single class of cyber attack despite encompassing multiple mechanisms for manipulating maritime situational awareness. Although these attacks share the objective of deceiving operators by presenting false or misleading maritime information, they differ substantially in how AIS data are generated, manipulated, and introduced into the maritime operational picture. These differences influence not only the technical implementation of the attacks but also their observable characteristics and operational manifestations.

To provide a structured understanding of AIS spoofing, this paper introduces a focused taxonomy comprising two primary categories: replay attacks and ghost ship attacks. Hybrid attacks are treated as a combination of these mechanisms rather than as a separate

primary category. The taxonomy classifies attacks according to the nature of the deception rather than the specific technical techniques used to execute them. It is not intended to provide an exhaustive classification of every form of AIS manipulation. Related threats that alter the live information of a genuine vessel without replaying historical data or creating a fabricated vessel are excluded from the taxonomy.

4.1.1. Replay Attacks

Replay attacks involve the capture and retransmission of previously legitimate AIS traffic. Instead of creating new vessel information, the attacker reuses authentic AIS messages outside their original temporal or geographical context. The transmitted information therefore remains syntactically correct and corresponds to a genuine vessel, but no longer accurately reflects the current operational situation.

Replay attacks may be realised through several mechanisms, including RF replay of captured AIS transmissions, replay of AIS packets within communication networks, replay of AIS feed information, or reintroduction of historical vessel tracks from databases. Regardless of the execution method, the objective is to manipulate maritime situational awareness by presenting legitimate but contextually incorrect information. Consequently, replay attacks primarily compromise information freshness and temporal consistency while preserving the authenticity of the underlying AIS messages.

4.1.2. Ghost Ship Attacks

Ghost ship attacks involve the generation of fabricated AIS information describing vessels that do not physically exist or involve the impersonation of legitimate vessels through fabricated information. Unlike replay attacks, these attacks require the creation of synthetic vessel identities, navigational data, and behavioural characteristics that appear plausible within the maritime environment.

Ghost ships may be introduced through RF-based message injection, network manipulation, compromised AIS feeds, or application and database manipulation. Depending on the attacker's objectives and level of access, fabricated targets may coexist with legitimate vessels or replace selected attributes of genuine maritime contacts. As a result, ghost ship attacks primarily challenge vessel authenticity, target verification, and the integrity of the recognised maritime picture.

4.1.3. Hybrid Attacks

Sophisticated adversaries may combine replay and ghost ship techniques to create highly realistic deception campaigns. Hybrid attacks blend legitimate historical AIS information with fabricated vessel data, producing manipulated maritime situations that appear operationally plausible over extended periods. For example, an attacker may reuse the historical trajectory of a genuine vessel while modifying its identity, timestamps, destination, position, or other voyage-related information. A hybrid attack may therefore exhibit the temporal characteristics of replay alongside the identity and authenticity inconsistencies associated with a ghost ship attack. By simultaneously exploiting the characteristics of replay and ghost ship attacks, hybrid attacks increase the complexity of detection and investigation and represent one of the most challenging categories of AIS spoofing.

4.2. Technical Execution Paths

While the proposed taxonomy classifies AIS spoofing according to the nature of the deception, each attack category may be realised through multiple technical execution paths depending on the attacker's objectives, technical capabilities, and level of access to the maritime environment. Analysis of existing literature identifies five principal execution paths that repeatedly appear across AIS spoofing campaigns. Although each execution path

exploits different components of the AIS ecosystem, they ultimately seek to manipulate the recognised maritime picture by altering the integrity, authenticity, or temporal validity of AIS information.

- **GNSS Manipulation.** GNSS spoofing indirectly manipulates AIS information by transmitting counterfeit satellite navigation signals to the target vessel [45,46]. The AIS transponder continues to operate normally but broadcasts erroneous position, speed, and course information derived from the falsified navigation data. Because the transmitted AIS messages remain protocol compliant, identifying these attacks relies on recognising inconsistencies between reported navigation data and independent positioning sources.
- **RF Layer Manipulation.** Software Defined Radio (SDR) technologies enable attackers to inject fabricated AIS messages or replay previously captured transmissions directly into the VHF communication channel [50]. Since manipulation occurs before information reaches shore-based processing systems, these attacks affect the integrity of the radio-frequency environment while preserving the standard operation of receiving AIS equipment.
- **Network Layer Manipulation.** AIS information distributed through IP-based communication networks may be manipulated using packet injection, interception, replay, redirection, or session hijacking techniques. Adversary-in-the-Middle (AiTM) attacks represent a particularly important execution path because they enable attackers to modify or replay AIS information exchanged between sensors, aggregation services, and command-and-control systems without altering the original RF transmission [51].
- **AIS Feed Manipulation.** AIS aggregation services represent attractive attack targets because a single compromise may simultaneously affect multiple downstream consumers [48]. By manipulating AIS feeds before distribution, attackers can introduce fabricated vessels, suppress legitimate targets, or replay historical traffic across several maritime systems.
- **Application and Database Manipulation.** Compromised AIS servers, databases, APIs, or visualisation platforms allow attackers to introduce fabricated vessels or replay historical traffic after data collection has already occurred [11]. Consequently, the operational picture presented to maritime personnel may differ significantly from the original sensor information despite no manipulation occurring within either the RF environment or the communication network.

Although these execution paths differ in their technical implementation, they all seek to manipulate maritime situational awareness by exploiting different stages of the AIS information lifecycle. Consequently, similar operational effects may arise from fundamentally different technical mechanisms. Understanding these execution paths therefore complements the proposed taxonomy and provides the basis for examining how replay and ghost ship attacks manifest across the different layers of the AIS ecosystem.

4.3. Layered Threat Analysis

The taxonomy presented in the previous sections distinguishes replay and ghost ship attacks as the two primary categories examined, while hybrid attacks combine characteristics of both mechanisms. The discussion related to the technical execution paths explains where manipulation may be introduced within the AIS ecosystem. However, these perspectives alone do not fully explain how the different attack categories manifest throughout the maritime operational environment or why they require different approaches to detection and analysis. The literature reviewed in Section 3 demonstrates that replay and ghost ship attacks have largely been investigated independently and predominantly from a detection perspective. While these studies provide valuable insights into individual attack mecha-

nisms, they do not offer a structured analytical framework for systematically comparing how the two attack categories propagate across the AIS ecosystem and influence maritime situational awareness.

To address this gap, the findings of the literature review were synthesised to develop a layered threat analysis that examines replay and ghost ship attacks across five complementary analytical layers: the Physical Layer, Network Infrastructure Layer, AIS Protocol Layer, AIS Message Validation Layer, and Operational Behaviour Layer. The layered analysis identifies the observable characteristics associated with each attack type across successive stages of AIS communication and processing. Figure 2 illustrates the proposed layered analytical approach.




Figure 2. Layered Threat Analysis Approach to Compare Replay and Ghost Ship Attacks.

Using a layered threat analysis approach, the characteristics of replay and ghost ship attacks identified throughout the literature were systematically compared across each analytical layer. The resulting synthesis is presented in Table 1, which highlights the principal distinctions between the two attack categories and demonstrates how apparently similar AIS spoofing attacks exhibit fundamentally different characteristics depending on the analytical perspective adopted.

Table 1. Layered Threat Analysis Comparing Replay and Ghost Ship Attacks.

Analysis Layer	Replay Attack Characteristics	Ghost Ship Attack Characteristics	Analytical Significance
Physical Layer	Legitimate AIS RF transmissions or previously decoded AIS messages are captured and retransmitted. The replayed signal originates from a different location or time, potentially introducing transmission delays and altered RF characteristics.	New AIS RF transmissions or fabricated AIS messages are generated directly by the attacker without requiring previously captured signals. The transmitted signal originates solely from the attacker's equipment while claiming a different vessel location.	Replay attacks are characterised by the reuse of existing transmissions, whereas ghost ship attacks involve the generation of entirely new transmissions. This distinction provides the first indication of the underlying deception mechanism.
Network Infrastructure	May require interception, storage, modification, and retransmission of legitimate AIS traffic. Replay attacks targeting IP-based infrastructures typically exploit compromised communication paths to reuse historical information.	Fabricated AIS information is injected through compromised servers, AIS aggregation platforms, APIs, or network infrastructure. Network compromise primarily serves as a delivery mechanism for newly generated information.	This layer distinguishes historical information reuse from synthetic information injection, providing insight into how deception propagates through maritime communication infrastructures.

Table 1. Cont.

Analysis Layer	Replay Attack Characteristics	Ghost Ship Attack Characteristics	Analytical Significance
AIS Protocol	Reuses legitimate AIS packets that naturally comply with message encoding, CRC validation, slot allocation, reporting intervals, and protocol timing requirements.	Requires attackers to generate protocol-compliant AIS packets. Although sophisticated generators produce valid messages, implementation errors or deviations in timing and slot allocation can expose the attack.	Replay attacks generally preserve protocol correctness because authentic messages are reused, whereas ghost ship attacks may reveal inconsistencies associated with message generation. Protocol analysis therefore provides stronger evidence for ghost ship attacks than for replay attacks.
AIS Message Validation	Broadcasts authentic vessel identities and navigational information. Manipulation is generally limited to message timing, geographical position, or selected dynamic fields while preserving message authenticity.	Creates synthetic or impersonated vessel identities together with fabricated static and dynamic information, including MMSI, IMO number, callsign, dimensions, voyage information, position, speed, and course.	Replay attacks primarily compromise information freshness, whereas ghost ship attacks compromise vessel authenticity and identity. This layer represents the fundamental distinction between the two attack categories.
Operational Behaviour	Reproduces authentic historical vessel trajectories and movement patterns. Vessel behaviour generally remains physically plausible but appears delayed, duplicated, or inconsistent with the current operational context.	Generates entirely synthetic vessel behaviour requiring realistic trajectories, speeds, headings, and navigation patterns. Inconsistencies frequently emerge through implausible manoeuvres, conflicting sensor observations, or unrealistic maritime behaviour.	Replay attacks primarily create uncertainty regarding temporal consistency and the reliability of operational information, whereas ghost ship attacks challenge the authenticity of the recognised maritime picture through fabricated maritime contacts.

The comparison reveals several important observations. At the Physical and Network Infrastructure layers, the principal distinction concerns the origin of the manipulated information. Replay attacks reuse legitimate AIS transmissions or previously recorded information, whereas ghost ship attacks require the generation and introduction of entirely synthetic maritime data. Although both execution paths ultimately distort the recognised maritime picture, they rely on fundamentally different deception mechanisms. Replay attacks exploit the credibility of authentic historical information, while ghost ship attacks depend on creating convincing but fabricated vessel identities and associated navigational data.

Further distinctions emerge at the AIS Protocol and AIS Message Validation layers. Replay attacks typically preserve protocol compliance because authentic AIS messages are reused without modification, making them difficult to distinguish using protocol analysis alone. In contrast, ghost ship attacks require the generation of protocol-compliant AIS messages and synthetic vessel attributes, introducing additional opportunities for detection through protocol inconsistencies, identity validation, behavioural plausibility assessment, and multisensor verification. Consequently, replay attacks primarily compromise information freshness and temporal validity, whereas ghost ship attacks primarily compromise vessel authenticity and identity.

These differences become most evident at the Operational Behaviour layer, where replay attacks reproduce genuine historical vessel movements outside their original temporal context, creating uncertainty regarding the current validity of maritime information. Ghost ship attacks, on the other hand, introduce entirely fabricated maritime contacts that challenge the credibility of the recognised maritime picture and require operators to determine whether the reported vessel actually exists. This demonstrates that replay and ghost ship attacks cannot be fully characterised solely by their technical implementation but

require a broader analytical perspective that considers how deception propagates across multiple layers of the AIS ecosystem.

Hybrid attacks may produce combined manifestations across these layers. Replayed messages may retain valid protocol structures and historically plausible movement, while fabricated identity, timing, or navigational attributes introduce inconsistencies at the message-validation and operational-behaviour layers. Their detection therefore requires operators to consider temporal validity and vessel authenticity simultaneously rather than treating the presence of one type of evidence as excluding the other.

The taxonomy, technical execution paths, and layered threat analysis provide a structured analytical perspective for distinguishing replay and ghost ship attacks beyond conventional detection-oriented classifications. By linking technical attack mechanisms with their operational manifestations, this study establishes a systematic basis for examining the differing operational implications associated with each attack category. These implications are explored in the following section, where the layered threat analysis is used to derive the operational capabilities required to detect and respond effectively to replay and ghost ship attacks and to recognise situations in which characteristics of both mechanisms coexist.

5. Operational Capability Requirements

The layered threat analysis presented in the previous section demonstrates that replay and ghost ship attacks differ not only in their technical implementation but also in the operational challenges they introduce for maritime operators. While both attack types manipulate AIS information, they do so through fundamentally different deception mechanisms. Replay attacks exploit the temporal validity of authentic AIS transmissions by reusing legitimate information outside its original operational context, whereas ghost ship attacks generate synthetic vessel information intended to create entirely fictitious maritime contacts. This differentiation demonstrates that the analytical requirements for identifying, interpreting, and responding to these attacks extend beyond technical detection and require distinct forms of operational reasoning. In hybrid attacks, these requirements coexist because operators must assess both the temporal validity and the authenticity of the reported information.

To systematically derive these requirements, the layered threat analysis was abstracted into a set of operational analytical capability requirements. The analysis considers how observable characteristics of AIS spoofing manifest across different analytical layers, the operational challenges these manifestations create, and the analytical capabilities required to interpret them. This progression provides a structured mechanism for translating technical attack characteristics into operational capability requirements that can inform maritime operations, capability development, and the design of training interventions. In this context, an operational analytical capability refers to the ability of maritime personnel to interpret technical evidence, assess its operational significance, and support informed decision-making during AIS spoofing incidents.

The analysis follows three successive levels of abstraction, illustrated in Table 2. The first level focuses on technical evidence, examining the physical, network, and protocol characteristics that provide indicators of AIS manipulation. The second level considers information validation, where the integrity, authenticity, and operational validity of reported maritime information are assessed. Finally, the third level addresses operational reasoning, examining how operators interpret the validated information to maintain situational awareness, support decision-making, and prioritise appropriate operational responses. These abstraction levels establish an analytical progression from low-level technical observations to high-level operational reasoning, providing the conceptual foundation for deriving operational capability requirements.

Table 2. Analytical Approach for Deriving Operational Capability Requirements.

Level of Abstraction	Analytical Layers	Primary Analytical Question	Purpose
Technical Evidence	Physical Layer, Network Infrastructure Layer, AIS Protocol Layer	What technical evidence indicates AIS manipulation?	Identify observable technical characteristics that distinguish replay and ghost ship attacks.
Information Validation	AIS Message Validation Layer	Can the reported maritime information be trusted?	Assess the authenticity, integrity and operational validity of AIS information.
Operational Reasoning	Operational Behaviour Layer–Operational Assessment/Operational Decision-Making	How should operators interpret the evidence and determine an appropriate response?	Transform validated technical evidence into operational understanding and informed decision-making.

Building upon the three levels of abstraction presented in Table 2, Table 3 synthesises the findings of the layered threat analysis into a structured set of operational capability requirements. The table systematically synthesises the layered threat analysis by mapping each operational capability to an observable characteristic, the distinct manifestations of replay and ghost ship attacks, and the corresponding operational challenge that maritime operators must resolve. This explicit mapping demonstrates how technical attack characteristics are systematically translated into distinct analytical requirements, providing a transparent rationale for capability development and the design of training interventions. Instead of treating replay and ghost ship attacks as variations in a common spoofing technique, the synthesis highlights the different reasoning processes required to identify, interpret, and respond to each attack type.

Although Table 3 focuses on comparing replay and ghost ship attacks, the identified capabilities also provide the basis for analysing hybrid attacks. In such cases, operators must combine temporal and behavioural analysis with identity validation, multisensor correlation, and information-provenance assessment. They must also avoid treating evidence of replay as excluding the simultaneous fabrication or modification of AIS information. The principal capability required for hybrid incidents is therefore the ability to integrate multiple forms of evidence, maintain competing explanations, and revise the initial attack classification as additional information becomes available.

Table 3. Operational capability requirements derived from the layered threat analysis.

Analytical Layer	Observable Characteristic	Replay Attack Manifestation	Ghost Ship Attack Manifestation	Operational Challenge	Required Operational Capability
Physical Layer	RF transmission origin	Legitimate RF transmissions are retransmitted outside their original spatial or temporal context.	New RF transmissions are generated to create synthetic AIS broadcasts.	Determine whether the received transmission originates from the claimed vessel.	Correlate RF observations with expected vessel position and transmission source.
	Physical vessel presence	Historical transmissions imply vessel presence where the vessel is no longer located.	AIS reports indicate vessels that have no corresponding physical presence.	Determine whether a reported vessel physically exists.	Validate reported vessel presence using independent maritime information sources.
	Geographical plausibility	Historical vessel locations are replayed outside their original operational context.	Synthetic vessels appear in implausible or unsupported locations.	Assess whether reported vessel activity is geographically credible.	Evaluate geographical consistency using contextual maritime information.

Table 3. Cont.

Analytical Layer	Observable Characteristic	Replay Attack Manifestation	Ghost Ship Attack Manifestation	Operational Challenge	Required Operational Capability
Network Infrastructure Layer	Information provenance	Legitimate AIS data are replayed through intermediate communication paths.	Fabricated AIS data are injected into communication or aggregation services.	Identify where information integrity may have been compromised.	Trace information provenance across the communication chain.
	Cross-source consistency	Different information sources provide inconsistent temporal views of legitimate traffic.	Independent information sources disagree because fabricated information is introduced.	Determine which operational picture is trustworthy.	Correlate information obtained from multiple independent sources.
AIS Protocol Layer	Protocol compliance	Messages remain protocol compliant because authentic transmissions are reused.	Synthetic messages attempt to emulate protocol-compliant behaviour.	Distinguish authentic protocol behaviour from fabricated communication.	Analyse protocol behaviour to identify indicators of message manipulation.
	Transmission timing	Authentic transmission intervals are preserved but correspond to historical activity.	Transmission timing is synthetically generated to appear credible.	Determine whether message timing reflects current vessel activity.	Assess temporal consistency of AIS transmissions.
	Message integrity	Original message structure remains unchanged.	Synthetic message construction may introduce inconsistencies or implementation artefacts.	Determine whether message construction indicates manipulation.	Validate message integrity and identify anomalies.
AIS Message Validation Layer	Vessel identity	Legitimate identities are reused with outdated operational information.	Identities may be fabricated, duplicated or impersonated.	Determine whether vessel identity is authentic.	Verify vessel identity using authoritative maritime information.
	Navigational information	Position, course and speed correspond to historical vessel movement.	Navigational data are artificially generated.	Determine whether reported navigation reflects genuine vessel movement.	Validate navigational information against independent evidence.
	Information validity	Authentic information is operationally outdated.	Newly generated information has no corresponding real-world counterpart.	Determine whether reported information remains operationally valid.	Assess information validity and operational relevance.
Operational Behaviour Layer—Operational Assessment	Temporal behaviour	Vessel behaviour reflects genuine historical activity displaced in time.	Vessel behaviour is entirely synthetic.	Determine whether observed behaviour reflects the current operational situation.	Reconstruct the temporal evolution of vessel activity.

Table 3. *Cont.*

Analytical Layer	Observable Characteristic	Replay Attack Manifestation	Ghost Ship Attack Manifestation	Operational Challenge	Required Operational Capability
Operational Behaviour Layer—Operational Assessment	Behavioural plausibility	Historical behaviour remains realistic but inconsistent with the present operational context.	Synthetic behaviour attempts to imitate realistic vessel operations but may violate expected patterns.	Determine whether observed behaviour is operationally plausible.	Evaluate behavioural patterns against expected maritime operations.
	Situational picture consistency	Replay attacks reduce confidence in the timeliness of the recognised maritime picture.	Ghost ship attacks reduce confidence in the authenticity of the recognised maritime picture.	Maintain an accurate and trustworthy recognised maritime picture.	Differentiate temporal inconsistencies from authenticity-related inconsistencies when assessing the operational picture.
Operational Behaviour Layer—Operational Decision-Making	Decision confidence	Operators must determine whether delayed legitimate information can still support operational decisions.	Operators must determine whether reported vessels actually exist before initiating an operational response.	Support confident decision-making under uncertainty.	Integrate multiple independent sources of evidence to support operational judgement.
	Response prioritisation	Response activities focus on validating information freshness and reconstructing historical vessel activity.	Response activities focus on verifying vessel authenticity and assessing the operational consequences of fabricated contacts.	Select response actions appropriate to the underlying deception mechanism.	Prioritise operational response according to the characteristics and impact of the spoofing attack, including incidents in which replay and fabrication indicators coexist.

6. Capability-Driven Design of Maritime Cyber Range Exercises

6.1. From Capability Requirements to Exercise Design

The operational capability requirements presented in the previous section establish the competencies required to identify, interpret, and respond to replay and ghost ship attacks. The next step is to translate these analytical requirements into an instructional design that enables maritime personnel to develop and demonstrate these capabilities within a realistic operational environment. The proposed approach adopts a capability-driven design philosophy in which exercise objectives are directly informed by the operational reasoning requirements identified through the layered threat analysis.

This design philosophy recognises that effective cyber range training extends beyond exposing participants to technical attack scenarios. Instead, learning activities should reproduce the cognitive processes required during real maritime operations, including information validation, evidence correlation, behavioural assessment, and decision-making under uncertainty. Consequently, replay and ghost ship attacks are not treated simply as technical events to be detected, but as operational problems requiring participants to interpret multiple sources of evidence, distinguish between different deception mechanisms, and justify their operational decisions. For hybrid attacks, participants must also determine whether temporal and authentication inconsistencies indicate the simultaneous use of replayed and fabricated information.

Figure 3 illustrates the proposed capability-driven exercise design approach. The presented approach establishes a structured progression from threat analysis and operational capability requirements to learning objectives specification, exercise implementation, and

assessment. In doing so, it ensures that the cyber range exercise will directly support the development and evaluation of the identified operational analytical capabilities.



Figure 3. Proposed Capability-Driven Exercise Design Approach.

6.2. Deriving Learning Objectives from Operational Capability Requirements

The operational capability requirements presented in Section 5 define the knowledge, reasoning processes, and decision-making abilities required to identify, interpret, and respond to AIS replay and ghost ship attacks. From an instructional design perspective, these capabilities represent complementary components of broader operational competencies rather than isolated technical skills. Consequently, the proposed exercise design synthesises related operational analytical capabilities into a set of integrated learning objectives that reflect the reasoning processes required during real-world maritime cyber incidents.

This capability-driven approach shifts the focus of cyber range training beyond technical attack detection towards the development of higher-order analytical competencies. Participants are expected to validate maritime information, interpret technical and operational evidence, assess the plausibility of vessel behaviour, and make informed operational decisions under conditions of uncertainty. Accordingly, the learning objectives are designed to cultivate the analytical reasoning required to integrate multiple sources of evidence, evaluate competing explanations, and justify operational responses, thereby reflecting the cognitive demands of maritime cybersecurity operations.

Table 4 provides an illustrative example of how learning objectives can be systematically derived from the operational capability domains identified in Section 5. Instead of establishing a one-to-one correspondence between individual capabilities and learning objectives, the proposed approach groups complementary capabilities into broader instructional domains that reflect integrated operational reasoning processes. The table is not intended to be exhaustive; instead, it illustrates how operational analytical capabilities can be translated into meaningful educational objectives for cyber range exercise design.

Table 4. Learning objectives derived from the operational capability domains.

Capability Domain	Operational Capabilities (Derived from Section 5)	Learning Objective (LO)
Information Validation	Validate vessel identity, verify navigational information, assess message integrity.	LO1 Validate the authenticity and integrity of AIS information using multiple independent maritime information sources.
Temporal and Behavioural Analysis	Assess temporal validity, evaluate behavioural plausibility, distinguish legitimate vessel activity from deceptive behaviour.	LO2 Differentiate replay and ghost ship attacks through temporal, behavioural, and contextual analysis.
Evidence Correlation and Situational Assessment	Correlate technical and operational evidence, integrate information from multiple heterogeneous sources, develop situational awareness.	LO3 Produce evidence-based situational assessments by integrating information from multiple maritime data sources.
Operational Decision-Making	Interpret validated information, determine appropriate operational responses, justify decisions under uncertainty.	LO4 Select and justify appropriate operational response actions based on validated evidence and the evolving operational context.

The illustrative learning objectives presented in Table 4 demonstrate how the proposed capability-driven approach can be applied to transform operational capability requirements into measurable educational outcomes. In practice, the specific learning objectives adopted for a cyber range exercise will depend on the intended audience, training context, and desired level of complexity. For advanced exercises involving hybrid attacks, the learning objectives should also require participants to maintain competing hypotheses and revise their classification when both replay and fabrication indicators are present.

6.3. Capability-Driven Scenario Design

The learning objectives presented in the previous section provide the foundation for the design of cyber range exercises. The proposed approach adopts a capability-driven design in which every stage of an exercise is intended to provide participants with opportunities to demonstrate specific operational analytical capabilities. This ensures that an exercise reflects the reasoning processes required during real maritime cyber incidents, where operators must interpret incomplete and evolving information, validate evidence from multiple sources, and make operational decisions under conditions of uncertainty.

Accordingly, the exercise scenarios are designed to progressively increase both the complexity of the operational situation and the analytical demands placed upon participants. Instead of immediately revealing the nature of the cyber attack, participants are initially exposed to operational anomalies that require investigation. As additional information becomes available through successive scenario injects, participants are expected to continuously reassess their hypotheses, correlate evidence from different maritime information sources, and determine whether the observed anomalies are consistent with legitimate operational conditions, replay attacks, or ghost ship attacks.

Hybrid scenarios are expected to deliberately combine temporal and authenticity-related inconsistencies. For example, a trainer may reuse a legitimate historical vessel trajectory while modifying selected identity, timestamp, destination, or navigational fields. Initial evidence may indicate a replay attack, while later registry, radar, communication, or network evidence reveals that fabricated information has also been introduced. Such scenarios require participants to avoid premature classification and determine whether the incident represents replay, fabrication, a hybrid attack, or a non-malicious data anomaly.

This progressive disclosure of information encourages continuous situational assessment instead of focusing on isolated technical analysis. Each inject is designed to provide additional operational evidence that supports one or more learning objectives, requiring participants to validate maritime information, assess temporal consistency, evaluate vessel behaviour, correlate evidence across heterogeneous information sources, or justify operational decisions. Consequently, the sequence of injects mirrors the progressive reasoning process employed during real maritime incident response, allowing participants to iteratively construct, refine, and validate their understanding of the evolving operational situation.

This capability-driven approach also supports flexibility in exercise development. By maintaining the learning objectives as the primary design driver, individual scenarios, injects, and operational artefacts can be adapted to different training audiences, levels of expertise, or operational contexts without altering the underlying educational approach. The following section illustrates how operational evidence and supporting artefacts can be incorporated into the exercise to facilitate the development and assessment of the intended operational analytical capabilities.

6.4. Operational Evidence and Exercise Artefacts

The capability-driven scenario design described in the previous section is realised through the systematic provision of operational evidence and supporting exercise artefacts that enable participants to achieve the intended learning objectives. An exercise should provide multiple sources of operational information enriched with representative technical information that must be analysed, validated, and interpreted throughout the scenario. This approach reflects the realities of maritime operations, where cybersecurity incidents are rarely identified through a single observable event but instead emerge through the correlation of multiple, often incomplete, sources of operational and technical evidence.

Accordingly, an exercise should incorporate a diverse set of operational artefacts that collectively represent the information environment available to maritime operators during an evolving AIS spoofing incident. These artefacts include maritime operational data, navigational information, system logs, communication records, and other supporting information sources, each containing representative technical information that participants must interpret within its operational context. The selection of these artefacts and the associated technical information is informed by the operational analytical capabilities identified in Section 5, the learning objectives derived in Section 6.2, and established literature on AIS spoofing, maritime situational awareness, and cyber range training. This ensures that every information source contributes to the intended educational outcomes while remaining grounded in realistic operational practice.

Table 5 presents illustrative examples of operational evidence, representative technical information, and their relationship to the learning objectives introduced in Section 6.2. The examples are not intended to be exhaustive; they demonstrate how operational artefacts can be enriched with realistic technical information and aligned with evidence from the literature to support the development of the intended operational analytical capabilities within a capability-driven exercise design.

Table 5. Illustrative Operational Evidence Supporting Capability-Driven Exercise Design.

Operational Evidence	Illustrative Technical Information	Associated Learning Objective(s)	Supporting Literature
Live AIS transmissions	MMSI consistency, SOG, COG, heading, message integrity	LO1	[4,12,17–19,22,38,39]
Historical AIS records	Timestamp continuity, repeated message patterns	LO2	[8,18,22,32,36,39,41]
Radar tracks	Position discrepancies, target correlation	LO3	[25,28,31,44,48]
Vessel registry information	MMSI registration, IMO number, call sign	LO1	[12,18,22,25,33,37,39]
GNSS information	Position mismatch, navigation discrepancies	LO3	[10,11,28,31,40,45,46,48]
Bridge observations and operator reports	Visual confirmation, traffic behaviour	LO2, LO3	[1,2,10,11,28,31,42,48]
Communication records	VHF reports, operational notifications	LO3	[1,10,11,28,42,48]
System and network logs	Authentication events, protocol anomalies	LO3, LO4	[4,15–17,27,30,38,42,43,48,52]

The operational evidence should be introduced progressively through a series of scenario injects that mirror the evolution of a real maritime cyber incident. Rather than providing all available information simultaneously, new artefacts and associated technical

information should be released as the operational situation develops, requiring participants to continuously update their situational understanding, reassess previously formed hypotheses, and justify their decisions using the evidence available at each stage of the exercise. This progressive disclosure reinforces analytical reasoning, promotes evidence-based decision-making, and reflects the dynamic nature of operational incident response within maritime environments.

6.5. Guidelines for Capability-Driven Assessment

This section provides guidelines for assessing participant performance in alignment with the operational capabilities and learning objectives defined in the preceding sections. Rather than measuring success solely by the successful completion of scenario tasks or the identification of an AIS spoofing attack, the proposed approach considers how participants apply relevant capabilities such as analytical reasoning, evidence interpretation, decision-making, and communication throughout the exercise. This ensures that assessment remains aligned with the operational capability requirements identified in Section 5 and the learning objectives derived in Section 6.2.

Assessment should be performed continuously during the execution of an exercise by combining instructor observations with objective evidence automatically generated by the cyber range environment. Relevant participant activities include interpreting operational artefacts, formulating and revising analytical hypotheses, correlating information obtained from multiple maritime sources, distinguishing among competing explanations, justifying operational decisions, and communicating findings throughout the incident response process. In hybrid scenarios, trainers should also examine whether participants recognise the coexistence of replay and fabrication indicators and revise their initial classification as additional evidence becomes available. This combination of qualitative observations and system-generated evidence enables a comprehensive evaluation of both technical proficiency and operational competence while maintaining consistency with realistic maritime operational practice.

Table 6 presents illustrative assessment areas corresponding to the operational capability domains introduced in this work. For each domain, it identifies the participant activities that may indicate the application of the relevant capability and the exercise records that may support trainer judgement. A scoring model is outside of the scope of this work. For this reason, Table 6 does not prescribe numerical scores, weighting, pass thresholds, or standardised performance levels. Future work will focus on designing, calibrating and validating a relevant scoring model.

Table 6. Illustrative Areas for Competency-Driven Assessment.

Capability Domain	Illustrative Assessment Criterion	Explanation	Examples of Assessment Evidence
Information Validation	Validates the authenticity and integrity of AIS information.	Validates AIS information across radar, GNSS and registry data.	Instructor observation, participant submissions, incident classification, AIS logs
Temporal & Behavioural Analysis	Differentiates replay and ghost ship attacks using temporal and behavioural evidence.	Identifies repeated messages, outdated vessel behaviour, timing inconsistencies, or implausible movement.	AIS logs, PCAP analysis, Wireshark captures, forensic artefacts
Evidence Correlation & Situational Assessment	Integrates heterogeneous operational information before reaching conclusions.	Develops and revises hypotheses using evidence from multiple sources.	Incident reports, screenshots, collected artefacts, SOC evidence repository

Table 6. *Cont.*

Capability Domain	Illustrative Assessment Criterion	Explanation	Examples of Assessment Evidence
Operational Decision-Making	Selects and justifies appropriate operational response actions.	Recommends proportionate response actions supported by available evidence.	Exercise timeline, instructor rubric, decision logs
Communication	Clearly communicates findings, recommendations and operational rationale.	Effectively briefs stakeholders and documents incident findings.	Exercise communications, chat logs, voice recordings, incident reports

The assessment process can follow four stages. First, trainers select the capability domains relevant to the exercise learning objectives. Second, they identify the participant activities through which these capabilities may become observable. Third, they collect the corresponding cyber-range records and operational artefacts. Finally, they use the collected evidence during assessment and debriefing to examine how participants interpreted the incident and justified their decisions. Evidence from multiple sources should be considered together, since a single action or artefact may not adequately represent the participant's analytical process.

The assessment outcomes can provide structured feedback during the post-exercise debriefing, enabling participants to reflect on both the decisions they made and the reasoning processes that informed those decisions. In doing so, the assessment supports reflective learning by identifying strengths, highlighting capability gaps, and encouraging participants to justify their analytical reasoning using the available operational evidence. By explicitly aligning assessment with the operational capabilities, learning objectives, and exercise design presented throughout this section, the proposed methodology establishes a coherent and transparent approach for evaluating the capabilities required to detect, analyse, and respond to AIS spoofing incidents within realistic maritime operational environments.

7. Cyber Range Implementation—Illustrative Example

The proposed capability-driven exercise design was implemented using a specialised Maritime Cyber Range [52], a virtualised environment designed to support realistic maritime cybersecurity training. While the design approach presented in the previous sections is platform-independent and may be implemented using different cyber range technologies, the Maritime Cyber Range presented in this paper serves as an illustrative implementation demonstrating how the derived operational capabilities, learning objectives, scenario design, operational evidence, and assessment strategy can be realised within a controlled training environment.

Figure 4 illustrates the architecture of the cyber range implementation. The environment integrates real and simulated AIS infrastructure with maritime surveillance and navigation components, including GNSS and radar data sources, AIS processing servers, command-and-control (C2) systems, and virtualised network infrastructure. The architecture also incorporates an adversary workstation used to emulate representative AIS spoofing techniques and an instructor-managed environment that enables the controlled execution of training scenarios. By combining real and synthetic maritime data, the cyber range provides a representative and configurable operational context in which replay and ghost ship attacks can be executed repeatedly, consistently, and safely while maintaining full control over scenario progression and evidence release.

The cyber range used in this study is designed to provide a representative operational environment rather than a complete physical replica of a ship bridge or maritime surveil-

lance centre. Its fidelity is therefore considered in relation to the information and analytical conditions required by the exercise learning objectives. In particular, the environment integrates AIS, radar, GNSS, network, and C2 information, enabling participants to analyse temporal and cross-source inconsistencies within a controlled operational setting. The representation of these heterogeneous information sources allows expected variations in sensor observations and system timing to be incorporated into scenario design rather than treating every cross-source inconsistency as evidence of malicious activity. These characteristics are particularly relevant to replay and ghost ship scenarios, where distinguishing malicious manipulation from normal operational or technical variation forms part of the required analytical reasoning.




Figure 4. Cyber Range Architecture.

Unlike conventional attack-centric cyber range exercises, the implementation is structured around the operational capabilities identified in Section 5 rather than the technical execution of individual attacks. Consequently, each scenario is designed to progressively expose participants to operational evidence that supports information validation, evidence correlation, situational assessment, and decision-making under uncertainty.

To illustrate how the proposed design approach is operationalised, Table 7 presents the progression of a representative ghost ship scenario. The exercise begins with participants monitoring routine maritime traffic using operational systems representative of those deployed in contemporary maritime organisations. During the initial phase, the recognised maritime picture is consistent across all available information sources. The instructor subsequently introduces a fabricated AIS target through the adversary environment, presenting a vessel with plausible navigational characteristics, including realistic identity, speed, heading, and voyage information. Rather than immediately revealing the deception, operational evidence is progressively disclosed throughout the exercise. Participants first observe the suspicious vessel on the AIS display before receiving additional information from independent sources, including radar observations, vessel registry information, communication records, and supporting operational documentation. As new evidence becomes available, participants are required to continuously reassess the evolving operational picture, validate the credibility of conflicting information, formulate evidence-based hypotheses, and determine whether the observed behaviour is consistent with a ghost ship attack. This progressive disclosure reflects real operational environments, where complete information is rarely available simultaneously, and directly operationalises the capability requirements and learning objectives derived in Sections 5 and 6.2.

Table 7. Progressive Execution of a Representative Ghost Ship Scenario.

Phase	Operational Activity	Operational Capability
Phase 1—Initial Maritime Picture	Participants monitor routine maritime traffic using AIS, radar and GNSS.	Establish operational baseline.
Phase 2—Suspicious AIS Activity	A fabricated vessel is introduced into the AIS environment with plausible navigational characteristics.	Recognise potential anomalies.
Phase 3—Multisource Validation	Additional evidence becomes available through radar, GNSS and other independent sources.	Validate information across multiple sources.
Phase 4—Evidence Correlation	Registry information, communications and operational reports are progressively released.	Correlate heterogeneous evidence and assess credibility.
Phase 5—Operational Decision	Participants determine whether the anomaly represents a ghost ship attack and justify the proposed response.	Situational assessment and evidence-based decision-making.

A similar implementation strategy is adopted for replay attack scenarios, although the required operational reasoning differs. Rather than detecting fabricated vessels, participants must recognise temporal inconsistencies introduced through the replay of previously legitimate AIS transmissions. Consequently, the exercise emphasises the validation of temporal information, identification of duplicated or outdated vessel behaviour, and correlation of historical and current operational observations. This distinction reflects the taxonomy and layered threat analysis presented earlier in the paper and demonstrates how different AIS spoofing mechanisms require different operational capabilities despite belonging to the same attack class.

Table 8 summarises how the replay, ghost ship, and hybrid scenarios can be translated into capability-driven cyber range exercises. The scenarios are deliberately designed to exercise different analytical and decision-making capabilities by varying the sequence of operational evidence and the reasoning required to identify each attack.

Table 8. Capability-Driven Comparison of Ghost Ship, Replay, and Hybrid Attack Scenarios.

Ghost Ship Scenario	Replay Attack Scenario	Hybrid Attack Scenario
Fabricated vessel introduced into the operational picture.	Previously legitimate AIS transmissions are replayed.	A legitimate historical trajectory is combined with modified identity, timing, or navigational information.
Radar fails to confirm vessel presence.	Temporal inconsistencies emerge between historical and current traffic.	Temporal inconsistencies coexist with identity or multisensor inconsistencies.
Registry and communication records reveal inconsistencies.	Duplicate vessel tracks or outdated navigation behaviour become evident.	Early evidence supports replay, while later evidence reveals fabricated or altered information.
Participants validate information across multiple independent sources.	Participants validate temporal consistency across multiple information sources.	Participants maintain competing hypotheses and reassess the classification as evidence develops.

Figure 5 illustrates components from the specialised maritime cyber range that simulates representative operational conditions during ghost ship and replay attacks.




Figure 5. Examples of Ghost Ship and Replay Attacks. (a) Ghost Ship attack creates AIS vessels near Port Facilities. (b) Replay attack execution using data from real sensor.

Throughout execution, instructors manage scenario progression through the controlled release of injects, operational evidence, and supporting artefacts while adjusting exercise complexity to the intended learning outcomes. Simultaneously, the cyber range records participant interactions, analytical decisions, communications, and system activities, generating objective evidence that supports structured debriefing and the capability-based assessment process presented in the previous section.

8. Limitations and Future Work

This work has three main limitations that should be considered when interpreting its contributions. First, the proposed taxonomy and layered threat analysis focus specifically on AIS replay and ghost ship attacks. Although hybrid attacks combining characteristics of both categories are considered, the taxonomy is not intended to provide an exhaustive classification of AIS manipulation or spoofing techniques. Other forms of manipulation, including attacks involving the falsification of live vessel data through GNSS manipulation, fall outside the scope of this study.

Second, the proposed capability-driven approach has been developed by translating the distinguishing technical and operational characteristics identified through the threat analysis into operational capability requirements, learning objectives, scenario design considerations, and assessment evidence. While the cyber range exercise design illustrates how this process can be operationalised, the proposed approach has not yet been empirically validated through its delivery to maritime practitioners. Consequently, this work does not make claims regarding training effectiveness or measurable capability improvement.

Third, the assessment guidelines identify relevant participant activities and sources of evidence that can support trainer judgement but do not prescribe a validated scoring model, performance thresholds, or standardised proficiency levels. The interpretation of collected evidence therefore remains dependent on trainer judgement and the specific objectives and context of an exercise.

Future work will focus on the empirical validation of the proposed approach through the delivery and evaluation of cyber range exercises involving maritime practitioners. This will include examining its effectiveness in developing the intended operational capabilities, evaluating participant performance across different experience levels, and refining the scenario design and assessment guidelines based on evidence collected during training activities. Further work will also focus on designing, calibrating, and validating a scoring model for capability-driven assessment and investigating opportunities for automated evidence collection and performance assessment within maritime cyber range environments. Finally, the applicability of the proposed approach to other maritime cyber threats will be

investigated to determine the extent to which the threat-analysis-to-capability-development process can support broader maritime cybersecurity training contexts.

9. Conclusions

This paper examined the distinct characteristics of replay and ghost ship attacks and their implications for capability-driven maritime cybersecurity training. A focused taxonomy and layered threat analysis were developed to differentiate these attacks according to their technical characteristics, operational manifestations, and defensive implications across the AIS ecosystem. The analysis highlights how replay and ghost ship attacks create different challenges for detection and operational reasoning, while also accounting for hybrid scenarios in which characteristics of both attack types may coexist.

Building upon this analysis, the paper identified the operational capabilities required to recognise, differentiate, and respond to these attacks and proposed a structured approach for translating them into learning objectives, differentiated scenario designs, operational evidence, and capability-driven assessment. The cyber range exercise design illustrates how this process can be applied in practice, providing trainers with a systematic approach for moving from threat analysis to the development and assessment of operational capabilities. Although demonstrated in the context of AIS replay and ghost ship attacks, the proposed approach provides a foundation that may be adapted to other maritime cyber threats where differences in attack characteristics have implications for operational reasoning and training design.

Ultimately, effective preparation for AIS spoofing incidents requires more than exposing participants to a generic spoofing scenario. It requires exercises that reflect the distinguishing characteristics of the threats participants may encounter and require them to interpret evidence, differentiate between competing explanations, and adapt their decisions as the operational picture evolves. By establishing an explicit connection between threat characteristics, operational capabilities, and exercise design, this work provides a foundation for developing maritime cybersecurity training that more closely reflects the analytical demands of real-world AIS spoofing incidents.

Author Contributions: Conceptualization, G.P., S.T., E.S. and S.S.; methodology, G.P., E.S. and S.S.; investigation, G.P.; resources, G.P. and S.T.; writing—original draft preparation, G.P.; writing—review and editing, E.S. and S.S.; visualization, G.P. and E.S.; supervision, E.S. and S.S.; project administration, S.S.; funding acquisition, S.S. All authors have read and agreed to the published version of the manuscript.

Funding: Funded by the European Union (Grant Agreement Nr. 101128049, SecAwarenessTruss Project). Views and opinions expressed are, however, those of the authors only and do not necessarily reflect those of the European Union or the Health and Digital Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: No new data were created or analyzed in this study. Data sharing is not applicable.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Stupak, T. Influence of automatic identification system on safety of navigation at sea. *TransNav* **2014**, *8*, 337–341. [CrossRef]
2. Kim, I.; Lee, S.; Youn, I. Adopting the audible alert system for the electronic chart display and information system for improvement of early navigational situation awareness. *J. Int. Marit. Saf. Environ. Aff. Shipp.* **2020**, *4*, 177–186. [CrossRef]

3. Emmens, T.; Amrit, C.; Abdi, A.; Ghosh, M. The promises and perils of Automatic Identification System data. *Expert Syst. Appl.* **2021**, *178*, 114975. [[CrossRef](#)]
4. Balduzzi, M.; Pasta, A.; Wilhoit, K. A security evaluation of AIS automated identification system. In Proceedings of the 30th Annual Computer Security Applications Conference, New Orleans, LA, USA, 8–12 December 2014; pp. 436–445.
5. Goudossis, A.; Katsikas, S.K. Towards a secure automatic identification system (AIS). *J. Mar. Sci. Technol.* **2019**, *24*, 410–423. [[CrossRef](#)]
6. Vancé, Z.; Renaud, A.; Napoli, A. Predicting AIS reception using tropospheric propagation forecast and machine learning. In Proceedings of the 2022 IEEE USNC-URSI Radio Science Meeting (Joint with AP-S Symposium), Denver, CO, USA, 10–15 July 2022; pp. 70–71.
7. Kirichevskaya, A. A Network Capacity Perspective on Digital Signatures for AIS. Diploma Thesis, Technische Universität Wien, Vienna, Austria, 2025.
8. Zhao, D.; Shi, Y.; Ding, S.X.; Li, Y.; Fu, F. Replay attack detection based on parity space method for cyber-physical systems. *IEEE Trans. Autom. Control* **2024**, *70*, 2390–2405. [[CrossRef](#)]
9. Cichocki, R.; Neumann, T. Cybersecurity challenges and vulnerabilities of the automatic identification system in maritime transport. *Arch. Transp.* **2026**, *77*, 27–43. [[CrossRef](#)]
10. Androjna, A.; Brčko, T.; Pavić, I.; Greidanus, H. Assessing cyber challenges of maritime navigation. *J. Mar. Sci. Eng.* **2020**, *8*, 776. [[CrossRef](#)]
11. Oruç, A.; Kavallieratos, G.; Gkioulos, V.; Katsikas, S. Perspectives on the cybersecurity of the integrated navigation system. *J. Mar. Sci. Eng.* **2025**, *13*, 1087. [[CrossRef](#)]
12. Wolsing, K.; Roepert, L.; Bauer, J.; Wehrle, K. Anomaly detection in maritime AIS tracks: A review of recent approaches. *J. Mar. Sci. Eng.* **2022**, *10*, 112. [[CrossRef](#)]
13. Duzenli, E.; Akyuz, E.; Kayisoglu, G.; Bolat, P. An integrated risk-based modelling approach for prediction of AIS spoofing attacks on-board ship. *Ocean Eng.* **2025**, *342*, 122895. [[CrossRef](#)]
14. Ashfaq Uz Zaman, S.M.; Qamar, F.; Mohd, M.; Suhaimi, N.H.S.; Khandakar, A. A review of Automatic Identification System approaches for maritime cyber security. *Secur. Commun. Netw.* **2026**, *2026*, 5898106. [[CrossRef](#)]
15. Leite Junior, W.C.; de Moraes, C.C.; de Albuquerque, C.E.P.; Machado, R.C.S.; de Sá, A.O. A triggering mechanism for cyber-attacks in naval sensors and systems. *Sensors* **2021**, *21*, 3195. [[CrossRef](#)] [[PubMed](#)]
16. Kessler, G.C.; Craiger, P.; Haass, J.C. A taxonomy framework for maritime cybersecurity: A demonstration using the Automatic Identification System. *TransNav* **2018**, *12*, 429–437. [[CrossRef](#)]
17. Sciancalepore, S.; Tedeschi, P.; Aziz, A.; Di Pietro, R. Auth-AIS: Secure, flexible, and backward-compatible authentication of vessels AIS broadcasts. *IEEE Trans. Dependable Secur. Comput.* **2022**, *19*, 2709–2726. [[CrossRef](#)]
18. Iphar, C.; Napoli, A.; Ray, C. Detection of false AIS messages for the improvement of maritime situational awareness. In Proceedings of the OCEANS 2015—MTS/IEEE Washington, Washington, DC, USA, 19–22 October 2015; pp. 1–7.
19. Androjna, A.; Perkovič, M.; Pavić, I.; Mišković, J. AIS data vulnerability indicated by a spoofing case-study. *Appl. Sci.* **2021**, *11*, 5015. [[CrossRef](#)]
20. Mazzarella, F.; Vespe, M.; Alessandrini, A.; Tarchi, D.; Aulicino, G.; Vollero, A. A novel anomaly detection approach to identify intentional AIS on-off switching. *Expert Syst. Appl.* **2017**, *78*, 110–123. [[CrossRef](#)]
21. Mednikarov, B.; Tsonev, Y.; Lazarov, A. Analysis of cybersecurity issues in the maritime industry. *Inf. Secur. Int. J.* **2020**, *47*, 27–43. [[CrossRef](#)]
22. Iphar, C.; Ray, C.; Napoli, A. Data integrity assessment for maritime anomaly detection. *Expert Syst. Appl.* **2020**, *147*, 113219. [[CrossRef](#)]
23. Murray, B.; Perera, L.P. An AIS-based deep learning framework for regional ship behavior prediction. *Reliab. Eng. Syst. Saf.* **2021**, *215*, 107819. [[CrossRef](#)]
24. Alessandrini, A.; Mazzarella, F.; Vespe, M.; Tarchi, D.; Aulicino, G.; Vollero, A. Mining vessel tracking data for maritime domain applications. In Proceedings of the 2016 IEEE 16th International Conference on Data Mining Workshops (ICDMW), Barcelona, Spain, 12–15 December 2016; pp. 361–367.
25. d’Afflisio, E.; Braca, P.; Chisci, L.; Battistelli, G.; Willett, P. Maritime anomaly detection of malicious data spoofing and stealth deviations from nominal route exploiting heterogeneous sources of information. In Proceedings of the 2021 IEEE 24th International Conference on Information Fusion (FUSION), Sun City, South Africa, 1–4 November 2021; pp. 1–7.
26. Volkov, V.V.; Ol’khovik, E.O.; Fedosenko, Y.S. Models of cyberattacks on the automatic identification system network. *Vestn. Gos. Univ. Morsk. Rechn. Flot. Im. Admirala S.O. Makarova* **2025**, *17*, 641–652. [[CrossRef](#)]
27. Goudosis, A.; Katsikas, S. Secure Automatic Identification System (SecAIS): Proof-of-concept implementation. *J. Mar. Sci. Eng.* **2022**, *10*, 805. [[CrossRef](#)]
28. Potamos, G.; Stavrou, E.; Stavrou, S. Enhancing maritime cybersecurity through operational technology sensor data fusion: A comprehensive survey and analysis. *Sensors* **2024**, *24*, 3458. [[CrossRef](#)] [[PubMed](#)]

29. Iphar, C.; Napoli, A.; Ray, C.; Alincourt, E.; Brosset, D. Risk analysis of falsified automatic identification system for the improvement of maritime traffic safety. In *Risk, Reliability and Safety: Innovating Theory and Practice*; Walls, L., Revie, M., Bedford, T., Eds.; CRC Press: Boca Raton, FL, USA, 2016; pp. 606–613.
30. Kessler, G.C.; Zorri, D.M. AIS spoofing: A tutorial for researchers. In Proceedings of the 2024 IEEE 49th Conference on Local Computer Networks (LCN), Caen, France, 8–10 October 2024; pp. 1–7.
31. Androjna, A.; Perkovič, M. Impact of spoofing of navigation systems on maritime situational awareness. *Trans. Marit. Sci.* **2021**, *10*, 361–373. [\[CrossRef\]](#)
32. Pallotta, G.; Vespe, M.; Bryan, K. Vessel pattern knowledge discovery from AIS data: A framework for anomaly detection and route prediction. *Entropy* **2013**, *15*, 2218–2245. [\[CrossRef\]](#)
33. Zheng, H.; Hu, Q.; Yang, C.; Mei, Q.; Wang, P.; Li, K. Identification of spoofing ships from Automatic Identification System data via trajectory segmentation and isolation forest. *J. Mar. Sci. Eng.* **2023**, *11*, 1516. [\[CrossRef\]](#)
34. Thomas, A.; Chiego, C. *Maritime Cybersecurity: AIS Manipulation Motivations in the Maritime Domain*; Center for International Maritime Security: Washington, DC, USA, 2022.
35. Li, M.; Zhou, J.; Chattopadhyay, S.; Goh, M. Maritime cybersecurity: A comprehensive review. *arXiv* **2024**, arXiv:2409.11417.
36. Arguedas, V.F.; Mazzarella, F.; Vespe, M. Spatio-temporal data mining for maritime situational awareness. In Proceedings of the OCEANS 2015—Genova, Genoa, Italy, 18–21 May 2015; pp. 1–8.
37. Jakovlev, S.; Daranda, A.; Voznak, M.; Lektuers, A.; Eglynas, T.; Juisis, M. Analysis of the possibility to detect fake vessels in the Automatic Identification System. In Proceedings of the 2020 61st International Scientific Conference on Information Technology and Management Science of Riga Technical University (ITMS), Riga, Latvia, 15–16 October 2020; pp. 1–5.
38. Louart, M.; Szkolnik, J.-J.; Boudraa, A.-O.; Le Lann, J.-C.; Le Roy, F. Detection of AIS messages falsifications and spoofing by checking messages compliance with TDMA protocol. *Digit. Signal Process.* **2023**, *136*, 103983. [\[CrossRef\]](#)
39. Zhang, T.; Zhao, S.; Cheng, B.; Chen, J. Detection of AIS closing behavior and MMSI spoofing behavior of ships based on spatiotemporal data. *Remote Sens.* **2020**, *12*, 702. [\[CrossRef\]](#)
40. Westbrook, T. Radiofrequency interference strategies targeting marine navigation systems: Political motives and consequences. *J. Balt. Secur.* **2023**, *9*, 69–97.
41. Bian, G.; Gao, X. Enhancing unmanned marine vehicle security: A periodic watermark-based detection of replay attacks. *Appl. Sci.* **2024**, *14*, 8298. [\[CrossRef\]](#)
42. Potamos, G.; Theodoulou, S.; Stavrou, E.; Stavrou, S. Maritime cyber threats detection framework: Building capabilities. In *Information Security Education—Adapting to the Fourth Industrial Revolution*; Drevin, L., Miloslavskaya, N., Leung, W.S., von Solms, S., Eds.; Springer International Publishing: Cham, Switzerland, 2022; pp. 107–129.
43. Bothur, D.; Zheng, G.; Valli, C. A critical analysis of security vulnerabilities and countermeasures in a smart ship system. In Proceedings of the 15th Australian Information Security Management Conference, Perth, Australia, 5–6 December 2017; pp. 81–87.
44. Bruin, E.R. On Propagation Effects in Maritime Situation Awareness: Modelling the Impact of North Sea Weather Conditions on the Performance of AIS and Coastal Radar Systems. Master's Thesis, Utrecht University, Utrecht, The Netherlands, 2016.
45. Park, S.; Cho, D.; Son, P.-W. Wide-area GNSS spoofing and jamming detection using AIS-derived spatiotemporal integrity monitoring. *arXiv* **2026**, arXiv:2603.11055.
46. Faragher, R.; MacDoran, P.F.; Mathews, M.B. Spoofing mitigation, robust collision avoidance, and opportunistic receiver localisation using a new signal processing scheme for ADS-B or AIS. In Proceedings of the 27th International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS+ 2014), Tampa, FL, USA, 8–12 September 2014; pp. 858–868.
47. Afenyo, M.; Caesar, L.D. Maritime cybersecurity threats: Gaps and directions for future research. *Ocean Coast. Manag.* **2023**, *236*, 106493. [\[CrossRef\]](#)
48. Vasan, D.; Hammoudeh, M.; Ahmed, A.F.; Naeem, H. Cyber-attacks: Securing ship navigation systems using multi-layer cross-validation defense. *Comput. Secur.* **2026**, *160*, 104706. [\[CrossRef\]](#)
49. Jones, K.D.; Tam, K.; Papadaki, M. Threats and impacts in maritime cyber security. *Eng. Technol. Ref.* **2015**, *1*, 1–5. [\[CrossRef\]](#)
50. Last, P.; Hering-Bertram, M.; Linsen, L. How Automatic Identification System (AIS) antenna setup affects AIS signal quality. *Ocean Eng.* **2015**, *100*, 83–89. [\[CrossRef\]](#)
51. Abdelmagid, A.M.; Javadnejad, F.; McShane, M.; Diaz, R.; Pinto, C.A. A new cyber risk identification and assessment approach of the maritime cyber risks. *Enterp. Inf. Syst.* **2025**, *19*, 2524848. [\[CrossRef\]](#)
52. Potamos, G.; Peratikou, A.; Stavrou, S. Towards a maritime cyber range training environment. In Proceedings of the 2021 IEEE International Conference on Cyber Security and Resilience (CSR), Rhodes, Greece, 26–28 July 2021; pp. 180–185.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.