Zero Trust at Sea: Security Gains and the Risk of Operational Disruption without Fallback Mechanisms

BJ Croteau, CB Landis, A Srinivasan

Cyber Science Department
United States Naval Academy
Annapolis, Maryland, United States of America

Email: croteau@usna.edu; clandis@usna.edu; srinivas@usna.edu

Abstract: *Zero-Trust Architecture (ZTA) has emerged as a foundational cybersecurity paradigm with wide adoption. However, assumptions made in shore-based enterprise networks pose challenges when applying zero-trust principles in maritime environments. In maritime systems, operations depend on intermittent, low-bandwidth, and often contested communications. Thus, strict zero trust enforcement can unintentionally disrupt mission-critical functions. This paper examines the benefits and limitations of applying ZTA in maritime contexts, with particular attention to safety-critical operations and military. This work outlines a maritime-adapted zero trust framework that integrates degraded-mode operation and fallback mechanisms to preserve operational continuity while maintaining security objectives and safe mission completion.*

Keywords: *Zero Trust, Architecture, Maritime, Framework, Tactical, Cybersecurity, Safety, Security, Assurance, Navy*

Introduction

Zero-Trust Architecture (ZTA) represents a paradigm shift in cybersecurity, necessitated by the erosion of traditional network perimeters. While its principles are being rapidly adopted across the defence, government, and critical infrastructure sectors, the assumptions underlying high-connectivity enterprise networks often fail when applied to the maritime domain. ZTA shifts the security focus from implicit trust, historically based on the 'air-gap' or physical location, to explicit distrust. This is enforced through continuous authentication, authorisation, and monitoring of every subject and device. Although ZTA offers a superior security paradigm, implementation without due diligence can result in inadvertent access denials, potentially compromising system availability, operational safety, or mission success.

Implementing ZTA aboard ships presents significant challenges, as ships' operations routinely experience intermittent low-bandwidth and contested communications. In such environments, rigid zero-trust enforcement can unintentionally disrupt safety-critical functions if connectivity to centralised, shore-side security services is degraded. This paper examines the benefits and limitations of applying ZTA in maritime contexts and can be generalised to other military and safety-critical operations. Building upon the tactical zero-trust architecture developed for the U.S. Marine Corps by Oshiro (2023), this work proposes a maritime ZTA framework (mZTAf) that integrates degrad-

ed-mode operations and fallback mechanisms to preserve operational continuity. While previous research has addressed tactical Zero Trust in isolated pockets, this paper proposes the mZTAf as a comprehensive architectural bridge, designed to unify the high-connectivity ambitions of modern Zero Trust with the disconnected, high-consequence realities of maritime operations.

Ships at sea operate without the benefit of reliable, high-speed Internet connectivity. Due to distance and mobility, these platforms rely on radio-frequency (RF) paths that are inherently vulnerable to weather, atmospheric attenuation, and saline conditions. Furthermore, line-of-sight satellite communication (SATCOM) is fragile; maritime-specific constraints, such as mast shadowing, vessel movement, and transitioning between satellite footprints, frequently interrupt signals. These environmental impairments threaten the persistent links to shore-side identity providers, necessitating a framework capable of functioning as an autonomous 'Local Trust Anchor'.

Furthermore, maritime environments feature a tight integration of Information Technology (IT) and Operational Technology (OT), encompassing navigation, engineering controls, and, on naval vessels, complex combat systems. A maritime ZTA must, therefore, explicitly address the protection and segmentation of OT assets alongside traditional IT resources. Because the mZTAf incorporates cyber-physical systems (CPS), digital policy decisions can have immediate and irreversible physical consequences, such as the loss of propulsion or steering, affecting the safety of the ship and its surroundings. Consequently, cyber-physical risk management within the mZTAf becomes a primary safety requirement rather than a mere extension of traditional IT security practices.

This work contributes the following:

1. An analysis of Tactical ZTA requirements within the maritime environment;
2. The development of a resilient ZTA framework tailored to maritime environments (mZTAf); and
3. A comprehensive use-case evaluation of the proposed framework.

The paper is organised as follows: the next section explores the existing body of research in this active field, providing the theoretical context for zero trust. This is followed by an evaluation of how a proposed 'Tactical ZTA' can be adapted to the unique rigours of the maritime environment. Subsequently, the authors outline the proposed mZTAf and analyse its effectiveness through several distinct use cases, highlighting both operational strengths and potential limitations. The paper concludes with a summary of findings and a discussion on future research directions.

Background

A ZTA relies on a mindset shift from traditional, static cybersecurity boundary defences toward a focus on users, assets, and resources. As such, ZTA is not strictly a technical architecture for organisations to adopt but an architecture of culture and paradigms concerning system design and operation. Although the overarching cybersecurity goal remains the same, that is, "to prevent unauthorised access to data and services", the ZTA method strives to achieve this goal by making "access control enforcement as granular as possible" (Rose *et al*. 2020, p. 4). This goal is also consistent with the International Maritime Organization's (IMO's) goal for maritime cyber risk management (International Maritime Organization 2025, §3.2).

Kindervag, Balaouras, and Coit (2010) first coined the term "zero trust" and proposed three concepts on which a ZTA must depend:

1. Ensure that all resources are accessed securely regardless of location.
2. Adopt a least privilege strategy and strictly enforce access control.
3. Inspect and log all traffic. (Kindervag, Balaouras, & Coit 2010, pp. 8-9)

Over time, researchers have continued developing these concepts, theoretically and practically (Kang *et al*. 2023). Today, many would point to the U.S. National Institute of Standards and Technology (NIST)'s Special Publication (SP) on the topic as a reliable source of ZTA concepts. Rose *et al*. (2020) defined a set of seven tenants of ZTAs that expand and elaborate on the initial three concepts:

1. All data sources and computing services are considered resources.
2. All communication is secured regardless of network location.
3. Access to individual enterprise resources is granted on a per-session basis.
4. Access to resources is determined by dynamic policy—including the observable state of client identity, application/service, and the requesting asset—and may include other behavioural and environmental attributes.
5. The enterprise monitors and measures the integrity and security posture of all owned and associated assets.
6. All resource authentication and authorisations are dynamic and strictly enforced before access is allowed.
7. The enterprise collects as much information as possible about the current state of assets, network infrastructure, and communications and uses it to improve its security posture. (Rose *et al*. 2020, pp. 6-7)

Various organisations, including NIST, have worked on developing ZTA concepts, including the U.S. Department of Defense (DoD) (Defense Information Systems Agency and National Security Agency Zero Trust Engineering Team 2022). The DoD Zero Trust Reference Architecture references NIST SP 800-207 throughout in developing a DoD approach to implementing ZTAs. To align better with DoD's military perspective, it organises the ZTA concepts and tenants into five tenants and seven pillars that focus on cybersecurity in more militaristic terms, like "Assume a Hostile Environment" (p. 21). Oshiro (2023) fittingly uses this DoD reference as a primary source in his analysis and for implementing a tactical ZTA—in the U.S. Marine Corps, these are known as 'tactical clouds'. The DoD Zero Trust Reference Architecture is among the most comprehensive ZTA references available publicly and, despite its use of military terms, could benefit many organisations that are developing and implementing their own ZTAs.

Related Works

To support this work's claim of presenting the first holistic ZTA framework specifically designed for the maritime environment, the current state of research in both Zero Trust and maritime cybersecurity is reviewed, with particular attention to the gap at their intersection.

Sunkara (2025) analysed and compared various ZTA standards, examining implementations across five large organisations in government, technology, and finance. The study highlighted how different solutions emphasise particular ZTA pillars and identified both benefits and persistent chal-

lenges. A critical challenge, particularly relevant to the maritime domain, is the integration of ZTA with legacy systems.

Google's internal Zero Trust implementation, BeyondCorp (Osborn *et al.* 2016), serves as a well-documented model in a large corporate environment. This approach leverages cloud-based technologies and reflects a broader trend in ZTA research toward Secure Access Service Edge frameworks (Aiello & Rimal 2023), which integrate network provisioning and information security, such as access control, in a unified manner. While highly effective for cloud-centric networks with diverse bring-your-own-device demands, this model does not translate easily to the dynamic maritime environment.

A recent survey of the maritime cyber ecosystem (Li 2024) indicates that most research remains focused on component-specific solutions, such as Electronic Chart Display and Information System (ECDIS) and Automatic Identification System (AIS) systems. Proposed holistic approaches include machine/deep learning-based defences, cryptographic schemes, blockchain-based defences, and network-level protections. Although ZTA falls conceptually under network-level approaches, its treatment in maritime research is often superficial. For example, the International Association of Ports and Harbors (IAPH) Cyber Resilience Guidelines mention Zero Trust only briefly under "Robust Authentication" (IAPH Data Collaboration Committee 2025), alongside standard recommendations like "change default passwords" and "multi-factor authentication" (Meier *et al.* 2025).

Research on maritime Zero Trust Architectures (ZTAs) has primarily focused on specific enabling capabilities rather than on providing a comprehensive maritime framework as proposed in this work. For example, Al-Khalidi *et al.* (2025) emphasised communications security, proposing certificateless digital signatures to ensure confidential and authenticated interactions. Dodd (2023) explored implementation aboard a U.S. Navy ship, leveraging a purchased software-defined network (SDN) for micro-segmentation and using keystroke dynamics for continuous authentication. Similarly, Stewart (2023) applied ZT principles to maritime OT, such as continuous trust monitoring, highlighting that even partial adoption of ZTA concepts strengthens system resilience against vulnerabilities.

Some initial work has begun to address holistic frameworks for maritime ZTAs. Bargh, Omar, & Choenni (2024) investigated inter-organisational communications in shipping, emphasising the need for ZTA resources to establish ad hoc trust relationships, including with external entities, to enable end-to-end smart shipping. They identified challenges, such as unreliable or disconnected communication paths, suggesting that future ZTAs should support graceful degradation of functions when connectivity to policy engines is limited. Collectively, these studies highlight that, while foundational elements exist, a fully integrated, maritime-specific ZTA framework remains an open research gap.

ZTA Evolution: From USMC Tactical Edge to Maritime Environment

In establishing a firm architectural foundation, Oshiro (2023) identified robust core requirements for implementing a ZTA in U.S. Marine Corps tactical cloud environments. Though originally designed for ground-based operations, these requirements serve as the structural foundation for a maritime framework, as they were designed specifically to survive the volatility of combat operations. The following analysis acknowledges the suitability and necessary adaptation of Oshiro's requirements to the specific rigours of a ship operating in the maritime environment.

As a foundation, Oshiro (2023) highlighted several security and operational requirements for a tactical ZTA (TZTA). These requirements were developed using the STRIDE methodology (Kohnfelder & Garg 1999) to ensure they address a comprehensive range of cyber threats, including spoofing, denial of service, and privilege elevation, which are as prevalent at sea as they are on land. By prioritising these requirements, the mZTAf can transition from a legacy air-gap defence to a dynamic, verifiable security posture.

The tactical ZTA framework proposed by Oshiro

Oshiro's (2023) framework starts by defining a "Tactical ZTA", which is a self-contained enclave separate from the "Enterprise ZTA". As shown in **Figure 1**, the TZTA has its own local policy engine and administrator and, importantly, local policy enforcement points that also can be integrated with 'legacy' equipment and systems. The TZTA communicates with the enterprise through gateways to exchange information about local behaviour and, possibly, to receive updates for the TZTA policy engine. While this makes sense for a Marine Corps deployed unit connecting with higher headquarters and joint and allied communications and information networks, this model does not necessarily work for a ship. The ideas and components, however, could be adapted for the maritime environment, as described in the next section.

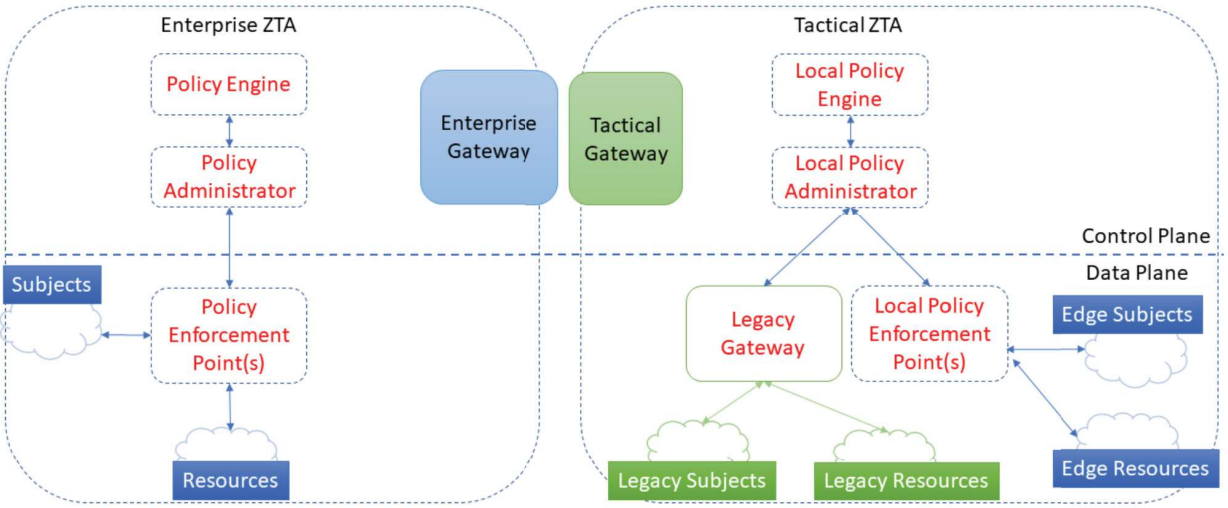

Figure 1: Zero trust architecture framework at the tactical edge (Oshiro 2023, p. 27)

Oshiro (2023) also described a trust framework where different networks can safely share information, consistent with ZTA principles, through a scoring system that identifies each system's security posture and data used and collected about users to perform internal authentication. Knowing the relative level of information the other system has about its users will help determine the most appropriate policy rules if information is to be shared. Deployed USMC units that work with allied military units or with civilian relief agencies would use this capability. Although a ship at sea may not require it, a ship in port or receiving a vendor's remote access to repair or to update specific equipment or systems may find this capability useful. These use cases (and more) will be analysed below.

Relevant technologies identified by Oshiro

Oshiro (2023) analysed several technologies for use in his TZTA. This subsection briefly addresses each technology's relevance to the maritime environment. Køien's (2021) Legacy Component Architecture that Oshiro adapted into his "Legacy Interface" could interface IT with maritime OT, including engineering industrial controllers and bridge systems like the AIS. The token-based scheme that incorporates blockchain methods (Hu 2022) is likely too heavyweight for most OT applications, but future advances in lightweight shared ledger technology may reduce this barrier. The Identify Federation Scheme from Hatakeyama, Kotani, and Okabe (2021) that Oshiro adopted into the partner trust framework is also less relevant due to the more isolated nature of maritime vessel operations. The Continuous Remote Attestation Framework for IoT (CRAFT), which is a software device attestation for legacy components, looks like a promising system to allow for some of the legacy systems described above to periodically assert their integrity in a scalable, low-overhead manner (Moreau, Conchon & Sauveron 2021). This system would likely need to be expanded or modified to include industrial control system (ICS) components, including programmable logic controllers (PLCs) or remote terminal units (RTUs), which may have less processing power and interfaces than the Internet of Things (IoT) systems envisioned in the original paper.

The network-based intrusion-detection system (IDS) proposed by Alalmaie, Nanda, and He (2022) employs multi-view feature set extraction and seems to be less relevant because sharing real-time data from a ship's network IDS with third parties ashore is an unreliable plan. The SysFlow framework (Hong *et al*. 2023), in which each host generates events based on kernel calls to system resources, could work within the subset of onboard systems having traditional central processing units and operating systems. While Oshiro (2023) does not identify a specific SDN technology, SDN serves as a foundational component within the mZTAf providing the granular controls necessary to enforce micro-segmentation and to uphold the principle of least privilege. By abstracting the control plane from the physical hardware, the SDN allows the Local Policy Administrator (LPA) to dynamically isolate the network traffic and to restrict access to only those resource segments essential for a specific mission or user role.

Security requirements: Alignment with the afloat context

While Oshiro (2023) focused on the broader tactical edge, this work analyses these requirements through the lens of a single vessel. Instead of using the denied, degraded, intermittent, and/or latent (DDIL) term, the communications-degraded or denied environment (CD2E) term more concisely captures typical maritime environmental conditions, thus more accurately describing the persistent connectivity challenges unique to afloat platforms.

1. **Mission Continuity in a CD2E**. The mZTAf must guarantee that mission-critical services, including propulsion, steering, and navigation, remain operational without shore-side connectivity. This necessitates local policy enforcement so that authorisations persist, even when external satellite communications (SATCOM) links are severed or internal damage physically isolates shipboard components.
2. **Granular Access Control with Continuous Multi-Factor Authentication**. The mZTAf must use contextual data, such as shipboard location and biometrics, to verify users and non-person entities continuously. These mechanisms must function independently of shore-side resources and be resilient enough to survive internal hull damage that may bisect the ship's network.
3. **Robust IDS and Metadata Tagging**. Rather than simple boundary defence, a robust IDS

must monitor for internal anomalies like latent logic bombs or unusual data combinations. It must identify disinformation, isolate malicious activity locally, and synchronise logs with the enterprise only after CD2E conditions end to refine global intelligence for the mZTAf.

Operational requirements: Viability in the maritime mission

The operational requirements address what and who the system must support to remain viable and effective in a maritime mission context.

1. **Support for Legacy Systems**. The mZTAf must secure legacy OT, such as PLCs that often lack modern hardware security like trusted platform modules and trusted execution environment. Wrapping critical, high-replacement-cost systems in a monitored security layer can protect the ship's engineering plant from cybersecurity intrusion.
2. **Ad Hoc Joint and Coalition Networking**. Ships require the ability to form temporary, secure interconnections with third parties, including equipment vendors and foreign partners. The mZTAf must facilitate secure collaboration for coordination through identity federation, ensuring mission partners can share data without compromising the ship's network integrity.

A Maritime Zero-Trust Architecture Framework

Having addressed the maritime-relevant aspects of the tactical ZTA, this section introduces a theoretical foundation of an mZTAf and offers an analysis of the mZTAf within this theory.

Theoretical underpinnings: The Purdue Model in a maritime context

The Purdue Enterprise Reference Architecture, commonly referred to as the 'Purdue Model', is a layered framework that structures and segments ICS and OT (Williams 1993). Specifically, the model organises industrial environments into hierarchical levels to improve reliability and safety. In a maritime context, these levels define the flow of data from physical machinery to shoreside enterprise systems:

- **Level 0 (Field Layer):** Physical process components, including sensors (for example, temperature gauges) and actuators (for example, valves and motors).
- **Level 1 (Control Layer):** The logic controllers, including PLCs and RTUs, that manage engine timing, fuel pumps, ballast, and so on.
- **Level 2 (Supervisory Layer):** Human-machine interfaces and engineering workstations where the crew monitors and controls machinery.
- **Level 3 (Operations Layer):** Integrated shipboard systems including navigation, radar, and terminal operating systems.
- **Level 4/5 (Enterprise Layer):** The ship's administrative IT network and the SATCOM links connecting the vessel to shoreside headquarters.

While the foundational requirements that Oshiro (2023) references provide a robust baseline for tactical environments, the maritime domain introduces unique environmental and connectivity constraints that necessitate a more resilient, platform-centric mZTAf. Unlike most terrestrial networks, maritime platforms rely on fragile RF and SATCOM paths vulnerable to weather, RF interference, and 'mast shadowing', which threaten the link between Purdue Level 4/5 and shoreside authorities. Consequently, the mZTAf must function as a fully autonomous local trust anchor,

capable of managing credential ageing and policy drift within Purdue Levels 0-3 during periods of isolation without needing to phone home. As such, the mZTAf systems aboard a ship must be capable of performing zero-trust authentication and policy changes independently of the enterprise network ashore.

The maritime zero-trust architecture framework (mZTAf) overview

The proposed mZTAf derives its foundation by federating the tactical approach, building a hierarchical system of enclaves. As illustrated in **Figure 2**, different networks or systems on a typical vessel comprise local, tactical ZTA enclaves. Each area represents a distinct system category, each serving as its own tactical ZTA enclave in that zone of the ship. These can make autonomous policy decisions and changes, even when disconnected from other systems on the ship. For example, the integrated navigation systems used on the ship's bridge (blue nodes) constitute one enclave while the engineering systems (red nodes) constitute another enclave that is likely to span multiple zones within the ship. These may be further subdivided into smaller enclaves for systems, such as electrical, steering, ballast, cargo, and firefighting. Yellow nodes represent the ship's business network that handles administrative and logistics aspects and the crew welfare network that can include crew members' personally-owned devices. On a cruise ship, the yellow nodes would incorporate a much larger network and, perhaps, additional nodes to accommodate the greater number of crew and passengers. A navy ship's mZTAf would incorporate an additional category of systems for its combat and weapons systems. The principles described here remain the same as one adapts the mZTAf to different types of ships.

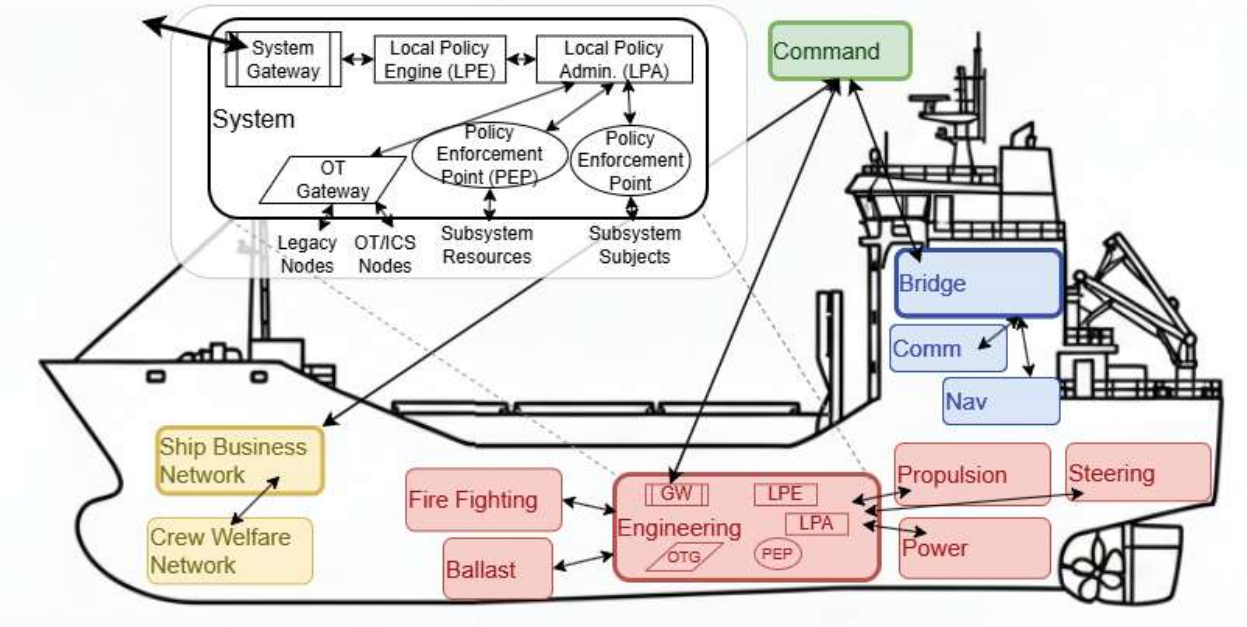

Figure 2: An overview of the Maritime Zero Trust Architecture Framework (mZTAf). The hierarchical and federated design of the system is shown. The single command node under the shipmaster's control is at the top of the tree structure. Major subsystems would each be a tactical ZTA-inspired enclave, which could be further subdivided into smaller subsystems. Each of the rounded-rectangle nodes consists of the components shown in the upper-left inset diagram.

The hierarchical nature of the proposed mZTAf allows for local, tactical ZTA nodes to operate in isolation when disconnected from other ZTA enclaves. Because each system has its own tactical

ZTA enclave, each enclave possesses its own Local Policy Engine (LPE), LPA, and Policy Enforcement Points (PEPs). The PEPs within the mZTAf enclave serving that zone of the ship connect with that enclave's LPA, which adjudicates the interactions between subjects and resources for the PEPs to enforce. Each enclave's LPA interacts with the enclave's LPE, which continually evaluates the defined policy versus the confidence level of each subject's authentication (Rose *et al*. 2020, pp. 9-10). As required in most shipboard instances, these enclaves utilise an OT Gateway (a type of PEP) to communicate with ICS and OT, for which Køien's (2021) Legacy Component Architecture could serve as a model. The term 'OT Gateway' is more representative of the systems aboard a ship than 'Legacy Gateway', which implies outdated technologies. Each system also maintains its own System Gateway to communicate with the System Gateways of other onboard systems and the company's enterprise ZTA at the headquarters facility ashore. Defining the communications involved with the OT and System Gateways could lead to new industry standards for maritime IoT/OT interoperability. All systems share the same SATCOM connection to connect with resources ashore, including the enterprise ZTA. In a ZTA, non-standard actions raise warnings that the nearest PEP must address with its LPA and LPE. The continued operation of the ship's systems is enabled by the mZTAf's distribution of these functions into ZTA enclaves throughout the ship. The mZTAf thus improves on a typical enterprise ZTA in which only one LPA exists, and this one LPA might not be reachable by the nearest PEP.

The framework specifically accounts for the maritime organisational structure where ultimate authority lies with the ship's master. The master is responsible for the confidentiality, integrity, and availability of all shipboard networks while adhering to IMO regulations and following guidelines, like the IMO's "Guidelines on Maritime Cyber Risk Management" that focuses on "safety and security management practices in the cyber domain" (IMO 2025, §2.3.1). To facilitate this, the mZTAf includes a 'command' enclave (the green node in **Figure 2**), which has the authority to affect the other LPEs' policy decisions and rules. To provide flexible and granular control of all connections within and between shipboard nodes, the mZTAf uses a tailored SDN infrastructure. Orchestration of this SDN is managed by each node's LPA. Within tactical ZTA enclaves comprising multiple sub-nodes, the parent LPA governs the connections between each sub-node gateway, ensuring that the overarching network control of the SDN remains centralised within the Command LPA.

In addition to the SDN, the mZTAf utilises Controlled Interfaces and Restricted Gateways to enforce boundaries between enclaves of varying security levels. These interfaces act as hardened Policy Enforcement Points (PEPs) that inspect and filter traffic at the application layer, rather than solely at the network layer. In a maritime environment, these restricted interfaces are vital for protecting Purdue Level 1 OT systems from lower-security administrative or crew networks. By implementing these strategic 'choke points', the LPA can physically and logically restrict communication paths. This ensures that, even if a device on a crew welfare network is compromised, the threat is prevented from moving laterally through an uncontrolled interface into the ship's critical navigation or propulsion control systems. This integrated architecture facilitates the dynamic delegation of authority, allowing the systems to grant access to personnel assuming non-traditional roles, which is a critical capability for ensuring mission continuity and completion during the unpredictable and high-stress events common to maritime operations.

The following procedure can be used to implement the mZTAf aboard a given ship:

1. Consider the requirements levied by organisations having jurisdiction over the ship, for example, the IMO, U.S. Coast Guard, and classification society.
2. Inventory the ship's networked assets and enumerate the interactions between them while grouping them into their respective zones and system categories.
3. Enumerate typical user requirements and interactions with these networked assets and systems.
4. Determine the multiple mechanisms through which the mZTAf will authenticate users periodically and continuously.
5. Microsegment each zone by restructuring the communications network between devices such that the PEPs and OT Gateways can enforce the access policies.
6. Test and evaluate the mZTAf's operation first by establishing a baseline before enabling its policy enforcement actions.
7. Configure inputs and conditions that allow the LPEs to update policy passed to the LPAs.

Continuous authentication

A practical ZTA requires methods to continually authenticate both humans and devices on the network, not just at the point of initial access, but at regular, short intervals to mitigate the risk of an adversary pivoting through the network. In a maritime environment, such rigid authentication poses a significant safety risk. To maintain accurate authentication without compromising the ship's safety (manoeuvrability, for example), the mZTAf must avoid over-reliance on any subset of factors. Tactical settings are highly susceptible to physical disruption; a crew member may suffer injuries rendering specific biometrics unusable, or environmental factors, such as cold-weather gear or protective goggles, may degrade sensor reliability. These environmental deviations increase the likelihood of false-negative results, which can degrade authentication reliability or trigger inadvertent access denials during critical operational windows. In a CD2E, where querying the enterprise ZTA ashore for credential resets is impossible, this local flexibility is vital for maintaining a self-sustaining trust posture.

To manage these complexities, the mZTAf integrates a hybrid access control model. While Role-Based Access Control (RBAC) provides the baseline for daily duties, the framework layers Attribute-Based Access Control (ABAC) to factor in the 'state' of the ship, such as geographic location or mission phase. This multi-modal approach ensures that the loss of a single factor does not deny access to critical Purdue Level 1 controls. During emergencies, the system can pivot between Mandatory Access Control (MAC), for strictly protecting sensitive combat data, and Discretionary Access Control (DAC), which allows the ship's master or a designated officer to manually override access denials. This arrangement enables LPAs to dynamically grant 'Emergency Responder' attributes to a junior technician, temporarily overriding standard role restrictions to allow for critical repairs.

Just as human operators may suffer physical casualties, shipboard devices and sensors at Purdue Levels 0 and 1 are susceptible to electronic corruption or adversarial spoofing. A rogue or compromised sensor providing false data, such as incorrect engine temperature or fuel levels, can lead to catastrophic kinetic failure. To maintain continuous device-level authentication without human intervention, the mZTAf can employ physical watermarking techniques (Satchidanandan & Kumar 2016) or by continuously verifying Physical Unclonable Functions (PUFs) (Aman, Basheer &

Sikdar 2019). The absence or distortion of this signature in the returning sensor data indicates an unauthorised or tampered device.

Rather than having a binary pass/fail authentication system, by utilising a well-designed set of heuristics to aggregate multiple authentication factors, an LPE (contained within the Policy Decision Point of **Figure 3**) calculates a confidence threshold that prioritises 'verified-safe operations' while still barring or challenging unauthorised or anomalous behaviour. Thus, the system evaluates the total context, combining available biometrics, device health, and mission state, to maintain operational momentum. Such a framework can also reduce 'authentication fatigue', ensuring that security measures do not become detractors from mission-critical work (Sasse 2014).

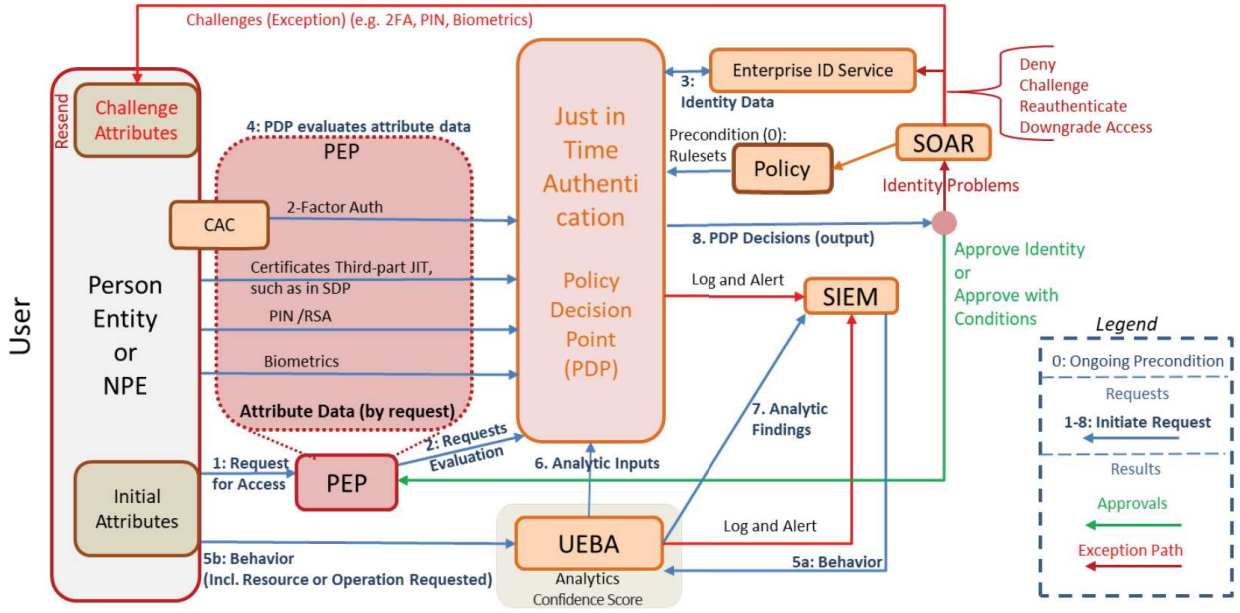

Figure 3: Dynamic, Continuous Authentication figure from the Department of Defense Zero Trust Reference Architecture document (Defense Information Systems Agency and National Security Agency Zero Trust Engineering Team 2022, p. 56)

Analysis of Proposed Solution

For the purposes of the case studies explored below, consider a well-known real-world ship MV Isla Bella (TOTE Maritime, IMO 9680841), a US-flagged 3,100 twenty-foot equivalent (TEU) container ship which has been operating between Florida and Puerto Rico since 2015 (Professional Mariner 2015). It has the following American Bureau of Shipping (ABS) classification: +A1 (E), Container Carrier, SH, SHCM, +AMS, GFS (DFD), +ACCU#1, CPS, UWILD, GP, TCM. The section below will illustrate the kinds of data to be gathered and decisions made to introduce the proposed mZTAf to this vessel.

The regulations that levy cybersecurity requirements include 33 CFR Part 104, which mandates a Vessel Security Assessment and a Vessel Security Plan that explain how vital systems are protected from tampering. The U.S. Coast Guard's "Cybersecurity in the MTS" Final Rule (effective July 16, 2025) requires cybersecurity awareness training and a cybersecurity plan. IMO Resolution MSC.428(98) requires that cybersecurity risk management must be integrated into the ship's Safe-

ty Management System. To keep its +ACCU (Automated Control) designation, the vessel must also adhere to the ABS Guide for Cybersecurity Events.

Inventorying and grouping the ship's networked assets would result in something similar to the following lists. Bridge: Yokogawa gyrocompass and steering control system, Furuno ECDS FMD-3300, Furuno radars, X-band and S-band, AIS. IT/Comms: Furuno FELCOM 500 Inmarsat FleetBroadband, Watchstander and Stateroom Computers. Engineering: MAN 8L70ME-C8.2GI main engine, MAN B&W ME-GI control system, MAN dual-fuel generators, Wärtsilä NACOS Platinum Integrated Automation System (IAS), Daewoo Shipbuilding & Marine Engineering (DSME) HiVAR fuel gas supply system, Cryos liquid natural gas (LNG) tanks, Techcross ballast water treatment system. Command: Master's computers in the Bridge and Cabin.

While considering typical user requirements and interactions with networked assets and systems to define RBAC rules, it would be helpful to start with a breakdown of a typical crew for a ship of this type. The Deck Division consists of the Master, three Mates, the Bosun, and 3-4 able-bodied seamen. The Engineering Division consists of the Chief Engineer, three Assistant Engineers, an Electro-Technical Officer, and 3-4 qualified members of the engineering division. The Stewards Division consists of a Chief Steward and two cooks. Occasionally, additional personnel may embark, including cadets, contractors, technicians, and auditors.

Some typical interactions the users would have could include the Master and Chief Engineer, who have superuser privileges on their respective systems. There are three watchstanding teams that typically rotate 4 hours on watch and 8 hours off watch. Regular ship maintenance typically happens during the day shift, 0800-1600, with engine systems testing between 0000-0400 and navigation system updates between 0800-1200.

One can propose that the crew members use passwords and PINs for primary authentication in the mZTAf, augmented with an authenticator app and company-issued smartcard for multi-factor authentication. For some sensitive systems, a fingerprint reader is used for biometric identification. The devices on the ship's network are authenticated through a variety of means, including network interface addresses and hardware Root of Trust mechanisms, such as Trusted Platform Modules (TPMs) or PUFs.

When considering how to microsegment the devices and assign them to specific PEPs and OT Gateways, one could devise a hierarchical set of enclaves that are grouped by function and location, similar to that shown in **Figure 2**. For example, an Engineering enclave would include the engineering control watchstander computers connected to PEPs, the MAN B&W ME-GI systems connected to an OT Gateway, and the Wärtsilä NACOS interfaces connected to either, depending on whether the device is IT or OT. Sub-enclaves that would be organised below include Power, Propulsion, and Steering, for example. The Bridge enclave would encompass the Integrated Bridge Systems, AIS Interfaces, Radar, and sub-enclaves for Communications (radios, weather, AIS, Global Positioning System) and Navigation (ECDIS, autosteering, Radar/Automatic Radar Plotting Aid).

A starting baseline policy could be generated by observing the normal operation of users and equipment for an extended period. This policy would allow for nominal interactions but would

generally block unobserved behaviour. To implement the dynamic mZTAf, the designers would consider different events that may trigger changes to the baseline policy. These would include phases of voyage—such as being docked, manoeuvring, underway at sea, conducting drills, and casualties. Other inputs to the LPEs could also include the Security Orchestration, Automation, and Response (SOAR) blended trust score (illustrated in **Figure 3**); top-down policy updates from parent LPE units; exception requests from lower-hierarchy units; and Manual/Safety overrides.

Use Case #1: Single ship in casualty (personnel or equipment) scenario

Based on the ZTA and maritime considerations described above, a major concern revealed for implementing ZTAs in maritime environments relates to ships at sea in casualty scenarios. In this context, 'casualty' refers to any form of unexpected equipment degradation or personnel injury that affects the safe, continued operation of the ship. The causes of such a casualty condition could include internal sources (for example, equipment or system failure, operator error, slips, trips, or falls) or external sources (for example, heavy seas or vessel collision or grounding). As an equipment example, a seawater pump could fail because of a clog, triggering the connected Wärtsilä NACOS to report the issue through the OT Gateway. The crew may then need to reconfigure the seawater cooling and fire suppression systems to a non-standard alignment, changing which signals must pass through PEPs, including OT Gateways. Thus, ABAC policies would trigger updates from the LPE to the LPA and, in turn, to the PEPs. As a personnel example, high seas could cause a mariner to fall down a ship's ladder, prompting another crew member to stand in for the injured member's duties. The LPE must recognise the watch-bill change such that the LPA can authorise the substitute crew member's actions. In each case, non-standard actions must occur to continue safely operating the ship (National Transportation Safety Board 2017).

As an extreme example, a ship can incur battle damage, exacerbating the casualty conditions in which the mZTAf must function. In battle, ships must be able to suffer significant damage and keep afloat. Perhaps one part of the ship is completely disconnected (from a computer network perspective) from another part that has the ZTA components necessary for monitoring and authorising actions. Perhaps the battle damage incapacitated a great number of crew members such that some crew members must fulfil roles they typically do not fill. ZTAs must account for these casualty conditions and the subsequent contingency actions to be viable in a maritime environment.

The mZTAf enclaves can each function in the above casualty scenarios, providing increased authentication scrutiny, as may be necessary, by the means available in that enclave, even when disconnected from the mZTAf command node. Otherwise, a casualty scenario could lead to dynamically inhibiting the necessary action, such as using a non-standard seawater system configuration or a crew member doing a job different from the member's normal duties.

Use Case #2: Remote Repair of Vendor-Locked Propulsion Controls (MAN B&W ME-GI)

To illustrate the mZTAf in action, consider a second scenario involving the *MV Isla Bella* while transiting between Florida and Puerto Rico. The vessel's MAN B&W ME-GI control system, critical for the operation of the dual-fuel main engine, requires an emergency software patch from a shore-side technician to resolve a fuel-injection anomaly. Per IMO Resolution MSC.428(98) and the ship's ABS +ACCU designation, this remote access must be strictly managed to prevent interference with the Safety Management System. In this scenario, the mZTAf facilitates this remote repair by tunnelling all external vendor traffic through a transient jump host acting as a highly scrutinised quarantine node. This architecture resolves the conflict between mission continuity

and granular access control by replacing permanent manufacturer 'backdoors' with just-in-time ephemeral identities managed by the LPA. Following the ad hoc trust principles of Bargh, Omar, & Choenni (2024), the vendor is isolated behind a Controlled Interface that performs deep application-layer inspection.

The framework employs ABAC, tagging the vendor's traffic with metadata such as Identity: MAN_Tech_Remote and Access_Target: Propulsion_Sub_Enclave. If the vendor's software, whether through malice or a configuration error, attempts to pivot from the MAN B&W controller to the Wärtsilä NACOS Platinum IAS or the Cryos LNG fuel gas supply system, the 'Restricted Gateway' for those enclaves identifies the unauthorised tag and drops the packets, preventing lateral movement into the broader OT network. The Hybrid Access Control logic further ensures that, even with valid repair credentials, the vendor is subject to MAC policies that automatically block write-access to engine parameters if the vessel is currently in a "Manoeuvring" phase (as determined by the LPE's voyage-phase input). If the Chief Engineer determines the repair is critical for safety during this phase, the Ship's Master can exercise DAC via the Command Node to issue a temporary, high-trust override token. This ensures that the *MV Isla Bella* maintains its security posture without a persistent link to shore, successfully neutralising third-party vulnerabilities while preserving the ship's kinetic safety.

Use Case #3: Secure Integration with Untrusted Port Infrastructure

In this scenario, the *MV Isla Bella* docks at a foreign, unfamiliar port. To reduce SATCOM costs and leverage higher bandwidth for crew welfare and logistical updates, the vessel connects to the port facility's terrestrial network, a system with which the ship's organisation has no prior trust relationship. The mZTAf manages this high-risk connection by strictly routing all port-provided backhaul through a dedicated 'Restricted Gateway' in the 'Communications Sub-Enclave'.

To ensure the integrity of the ship's Purdue Level 0 and 1 system, the LPA enforces ABAC policies that factor in the vessel's current connection state. Because the 'Connection_Source' attribute is set to 'Untrusted_External', the framework automatically assigns a low-trust metadata tag to all incoming packets, such as Source: Untrusted_Port_Network and Access_Level: Welfare_Only. While a VPN tunnel within this path allows the Stewards Division and off-watch crew to conduct high-bandwidth video calls, the mZTAf utilises MAC to ensure these packets are physically and logically incapable of reaching other components, including the Wärtsilä IAS or the Yokogawa steering controls.

The framework employs Controlled Interfaces acting as digital 'data diodes' for this specific connection. If an adversary, having compromised the port's cargo office or the ship's welfare network, attempts to pivot toward the engineering enclave or bridge enclave, the LPE identifies the unauthorised Welfare_Only attribute and immediately drops the traffic. Automated alarms are triggered at the Command Node, notifying the Master and the designated Engineering Officer of the lateral movement attempt. This ensures, that even when utilising untrusted shore-side backhaul at a foreign pier, the mZTAf maintains a hardened security posture, protecting critical kinetic systems from external vulnerabilities.

Use Case #4: Ship hijack scenario

Piracy is a longstanding characteristic of the maritime environment. When attempting to hijack a ship having an implemented mZTAf, its policies could prevent pirates' ability to control ship

systems. This scenario points to multiple ramifications and design considerations. Recognising piracy as illegal activity and having a desire to limit such activity, an mZTAf preventing pirates from controlling the ship can help keep the ship on its intended track; however, such a condition could frustrate the pirates and could increase the potential for them to harm the ship's crew unless it'unlocks' the system. Cases like this point to the design recommendation for ships to have a 'citadel' into which the crew can retreat during a hijacking event and, ideally, could control the ship from there (Carral *et al.* 2015). If a ship has a citadel, that would be the best place for the mZTAf command node from which the Master could implement a new DAC to prohibit all others access so that the crew could continue to control the ship while in the citadel. Hijack scenarios further justify the necessity to enable local, tactical ZTAs within the mZTAf to continue functioning such that no need exists for a physical manual override switch to bypass the ZTA because the hijackers could also use that switch to override the crew's control of the ship from the citadel.

As a counterpoint, consider that some ship 'hijack' scenarios are sanctioned, either through privateering or law-enforcement operations (Soons 2004). As a coast guard boards and seizes a ship, the mZTAf could hinder law-enforcement activities if the boarding team cannot take control of the ship involved in illegal activity. Future research could explore the benefits and risks of the malicious use of cybersecurity measures to protect illegal activity in the maritime environment. Such research could examine the risks involved with international maritime authorities having an ability to 'unlock' vessels, perhaps by comparing the concept with firefighters' access to a building's control systems versus having a 'Clipper Chip' capability (Denning & Barlow 1996) for commercial ships' systems. Also, could a ship's organisational headquarters control the ship remotely when the ship's crew is under duress? Future research on the legalities involved with operating such a command node is also warranted, especially regarding the liability incurred by the physical effects of operating the ship from a command node in the ship's citadel or in the organisation's headquarters ashore.

Conclusion

This work investigated the very important ideas surrounding increasing the security of maritime vessels via the popular zero-trust principles. Because the shipboard environment is much different than that of traditional IT networks, the same approaches that are successful on static land-based systems may run into issues when deployed at sea, where network connections are often degraded and intermittent and there is a higher likelihood of casualties that require non-standard responses. Inspired by a work that explored implementing a tactical ZTA for a deployed U.S. Marine Corps unit, this work analysed the assumptions, goals, and implementation technology for the maritime environment. This work proposed a maritime zero-trust architecture framework (mZTAf) that addresses many of the challenges unique to the maritime environment. Finally, this architecture framework was conceptually validated by analysing four use cases where strengths and weaknesses of the proposed approach were discussed.

While many challenges to implementing this approach exist beyond this introductory work, there is promise that such an effort will be successful. Further research is needed to more fully develop the rules of the local policy engines and how logging and exceptions are communicated back to central nodes when connections exist or are re-established. The software-defined network infrastructure that allows for granular control of connections also needs further development. Finally, the entire framework would be improved by a complete implementation on a simulated or digital

twin of a typical ship where higher-fidelity validation can occur and where researchers can use Artificial Intelligence and machine learning to detect anomalous and malicious activity over time.

Acknowledgments
The authors thank Gurminder Singh and Alan Shaffer of the U.S. Naval Postgraduate School for their input.

Disclaimer
The views expressed in this article are those of the author(s) and do not reflect the official policy or position of the U.S. Naval Academy, the Department of the Navy, the Department of War, or the U.S. Government.

References
Aiello, S & Rimal, BP 2023, 'Secure access service edge convergence: Recent progress and open issues', *IEEE Security & Privacy*, vol. 22, no. 2, pp. 8-16, <https://doi.org/10.1109/MSEC.2023.3326811>.

Alalmaie, AZ, Nanda, P & He, X 2022, 'Zero trust-NIDS: Extended multi-view approach for network trace anonymization and auto-encoder CNN for network intrusion detection', *IEEE International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom)*, pp. 449-56, <https://doi.org/10.1109/TrustCom56396.2022.00069>.

Al-Khalidi, M, Al-Zaidi, R, Ali, T, Khan, S & Bashir, AK 2025, 'AI-optimized elliptic curve with certificate-less digital signature for zero trust maritime security', *Ad Hoc Networks*, vol. 166, p. 103669, <https://doi.org/10.1016/j.adhoc.2024.103669>.

Aman, MN, Basheer, MH & Sikdar, B 2019, 'Two-factor authentication for IoT with location information', *IEEE Internet of Things Journal*, vol. 6, no. 2, pp. 3335-51, <https://doi.org/10.1109/JIOT.2018.2882610>.

Bargh, MS, Omar, A & Choenni, S 2024 'Zero-trust security model applied to smart shipping: Towards a feasible architecture', *Advances in Knowledge-Based Systems, Data Science, and Cybersecurity*, vol. 1, pp. 78-107, <https://doi.org/10.54364/cybersecurityjournal.2024.1105>.

Carral, L, Fernández-Garrido, C, de Troya, JJ & Fraguela, JÁ, 2015, 'Considering anti-piracy ship security: Citadel design and use', *Shipbuilding*, vol. 66, no. 3, pp. 75-90, <https://hrcak.srce.hr/146159>.

Defense Information Systems Agency (DISA) and National Security Agency (NSA) Zero Trust Engineering Team 2022, *Department of Defense (DoD) Zero Trust Reference Architecture*, Version 2.0, July, viewed 10 February 2026, <https://dodcio.defense.gov/Portals/0/Documents/Library/(U)ZT_RA_v2.0(U)_Sep22.pdf>.

Denning, DE & Barlow, JP 1996, 'The Denning-Barlow Clipper Chip debate', *High Noon on the Electronic Frontier: Conceptual Issues in Cyberspace*, ed. P Ludlow, MIT Press, Cambridge, MA, US, pp. 217-24.

Dodd, AJ 2023, 'Continuous authentication in zero trust architecture for afloat platforms', Thesis, Naval Postgraduate School, Monterey, CA, US, restricted distribution.

Hatakeyama, K, Kotani, D & Okabe, Y 2021, 'Zero trust federation: Sharing context under user control towards zero trust in identity federation', *2021 IEEE International Conference on Pervasive Computing and Communications Workshops and other Affiliated Events (PerCom Workshops)*, pp. 514-9, <https://doi.org/10.1109/PerComWorkshops51409.2021.9431116>.

Hong, S, Xu, L, Huang, J, Li, H, Hu, H & Gu, G 2023, 'SysFlow: Toward a programmable zero trust framework for system security', *IEEE Transactions on Information Forensics and Security*, vol. 18, pp. 2794-2809, <https://doi.org/10.1109/TIFS.2023.3264152>.

Hu, V 2022, *Blockchain for access control systems, interagency or internal report 8403*, National Institute of Standards and Technology, Gaithersburg, MD, US, <https://doi.org/10.6028/NIST.IR.8403>.

International Association of Ports and Harbors (IAPH) Data Collaboration Committee 2025, *Cyber Resilience Guidelines for Emerging Technologies in the Maritime Supply Chain*, viewed 30 March 2026, <https://wwwcdn.imo.org/localresources/en/OurWork/Security/Documents/IAPH%20Cyber%20resilience%20guidelines%20for%20emerging%20technologies%20in%20the%20maritime%20supply%20chain%20ENG.pdf>.

International Maritime Organization (IMO) 2025, *Guidelines on maritime cyber risk management, MSC-FAL.1/Circ.3/Rev.3*, 4 April 2025, viewed 26 February 2026, <https://wwwcdn.imo.org/localresources/en/OurWork/Security/Documents/MSC-FAL.1-Circ.3-Rev.3.pdf>.

Kang, H, Liu, G, Wang, Q, Meng, L & Liu, J 2023, 'Theory and application of zero trust security: A brief survey', *Entropy*, vol. 25, no. 12, p. 1595, <https://doi.org/10.3390/e25121595>.

Kindervag, J, Balaouras, S & Coit, L 2010, 'No more chewy centers: Introducing the zero trust model of information security', *Forrester Research*, vol. 3, pp. 1-10, viewed 21 January 2026, <https://media.paloaltonetworks.com/documents/Forrester-No-More-Chewy-Centers.pdf>.

Kohnfelder, L & Garg, P 1999, 'The threats to our products', *Microsoft Interface,* Microsoft Corporation*,* viewed 30 March 2026, <https://shostack.org/files/microsoft/The-Threats-To-Our-Products.docx>.

Køien, GM 2021, 'Zero-trust principles for legacy components: 12 rules for legacy devices: An antidote to chaos', *Wireless Personal Communications*, vol. 121, pp. 1169-86, <https://doi.org/10.1007/s11277-021-09055-1>.

Moreau, L, Conchon, E & Sauveron, D 2021, 'CRAFT: A continuous remote attestation framework for IoT', *IEEE Access*, vol. 9, pp. 46430-47, <https://doi.org/10.1109/ACCESS.2021.3067697>.

National Transportation Safety Board (NTSB) 2017, '*Collision between US Navy Destroyer* John S McCain *and Tanker* Alnic MC*, Singapore Strait, 5 Miles Northeast of Horsburgh Lighthouse*',

Marine Accident Report NTSB/MAR-19/01, Washington, D.C., US, viewed 11 February 2026, <https://www.ntsb.gov/investigations/accidentreports/reports/mar1901.pdf>.

Osborn, B, McWilliams, J, Beyer, B & Saltonstall, M 2016, 'BeyondCorp: Design to deployment at Google,' *USENIX ;login:*, vol. 41, no. 1, pp. 28-35, viewed 30 March 2026, <https://www.usenix.org/publications/login/spring2016/osborn>.

Oshiro, DM 2023, 'Zero trust architecture implementation for the Marine Corps tactical cloud', Thesis, Naval Postgraduate School, Monterey, CA, US, <https://hdl.handle.net/10945/72380>.

Professional Mariner 2015, *Isla Bella specifications*, 9 November, viewed 30 March 2026, <https://professionalmariner.com/isla-bella-specifications/>.

Rose, S, Borchert, O, Mitchell, S & Connelly, S 2020, *Zero Trust Architecture, Special Publication 800-207*, National Institute of Standards and Technology, Gaithersburg, MD, US, <https://doi.org/10.6028/NIST.SP.800-207>.

Sasse, MA 2014, '"Technology should be smarter than this!": A vision for overcoming the great authentication fatigue', *Secure Data Management (SDM 2013)*, eds. W Jonker & M. Petković, LNCS 8425, Springer, Cham, Switzerland, pp. 33-6, <https://doi.org/10.1007/978-3-319-06811-4_7>.

Satchidanandan, B & Kumar, PR 2016, 'Dynamic watermarking: Active defense of networked cyber-physical systems', *Proceedings of the IEEE*, vol. 105, no. 2, pp. 219-40, <https://doi.org/10.1109/JPROC.2016.2575064>.

Soons, AHA 2004, 'Law enforcement in the ocean—An overview', *WMU Journal of Maritime Affairs*, vol. 3, pp. 3-16, <https://doi.org/10.1007/BF03195046>.

Stewart, AD 2023, 'Cybersecurity of advanced machinery systems: The foundation of cyber resilient ships', *Naval Engineers Journal*, vol. 135, no. 2, pp. 59-69, <https://www.ingentaconnect.com/content/asne/nej/2023/00000135/00000002/art00017>.

Sunkara, G 2025, 'Implementing zero trust architecture in modern enterprise networks', *SAMRIDDHI: A Journal of Physical Sciences, Engineering and Technology*, vol. 17, no. 3, pp. 1-11, <https://doi.org/10.18090/samriddhi.v17i03.01>.

Williams, TJ 1993, 'The Purdue enterprise reference architecture', *IFAC Proceedings*, vol. 26, no. 2, pp. 559-64, <https://doi.org/10.1016/S1474-6670(17)48532-6>.